

NBER WORKING PAPER SERIES

INCOMPLETE INFORMATION IN MACROECONOMICS:
ACCOMMODATING FRICTIONS IN COORDINATION

George-Marios Angeletos
Chen Lian

Working Paper 22297
<http://www.nber.org/papers/w22297>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2016

Prepared for the Handbook of Macroeconomics, vol 2A and 2B. For helpful comments and suggestions, we thank the editors, John Taylor and Harald Uhlig; our discussant, Alp Simsek; various participants at the Handbook Conference hosted by the Becker Friedman Institute in Chicago; and especially Harris Dellas. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 Elsevier, circulated with permission. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination
George-Marios Angeletos and Chen Lian
NBER Working Paper No. 22297
June 2016
JEL No. C7,D8,E01,E3,E40,G1

ABSTRACT

This chapter studies how incomplete information helps accommodate frictions in coordination, leading to novel insights on the joint determination of expectations and macroeconomic outcomes. We review and synthesize recent work on global games, beauty contests, and their applications. We elaborate on the distinct effects of strategic uncertainty relative to fundamental uncertainty. We demonstrate the potential fragility of workhorse macroeconomic models to relaxations of common knowledge; the possibility of operationalizing the notions of “coordination failure” and “animal spirits” in a manner that unifies unique- and multiple-equilibrium models; and the ability of incomplete information to offer a parsimonious explanation of important empirical regularities. We provide a general treatment of these ideas, as well as specific applications in the context of business cycles, financial crises, and asset pricing.

George-Marios Angeletos
Department of Economics, E52-530
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
angelet@mit.edu

Chen Lian
Department of Economics
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
lianchen@mit.edu

Contents

| | | |
|----------|---------------------------------------------------------------------------------|-----------|
| 1 | Introduction | 4 |
| 2 | Framework | 10 |
| 2.1 | Actions and Payoffs | 10 |
| 2.2 | Examples | 11 |
| 2.3 | Shocks and Information | 13 |
| 2.4 | Equilibrium | 16 |
| 3 | Imperfect Coordination | 17 |
| 3.1 | Some Key Notions | 17 |
| 3.2 | Informational Friction | 21 |
| 3.3 | Imperfect Coordination Equals Incomplete Information | 21 |
| 3.4 | Complete-Information Benchmark and Coordination Motives | 23 |
| 4 | Global Games: Theory | 27 |
| 4.1 | Set Up | 28 |
| 4.2 | Interpretation | 28 |
| 4.3 | Complete Information and Multiple Equilibria | 29 |
| 4.4 | Incomplete Information and Unique Equilibrium | 30 |
| 4.5 | Proof of Equilibrium Uniqueness | 31 |
| 4.6 | The Role of Public Information | 33 |
| 4.7 | Intuition and Take-home Lessons | 34 |
| 4.8 | Extensions and Additional Lessons | 36 |
| 5 | Global Games: Applications | 37 |
| 5.1 | Currency Crises and Bank Runs | 37 |
| 5.2 | Big Players and Lenders of Last Resort | 38 |
| 5.3 | Debt runs, Default Premia, and Credit Ratings | 39 |
| 5.4 | “Liquidity Black Holes” and “Market Freezes” | 40 |
| 5.5 | Safe Assets | 41 |
| 5.6 | Business Cycles and the Great Recession | 41 |
| 5.7 | Prices and Other Endogenous Public Signals | 42 |
| 5.8 | Policy Interventions | 45 |
| 5.9 | Dynamics | 46 |
| 5.10 | Other Related Research | 47 |
| 6 | Coordination and Synchronization | 47 |
| 6.1 | The Calvo friction, Asynchronous Choice, and Coordination | 48 |
| 6.2 | Asynchronous Awareness and Coordination (with Application to Bubbles) | 53 |
| 7 | Beauty Contests: Theory | 55 |
| 7.1 | Set Up | 55 |
| 7.2 | Complete-Information Benchmark | 56 |
| 7.3 | Equilibrium with Incomplete Information | 56 |

| | | |
|-----------|---------------------------------------------------------------------------------|------------|
| 7.4 | Simplifying Assumptions | 57 |
| 7.5 | The Structure of Higher-order Beliefs | 59 |
| 7.6 | Positive Implications | 59 |
| 7.7 | Complementarity, Information, and Volatility | 61 |
| 7.8 | Dynamics and Learning | 62 |
| 7.9 | Endogenous Information Acquisition | 67 |
| 7.10 | A Note on Solution Methods | 70 |
| 8 | Beauty Contests: Applications | 72 |
| 8.1 | Real Rigidity and Technology Shocks | 72 |
| 8.2 | General-Equilibrium Dampening, and Rigidity vs Overshooting | 78 |
| 8.3 | Nominal Rigidity and Monetary Shocks | 78 |
| 8.4 | Parenthesis: Lucas (1972) | 83 |
| 8.5 | Sticky Information and Rational Inattention | 86 |
| 8.6 | Survey Evidence on Informational Frictions | 89 |
| 8.7 | Demand-Driven Fluctuations | 92 |
| 8.8 | Financial Markets | 97 |
| 9 | Efficiency and Policy Implications | 105 |
| 9.1 | Constrained Efficiency | 105 |
| 9.2 | Efficient Coordination in Beauty Contests | 108 |
| 9.3 | Policy Implications | 109 |
| 9.4 | Welfare Effects of Public Information (and Central-Bank Transparency) | 110 |
| 10 | Conclusion | 112 |
| | Appendix: Proofs | 114 |
| | References | 131 |

1 Introduction

Modern economies are vastly complex networks, in which the decisions of any given agent are highly dependent on expectations of the decisions of many other agents. For instance, how much a firm wants to hire and invest depends on the firm's expectations of the future demand for its product, which in turn depends on the spending plans of the consumers who are likely to buy the product of that firm; but consumer plans themselves depend on expectations of income and labor-market conditions, which in turn depend on the decisions of other firms and other consumers, and so on. Markets, the government, and other social institutions help facilitate the coordination of these expectations, and of the associated decisions, but only up to a point.

The limits to such coordination are rarely acknowledged in macroeconomics. Workhorse models are constructed as if all economic agents could confer with one another, reach a unanimous consensus on the current state and future prospects of the economy, and effortlessly coordinate on a commonly known profile of actions, or allocation, at all times.¹

How would the predictions of these models differ if there were non-trivial frictions in the agents' ability to reach such a consensus and to coordinate their actions? What are useful ways for formalizing and operationalizing such frictions in the first place? And how does the accommodation of such frictions affect our interpretation of the observed phenomena and the guidance we may offer to policy makers?

This chapter reviews and synthesizes recent advances made in addressing these questions. We argue that incomplete information offers a useful method for introducing frictions in coordination and for enriching the dynamics of expectations in macroeconomic models. This enrichment leads to the questioning of existing interpretations of, and helps shed new light on, important phenomena such as business cycles and crises.

Background and Key Ideas

Before reviewing the contents of this chapter, we need to clarify the benchmark from which we depart, and the nature of the departure.

It is useful to start with the textbook RBC model. This model, like other workhorse models in macroeconomics and finance, is based on the Arrow-Debreu paradigm. In this paradigm, agents attain perfect communication: the general equilibrium is akin to the Nash equilibrium of a game of complete information, in which the players attain common knowledge, not only about exogenous fundamentals such as preferences and technology, but also about endogenous outcomes such as employment, output, and asset prices.²

While modern macroeconomic models often depart from the Arrow-Debreu paradigm by allowing for monopolistic power, sticky prices, credit constraints, search frictions, and so forth, they typically impose that all agents share the same information. This assumption, together with standard equilibrium concepts, tends to trivialize the problem of forecasting the actions of other agents and the associated macroeconomic outcomes: in equilibrium, all agents share the same "state of mind" about where the economy is heading to. It is in this sense that perfect coordination of beliefs and actions is imposed.

This property is evident in unique-equilibrium DSGE models that pin down equilibrium outcomes as functions of commonly known fundamentals such as preferences and technologies. But it also applies to certain multiple-equilibrium models that aimed to capture the idea that recessions, bank runs, and liquidity crises are the product of "coordination failures", such as Diamond (1982), Bryant (1983), Cooper and John (1988), Diamond and Dybvig (1983), Calvo (1988), and Obstfeld (1986, 1996).

¹This allocation need not be a socially desirable one; it only has to be consistent with equilibrium.

²Whenever we say "fundamentals", the game theorist can read "payoff types"; and whenever we say "endogenous outcomes" or "economic activity", the game theorist can read "actions".

Let us explain. The models in the aforementioned papers boil down to static coordination games that admit multiple equilibria. Some of these equilibria can be interpreted as coordination failures because they are welfare-inferior to other equilibria. Nevertheless, the following property holds in *any* of the equilibria of these models: the agents know, not only which strategy profile has been selected, but also the particular actions played by all other agents. This represents a certain oxymoron: although a coordination failure obtains *across* equilibria, coordination remains flawless *along* any given equilibrium.³

The oxymoron gets worse: once we allow the coordination that obtains along any given equilibrium to be imperfect, in the manner that we formalize in the sequel, the multiplicity of equilibria may vanish; and when this happens, the conventional formalizations of “coordination failures” and “animal spirits” fall apart. This underscores the potential fragility of the predictions of standard macroeconomic models—and of the associated interpretations of macroeconomic phenomena—to the kind of friction we are after.

Nonetheless, we also show that this friction offers a new, refined, version of the aforementioned concepts. Interestingly, this new version helps unify unique- and multiple-equilibria models. Most importantly, it leads to new insights into the nature of important phenomena, including financial crises, business cycles, and asset-price fluctuations; and to new lessons for policy.

But how can one operationalize the sought-after friction in coordination? We propose that incomplete information offers a useful method for doing so.⁴

We thus depart from standard macroeconomic models in one crucial respect: we let agents have differential information about the aggregate shocks hitting the economy. This transforms the macroeconomic models of interest into games of incomplete information, thereby also opening the door to rich higher-order beliefs.

Once we look at the equilibria of the modified models, we find that there is no more a unanimous consensus about where the economy is heading: agents fail to reach the same equilibrium beliefs about one another’s actions and the associated macroeconomic outcomes. In this sense, coordination has become imperfect.

Shedding light on how this imperfection affects the predictions of macroeconomic models and our interpretation of macroeconomic phenomena is the goal of this chapter.

Preview and Main Lessons

Our analysis alternates between two modes: one abstract, seeking to highlight general properties; and another applied, spelling out concrete lessons in the context of specific applications.

In Section 2, we set up our baseline framework. The framework is highly stylized: it boils down to a static game with a continuum of agents. Yet, the framework is sufficient for our purposes. First, it encapsulates the fixed-point relation between expectations of economic outcomes and actual behavior that is at the core of most macroeconomic models. Second, it facilitates a sharp analysis of how the theoretical properties of this fixed point, and the implied restrictions on the observables of the model, vary with the information structure. Last but not least, it helps clarify the language employed throughout this chapter.

In Section 3, we offer our preferred definitions of key concepts such as “fundamentals,” “incomplete information”, and “imperfect coordination”. We next distinguish the uncertainty that the agents face about the behavior of others and the associated macroeconomic outcomes, hereafter referred to as “strategic uncertainty”, from the uncertainty they face about fundamentals such as preferences and technologies. We finally

³These points are exact in static coordination models, such as those studied in the aforementioned works, but become fussy in dynamic sunspot models. See the remark at the end of Subsection 3.4; see also Aumann, Peck and Shell (1988) and Peck and Shell (1991) on the relation between sunspot equilibria in dynamic macroeconomic models and correlated equilibria in games.

⁴By “incomplete information” we refer to situations in which agents have dispersed private information about, and lack common knowledge of, aggregate shocks. Our preferred definitions of this and other key concepts are provided in Section 3, after the introduction of our framework.

elaborate on how incomplete information helps disentangle these two types of uncertainty and how it helps formalize the sought-after friction in coordination.

For the remainder of the chapter, we concentrate on two special cases of our framework, each of which corresponds to a different class of models. The first class, which is introduced in Section 4, is known in the literature as “global games”. The second class, which is introduced in Section 7, is often referred to as “beauty contests”. These classes exhibit strategic complementarity and use incomplete information to accommodate a friction in coordination; what differentiates them is the strength of the coordination motives and the potential impact of the friction on the determinacy of the equilibrium.

In Section 4, we thus review the following result from the literature on global games, which was initiated by Carlsson and Van Damme (1993a, 1993b) and was further advanced by Morris and Shin (1998, 2001, 2003). Take a classic coordination model, as in Diamond (1982), Diamond and Dybvig (1983), Calvo (1988), and Obstfeld (1986, 1996). Modify it by introducing idiosyncratic noise in the agents’ observation of the fundamentals. In the limit as this noise vanishes, the modified model is indistinguishable from the original model in the sense that every agent has nearly perfect knowledge of the underlying fundamentals. And yet, the modified model admits a unique equilibrium.

At the core of this result is the subtle difference between perturbations of mutual knowledge (“everybody knows A”) and perturbations of common knowledge (“everybody knows A,” “everybody knows that everybody knows A,” ... ad infinitum). The friction may be small in the sense that every agent faces little uncertainty about the underlying fundamentals; and yet it can have a sizable effect on equilibrium outcomes insofar as it shatters common knowledge.

Notwithstanding this subtlety, the result reveals a discontinuity in the predictions of a large body of applied work to perturbations of the informational assumptions. At first glance, this discontinuity is quite troubling: the interpretations of certain phenomena and the related policy implications that macroeconomists have developed on the basis of multiple-equilibria models appear to be in jeopardy.

We argue that this worry is largely not warranted.

First, equilibrium multiplicity is no more needed in order for economic outcomes to be fragile, for self-fulfilling beliefs to matter, and for coordination failures to occur. For instance, we show that small exogenous shocks can still trigger large changes in equilibrium outcomes because, and only because, they trigger large changes in the equilibrium expectations that agents hold about one another’s choices.

Second, the determinacy of the equilibrium under incomplete information is sensitive to the particular way that the information structure is modeled. For instance, if market signals such as prices create sufficient correlation in beliefs, or approximate common knowledge, multiplicity may survive the kind of perturbations that would have induced equilibrium uniqueness in the absence of such endogenous signals.

These observations help resuscitate the spirit of the earlier, complete-information models. But this is not the endgame. By offering a more structured way for modeling expectations, the global-games methodology qualifies existing applied lessons and leads to new ones.

In Section 5, we expand on these lessons within the context of several applications. Some of these applications regard bank runs, currency attacks, and sovereign debt crises; others are in the context of business cycles, market freezes, and bubbles. Inter alia, we highlight a number of novel lessons regarding the adverse effects of big speculators, the effectiveness of lenders of last resort such as the IMF, and the role of bank regulation in preventing coordination failures.

We also discuss a few extensions that endogenize the information structure in global games. This permits us to elaborate on the robustness and the empirical content of the uniqueness result. More importantly, it

leads to new applied lessons, such as the possibility that an informed policy maker finds herself trapped in a situation where her policy actions are shaped by market expectations rather than the other way around.

Summing up, although the global-game methodology does not offer a panacea for equilibrium uniqueness, it helps open “the black box” of expectations and coordination. This leads to useful applied lessons that were not possible in the context of the earlier, complete-information, literature on coordination failures.

In Section 6, we review two contributions that shift attention to a dynamic issue, the ability of the agents to *synchronize* their choices. The first contribution, which is by Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2001), uses a friction as in Calvo (1983) to introduce asynchronous choice in a dynamic coordination game and shows that this helps select a unique equilibrium, even when the Calvo-like friction is vanishingly small. The second contribution, which is by Abreu and Brunnermeier (2003), is also concerned with a dynamic coordination game; but instead of adding a Calvo friction, it introduces asynchronous awareness of a shock in the environment; it is then shown that such asynchronous awareness results, not only in asynchronous responses to the aforementioned shock, but also in significant delay of these responses. We discuss how these contributions connect to the global-games literature and how they shed light on the subtle relation between synchronization and coordination.

In Section 7, we shift gears and study a class of models that feature a weaker form of coordination motives than that in the global-games literature. These games, which are often described as beauty contests, have linear best-response conditions and admit a unique equilibrium, features shared by the type of (log)linear DSGE models commonly used to study business cycles and macroeconomic policy, as well as by certain asset-pricing models.

In this class of models, equilibrium determinacy is no longer an issue and the key preoccupation is with the stochastic properties of the equilibrium. Accordingly, the questions we pose are the following: How does an economy respond to innovations to fundamentals? How does this response depend on strategic uncertainty as opposed to fundamental uncertainty? Do coordination motives matter for level and the nature of aggregate volatility? Can one discern a role for “animal spirits” and “self-fulfilling beliefs”?

The core of Section 7 is devoted to answering the above questions in a relatively abstract manner. The analysis builds heavily on Morris and Shin (2002), Angeletos and Pavan (2007), and Bergemann and Morris (2013), although we adapt and extend the results of those papers in ways that suit our purposes. Inter alia, we establish the following three general points.

First, incomplete information can help generate significant rigidity, or inertia, in the response of macroeconomic outcomes to aggregate shocks to fundamentals. This inertia can be quantitatively significant even if the noise in the observation of the fundamentals is small, a finding that underscores the distinct role played by strategic uncertainty, or imperfect coordination, relative to fundamental uncertainty.

Second, inertia may obtain at the macro level even if it is largely absent at the micro level. This finding may help liberate macroeconomists from ad-hoc adjustment costs that DSGE models have to use in order to match the inertia in the estimated responses of macroeconomic variables to identified shocks to fundamentals. It also highlights a subtle difference between incomplete information and the formalizations of “inattention” found in Sims (2003), Reis (2006), and Gabaix (2014): whereas these formalizations are designed so as to introduce inertia at the microeconomic level, the friction we accommodate in this chapter is uniquely positioned to generate inertia at the macroeconomic level.

Third, this friction can also manifest as extrinsic volatility, that is, as volatility in equilibrium outcomes that is not spanned by the volatility in either the fundamentals or beliefs thereof. It follows that the traditionally dichotomy between models that admit multiple equilibria and those that do not may no longer

be central to the understanding of “animal spirits” and “self-fulfilling beliefs”: the same type of volatility can now obtain in either class of models.

In Section 8, we consider several applications. In our first application, which is based on Angeletos and La’O (2010), we show how incomplete information can induce rigidity in the response of employment and output to technology shocks. This helps reconcile the RBC paradigm with facts that have so far been considered prima-facia evidence against that model and in favor of the New-Keynesian model (Gali, 1999). We also discuss other works that use informational frictions to generate real rigidities in other contexts.

Our second application concerns nominal rigidity. In particular, we review Woodford (2003), who shows how inertia in higher-order beliefs can be a potent, if not superior, substitute to sticky prices. We also discuss the connection to the complementary works of Mankiw and Reis (2002) and Mackowiak and Wiederholt (2009), as well as to the seminal contribution by Lucas (1972).

We proceed to review some recent empirical work that uses surveys of economic forecasts to measure the role of frictions in information, most notably by Coibion and Gorodnichenko (2012). We argue that this work provides important evidence in support for the main ideas reviewed in this chapter; but we also highlight certain caveats in the mapping from this evidence to the theory.

In another application, we argue that incomplete information helps accommodate the notion of demand-driven fluctuations without the need for either nominal rigidity or constraints on monetary policy. We also discuss how the incorporation of plausible higher-order belief dynamics in DSGE models can help match salient features of the business-cycle data.

We then turn to a review of applications in finance, which have indicated how incomplete information can also offer a parsimonious explanation to asset-pricing puzzles such as momentum, excess volatility, and the disconnect between exchanges rates and macroeconomic fundamentals.

Section 9 concludes our analysis by touching upon certain normative questions. We first review a notion of constrained efficiency that helps dissect the normative properties of the equilibrium in the class of incomplete-information models we are concerned with. We then discuss two sets of applications: a literature that characterizes optimal monetary policy in the presence of informational frictions; and a literature that studies the welfare effects of the information disseminated by markets, policy makers, or the media.

Additional points and position in the literature

We now offer a few additional remarks about the theme of our chapter and its position in the literature.

1. The hallmark of the friction we accommodate in this chapter is the inability of agents to reach a consensus about the state of the economy. Along with it comes a certain disentanglement of the uncertainty agents face about endogenous economic outcomes such as employment, output, and asset prices, from their uncertainty about exogenous fundamentals such as preferences and technologies. This disentanglement is made possible with the help of certain types of private information, which permit higher-order beliefs of fundamentals to diverge from the corresponding first-order beliefs. This divergence, in turn, is instrumental to understanding the equilibrium properties of the models we study. However, once equilibrium is imposed, higher-order beliefs become a sideshow: all that matters for the behavior of an agent is the equilibrium expectations (first-order beliefs) she forms about the relevant economic outcomes. Furthermore, for the applications we have in mind, it seems easier to detect and quantify this kind of expectations, as opposed to the underlying belief hierarchies. These observations explain why the focal point of our analysis is the joint determination of equilibrium outcomes and expectations thereof: higher-order beliefs is the machinery, not the essence, of what we are after.

2. The uncertainty agents face about economic outcomes is tightly connected to the one they face about the behavior of other agents, which is often referred to as “strategic uncertainty” in game theory. Starting from a complete-information model, one notes that the equilibria of such a model rule out any uncertainty about the behavior of others conditional on payoffs.⁵ One may then seek to accommodate such uncertainty by relaxing the solution concept. Alternatively, one may maintain the solution concept and instead engineer the sought-after uncertainty with the help of incomplete information. The approach taken in this chapter is the latter. Whenever we talk about strategic uncertainty in the sequel, we therefore refer to the uncertainty that agents face about the actions of others *on* equilibrium, due to incompleteness of information.

3. The frictionless coordination of beliefs and actions that is embedded in standard macroeconomic models hinges on the rational-expectations equilibrium concept just as much as it hinges on the conventional assumption of complete information. It follows that the relaxation of the latter assumption—the approach favored in this chapter—can also be seen as a substitute for relaxing the solution concept. We elaborate on this point in due course.

4. The preceding remarks help distinguish our chapter from the recent literature on “news” and “noise shocks” that has followed Beaudry and Portier (2006).⁶ This literature extends workhorse models by letting agents observe noisy signals of future fundamentals, but typically maintains the assumption that all agents share the same information. In so doing, it fails to accommodate the type of friction we are after. In short, it enriches the stochasticity in expectations of fundamentals, but does not disentangle uncertainty of endogenous economic outcomes from uncertainty of exogenous fundamentals. The same point applies to Bloom (2009) and the subsequent literature on “uncertainty shocks”.

5. Similar observations help position our chapter also vis-a-vis a strand of the literature that studies various forms of inattention. Consider, in particular, Sims (2003), Reis (2006), and Gabaix (2014). These papers study single-agent problems in which it is costly for an agent to acquire information about an exogenous payoff-relevant variable (or otherwise make her action covary appropriately with that variable). In so doing, these papers provide useful decision-theoretic foundations of informational frictions. But they also bypass the distinction between fundamental and strategic uncertainty, for such a distinction is meaningful only in settings in which agents interact with one another.

6. Applied work often confounds the aforementioned two types of uncertainty. To some extent, this is unavoidable: in many applications, accommodating uncertainty about the actions of others *requires* uncertainty about fundamentals. Yet, not only are the two types of uncertainty conceptually distinct, but they can also have different implications. For instance, in the class of global games we study in Sections 4–5, what matters for equilibrium determinacy is the strategic uncertainty that is induced by incomplete information, not the underlying fundamental uncertainty. Furthermore, as we elaborate in Sections 7–8, incomplete information and the resulting strategic uncertainty can help operationalize the notions of coordination failure and animals spirits within unique-equilibrium models, in a manner that no kind of fundamental uncertainty alone can do. An integral part of our analysis is therefore to disentangle the two types of uncertainty as much as possible. We hope that this will help clarify, not only the theoretical underpinnings, but also the empirical implications of the blossoming macroeconomics literature on informational frictions.

⁵Strictly speaking, this statement applies only to Nash equilibria in pure strategies. However, mixed strategies are not relevant in our context, because we study settings in which there is no loss in ruling out mixed strategies. Also, the statement need not apply to the broader class of *correlated* equilibria, but this precisely because this solution concept departs from complete information.

⁶See, inter alia, Barsky and Sims (2011, 2012), Blanchard et al. (2013), Christiano et al. (2008), Jaimovich and Rebelo (2009), and the baseline model in Lorenzoni (2009). See also the related work by Collard and Dellas (2009) on imperfect observability of current fundamentals.

7. The aforementioned literature contains many strands. This chapter occupies only the sub-space in this literature that relates to coordination, higher-order beliefs, and strategic uncertainty. For complementary reviews of the literature, see Sims (2011), Mankiw and Reis (2011), and Veldkamp (2011).

8. Models with search and trading frictions are occasionally interpreted as models with imperfect coordination; see, e.g., Diamond (1982) and Shimer (2005). This is not what *we* have in mind. To the extent that these models maintain common knowledge of aggregate fundamentals and aggregate outcomes, for our purposes they impose the same level of coordination as the Arrow-Debreu framework.⁷

9. We are more sympathetic to the notion that asynchronous choice impedes coordination. We expand on this point in Section 6 by uncovering a subtle connection between the role of incomplete information and that of a Calvo-like friction in settings with strategic complementarities.

10. Our treatment of incomplete information is consistent with standard treatments in Bayesian games. In particular, our baseline framework specifies information in terms of Harsanyi types: the signal of an agent encodes, not only her beliefs about payoffs (fundamentals), but also her beliefs, or “state of mind”, regarding the types of other agents. We nevertheless adjust some of the definitions in manners that suit our applied purposes. We also bypass a number of deeper theoretical issues that are beyond the scope of this chapter, such as the equivalence between Harsanyi types and belief hierarchies (Mertens and Zamir, 1985) and the approximation of common knowledge with common beliefs (Monderer and Samet, 1989, Morris and Shin, 1997), as well as all the concerns of epistemic game theory (Dekel and Siniscalchi, 2015).

11. In our baseline framework, as well as in various applications, we favor a broad interpretation of incomplete information: signals are allowed to capture subjective states of mind about the beliefs and actions of others. For certain questions, however, a more narrow interpretation may be appropriate: when we study the aggregation of information through markets or other channels, or when we touch upon the desirability of central-bank communication, it is safer to assume that signals represent hard information.

2 Framework

In this section we introduce our main framework. Although it is static and abstract, the framework is quite flexible and stylizes role of general-equilibrium interactions, and of coordination, in a variety of applications. Furthermore, it facilitates the formalization of a number of key concepts in the next section; it helps illustrate how the predictions of standard macroeconomic models hinge on strong assumptions about the determination of beliefs; and it paves the way to a number of applications presented later on.

2.1 Actions and Payoffs

Consider a one-shot game, or “economy”, with a large number of players, or “agents”, indexed by $i \in [0, 1]$. Each agent i chooses an action $k_i \in D_k \subseteq \mathbb{R}^n$, with $n \geq 1$. His payoff is given by

$$u_i = U(k_i, \mathbf{K}, \theta_i), \tag{1}$$

where $\mathbf{K} \in D_{\mathbf{K}}$ denotes the distribution of actions in the cross section of the population (“the aggregate economic activity”), $\theta_i \in D_{\theta} \subseteq \mathbb{R}^m$, $m \geq 1$, summarizes any exogenous variable that is payoff-relevant to agent i (“her fundamental”), and U is a function, which will be further specified as we proceed.

⁷That said, trading frictions can be instrumental for sustaining lack of common knowledge: accommodating incomplete information *requires* a departure from the Arrow-Debreu framework.

As we move on, we will discuss a number of applications that can be nested either directly in the above framework or in appropriate extensions of it. In some of the applications, the players are speculators deciding whether to attack a currency, or depositors deciding whether to run against a bank. In others, they are firms making production and pricing choices, or consumers choosing how much to spend. Also note that the specification assumed in condition (1) allows an agent’s payoff to depend on the entire distribution of actions in the cross-section of the population. In applications, however, the first and second moments (i.e., the mean and the dispersion) often suffice. Furthermore, although we momentarily allow for both the actions and the fundamentals to be multi-dimensional, we will soon restrict attention to settings where $n = m = 1$. In such a case, we let $K \equiv \int k d\mathbf{K}(k)$ and $\sigma_k \equiv \sqrt{\int (k - K)^2 d\mathbf{K}(k)}$ denote, respectively, the mean action and the cross-sectional standard deviation of actions.

Remark 1. One may find it more appealing to represent the macroeconomy as a network where each agent is connected to a small set of other agents as opposed to the entire population. The ideas we develop in this chapter apply in such settings as well. For our purposes, the fact that the economy is modeled as a complete and symmetric network is mostly a simplification. For other purposes, however, it may be important to abandon this simplification. For instance, an important question that we will not address is how the effects of strategic uncertainty on observables depend on the network structure.

Remark 2. Recall that the question of interest for us is the equilibrium expectations of economic outcomes. In the static framework we introduce in this section, this question is reduced to the question of the equilibrium expectations of the *contemporaneous* actions of others. In dynamic applications in which decisions are forward-looking, the relevant expectations also regard the actions of others in *future* periods.

Remark 3. Notwithstanding the previous remark, there is an important multi-period setting that defies the aforementioned distinction and can be readily nested in our static framework: an Arrow-Debreu economy, with a complete market for time- and state-contingent goods. In such an economy, θ_i captures the preferences and endowments of agent i , k_i captures her net demand of the various goods, and the dependence of the agent’s utility on \mathbf{K} emerges from the dependence of equilibrium prices on the net demands of other agents. Importantly, because in a complete Arrow-Debreu market all agents get to observe all the relevant prices, it is *as if* they also observe the actions of other agents. This underscores that the Arrow-Debreu framework and many workhorse macroeconomic models that are based on it are, in effect, a particular class of static, complete-information games. Conversely, accommodating incomplete information in macroeconomic models requires some kind of market incompleteness (such as missing forward markets).

2.2 Examples

We now use two simple examples to illustrate how the abstract payoff structure assumed above corresponds to a reduced-form representation of more concrete applications. The first example is a neoclassical economy similar to those studied in Angeletos and La’O (2010, 2013), Benhabib, Wang, and Wen (2015), Sockin and Xiong (2015), and Huo and Takayama (2015a). The second is a monetary economy similar to those studied in Woodford (2003), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009). Later on, we will explore the implications of a certain kind of informational friction in these and other examples. For now, we abstract for the informational friction and focus on demonstrating how these examples can be nested in our framework. Apart from providing concrete interpretations of the framework, these examples help clarify that, for our purposes, strategic interaction is often synonymous to market-based general-equilibrium effects.

An neoclassical economy. There is a continuum of “farmers”, which can be interpreted as both firms and consumers, and a continuum of differentiated consumption goods. Each farmer specializes in the

production of a single good but consumes all the goods in the economy. We let $i \in [0, 1]$ index both the farmer and the good she produces. There are two stages. Production takes places in stage 1; trading and consumption take place in stage 2. Preferences are given by

$$u_i = v(c_i) - n_i, \quad (2)$$

where v is a strictly increasing and strictly concave function, n_i denotes labor, and c_i is a CES aggregator of the farmer's consumption of all the goods. That is,

$$c_i = \left[\int c_{ij}^{1-\eta} dj \right]^{\frac{1}{1-\eta}}$$

where c_{ij} denotes the consumption of good j by household i and $\frac{1}{\eta}$ is the elasticity of substitution across goods. The budget constraint of any farmer i is given by $\int p_j c_{ij} dj = p_i q_i$, where p_j denotes the (nominal) price of good j and q_i the farmer's output. Finally, production is given by

$$q_i = A_i n_i,$$

where A_i is the farmer's exogenous productivity.

As is well known, the CES specification implies that the optimal consumption bundle of farmer i satisfies

$$\int p_j c_{ij} dj = P c_i \quad \text{and} \quad \frac{c_{ij}}{c_i} = \left(\frac{p_j}{P} \right)^{-1/\eta} \quad \forall j,$$

where $P \equiv \left[\int p_j^{1-1/\eta} dj \right]^{\frac{1}{1-1/\eta}}$ is the ideal price index. Market clearing imposes $\int c_{ji} dj = q_i \quad \forall i$. Using this yields the following relation between the market-clearing prices and the quantities produced:

$$p_i = P \left(\frac{q_i}{Q} \right)^{-\eta} \quad \forall j, \quad (3)$$

where $Q \equiv \left[\int q_j^{1-\eta} dj \right]^{\frac{1}{1-\eta}}$ measures aggregate output. The budget constraint, on the other hand, gives $c_i = \frac{p_i q_i}{P} = Q^\eta q_i^{1-\eta}$. It follows that the utility of farmer i reduces to the following:

$$u_i = v(c_i) - n_i = v \left(Q^\eta q_i^{1-\eta} \right) - \frac{q_i}{A_i}.$$

We conclude that the example we have introduced here can be readily nested in our framework if we let

$$U(k_i, \mathbf{K}, \theta_i) = v \left(\left[\int \exp(x)^{1-\eta} d\mathbf{K}(x) \right]^{\frac{\eta}{1-\eta}} \exp(k_i)^{1-\eta} \right) - \exp(k_i - \theta_i), \quad (4)$$

with $k_i \equiv \log q_i$ and $\theta_i \equiv \log A_i$. That is, the economy under consideration can be interpreted as a game in which the players are the farmers, the actions are the quantities these farmers produce, the fundamentals are their exogenous productivities, and the payoffs are given by (4).⁸

⁸Note that k_i and θ_i could have been defined as the absolute levels of q_i and A_i , rather than their logarithms. We opt for the logarithmic transformation because of the following reason: when we assume that productivities are log-normally distributed and that the available information is Gaussian, the power specification of preferences and technologies in this example guarantees the existence of a unique equilibrium in which $k_i \equiv \log q_i$ can be expressed as a linear function of the firm's expectation of

A monetary economy. The structure of the economy is as above, except for the following two modifications. First, the farmers set nominal prices in stage 1 and commit to accommodate any quantity demanded in stage 2 at these prices. Second, the level of nominal GDP is given by

$$PQ = M, \tag{5}$$

where M is an exogenous variable. In line with much of the related literature, one may interpret the variation in M as monetary shocks.

Following similar steps to those above, we can solve for quantities as functions of prices (rather than the other way around). In particular, using (3), (5), and the budget constraint, we get

$$q_i = Q \left(\frac{p_i}{P} \right)^{-1/\eta} = MP^{1/\eta-1} p_i^{-1/\eta} \quad \text{and} \quad c_i = \frac{p_i q_i}{P} = MP^{1/\eta-2} p_i^{1-1/\eta}.$$

The present example is nested in our framework by letting $k_i \equiv \log p_i$, $\theta_i \equiv (\theta_{i1}, \theta_2) \equiv (\log A_i, \log M)$, and⁹

$$U(k_i, \mathbf{K}, (\theta_{i1}, \theta_2)) = v \left(\exp(\theta_2) \left[\int \exp(x)^{1-1/\eta} d\mathbf{K}(x) \right]^{\frac{1/\eta-2}{1-1/\eta}} \exp(k_i)^{1-1/\eta} \right) - \exp(\theta_2 - \theta_{i1} - \frac{1}{\eta} k_i) \left[\int \exp(x)^{1-1/\eta} d\mathbf{K}(x) \right]^{-1}. \tag{6}$$

That is, the actions now are the nominal prices set by the farmers and the fundamentals are the productivity and monetary shocks. Note that in this example, productivity is allowed to have both an idiosyncratic and an aggregate component, whereas the monetary shock has only an aggregate component. Woodford (2003) and Mankiw and Reis (2002) are then nested by shutting down all the productivity shocks, whereas Mackowiak and Wiederholt (2009) is nested by letting the productivity shocks be purely idiosyncratic.

Remark 1. The above two examples can be thought as representing two diametrically opposite cases: in the first, firms set quantities and let prices adjust; in the second, firms set prices and let quantities adjust. This difference is inconsequential when the firms face no uncertainty, but matters for both positive and normative issues if the firms are uncertain about either the relevant exogenous fundamentals or the choices of other firms. We elaborate on these points in Sections 8 and 9.

Remark 2. Although our examples are populated by “farmers”, one should of course not take this too literally. Starting with Lucas (1972), various authors have considered models that feature more appealing micro-foundations along with informational frictions. The key question for our purposes is whether these models make room for non-trivial strategic, or higher-order, uncertainty. This is what distinguishes the more recent works mentioned above from Lucas’ earlier contribution: as further explained in Subsection 8.3, the specific model used in Lucas (1972) shuts down the effects we are after in this chapter.

2.3 Shocks and Information

So far, we have specified actions and payoffs. To complete the description of the framework, we must specify the exogenous stochastic structure (shocks and information) and must also pick a solution concept. We complete these two tasks in, respectively, the present and the next subsection.

$\theta_i \equiv \log A_i$ and K , with K itself being equal to $\log Q$ plus a constant. This in turn means that the neoclassical economy under consideration maps into the specific class of games we study in Sections 7 and 8.

⁹The logarithmic transformation is used for the same reason as in footnote 8.

We let $\omega_i \in D_\omega \subseteq \mathbb{R}^l$ denote the signal received by agent i , where by “signal” we mean the entire information set of the agent, or equivalently her Harsanyi type. We denote by Ω the distribution of ω_i in the cross-section of the population, and by Θ the cross-sectional distribution of θ_i . We let D_Ω and D_Θ denote the sets of the possible values for, respectively, Ω and Θ . We let \mathbf{S} denote the distribution that Nature uses to draw, in an i.i.d. fashion, a pair $s_i = (\omega_i, \theta_i)$ for each i , and let $D_{\mathbf{S}}$ denote the set of possible values for \mathbf{S} . We assume that a version of the law of large numbers applies, so that \mathbf{S} is also the cross-sectional distribution of s_i in the population. It follows that Θ and Ω are also the marginals of \mathbf{S} in, respectively, the θ and ω direction. We introduce aggregate uncertainty by allowing \mathbf{S} (and hence also Θ and Ω) to be random, drawn from the set $D_{\mathbf{S}}$ according to some fixed distribution \mathcal{P} , which constitutes the common prior.

One can thus describe the “moves of Nature” as follows. First, Nature draws \mathbf{S} according to \mathcal{P} . Next, Nature uses \mathbf{S} to draw, in an i.i.d. manner, a pair $s_i = (\theta_i, \omega_i)$ for each agent i . Finally, Nature reveals ω_i (and only that) to agent i .

The objects $\{U, D_k, D_{\mathbf{K}}, \mathcal{P}, D_\theta, D_\Theta, D_\omega, D_\Omega, D_{\mathbf{S}}\}$ and all the aforementioned facts are common knowledge. What may not be common knowledge is the realization of \mathbf{S} , and therefore also the realizations of Θ and Ω . Different specifications of the stochastic structure—that is, of the prior \mathcal{P} and the associated domains for the fundamentals and the signals—can accommodate different scenarios about how much each agent knows, not only about her own fundamentals, but also about the fundamentals and the information of other agents. Importantly, we can accommodate aggregate and idiosyncratic shocks, not only to fundamentals, but also to information.

To give a concrete example, consider the neoclassical economy introduced earlier on and let $\theta_i \equiv \log A_i = \bar{\theta} + \xi_i$, where $\bar{\theta}$ and ξ_i are independent Normally distributed variables with mean zero and variances, respectively, $\sigma_{\bar{\theta}}^2$ and σ_{ξ}^2 , and ξ_i is i.i.d. across i . Then, variation in $\bar{\theta}$ represents aggregate TFP shocks, whereas variation in ξ_i represents idiosyncratic TFP shocks. Concerning the information structure, we could then let $\omega_i = (x_i, z)$, where $x_i = \theta_i + \epsilon_i$, $z = \bar{\theta} + \zeta$, and ϵ_i and ζ are Normally distributed, independent of one another, of $\bar{\theta}$, and of ξ_i for all i , with mean zero and fixed variances σ_ϵ^2 and σ_ζ^2 . This would mean that the information set of an agent is given by the pair of two signals, a private signal of her own TFP and a public signal of the aggregate TFP. Finally, because all the relevant cross-sectional distributions, namely Θ , Ω , and \mathbf{S} , are now Normal distributions with fixed variances, we can think of the aggregate shocks in terms of the pair $(\bar{\theta}, \zeta) \in \mathbb{R}^2$ rather than of the high-dimensional objects (Θ, Ω) or \mathbf{S} .

In the example described above, x_i is a signal, not only of θ_i , but also of $\bar{\theta}$; and since $\bar{\theta}$ is the mean of the realized distribution of x_j for $j \neq i$, it follows that x_i is also signal of the signals of others. More generally, note that different realizations of \mathbf{S} correspond to different conditional distributions for ω_i , which means that ω_i is a signal of \mathbf{S} and hence also of Ω .¹⁰ This underscores that ω_i shapes the agent’s belief, not only about her own payoff, but also about the *information* of others. Once we impose equilibrium, the latter will mean that ω_i shapes the beliefs about the *actions* of others.

We can now offer the following definitions, which help fix the language we use in this chapter.

Definition 1. *The realized **fundamentals** are given by Θ , the distribution of the exogenous payoff characteristic θ_i in the cross-section of the population. The realized **information** is given by Ω , the corresponding distribution of the signal ω_i . The **state of Nature** is given by \mathbf{S} , the joint distribution of (ω_i, θ_i) .*¹¹

¹⁰Note that ω_i is also a signal of Θ , the fundamentals of others. This, however, is not relevant per se. An agent cares only about her own fundamental and the actions of others. In equilibrium, the latter are pinned down by their information. What is relevant is therefore only the information that ω_i contains about the agent’s own fundamental and the information of others.

¹¹Note that \mathbf{S} is the joint distribution of $s_i = (\theta_i, \omega_i)$, not just the pair of the corresponding marginals: \mathbf{S} contains more information than the pair (Θ, Ω) .

Note that the above objects, as many others we introduce below, refer to the cross-section of the population, as opposed to any specific agent. This reflects merely our focus on macroeconomic outcomes, as opposed to the behavior of each individual player.

We now introduce our notation for first- and higher-order beliefs. Let b_i denote the belief of agent i about her own fundamental and the aggregate fundamentals, that is, the posterior of (θ_i, Θ) conditional on ω_i . This belief encapsulates the agent’s information about the payoff structure of the environment and is also known as her *first-order* belief (of the underlying fundamental). Clearly, b_i is a function of ω_i . Let \mathbf{B} denote the realized distribution of b_j in the cross-section of the population; this is a function of Ω . Define the *second-order* belief of agent i as her posterior about \mathbf{B} and denote it by b_i^2 . Since \mathbf{B} is a function of Ω and since ω_i pins down i ’s belief about Ω , ω_i pins down b_i^2 as well. Let \mathbf{B}^2 denote the cross-sectional distribution of second-order beliefs. We can then iteratively define, for any $h \geq 3$, the h -th order belief of i , b_i^h , as her belief about \mathbf{B}^{h-1} ; and the object \mathbf{B}^h as the cross-sectional distribution of b_j^h . To ease notation, we also let $b_i^1 \equiv b_i$ and $\mathbf{B}^1 \equiv \mathbf{B}$.

Definition 2. *The (aggregate) beliefs of fundamentals are given by \mathbf{B} , the cross section of the first-order beliefs of own and aggregate fundamentals. The corresponding belief hierarchy is given by $\{\mathbf{B}^h\}_{h=1}^\infty$. An agent faces **higher-order uncertainty** (about the fundamentals) if she is uncertain about $\{\mathbf{B}^h\}_{h=1}^\infty$.*

Because $\{\mathbf{B}^h\}_{h=1}^\infty$ is pinned down by Ω , we can think of Ω interchangeably either as the cross-sectional profile of information or as a summary statistic of the belief hierarchy. If Ω is known to an agent, then so is $\{\mathbf{B}^h\}_{h=1}^\infty$. Conversely, uncertainty about Ω helps accommodate higher-order uncertainty about fundamentals. Importantly, once we impose equilibrium, such higher-order uncertainty helps induce *first-order* uncertainty about the actions of others, ultimately capturing the coordination friction we are after.

Remark 1. For the rest of this chapter, and unless otherwise stated, the terms “higher-order beliefs” and “higher-order uncertainty” refer to the kind of higher-order beliefs of fundamentals defined above, as opposed to higher-order beliefs of either Harsanyi types (beliefs of beliefs of... Ω) or actions (beliefs of beliefs of... \mathbf{K}). In general, these three types of higher-order beliefs are distinct from one another. For instance, in Aumann (1974, 1987), private information about payoff-irrelevant variables (correlation devices) helps sustain higher-order uncertainty about the information and the actions of others, while maintaining common knowledge of fundamentals.¹² Furthermore, the distinction is central to epistemic game theory. For the purposes of this chapter, however, one can think of the three types of higher-order uncertainty as different facets of one and the same departure from workhorse macroeconomic models. Indeed, in the applications we study in this chapter, and insofar as the equilibrium is unique, the equilibrium beliefs of actions is pinned down by the hierarchy of beliefs of fundamentals. It follows that the friction we are after can be accommodated only by introducing higher-order uncertainty about fundamentals.

Remark 2. Whenever we use the terms “expectations”, “beliefs”, or “uncertainty” without the explicit qualifier “higher-order” in front of these terms, we refer to first-order beliefs.

Remark 3. In game theory, it is a standard practice to write the payoff of a player as function of her Harsanyi type. This could be done here by recasting payoffs as $V(k_i, \mathbf{K}, \omega_i) \equiv \int U(k_i, \mathbf{K}, \theta_i) db(\theta_i|\omega_i)$, where $b(\theta_i|\omega_i)$ is the posterior belief about θ_i conditional on ω_i . We have opted to express payoffs as functions of θ_i rather than ω_i in order to disentangle fundamentals from information sets and to accommodate the possibility that the econometrician observes θ_i but not ω_i .

¹²With our notation, this corresponds to situations where \mathbf{B} is commonly known, but Ω is not. This is because Ω contains correlation devices about which agents have private information. By the same token, although higher-order beliefs collapse to first-order beliefs in the case of fundamentals, the same is not true in the case of actions.

2.4 Equilibrium

Let us now turn to the solution concept. In line with the vast majority of macroeconomic research, we assume that agents play according to Rational-Expectations Equilibrium. In the present framework, this means the following.

Definition 3. A *rational-expectations equilibrium* (or, simply, an *equilibrium*) is a strategy $k^* \in D_\omega \rightarrow D_k$ and a mapping $\mathbf{K}^* : D_\Omega \rightarrow D_{\mathbf{K}}$ such that:

(i) the strategy k^* is the best response to the mapping \mathbf{K}^* , that is,

$$k^*(\omega) \in \arg \max_{k \in D_k} \mathbb{E}[U(k, \mathbf{K}^*(\Omega), \theta)|\omega] \quad \forall \omega; \quad (7)$$

(ii) for any Ω , $\mathbf{K}^*(\Omega)$ is the distribution of k that obtains when the distribution of ω is Ω and $k = k^*(\omega)$.

This definition is essentially the same as the definition of Bayesian-Nash Equilibrium in games. This, however, does not mean that the agents themselves engage in strategic reasoning. The context we have in mind is a large market-based economy, in which there is myriad of firms and consumers, each one trying to predict economic outcomes such as aggregate employment and output. This is quite different from an arms race or other contexts in which it may be more appealing to think that each player is explicitly trying to second guess the moves of other players. We thus envision that each agent treats \mathbf{K} as an exogenous stochastic process and we seek to understand the predictions that any given model makes about this process.

Under this interpretation, the adopted solution concept imposes that all agents perceive the same stochastic process for \mathbf{K} , as well as that this commonly perceived process coincides with the actual one generated by the joint behavior of all agents. Furthermore, ω_i becomes a signal of the realized value of \mathbf{K} . By the same token, we can think of ω_i as the agent’s “state of mind” about what is going on in the economy. Importantly, insofar as different agents do not share the same ω_i , they also do not need to share the same expectations about the endogenous outcomes of interest.

Summing up, we have that ω_i can play any subset of the following three modeling roles. First, it shapes the agent’s beliefs about her own fundamental. Second, it shapes the agent’s beliefs about the information of others and thereby also her higher-order beliefs of the fundamentals. Third, it shapes the agent’s equilibrium expectations about the actions of others and thereby about the endogenous outcomes of interest. The first two modeling roles are by construction; the third rests on imposing the adopted solution concept.

Some of the applied literature on informational frictions relies on the first modeling role: it removes knowledge of θ_i and lets ω_i be a noisy signal of θ_i , often in settings that rule out strategic interactions. In this chapter, by contrast, we are primarily concerned with the other two modeling roles, which are themselves relevant only in the presence of strategic interactions.¹³

Remark 1. Once the equilibrium concept has been imposed, higher-order beliefs become a sideshow: in equilibrium, an agent’s signal ω_i pins down her posterior about the joint distribution of θ_i and K , which is all that matter for her behavior. Therefore, the equilibrium can be characterized without ever invoking higher-order beliefs. However, understanding the structure of higher-order beliefs turns out to be instrumental to understanding the structure of the entire equilibrium set and its sensitivity to variations of the information structure. This will become evident as we proceed.

Remark 2. When we move on to dynamic settings, the equilibrium concept we use is essentially the same as Perfect Bayesian Equilibrium. The only difference—and a useful simplification—is that we do not

¹³With this point in mind, in the sequel we will often consider specifications of the information structure that help isolate the stochastic variation in the equilibrium expectations of economic outcomes from the stochastic variation in beliefs of fundamentals.

always need to specify out-of-equilibrium beliefs. This is because, in the typical macroeconomic setting, each private agent is too small to trigger deviations that are detectable by, and payoff-relevant to, any other agent. An important exemption to this rule, however, is settings with big players, such as the government.

Remark 3. While we have introduced the notation and the concepts needed for our purposes, we have not spelled out the mathematical structure of the underlying probability space and have swept under the rug technical issues regarding measurability and existence. We hope that this simplifies the exposition without obstructing the understanding of the more substantive issues. If the reader is annoyed by this imprecision, he/she can read the next section as if the sets of actions, fundamentals, and signals were finite. In subsequent sections, we will impose enough assumptions on the stochastic structure and the function U to guarantee the existence of an equilibrium, as well as its tractable characterization.

3 Imperfect Coordination

In this section, we use our framework to define certain key notions and to elaborate on the central theme of this paper. We also discuss how equilibrium expectations and outcomes are determined under incomplete information, and contrast them to their complete-information counterparts. The complete-information case defines the benchmark relative to which equilibrium expectations are enriched.

3.1 Some Key Notions

Macroeconomic models are typically used to provide a structural interpretation of the data, or to deliver counterfactuals that can help guide policy. To clarify what this means within our framework, we start by distinguishing the exogenous objects that the theorist can specify “at will” from the endogenous objects that she wishes to predict. The exogenous objects are the payoff function U together with the stochastic structure of the fundamentals and the information sets introduced in the previous subsection. The endogenous objects are the actions of the agents and the realized payoffs. Because the payoffs are themselves pinned down by the fundamentals and the actions, the only “true” endogenous objects are the actions. With this in mind, we adopt the following definition.

Definition 4. *The **economy’s outcome**, or the economy’s endogenous state, is given by \mathbf{K} , the distribution of actions in the cross-section of the population.*

In the context of an application, \mathbf{K} may represent the hiring choices of the firms, the spending choices of consumers, or the capital stock of the economy (which is itself the product of past investment choices by firms and consumers). What a model ultimately does is to impose certain restrictions on the joint distribution of these endogenous objects with the exogenous shocks to fundamentals and information sets.

Suppose that the data that are available for testing or quantifying the model regard *at most* the following objects: the fundamentals, the agents’ expectations of their fundamentals, their expectations of the endogenous state, and their actions. Then, the testable implications, or predictions, of the model are the restrictions that it imposes on the joint distribution of these objects. Of course, one can expand this list of objects by including the information sets and/or the higher-order beliefs. For applied purposes, however, this seems redundant, because data on information sets and higher-order beliefs are hardly available, in contrast to data on choices and on forecasts of economic outcomes. With these points in mind, we let π_i denote the expectation of agent i about \mathbf{K} ; we let $\mathbf{\Pi}$ denote the distribution of such expectations in the population; and we adopt the following definition.

Definition 5. The *predictions* of the model are the restrictions it imposes on the joint distribution of $(\theta_i, b_i, \pi_i, k_i, \Theta, \mathbf{B}, \mathbf{\Pi}, \mathbf{K})$.

One can thus think of any model as a “box” that takes the joint distribution of (s_i, \mathbf{S}) as an input and delivers a joint distribution for $(\theta_i, b_i, \pi_i, k_i, \Theta, \mathbf{B}, \mathbf{\Pi}, \mathbf{K})$ as an output. Depending on the question of interest and/or the available data, one can then look at certain conditionals or marginals of the aforementioned distribution, or certain moments of it. One can thus also answer the following type of applied questions: How much does individual activity respond to idiosyncratic shocks to fundamentals? What is the corresponding response of aggregate activity to aggregate shocks to fundamentals? Are expectations of economic outcomes pinned down by expectations of fundamentals? What is the volatility in aggregate economic activity and how much of it is explained by volatility in fundamentals?

To illustrate, consider our earlier neoclassical economy from Subsection 2.2, in which case θ_i captures the exogenous TFP of an agent (farmer, firm, or island) and k_i captures her production level. Assume that TFP is log-normally distributed in the cross section, let $\bar{\theta}$ denote the aggregate log-TFP, and assume $\omega_i = (\theta_i, \bar{\theta})$; the latter means that every agent knows perfectly both her own TFP level and the aggregate TFP level. These assumptions, in conjunction with the power-form specification of preferences and technologies, imply that the equilibrium levels of output at the individual and the aggregate level are given by, respectively,

$$\log q_i \equiv k_i = \kappa_1(\theta_i - \bar{\theta}) + \kappa_2\bar{\theta} \quad \text{and} \quad \log Q \equiv K = \kappa_2\bar{\theta},$$

where κ_1 and κ_2 are fixed scalars, pinned down by the underlying preference and technology parameters. Equilibrium outcomes are therefore pinned down by fundamentals—and so are equilibrium expectations. Furthermore, the joint distribution of $(\theta_i, b_i, \pi_i, k_i, \Theta, \mathbf{B}, \mathbf{\Pi}, \mathbf{K})$ is now conveniently summarized by the distribution of $(\theta_i, k_i, \bar{\theta}, K)$. Turning to some of the applied questions raised above, note that the “micro” elasticity of the response of individual output to idiosyncratic TFP shocks is given by κ_1 , whereas the corresponding “macro” elasticity is given by κ_2 . Finally, note that

$$\kappa_1 = \frac{Cov(k_i, \theta_i | \bar{\theta})}{Var(\theta_i | \bar{\theta})} \quad \text{and} \quad \kappa_2 = \frac{Cov(K, \bar{\theta})}{Var(\bar{\theta})},$$

which illustrates how the joint distribution of $(\theta_i, k_i, \bar{\theta}, K)$ contains the model’s predictions about the two elasticities of interest.¹⁴

Later on, we will explore how this kind of predictions change as one departs from standard informational assumptions. Importantly, we will draw a distinction between two kinds of departure: those that regard the information, or beliefs, agents have about their fundamentals; and those regard the information, or beliefs, agents have about one another’s beliefs and the associated friction in coordination (or lack thereof).

To do so, we need to introduce a few additional definitions. We start by distinguishing the two kinds of uncertainty mentioned in the Introduction.

Definition 6. An agent faces *fundamental uncertainty* if and only if, conditional on her information, she is uncertain about the value of (θ_i, Θ) .

¹⁴Macroeconomists are primarily interested in the macro elasticity, κ_2 . Empirical studies that use macro data, such as the paper by Galí (1999) reviewed in Subsection 8.1, seek to estimate this elasticity directly. By contrast, studies based on cross-sectional data concentrate on the identification of the micro-level elasticity, κ_1 (where “micro” may mean either at the level of an individual firm/household, or at the level of a region). The difference between the two types of elasticities reflects general-equilibrium effects that operate at the aggregate level and can thus not be identified from the cross section. See, however, Subsection 8.2 for a discussion of how incomplete information can help reduce this difference in the short run.

Definition 7. *An agent faces **uncertainty about the actions of others**, or about the economy’s outcome, if and only if, conditional on her information, she is uncertain about the value of \mathbf{K} .*

As noted in the Introduction, the latter type of uncertainty is connected to the notion of strategic uncertainty in games. In what follows, we use the term “strategic uncertainty” to refer to the type of uncertainty defined in Definition 7 and proceed to study how incomplete information helps accommodate such uncertainty in equilibrium. This is consistent with Morris and Shin (2002, 2003). Note, however, that the same term is often used in game theory to refer to a distinct object, namely the uncertainty that agents may face about the strategies of others outside equilibrium.

Also note that Definitions 6 and 7 refer to the uncertainty the agents face in the interim stage, after they have received their information, as opposed to the stochasticity that exists at the ex ante stage. For instance, consider the neoclassical economy introduced in Subsection 2.2 and suppose that TFP is random but perfectly observed. We would say that the agents in this economy face no fundamental uncertainty.¹⁵

The aforementioned definitions invite us to put ourselves in the shoes of a particular agent and to examine the uncertainty the agent faces conditional on her information. The next definition, instead, invites us to inspect the cross-section of agents conditional on the state of nature. It then asks the following question: have the agents been able to reach a common belief about the relevant economic outcomes?

Definition 8. ***Coordination is imperfect** in state \mathbf{S} if and only if, in that state, a positive measure of agents fail to reach the same belief about \mathbf{K} .*

To see what motivates this definition, consider the following hypothetical scenario. Prior to playing an equilibrium, the agents get together in the same room and communicate with one another until they reach common knowledge, not only of the fundamentals, but also of their intended actions and the relevant economic outcomes (i.e., of \mathbf{K} in our setting). We think of this situation as being conducive to “perfect coordination”. Conversely, we say that there is a friction in coordination if agents are unable to reach a common belief about one another’s actions and thereby about the relevant economic outcomes.

The following remark helps refine the friction we are after. In our framework, once we impose equilibrium, reaching the same belief about the choices of others means facing no uncertainty about them. The last property, however, is an artifact of the static nature of our framework. In dynamic settings, agents may face uncertainty about the *future* choices of others, even if they do not face any uncertainty about the *current* choices of others. Reaching a common belief about the relevant economic outcomes is therefore distinct from facing no uncertainty about them. In particular, standard macroeconomic models accommodate for uncertainty in future economic activity, due to uncertainty in future fundamentals; they nevertheless impose that agents always hold exactly the same equilibrium beliefs about all future economic outcomes. It is this restriction—the unanimous consensus about the current state and the future prospects of the economy—that we interpret as perfect coordination and that we seek to depart from.

We can now also elaborate on our earlier claim that some of the related applied literature either confounds the roles of payoff and strategic uncertainty, or is exclusively concerned with fundamental uncertainty. When Sims (2003), Reis (2006), and Gabaix (2014) introduce different (but also complementary) formalizations of “inattention”, they focus on single-agent decision problems, thus precluding a meaningful distinction between the two types of uncertainty.¹⁶ When Bloom (2009), Bloom et al. (2014), and Arellano, Bai and Kehoe

¹⁵One has to be careful with the extension of this notion to dynamic applications: in the RBC model (with capital), we would say that fundamental uncertainty is absent when the agents know, not on current TFP, but also future TFP.

¹⁶The same applies to many of the recent applications of rational inattention, including Luo (2008), Matejka (2015a, 2015b), and Matejka et al (2016); but not to those that let rational inattention be, in effect, a specific micro-foundation of strategic uncertainty, such as Mackowiak and Wiederholt (2009, 2015). See Subsection 8.5 for further discussion.

(2012) study the role of “uncertainty shocks”, or when Jaimovich and Rebelo (2009), Christiano et al. (2008), Beaudry and Portier (2006), Barsky and Sims (2011, 2012), and Blanchard et al. (2013) study the role of “news” and “noise shocks”, they consider different facets of fundamental uncertainty, but do not accommodate the kind of strategic uncertainty, or the friction in coordination, that we are after in this chapter.¹⁷

By contrast, when Aumann (1974, 1987) introduced the concept of correlated equilibrium and formalized its connection to Bayesian rationality in games, he ruled out fundamental uncertainty and used incomplete information to model *exclusively* uncertainty about the actions of others. A similar point applies to Rubinstein’s (1989) “email game”: an imperfection in communication was introduced in order to inhibit coordination, not to add uncertainty in payoffs.

The preceding observations help clarify the theme of this chapter. However, as noted in the Introduction, a sharp separation between fundamental and strategic uncertainty may not always be possible, especially once we impose equilibrium.

For instance, in the models we study in Sections 7–8, the equilibrium is a unique irrespective of the information structure and the action of each agent is pinned down by her hierarchy of beliefs about the fundamentals. It follows that uncertainty about the equilibrium value of \mathbf{K} can obtain in any particular state if and only if the agents lack common knowledge of fundamentals in that state, which in turns means that at least some of them face fundamental uncertainty in some state of nature.

What is more, suppose that all agents share the same fundamental, namely $\theta_i = \bar{\theta}$ for all i , and that the information of each agent i consists of a single private signal $x_i = \bar{\theta} + \epsilon_i$, where ϵ_i is idiosyncratic noise that is drawn from a commonly known distribution and that washes out at the aggregate level. Under these assumptions, $\bar{\theta}$ becomes a sufficient statistic for Ω , the profile of information in the cross section, and thereby also for the equilibrium profile of actions. Once we impose equilibrium, no *apparent* difference therefore remains between the two types of uncertainty: predicting \mathbf{K} is the same as predicting $\bar{\theta}$.

To better appreciate the role of strategic uncertainty, it is therefore useful to do one, or both, of the following: (i) consider sufficient rich information structures that make sure that the uncertainty agents face about \mathbf{K} is not spanned by their uncertainty about fundamentals; (ii) momentarily drop equilibrium and, instead, contemplate the kind of higher-order reasoning that may justify certain forecasts of \mathbf{K} on the basis of the common knowledge of the environment and of the rationality of the agents. Both of these paths will be explored throughout this chapter.

Remark. No matter the information structure, once equilibrium is imposed, the distinction between the two forms of uncertainty becomes blurred in the eyes of the agents inside the model: both the fundamentals and the endogenous outcomes are treated by each agent as exogenous stochastic processes. In the eyes of the outside observer, however, the two forms of uncertainty can be distinguished by the different marks they leave on the observables of the model. Dissecting this difference is an integral part of this chapter.¹⁸

¹⁷The same statement applies to the baseline model in Lorenzoni (2009): the key mechanism in that paper rests on the uncertainty consumers face about future fundamentals (TFP) and the deviation of monetary policy from replicating flexible-price allocations. The paper contains also an extension that features dispersed information. This extension is used to justify higher fundamental uncertainty and to raise the quantitative potential of the aforementioned mechanism. It also makes a contribution to the literature we are concerned with, because of the elegant ways in which the author deals with the geography of information and the dynamics of learning. Yet, the key mechanism does not rest on the dispersion of information, and is therefore outside the scope of our chapter.

¹⁸This remark is consistent with the Bayesian approach to games (e.g., Harsanyi, 1967–1968, Aumann, 1987, Aumann and Heifetz, 2002); but it also reinforces why we treat incomplete information and higher-order beliefs as instruments for generating first-order uncertainty about equilibrium outcomes, rather than objects of independent interest.

3.2 Informational Friction

As noted in the Introduction, the present chapter views incomplete information as a device for introducing frictions in coordination, rather than as a device for introducing uncertainty about payoff-relevant fundamentals such as technology and preferences. In this context, we now proceed to clarify two distinct forms of informational friction and to point out which one is more relevant for our purposes.

Definition 9. *Information is said to be **perfect** if (θ_i, \mathbf{S}) is known to all agents in all states \mathbf{S} . If the converse is true, information is said to be **imperfect**.*

Definition 10. *Information is said to be **complete** if Ω is known to all agents (or at least to all but a zero-measure set of agents) in all states of nature. If the converse is true, information is said to be **incomplete**.*

This terminology differs somewhat from the one typically found in the related applied literature, where terms “informational friction” and “imperfect”, “dispersed”, or “heterogeneous” information are often used to refer to either of the two concepts defined above. The problem is that, although many papers depart from the clearly defined benchmark of perfect information, it is not always clear whether their key departure is (i) the introduction of some type of fundamental uncertainty or (ii) the relaxation of common knowledge and the associated strategic uncertainty. With the above two definitions, we seek to clarify the difference between two important types of informational frictions.¹⁹

Note, in particular, that perfect information imposes common knowledge of both Ω and Θ , whereas complete information imposes only common knowledge of Ω .²⁰ For example, suppose that Ω is commonly known, but Θ is not. Then, information is *imperfect* in the sense that the agents are uncertain about one another’s fundamentals, and yet information is *complete* in the sense that the agents face no uncertainty about one another’s information sets. Finally, note that our definition of complete information rules out private information about aggregate shocks, but allows for private information about idiosyncratic shocks; that is, complete information imposes that all agents share the same beliefs about both Θ and more generally about \mathbf{S} , but allows each agent to know more than others about the idiosyncratic component of her fundamental.²¹

3.3 Imperfect Coordination Equals Incomplete Information

So far, the notions of imperfect coordination, strategic uncertainty, and incomplete information have been defined in a manner that allows them to be disconnected from one another. However, once equilibrium is imposed, these notions become essentially the same.

Proposition 1 (Information and Coordination). *The following are true only when information is incomplete:*

- (a) *higher-order beliefs of fundamentals diverge from first-order beliefs;*
- (b) *agents face uncertainty about the actions of others in any given equilibrium;*
- (c) *coordination is imperfect in any given equilibrium.*

¹⁹Note that our notion of incomplete information is the same as the one found in game theory, but our notion of imperfect information is distinct. Although we regret this discrepancy, we believe that the definitions we adopt here help clarify both the language used in macroeconomics and the mechanisms featured in different applications.

²⁰To be precise, our definition of complete information requires only *mutual* knowledge of Ω . But since Ω is the entire profile of Harsanyi, mutual knowledge of Ω implies common knowledge of Ω .

²¹The last point underscores that the heterogeneity, or dispersion, of information that is relevant for our purposes is the one that regards aggregate shocks, as opposed to the one that is at the centre of the the Mirrlees literature (aka New Public Finance): that literature introduces private information about idiosyncratic fundamentals in order to accommodate certain incentive problems, but maintains the assumption that all aggregate shocks (including the cross-sectional distribution of information sets) are commonly known, thus ruling out all the effects that are of interest to us in this chapter.

We prove this by showing that complete information rules out (a), (b), and (c). Recall first that first-order beliefs \mathbf{B} are pinned down by Ω . If information is complete, meaning that Ω is commonly known, then so is \mathbf{B} . The common-prior assumption then imposes that higher-order beliefs collapse to first-order beliefs. Next, note that, in any given equilibrium, the endogenous outcome \mathbf{K} is given by a known function \mathbf{K}^* of the realized Ω . If the latter is commonly known, then so is \mathbf{K} .²² That is, complete information rules out strategic uncertainty in equilibrium. Finally, note that the equilibrium beliefs of \mathbf{K} are pinned down by the beliefs of Ω , which themselves collapse to a Dirac on the true Ω when information is complete. It follows that all agents share the same belief about \mathbf{K} , which means the friction in coordination vanishes.

These arguments prove that incomplete information is necessary for obtaining (a), (b), and (c) in Proposition 1. Sufficiency is also true, provided we consider non-trivial forms of incomplete information; this will become clear as we proceed.²³ Furthermore, one can strengthen part (b) to read as follows: a agent can face uncertainty about \mathbf{K} even if she knows Θ , or even if she knows both Θ and $\{\mathbf{B}^h\}_{h=1}^H$ up to any finite H . This illustrates the richness of the uncertainty that agent can face about one another's equilibrium actions and the associated economic outcomes when, and only when, information is incomplete.

Apart from clarifying how complete information rules out the friction we are after in this chapter, the above result also highlights the following point.

Fact. *Workhorse macroeconomic models typically rule out imperfect coordination, not only because they assume complete information, but also because they impose a strong solution concept.*

This is due to the fact that the Rational-Expectations Equilibrium concept *itself* rules out uncertainty about the *strategies* of other agents, regardless of the information structure and the number of equilibria. It follows then that the agents can face uncertainty about the actions of others in equilibrium if and only if they do not share the same information, which is how incomplete information is defined.

Yet, this does not mean that incomplete information must be taken literally: as hinted in the Introduction, one can view incomplete information also as a substitute for relaxing the equilibrium concept.

Let us elaborate on what we have in mind. Fix a standard macroeconomic model, such as the textbook RBC model or its New-Keynesian sibling. These models deliver certain predictions about the nature of the business cycle, namely about the co-movement of key economic outcomes such as employment, consumption, and investment with one another, as well as with the underlying shocks to TFP or other fundamentals. These predictions rely, not only on features of the micro-foundations such as the Frisch elasticity of labor supply and the degree of price stickiness, but also on the combination of a strong solution concept with strong informational assumptions—a combination that imposes flawless coordination in beliefs.

Macroeconomists have long debated about the elasticity of labor supply, the degree of nominal rigidity, and numerous other aspects of the micro-foundations of macroeconomic models, but have paid relatively little attention to the bite that the aforementioned combination has on the predictions of their models, and thereby on their interpretation of the data. By contrast, it is precisely this combination which we are uncomfortable with and whose bite we would like to relax. Incomplete information is a vehicle for doing so and, in this sense, it is also as a substitute for relaxing the solution concept.²⁴

²²If there are multiple equilibria, the function \mathbf{K}^* varies across equilibria, but it is commonly known to the agents in any given equilibrium. This is true even in the case of sunspot equilibria; this case is nested by allowing ω_i to contain a publicly observed sunspot. Finally, if we consider correlated equilibria with imperfect correlation, the function \mathbf{K}^* is still commonly known in any equilibrium, but now different agents have different beliefs about the realized \mathbf{K} , simply because they have private information about the underlying correlation devices and therefore also about Ω .

²³See also Weinstein and Yildiz (2007a).

²⁴In this regard, the approach we take in this chapter can be viewed as complementary to, albeit distinct from, the approaches taken by Evans and Honkapohja (1999, 2009), Guesnerie (1992, 2008), Fuster, Laibson and Mendel (2010), Kurz (1994, 2012),

Summing up, we have now formalized the following point, which was anticipated in the Introduction.

Fact. *Incomplete information is a modeling device that expands the sources of uncertainty the economic agents face about endogenous economic outcomes and that allows the expectations of these outcomes to deviate from the expectations of fundamentals.*

We conclude this subsection by noting that, for our purposes, the uncertainty agents may face about their *own* fundamentals is of little consequence on its own right. To formalize this point, let

$$\hat{U}(k_i, \mathbf{K}, b_i) \equiv \int U(k_i, \mathbf{K}, \theta_i) db(\theta_i, \Theta),$$

where b_i is the agent’s first-order belief.²⁵ The following is then immediate.

Proposition 2. *The equilibrium set of the original model coincides with the equilibrium set of a variant model in which the agents are always perfectly informed about their own fundamentals. In this variant games, payoffs are given by $\hat{U}(k, \mathbf{K}, \hat{\theta}_i)$, individual fundamentals by $\hat{\theta}_i \equiv b_i$, and aggregate fundamentals by $\hat{\Theta} \equiv \mathbf{B}$.*

In this sense, we can always recast any given model as one in which the agents are perfectly informed about their own fundamentals. Of course, this does not mean that uncertainty about fundamentals is irrelevant. There is a long tradition in macroeconomics that studies different kinds of fundamental uncertainty within complete-information, and often representative-agent, models. Examples include a voluminous literature on the equity premium, as well as the recent literatures on news and uncertainty shocks. The sole role of the above result is to underscore, once more, the type of uncertainty that we are interested in. For our purposes, it is not relevant per se whether each agent knows her *own* fundamental; what is relevant is whether agents have common knowledge of *aggregate* fundamentals. To isolate the latter channel, we will therefore occasionally consider examples in which every agent i knows her own θ_i perfectly, but not necessarily Θ . By contrast, the leading example in the literature shuts down heterogeneity ($\theta_i = \theta$ for all i) and confounds the lack of common knowledge of aggregate fundamentals with uncertainty about own fundamentals.

3.4 Complete-Information Benchmark and Coordination Motives

The preceding subsection clarified the sense in which incomplete information introduces a friction in coordination. To elaborate on the observable implications of this friction and the applied value of the approach taken in this chapter, we will have to get more concrete. This is what we do in Sections 4–8. As a prelude to this, however, it is useful to work out the complete-information benchmark of our abstract framework. We do so in this subsection with the help of additional, simplifying, assumptions on the action space D_k and the payoff function U . We then use this benchmark to identify a feature that distinguishes the type of models studied in Sections 4–6 (“global games” and their applications) from those studied in Sections 7–8 (“beauty contests” and their applications), namely a feature that regards the strength of coordination motives.

With the function \hat{U} defined as before, we impose the following:

Assumption 1. *D_k is a closed interval in \mathbb{R} and \hat{U} is strictly concave and twice differentiable in k .*

and Sargent (2008). See also the review in Woodford (2013). What is common between our approach and these alternative approaches is the desire to relax, or enrich, the stochastic structure of beliefs.

²⁵Note that the definition of \hat{U} treats \mathbf{K} as a deterministic variable: the possibility that \mathbf{K} correlates with Θ in equilibrium is not relevant here. By the same token, \hat{U} depends on b_i only through the marginal of b_i over θ_i ; that is, it depends on the agent’s belief about her *own* fundamental, but is otherwise invariant to her belief about the *aggregate* fundamentals.

This simplifies the analysis by letting the action be uni-dimensional and by guaranteeing both the existence and the uniqueness of a solution to the individual's decision problem, that is, to the choice of k_i that is optimal for given beliefs of fundamentals and of the actions of others. We express this solution as follows:

$$k_i = \gamma(\mathbf{K}, b_i) \equiv \arg \max_k \hat{U}(k, \mathbf{K}, b_i).$$

Note here that agent i faces no uncertainty about \mathbf{K} , not only because information is complete, but also because she knows that others play a given equilibrium. That is, as anticipated, strategic uncertainty has been ruled out by the combination of a strong informational assumption with a strong solution concept.

With this in mind, we have the following result.

Proposition 3 (Complete Info). *Suppose information is complete. There exists a function Γ , which depends only on the function \hat{U} , such that the following is true: in any equilibrium and any state of nature, \mathbf{K} solves*

$$\mathbf{K} = \Gamma(\mathbf{K}, \mathbf{B}). \tag{8}$$

This result follows trivially from letting $\Gamma(\mathbf{K}, \mathbf{B})$ be the distribution of k that obtains when $k = \gamma(\mathbf{K}, b)$ and by noting that the function γ is essentially exogenous, in the sense that it is pinned down by the function \hat{U} . When the equilibrium is unique, this result means that equilibrium actions are pinned down *exclusively* by first-order beliefs of fundamentals. When, instead, there are multiple equilibria, the theory makes room for sunspots, that is, for random selections across the multiple solutions of condition (8). For any such selection, however, the equilibrium value of \mathbf{K} remains tied to \mathbf{B} through condition (8). Furthermore, whether the equilibrium is unique or not, the heterogeneity in the cross-section of actions is pinned down by the heterogeneity in beliefs of own fundamentals.²⁶ Finally, if information is perfect (which, recall, is a stronger requirement than the requirement of complete information), nothing essential changes: we merely need to replace b_i with θ_i and \mathbf{B} with Θ .

By contrast, these restrictions are relaxed when information is incomplete: equilibrium outcomes, whether at the aggregate level or in the cross section, may depend on higher-order beliefs, which themselves may differ from first-order beliefs. What exactly this implies for the observables of any given model will be explored in the subsequent sections of this chapter. For now, we continue with the complete-information benchmark.

To further sharpen the analysis, we next impose the following:

Assumption 2. \hat{U}_k depends on \mathbf{K} only through the mean action K and is differentiable in it.

We can then rewrite the best response of agent i as $k_i = g(K, b_i)$, where g is a function that is differentiable in K and that is implicitly defined by the solution to $\hat{U}_k(g(K, b), K, b) = 0$. Aggregating gives

$$K = G(K, \mathbf{B}),$$

where the function $G : D_k \times D_{\mathbf{B}} \rightarrow D_k$ is defined by $G(K, \mathbf{B}) \equiv \int g(K, b) d\mathbf{B}(b)$. This is essentially the same as Proposition 3 above, except for the fact that we now have a fixed-point relation in the mean action K alone as opposed to the entire distribution \mathbf{K} .

Since g is differentiable in K , G is also differentiable in K (as well as continuous in it). This, together with the compactness of D_k , guarantees the existence of equilibrium. To characterize the determinacy of the equilibrium, we introduce the following notions.

²⁶To see this, note that $k_i = \gamma(\mathbf{K}, b_i)$ can differ across two agents only if these agents have a different b_i ; and because all agents have the same beliefs about Θ , different b_i means different beliefs about θ_i .




Figure 1: Best responses with weak and strong complementarity

Definition 11. *The economy exhibits **strategic substitutability** if and only if $G_K(K, \mathbf{B}) < 0$ for all (K, \mathbf{B}) ; and **strategic complementarity** if and only if $G_K(K, \mathbf{B}) > 0$ for all (K, \mathbf{B}) .*

Definition 12. *Strategic complementarity is of the **weak** form if $G_K(K, \mathbf{B}) \in (0, 1)$ for all (K, \mathbf{B}) , and of the **strong** form if $G_K(K, \mathbf{B}) > 1$ for some (K, \mathbf{B}) such that $K = G(K, \mathbf{B})$ and K is in the interior of D_k .*

We have the following result.

Proposition 4 (Equilibrium Determinacy). *Suppose information is complete.*

- (i) *An equilibrium always exists.*
- (ii) *The equilibrium is unique if the economy exhibits either strategic substitutability or weak strategic complementarity.*
- (iii) *There exist multiple equilibria if the economy exhibits strong strategic complementarity.*

We illustrate these possibilities in Figure 1 for two examples. In both examples, we fix a particular realization of \mathbf{B} and draw $G(K, \mathbf{B})$ against K . The case of weak complementarity corresponds to the solid curve, which intersects only once with the 45-degree line; point B then gives the unique equilibrium value of K , for a given value of Θ . The case of substitutability is similar and is thus omitted. The case of strong complementarity, on the other hand, corresponds to the dotted curve, which intersects three times with the 45-degree line; points A and C then give the two stable equilibrium outcomes of the economy, whereas point B represents an unstable equilibrium outcome.²⁷

The above categorization anticipates the two classes of models we study in the sequel: Section 4 introduces a class of models that feature strong complementarity and that have been used to study self-fulfilling crises; Section 7 turns attention to a class of models that feature weak complementarity and that are commonly used to study business cycles and asset prices. These two classes of models behave quite differently, not only under complete information, but also under incomplete information. This is because the impact of higher-order uncertainty depends on how essential coordination is in the first place. As a result, higher-order uncertainty has a more pronounced effect in the former class than in the latter one.²⁸

²⁷The notion of stability used here is that of iterated best responses.

²⁸In particular, the beauty contests studied in Section 7 pass Weinstein and Yildiz's (2007b) criterion of "global stability under uncertainty," which means that beliefs of sufficiently high order have a vanishing effect. This is not the case with either the static global games of Section 4, or the related dynamic games of Section 6.

We conclude this section by noting how complete information ties “animal spirits” to equilibrium indeterminacy. To formalize this point, we define the notion of animal spirits as follows.

Definition 13. *The economy features **animal spirits** if there are two states $(\mathbf{S}_1, \mathbf{S}_2)$, with respective beliefs of fundamentals $(\mathbf{B}_1, \mathbf{B}_2)$ and respective outcomes $(\mathbf{K}_1, \mathbf{K}_2)$, such that $\mathbf{B}_1 = \mathbf{B}_2$ and $\mathbf{K}_1 \neq \mathbf{K}_2$.*

This definition is consistent with those found in Cass and Shell (1983) and the extant macroeconomic literature on sunspot fluctuations: it simply says that equilibrium actions can vary without any variation in the agent’s beliefs of the fundamentals. For instance, if we consider the textbook RBC model, animal spirits would require that employment, investment, and consumer spending vary without any commensurate variation in the firms’ and the consumers’ beliefs of preferences and technologies.²⁹

Note now that the action of any given agent is always pinned down by her belief of her fundamental and her belief of the actions of others. It follows that two states can be associated with different equilibrium outcomes even when they contain the same beliefs about the fundamentals only if the two states are associated with different equilibrium beliefs about the endogenous outcomes themselves. In this sense, the notion of animal spirits is inherently tied to the notion of self-fulfilling beliefs. Finally, the following is true.

Corollary. *Suppose information is complete. The economy can feature animal spirits if and only if it admits multiple equilibria.*

This result highlights more generally the tendency in the profession to associate the notion of animal spirits with models that admit multiple equilibria, and to treat this notion as inconsistent with the class of unique-equilibrium DSGE models often used in the study of business cycles and monetary policy. In this context, a key point of our analysis will be to show how incomplete information blurs the distinction between multiple- and unique-equilibrium models in two complementary ways: first, by inducing equilibrium uniqueness in models that admit multiple equilibria under complete information; and second, by allowing animal spirits to obtain despite a unique equilibrium.³⁰

Remark. Recall that our definition of complete information requires that agents share the same information, not only with regard to the underlying fundamentals, but also with regard to *everything* else. This allows for public sunspots but rules out *imperfect* correlation devices as in Aumann’s (1974, 1987) work on correlated equilibrium. If, instead, one imposes common knowledge of fundamentals but allows private information about payoff-irrelevant variables, then one can devise examples in which (i) there is a unique correlated equilibrium and (ii) this equilibrium features animal spirits because agents respond to their private information about the underlying correlation devices in this equilibrium. To see this, consider a three-player variant of the “matching pennies” game in which there is no (Nash) equilibrium in pure strategies, but there is a correlated equilibrium in which the two players coordinate on playing a jointly mixed strategy against a mixed strategy of the third player. Although this example is too contrived to be of any practical relevance, it clarifies the following point: in general, incomplete information can help sustain “animal spirits”, not only by introducing higher-order uncertainty about the fundamentals, but also by sustaining strategic uncertainty even in settings that maintain common knowledge of fundamentals.

Consider now an overlapping-generations macroeconomic model in which the fundamentals are common knowledge and remain constant over time, but a different sunspot is realized in every period. Even if the

²⁹Our notion of animal spirits should therefore not be confused with the less-common notion adopted in Lorenzoni (2009) and Barksy and Sims (2012): in these papers, “animal spirits” are defined as the noise in a public signal of future TFP, that is, as a particular shock to beliefs of fundamentals.

³⁰An alternative approach, which will not be considered in this paper, is to formalize “animal spirits” as a deviation from individual rationality. See, e.g., Akerlof and Shiller (2010).

sunspot is publicly observed in the period that is realized, the model is akin to a game with imperfect correlation devices, because the generation of agents that acts in any given period does not see the sunspot upon which future generations can condition their choices. This example underscores that the concept of sunspot equilibria in *dynamic* macroeconomic models is closely related to the concept of correlated equilibria, which itself allows for incomplete information. As a consequence, the sunspot volatility that has been documented in prior work does not *strictly* require equilibrium multiplicity. For a more detailed discussion of these issues, we refer the reader to Aumann, Peck and Shell (1988) and Peck and Shell (1991); see also Jackson and Peck (1991) and Solom (2003) for two applications of these ideas.

4 Global Games: Theory

In this section, we restrict attention to a special case of our framework that imposes strong complementarity. We view this case as representative of a variety of multiple-equilibrium models that were developed in order to formalize the idea that macroeconomic outcomes are driven by coordination failures, self-fulfilling beliefs, and animal spirits.³¹ Of course, our framework is too stylized to capture either the rich micro-foundations or the intricate dynamics of some of these models. Nonetheless, the case studied contains two key features from the related literature: the role of coordination and the existence of multiple self-fulfilling equilibria under complete information.

The main theme of this section is to show how the introduction of incomplete information in such settings can induce a unique equilibrium, possibly one in which outcomes are pinned down merely by the underlying fundamentals. In addition, we establish the existence of a certain type of discontinuity in equilibrium outcomes as we move from the complete-information benchmark to a perturbation with arbitrarily small noise in the agent's information of the fundamentals. Combined, these results provide a sharp illustration of the potential sensitivity of the predictions of the theory to the kind of friction we study in this chapter. They also highlight that the driving force is the uncertainty faced by agents about one another's actions, rather than their uncertainty about the underlying fundamentals. With this in mind, we then proceed to describe the distinct role that private and public information play in shaping the level of strategic uncertainty.

The analysis in this section builds heavily on the influential contributions of Carlsson and Van Damme (1993a, 1993b) and Morris and Shin (1998, 2001, 2003). Carlsson and Van Damme's work highlighted the possible fragility of multiple equilibria by showing how small perturbations of the payoff and information structures can guarantee a unique rationalizable outcome within arbitrary two-by-two games (two players, two actions). This outcome was also shown to coincide with the one selected by the risk-dominance criterion proposed by Harsanyi and Selten (1988). Morris and Shin's work extended this type of result to models that were better suited for macroeconomics and finance, thus paving the way to a large applied literature; it also expanded the theoretical foundations.

Remark 1. The closely related contributions of Chamley (1999), Frankel and Pauzner (2000), and Burdzy, Frankel and Pauzner (2001) were made at roughly the same time as those of Morris and Shin. An important precedent to all these works is Rubinstein (1989), which first highlighted the fragility of coordination to perturbations of common knowledge.

Remark 2. The term "global game" was introduced by Carlsson and Van Damme (1993b) within the

³¹See, inter alia, Azariadis (1981), Benhabib and Farmer (1994, 1999), Cass and Shell (1983), Cooper and John (1988), Cooper (1999), Diamond and Dybvig (1983), Farmer (1996), Farmer and Woodford (1997), Guesnerie and Woodford (1992), Howitt and McAfee (1992), Murphy, Shleifer, and Vishny (1989), Matsuyama (1991, 1995), Obstfeld (1986, 1996), Shell (1977, 1989), and Woodford (1986, 1991).

context of two-by-two games that had the following key features: first, information was incomplete; second, types could be ordered in such a way that one action is dominant for sufficiently low types and another action is dominant for sufficiently high types. In the subsequent literature, the number of players and actions is often larger, and strategic complementarity is often imposed, yet appropriate versions of the aforementioned key features are maintained. For the purposes of our paper, we restrict attention to a particular class of games that are often referred to in the literature as “games of regime change”.

4.1 Set Up

Throughout this section, we impose that actions are binary, that there is no heterogeneity in fundamentals, and that payoffs take a particularly simple form.

Assumption 3 (Global Game). $D_k = \{0, 1\}$ and $D_\theta = \mathbb{R}$. Furthermore, θ_i is identical across i and is henceforth denoted simply by θ . Finally,

$$u_i = U(k_i, K, \theta) = \begin{cases} k_i b & \text{if } K \geq \theta \\ -k_i c & \text{if } K < \theta \end{cases}$$

where K is the average action (equivalently, the fraction of the population that chooses $k = 1$) and where b and c are known positive scalars.

The key feature of this specification is that it represents environments with strong strategic complementarity. As long as the agent faces no uncertainty about either K or θ , her best response is given by $k = G(K, \theta)$, where

$$G(K, \theta) \equiv \arg \max_{k \in \{0, 1\}} U(k, K, \theta) = \begin{cases} 1 & \text{if } K \geq \theta \\ 0 & \text{if } K < \theta \end{cases}$$

It is therefore as if the best response of an agent has infinite slope in a neighborhood of $K = \theta$.

Remark. The definition of strong complementarity that we introduced in Subsection 3.4 requires that G is continuous and that it has slope higher than one in a point where it crosses the 45° line. Here, G is discontinuous, so the earlier definition does not apply anymore, but the essence remains the same. Also, note the following slight abuse of notation: before, G was a function of \mathbf{B} , the entire cross-sectional distribution of first-order beliefs; now, it is a function of θ , because the latter happens to be sufficient statistic for \mathbf{B} under the specification we consider in this section.

4.2 Interpretation

To ease the transition to applications, and following Angeletos, Hellwig, and Pavan (2007), we interpret the aforementioned payoff specification as a “game of regime change”.

There are two possible regimes, the status quo and an alternative. Each agent can choose between an action that is favorable to the alternative regime and an action that is favorable to the status quo. We henceforth refer to these actions as, respectively, “attack” and “not attack”. We denote the regime outcome with $R \in \{0, 1\}$, where $R = 0$ represents the survival of the status quo and $R = 1$ represents its collapse. We similarly denote the action of an agent with $k_i \in \{0, 1\}$, where $k_i = 0$ represents “not attack” and $k_i = 1$ represents “attack”. Next, we normalize the payoff from not attacking to zero and let the payoff from attacking be $b > 0$ if the status quo is abandoned and $-c < 0$ otherwise. This means that it is individually

optimal to attack the status quo if and only if the latter is expected to collapse. Finally, we assume the status quo is abandoned ($R = 1$) if and only if

$$K \geq \theta,$$

which means that θ identifies the minimal size of attack that is needed for the status quo to be overturned. It then follows that the payoff of the agent can be expressed as in Assumption 3.

One may wish to interpret the above game as a model of revolutions and political change.³² Perhaps more interestingly for macroeconomics, the above game can be used to capture the role of coordination in the context of financial crises.

Consider, in particular, the literature on self-fulfilling currency attacks, such as Obstfeld (1996). In this context, the players can be interpreted as speculators choosing whether to attack a currency peg or not; θ represents the resources (e.g., reserves) that are available to the central bank for defending the peg, or more generally, the ability and the willingness of the policy maker to withstand a speculative attack; c represents the interest-rate differential between domestic and foreign assets, or other costs suffered by a speculator when short-selling the domestic currency; b represents the gains of an attacking speculator in the event of devaluation; and a “regime change” occurs when a sufficiently large mass of speculators attacks the currency, forcing the central bank to abandon the peg.

Similarly, in models of self-fulfilling bank runs, such as Diamond and Dybvig (1983), θ may represent the liquidity of a bank or of the financial system at large. A “regime change” occurs once a sufficiently large number of depositors decides to withdraw their deposits, forcing the bank to suspend its payments.

In models of self-fulfilling debt crises, such as Calvo (1988), θ may represent the long-run profitability of a firm or the growth potential of a country, and “regime change” occurs when the firm’s/country’s creditors decline to roll over its outstanding short-term debt because they fear that other creditors will refuse to roll over, forcing the firm/country into default.

Finally, regime change can also mean the “big push” in Murphy, Shleifer and Vishny (1989), that is, a situation where the incentives of an agent to switch from one technology to another, or to leave the village for the city, or to enter a particular industry, depend crucially on how many other agents are doing the same.

Remark. The aforementioned papers assume complete information. The more recent literature that incorporates incomplete information in such applications is discussed in the next section. Also, the applications we have in mind are intrinsically dynamic. One may thus question whether it is appropriate to model them as one-shot games. We will revisit this issue at the end of the next section.

4.3 Complete Information and Multiple Equilibria

Suppose for the moment that information is complete. Without any serious loss, suppose further that information is perfect, meaning that the true realization of θ is known to all agents in all states of Nature.

Let $\theta_L \equiv 0$ and $\theta_H \equiv 1$. For $\theta \leq \theta_L$, the fundamentals are so weak that the regime is doomed with certainty and the unique equilibrium has every agent attacking. For $\theta > \theta_H$, the fundamentals are so strong that the regime can survive an attack of any size and the unique equilibrium has every agent not attacking. For intermediate values, $\theta \in (\theta_L, \theta_H]$, the regime is sound but vulnerable to a sufficiently large attack and there are multiple equilibria sustained by self-fulfilling expectations. In one equilibrium, individuals expect everyone else to attack, so they then find it individually optimal to attack, the status quo is abandoned

³²See, e.g., Atkeson (2000) and Edmond (2013) for related interpretations.

and expectations are vindicated. In another, individuals expect no one else to attack, they thus find it individually optimal not to attack, the status quo is spared and expectations are again fulfilled.

Proposition 5 (Multiple Equilibria). *Suppose information is complete (θ is common knowledge). When $\theta \leq \theta_L$, the unique equilibrium outcome is $K = R = 1$. When $\theta > \theta_H$, the unique equilibrium outcome is $K = R = 0$. And when $\theta \in (\theta_L, \theta_H]$, both $K = R = 0$ and $K = R = 1$ are equilibrium outcomes.*

The interval $(\theta_L, \theta_H]$ thus represents the set of “critical fundamentals” for which multiple equilibria are possible under complete information. Each equilibrium is sustained by different self-fulfilling expectations about what other agents do.

Under the assumed payoff function, the two equilibria can be readily ranked: the equilibrium with $K = R = 1$ is Pareto superior to that with $K = R = 0$. One can thus associate the latter equilibrium with a “coordination failure”. More generally, which equilibrium represents a coordination failure varies depending on the context of interest. In the context of self-fulfilling currency attacks, for example, a coordination failure is the no-attack equilibrium from the perspective of foreign speculators, whereas it is the attack equilibrium from the perspective of domestic agents. Regardless of which perspective is adopted, however, the notion of coordination failure is associated with a particular selection across multiple equilibria.

In the next two subsections, we will show how allowing information to be incomplete can remove the indeterminacy of the equilibrium. In fact, we will show that this can be true even if the noise in the available information is arbitrarily small. While this implies that a realistic perturbation of standard models can potentially destroy the *traditional* view of coordination failures and animal spirits, we will argue that these notions can be meaningfully resurrected even in models with a unique equilibrium.

4.4 Incomplete Information and Unique Equilibrium

Following Morris and Shin (1998, 2001), we restrict the stochastic structure as follows. Nature draws θ from a uniform distribution over the entire real line.³³ Conditioning on θ , Nature draws a private signal for each agent i . This signal is given by

$$x_i = \theta + \sigma \epsilon_i,$$

where ϵ_i is an idiosyncratic noise term and $\sigma > 0$ parameterizes the level of noise. The noise is independent of θ and is drawn from a smooth distribution over the real line, with a strictly increasing c.d.f. denoted by Φ . These facts are common knowledge, but the realizations of θ and the signals are not; instead, the information set, ω_i , of agent i is simply the signal x_i . Without any loss, we will assume that an agent attacks whenever she is indifferent between attacking and not attacking.³⁴

Note that the scalar σ parameterizes how informed each agent is about the realization of θ . When σ is *exactly* zero, the model reduces to the complete-information benchmark. When, instead, σ is positive but small enough, every agent is nearly perfectly informed, and we have a seemingly tiny perturbation in the exogenous primitives of the environment. One may have expected that such a tiny change in the assumptions of the model would imply a tiny difference in its predictions. This turns out not to be the case: the predictions of the theory are discontinuous at $\sigma = 0$.

³³This specification of the common prior about θ is only for simplicity. For the results that follow, it suffices to consider any smooth prior over an interval that strictly contains the critical region and to let σ , the noise in the signal, be small enough. What the assumption of a uniform (or uninformative) prior does is to guarantee that uniqueness obtains for any σ , not just for σ small enough. The role of precise priors, or the related role of public information, is discussed later.

³⁴This assumption resolves an indeterminacy that obtains in a zero-probability event, namely when the private signal x takes the particular value x^* characterized in the proof of Proposition 6 below.

Proposition 6 (Morris-Shin). *There is a unique equilibrium for all $\sigma > 0$. In this equilibrium, the size of the attack K is monotonically decreasing in θ , and regime change occurs ($R = 1$) if and only if $\theta < \theta^*$, where $\theta^* = 1 - b/c \in (0, 1)$.*

This result highlights the sharp discontinuity mentioned above: a seemingly tiny perturbation in a model's assumptions regarding the agents' information implies a huge difference in the model's predictions. This discontinuity seems troubling as it calls into question the insights and policy lessons obtained in the large multiple-equilibria literature of the 80s and 90s.

We prove this result in the next subsection. We then proceed to elaborate on its robustness, its theoretical underpinnings, and its implications for applied work.

4.5 Proof of Equilibrium Uniqueness

The proof is based on the procedure of iterated deletion of dominated strategies. For any $\hat{x} \in [-\infty, +\infty]$, let $K_{\hat{x}}(\theta)$ denote the size of aggregate attack when every (or almost every) agent attacks if and only if $x \leq \hat{x}$. Next, define the function

$$V(x, \hat{x}) = \mathbb{E} [U(1, K_{\hat{x}}(\theta), \theta) - U(0, K_{\hat{x}}(\theta), \theta) \mid x].$$

This is the difference in utility between attacking and not attacking for an agent who has private information x and expects the other agents to attack if and only if their signals fall below \hat{x} .

Let us determine V . First, note that, when other agents follow a threshold strategy with threshold \hat{x} , the resulting size of the attack is given by the following:

$$K_{\hat{x}}(\theta) = \text{Prob}(x \leq \hat{x} | \theta) = \text{Prob}(\theta + \sigma \epsilon \leq \hat{x} | \theta) = \text{Prob}(\epsilon \leq \frac{1}{\sigma} (\hat{x} - \theta)) = \Phi\left(\frac{1}{\sigma} (\hat{x} - \theta)\right).$$

By implication, the regime collapses ($R = 1$) if and only if $\theta < \hat{\theta}$, where $\hat{\theta} = \hat{\theta}(\hat{x})$ is the unique solution to $K_{\hat{x}}(\hat{\theta}) = \hat{\theta}$, or equivalently the inverse of the following:

$$\hat{x} = \hat{\theta} + \sigma \Phi^{-1}(\hat{\theta}).$$

It follows that the agent with signal x attaches the following probability to the event of regime change:

$$\text{Prob}(R = 1 | x) = \text{Prob}(\theta \leq \hat{\theta} | x) = \text{Prob}(\theta - x \leq \hat{\theta} - x) = \text{Prob}\left(\epsilon \geq \frac{1}{\sigma} (x - \hat{\theta})\right) = 1 - \Phi\left(\frac{1}{\sigma} (x - \hat{\theta})\right).$$

We can express the payoff V as follows:

$$V(x, \hat{x}) = b - b\Phi\left(\frac{1}{\sigma} [x - \hat{\theta}(\hat{x})]\right) - c,$$

where $\hat{\theta} = \hat{\theta}(\hat{x})$ is the aforementioned unique solution to $K_{\hat{x}}(\hat{\theta}) = \hat{\theta}$.

Before we proceed, we wish to emphasize that the specific functional form of V is not essential. As it will become clear below, the result is driven by the monotonicity and continuity properties of the function V and of the related function h that we introduce below.

With this point in mind, note $\hat{\theta}$ is increasing in \hat{x} . It follows that $V(x, \hat{x})$ is increasing in \hat{x} : The more aggressive the other agents are, the higher the expected payoff from attacking. Moreover, $V(x, \hat{x})$ is decreasing in x : The higher the value of the private signal, the lower the expected payoff from attacking.

Next, note that $V(x, \hat{x})$ is continuous in x and satisfies $V(x, \hat{x}) \rightarrow b - c > 0$ as $x \rightarrow -\infty$ and $V(x, \hat{x}) \rightarrow -c < 0$ as $x \rightarrow +\infty$. We can thus define a function h such that $x = h(\hat{x})$ is the unique solution to $V(x, \hat{x}) = 0$. Because $V(x, \hat{x})$ is continuous in both arguments, decreasing in x , and increasing in \hat{x} , $h(\hat{x})$ is continuous and increasing in \hat{x} .

The function h defined above summarizes best responses within the set of monotone strategies: assuming that agents $j \neq i$ attack if and only if $x_j \leq \hat{x}$, agent i finds it optimal to attack if and only if $x_i \leq h(\hat{x})$. By the same token, the fixed points of h identify the set of monotone equilibria: if there is an equilibrium in which an agent attacks if and only if $x < x^*$, then x^* solves $x^* = h(x^*)$.

Finally, note that $x^* = h(x^*)$ if and only if $V(x^*, x^*) = 0$. Using the definition of V , and letting $\theta^* = \hat{\theta}(x^*)$, we have that $V(x^*, x^*) = 0$ if and only if $\theta^* = 1 - c/b$. It follows that there exists a *unique* threshold x^* such that $x^* = h(x^*)$. This proves that there exists a unique equilibrium in monotone strategies.

We next prove that there is no other equilibrium. In fact, we prove a stronger result: the only strategy that survives iterated deletion of dominated strategies is the monotone equilibrium strategy identified by the aforementioned threshold.

Construct a sequence $\{\underline{x}_j\}_{j=0}^\infty$ by $\underline{x}_0 = -\infty$ and $\underline{x}_j = h(\underline{x}_{j-1})$ for all $j \geq 1$. In particular, letting $\underline{\theta}_{j-1}$ be the solution to

$$\underline{x}_{j-1} = \underline{\theta}_{j-1} + \sigma\Phi^{-1}(\underline{\theta}_{j-1}), \quad (9)$$

we have

$$V(x, \underline{x}_{j-1}) = b - b\Phi\left(\frac{1}{\sigma}(x - \underline{\theta}_{j-1})\right) - c$$

and thus

$$\underline{x}_j = \underline{\theta}_{j-1} + \sigma\Phi^{-1}\left(1 - \frac{c}{b}\right). \quad (10)$$

Hence, $\underline{x}_0 = -\infty$, $\underline{\theta}_0 = 0$, $\underline{x}_1 = \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}\left(\frac{b-c}{b}\right)$, and so on. Clearly, the sequence $\{\underline{x}_j\}_{j=0}^\infty$ is increasing and bounded above by x^* . Hence, the sequence $\{\underline{x}_j\}_{j=0}^\infty$ converges to some \underline{x} . By continuity of h , the limit \underline{x} must be a fixed point of h . But we already proved that h has a unique fixed point. Hence, $\underline{x} = x^*$.

Next, construct a sequence $\{\bar{x}_j\}_{j=0}^\infty$ by $\bar{x}_0 = +\infty$ and $\bar{x}_j = h(\bar{x}_{j-1})$ for all $j \geq 1$. Note that this sequence is decreasing and bounded below by x^* . Hence, the sequence $\{\bar{x}_j\}_{j=0}^\infty$ converges to some \bar{x} . By continuity of h , \bar{x} must be a fixed point of h . But we already proved that h has a unique fixed point. Hence, $\bar{x} = x^*$.

What is the meaning of these sequences?

Consider \underline{x}_1 . If nobody else attacks, the agent finds it optimal to attack if and only if $x \leq \underline{x}_1$. By complementarity, if some people attack, the agent finds it optimal to attack *at least* for $x \leq \underline{x}_1$. That is, for $x \leq \underline{x}_1$, attacking is *dominant*. Next, consider \underline{x}_2 . When other agents attack if and only if it is dominant for them to do so, that is, if and only if $x \leq \underline{x}_1$, then it is optimal to attack if and only if $x \leq \underline{x}_2$. By complementarity, if other agents attack at least for $x \leq \underline{x}_1$, then it is optimal to attack *at least* for $x \leq \underline{x}_2$. That is, attacking becomes dominant for $x \leq \underline{x}_2$ after the second round of deletion of dominated strategies.

More generally, for any $j \geq 1$, we have that attacking becomes dominant $x \leq \underline{x}_j$ after the j round of deletion of dominated strategies. Hence, $\{\underline{x}_j\}_{j=0}^\infty$ represents iterated deletion of dominated strategies “from below”. Symmetrically, $\{\bar{x}_j\}_{j=0}^\infty$ represents iterated deletion of dominated strategies “from above”.

To recap, the only strategies that survive j rounds of iterated deletion of dominated strategies are functions k such that $k(x) = 1$ for all $x \leq \underline{x}_j$ and $k(x) = 0$ for $x > \bar{x}_j$. The value of $k(x)$ is still “free” at the j -th round only for $x \in (\underline{x}_j, \bar{x}_j)$. But we already proved that both \underline{x}_j and \bar{x}_j converge to x^* as $j \rightarrow \infty$. Hence, in the limit, the only strategy that survives is the function k such that $k(x) = 1$ for $x \leq x^*$ and $k(x) = 0$ for $x > x^*$.

We have thus proved that there exists a unique rationalizable strategy, and hence a fortiori also a unique equilibrium. In this equilibrium, an agent attacks if and only if $x \leq x^*$, and the status quo collapses if and only if $\theta \leq \theta^*$, where the pair (x^*, θ^*) is the unique fixed point to the iteration defined in (9) and (10). From this last fact, it is then immediate that $\theta^* = \frac{b-c}{b}$.

Remark. The property that the best-response function h is increasing and admits a unique fixed point means, in effect, that the incomplete-information game under consideration features weak complementarity. But recall that the complete-information counterpart features strong complementarity. It follows that the introduction of incomplete information has effectively transformed the game from one of strong complementarity to one of weak complementarity, which in turn helps explain the uniqueness result. See Vives (2005), Vives and Van Zandt (2007), and Methevet (2010) for a further discussion of this point, and for complementary analyses of how the uniqueness result can be understood via either super-modular methods (Vives, 2005, Vives and Van Zandt, 2007) or a contraction mapping argument (Methevet, 2010).

4.6 The Role of Public Information

In the preceding analysis, we assumed that the common prior about θ was uninformative (i.e., an improper uniform over entire real line). If the prior is not uniform, uniqueness obtains for σ small enough, but may not obtain otherwise. A similar property applies if we introduce a public signal: to the extent that there is sufficient public information, multiplicity survives the introduction of dispersed private information.³⁵

To illustrate this point consider the following modification of the information structure. The private signal is now given by

$$x_i = \theta + \sigma_\epsilon \epsilon_i,$$

where ϵ_i is drawn from a standardized Normal distribution (with mean 0 and variance 1). Similarly, the public signal is given by

$$z = \theta + \sigma_\zeta \zeta,$$

where ζ is an aggregate noise term that is also drawn from a standardized Normal distribution. The scalars σ_ϵ and σ_ζ parameterize the level of noise in the two signals.

One can establish the following:

Proposition 7 (Public Info). *The equilibrium is unique if and only if $\sigma_\epsilon \leq \sqrt{2\pi}\sigma_\zeta^2$.*

The distinct nature of the two types of information can be further illustrated by comparing the limit as $\sigma_\epsilon \rightarrow 0$ (for fixed $\sigma_\zeta > 0$) to that as $\sigma_\zeta \rightarrow 0$ (for fixed $\sigma_\epsilon > 0$). In either limit, the beliefs of every agent about θ converge in probability to the true θ . Both limits therefore approach “perfect knowledge” with regard to fundamentals, and are indistinguishable from the perspective of how informed the agents are about the fundamentals. Yet, the two limits look very different from the perspective of how well the agents can predict one another’s actions. Because of this, they also make very different predictions about economic behavior.

Proposition 8 (Limits). *(i) In the limit as $\sigma_\epsilon \rightarrow 0$ for given $\sigma_\zeta > 0$, the probability of regime change converges to 1 for all $\theta < \theta^*$ and to zero for all $\theta > \theta^*$, where $\theta^* = 1 - c/b$.*

(ii) Pick an arbitrary compact subset of the critical region, $A \subset (\theta_L, \theta_H]$. In the limit as $\sigma_\zeta \rightarrow 0$ for given $\sigma_\epsilon > 0$, there is an equilibrium in which the probability of regime change converges to 1 for all $\theta \in A$, as well as an equilibrium in which the probability of regime change converges to 0 for all $\theta \in A$.

³⁵In terms of strategic uncertainty, there is no difference between a public signal and the common prior: when studying the hierarchy of beliefs and/or solving for the equilibrium strategies, one can always re-cast a public signal as part of the common prior. The distinction between the prior and a public signal is therefore useful only for applied purposes.

Part (i) recasts Proposition 7 as a particular limit in which private information is infinitely precise *relative* to public information. Part (ii), by contrast, recasts the complete-information benchmark as the limit in which it is the relative precision of public information that is infinite. In combination, these two results therefore underscore the distinct role played by the two types of information in shaping the equilibrium beliefs that agents form about one another’s actions and in determining their ability to coordinate.

Remark. Recall the function h from the proof of Proposition 6 in the previous subsection; that function described best responses within the set of monotone strategies. In that proof, we saw that h admitted a unique fixed point. We then commented that this property meant that the introduction of private information had transformed the game from one of strong complementarity to one of weak complementary. The role of public information is the exact opposite: when public information is sufficiently precise relative to private information, the function h admits multiple fixed points. That is, sufficiently precise public information brings back strong complementarity.³⁶

4.7 Intuition and Take-home Lessons

A deeper understanding of the precise logic behind the preceding results requires a review of important advances in game theory, an endeavor which is beyond the scope of this paper. We refer the interested reader to Rubinstein (1989) for an early seminal contribution that highlighted the fragility of coordination to perturbations of common knowledge; to Monderer and Samet (1989, 1996) and Kajii and Morris (1997a, 1997b) for what it means to have “approximate” common knowledge and the related robustness of equilibria to incomplete information; to Morris, Postlewaite, and Shin (1995) for the contagious effects of higher-order uncertainty; to Morris and Shin (1997, 2003), Frankel, Morris and Pauzner (2003), and Mathevet and Steiner (2013) for extensions of the global-games uniqueness result to richer settings than the simple one we have studied in this section; to Mathevet (2010) for a variant proof based on a contraction mapping argument; to Vives (2005) and Vives and Van Zandt (2007) for related techniques from supermodular games; to Weinstein and Yildiz (2007a) for a powerful result that we discuss briefly in the next subsection; and to Morris, Shin, and Yildiz (2016) for the importance of “rank beliefs” (the probability the players assign to their signal being higher than that of their opponents) and for the common belief foundations of global games. For our purposes, it suffices to emphasize the following key intuitions.

First, note that equilibrium imposes that agents know one another’s *strategies*, that is, they know the mappings from their information sets (or Harsanyi types) to their actions. If we assume that all agents share the same information, then this imposes that all agents face no uncertainty about their actions. The absence of this kind of strategic uncertainty is conducive to multiple equilibria: it is “easy” to coordinate on one of many equilibrium actions when the agents are confident that other agents will do the same. But once information is incomplete, the agents may face uncertainty about one another’s actions, and this type of uncertainty can hinder coordination. It follows that the determinacy of the equilibrium hinges on the level of strategic uncertainty: the higher the level of strategic uncertainty, the harder to sustain multiple equilibria.

Next, note that the level of strategic uncertainty is not necessarily tied to the level of fundamental uncertainty: the uncertainty an agent faces about the beliefs and actions of other agents has to do more with the heterogeneity of the information and the associated higher-order uncertainty, and less with the overall level of noise in the observation of the fundamentals. In fact, when private information becomes more

³⁶These points can be inferred from the proof of Proposition 7 in the Appendix, noting that the function G in that proof is a transformation of the function h from the space of strategies (as indexed by the threshold x^*) to the space of regime outcomes (as indexed by θ^*).

precise, the uncertainty that an agent i faces about the fundamentals necessarily decreases, yet it is possible that her uncertainty about beliefs and actions of any other agent j increases. This is because an increase in the precision of private information implies that the beliefs and actions of agent j become more anchored to her own private information, which is itself unknown to agent i . This anchoring effect in turn explains why private information can exacerbate higher-order uncertainty and thereby hinder coordination.

This intuition can be formalized as follows. First, note that, when information is complete, the equilibrium belief of any agent about K is a direct measure of the realized value of K . That is, agents are perfectly informed about the size of the attack in equilibrium, irrespective of the equilibrium selected. Next, note that, in the diametrically opposite scenario where an agent is completely agnostic about the size of the attack, her belief about K is uniform over the $[0, 1]$ interval. Finally, consider what happens under incomplete information. Let σ_ϵ be small enough so that the equilibrium is unique and consider the “marginal” agent, that is, the type who is indifferent between attacking and not attacking in equilibrium. Morris and Shin (2003) show that the following is true.

Proposition 9. *In the limit as private information becomes infinitely precise (namely, $\sigma_\epsilon \rightarrow 0$ for given $\sigma_\zeta > 0$), the marginal agent’s belief about K converges to a uniform distribution over $[0, 1]$. That is, the marginal agent learns perfectly the fundamental in the limit, and yet she remains completely uninformed about the actions of others.*

The sharpness of this last result hinges on focusing on the beliefs of the marginal agent, or of agents whose distance from the marginal agent vanishes as $\sigma_\epsilon \rightarrow 0$, as opposed to agents sufficiently far from the threshold x^* . Nevertheless, the combination of this result with the preceding observations helps explain why private information contributes to equilibrium uniqueness, whereas public information is conducive to multiplicity.

To summarize, in settings with coordination motives, information plays a dual role: it shapes the beliefs of each agent, not only about the exogenous payoff-relevant fundamentals, but also about the endogenous actions of the others. In the preceding result, it is the second channel that matters most. Private and public information are similar vis-a-vis the first function, but are distinct vis-a-vis the second function.

We close this section by discussing the nature of private and public signals.

For certain applied purposes, one may seek a rather literal interpretation of these signals. For instance, the private signal x_i may correspond to the proprietary information of a hedge fund, the “local” knowledge that a bank may have about the entire financial network, or the information a firm or consumer extracts about the aggregate economy from its own transactions. Similarly, the public signal z may be a proxy for financial news in the media, announcement or choices made by policy makers, or market signals such as prices. We discuss some of these interpretations, and the additional lessons they can lead to, in Section 5.

A literal interpretation of the signals is therefore useful to the extent that the researcher can envision empirical counterparts to them, and potentially even measure them. It is also useful if one wishes to study questions relating either to the collection and aggregation of information or to the welfare effects of the signals disseminated by markets, the media, and policy makers—issues that we address latter on.

That said, the mapping from the theory to the real world may not be as easy as suggested by the sharp dichotomy between private and public information employed above. For instance, suppose that the noise in the private signals happens to be the sum of two components, one idiosyncratic and one aggregate: $x_i = \theta + u + \epsilon_i$, where u is the aggregate noise and ϵ_i is the idiosyncratic noise. Suppose further that we introduce a public signal of the aggregate noise u , as opposed to a public signal of the fundamental θ . Then, this public signal will only increase the reliance of individual decisions on the private signals, which in turn may contribute to higher strategic uncertainty.

The above example may appear to be esoteric, but once one starts thinking about applications, the distinction between “fundamentals” and “noise” can get fussy. For instance, in the context of asset markets, should one think of shocks to discount factors as “noise” or as “fundamentals”? This question that can be meaningfully addressed only within the context of a fully-specified micro-founded model.

Finally, a literal interpretation of the signals is too narrow for some of the issues we have in mind. In the real world, and especially in the context of macroeconomic phenomena, it seems hard to measure, or even comprehend, all the ways in which economic agents collect and exchange information about either their idiosyncratic circumstances or the overall state of the economy. More generally, it is unclear, at least to us, why the most useful explanations of real-world phenomena are those based on models that rule out any friction either in how agents form beliefs about the relevant economic outcomes or in how they coordinate their behavior—which is what workhorse macroeconomic models typically do. From this perspective, the signal structures we employ throughout this chapter are ultimately modeling devices that permit the researcher to “open up the black box” of how agents form expectations about endogenous economic outcomes and how they coordinate their behavior. In our view, this basic point is central to understanding the precise applied value, and the testable implications, of incorporating incomplete information in macroeconomics.³⁷

4.8 Extensions and Additional Lessons

The preceding analysis has focused on a rather narrow class of binary-action games that can be described as games of regime change. Morris and Shin (2003) extend the global-games uniqueness result to a broader class of binary-action games, allowing for a more flexible form of payoff interdependence. Frankel, Morris, and Pauzner (2003) provide a generalization of the global-games uniqueness result to settings where the action belongs to an arbitrary, finite set, while maintaining strategic complementarity (appropriately defined in terms of supermodularity). Guimaraes and Morris (2007) consider an application with a continuous action (a portfolio choice). Goldstein and Pauzner (2005) consider an example that relaxes the assumption of strategic complementarity, in exchange for a weaker single-crossing condition³⁸ and the assumption that both the prior about the fundamentals and the distribution of the idiosyncratic noise are uniform.

Morris and Shin (2003) also establish the result in Proposition 9, namely that the marginal agent is agnostic about the size of the attack in the limit as private information becomes infinitely precise—a property to which they refer as “Laplacian belief”. Whether this property is realistic or not is open for debate. Morris and Shin (2003) argue that this is a “plausible” restriction the theorist can impose in order to sharpen the predictions he can deliver. By contrast, Atkeson (2000), Angeletos and Werning (2006), Angeletos, Hellwig, and Pavan (2006, 2007), and others have argued that, for many applications of interest, the endogeneity of the available information may lead to situations where this property is violated: market signals such as prices, and past economic outcomes, may endogenously limit the uncertainty that agents, including the marginal agent, face about one another’s actions. More on this in the next section.

The aforementioned works, as well as the applications considered in the next section, limit attention to specific payoff and information structures. Weinstein and Yildiz (2007a) aim for greater generality and obtain a revealing result: for any given game and any given rationalizable action a of any given type in the

³⁷Bergemann and Morris (2013) corroborate this point in the related class of games studied in Section 7, by showing that the equilibrium outcomes that can be sustained by any Gaussian information structure can be replicated by those that are sustained under the assumed private and public signals. For two alternative approaches within the global-games literature, see Cornand and Heinemann (2009) and Izmalkov and Yildiz (2010). The first paper studies a model with multiple group-specific signals; the second uses a heterogeneous-prior specification of the belief hierarchy.

³⁸This single-crossing condition is the following: $u(K, x)$ crosses zero once, going from negative to positive, as K varies from 0 to 1, where $u(K, x)$ is the net payoff from attacking when the agent receives signal x and knows K .

Universal Type Space, there is a nearby game in which the given action a becomes the unique rationalizable action for a perturbed version of the original type.³⁹ A corollary of this is that equilibrium multiplicity can be viewed as a knife-edge situation, which can always be eliminated by considering an arbitrarily small perturbation of the payoff and information structures.

At first sight, Weinstein and Yildiz’s result may appear to deliver a deathblow to multiple-equilibria models and their applications: multiplicity is degenerate! However, the actual meaning of the result is different. The result applies to *every* rationalizable action of the original game. It follows that, whenever we have a model with multiple equilibria, we can find small perturbations—indeed, open sets of such perturbations—that select *any* of the original equilibrium outcomes as the unique rationalizable outcome.

For applied purposes, this means the following. In models that admit multiple equilibria, the inability of theorist to tell which equilibrium is selected should be interpreted as follows: the theorist is unable to reach sufficiently sharp predictions on the basis of the assumptions he has made in his model (that is, on the basis of his prior knowledge about the environment). While Weinstein and Yildiz (2007a) result suggests that this type of multiplicity is fragile, the inability to make sufficiently sharp prediction about the equilibrium remains: but it now concerns lack of knowledge about the details of the information structure rather than lack of knowledge about which equilibrium is played.⁴⁰

What lesson can be drawn? The global-games methodology is not a panacea for getting rid of multiple equilibria, or for giving policy makers the satisfaction of sharp policy advice. Instead, the applied value of the global-games uniqueness result rests on elucidating the mechanics of coordination and on highlighting the importance of information, and communication, for the questions of interest. We try to give concrete examples of this value later in Section 5.

5 Global Games: Applications

In this section, we review recent applications of global games to macroeconomics and finance. Many of these applications can be classified into two broad categories. The first one treats the information structure as exogenous, makes assumptions that guarantee equilibrium uniqueness, and uses this to shed new light on phenomena such as bank runs, financial crises, and business cycles and on related policy questions. The second one endogenizes the information structure in manners that seem empirically relevant, studies how this can bring back multiple equilibria, and sheds further light on applied questions.

Subsections 5.1-5.6 review the first set of applications. Subsections 5.7-5.9 turn to the second set. Subsection 5.10 mentions additional research that does not necessarily fit in the previous two categories.

5.1 Currency Crises and Bank Runs

In the paper that popularized the global-games approach, Morris and Shin (1998) introduced incomplete information in the currency-attack model of Obstfeld (1996). Although multiple equilibria exist when the fundamentals (such as reserves) are common knowledge, a unique equilibrium is obtained by adding a small amount of idiosyncratic noise in the speculators’ information about the fundamentals. The size of the attack and the devaluation outcome no longer depend on sunspots but may exhibit a strong non-linearity

³⁹What “nearby” and “perturbed” mean can be delicate; see Weinstein and Yildiz (2007a) for details.

⁴⁰To put it differently, the common thread is that the variation in the observable outcomes of the model is not panned by either the underlying payoff relevant fundamentals or the agent’s beliefs about them. In the original, multiple-equilibrium, model, the residual variation is attributed to a pure sunspot. In the perturbed, unique-equilibrium, model, the residual variation is attributed to higher-order uncertainty and rich “noise” in the information structure.

(or near-discontinuity) with respect to the fundamentals. A large attack can thus be triggered by small changes in the underlying fundamentals, helping reconcile the uniqueness of the equilibrium with the fact that speculative attacks are abrupt and often happen without a significant change in measured fundamentals. Moreover, policy analysis can now be conducted without the difficulties associated with multiple equilibria and arbitrary, ad hoc, selections. The paper shows how a marginal increase in a “tax” on capital outflows or in domestic interest rates can curtail a speculative attack, a policy conclusion that would not have been obtained in the presence of multiple equilibria.⁴¹

Turning to the context of bank runs, Goldstein and Pauzner (2005) “globalize” Diamond and Dybvig (1983) and proceed to study the implications for the design of the optimal deposit contract. To understand the contribution of this paper, note first that the characterization of optimal contract in Diamond and Dybvig (1983) is based on the assumption that a self-fulfilling run never takes place. This entails a particular equilibrium selection, which is possible in the original, complete-information, setting of Diamond and Dybvig, but not once incomplete information is introduced. Instead, the unique equilibrium now makes self-fulfilling runs *inevitable* for sufficiently low bank fundamentals. Goldstein and Pauzner’s key contribution is therefore to study how this inevitability reshapes the design of the optimal design contract. The paper shows that the optimal contract under incomplete information penalizes early withdrawals relative to the complete-information benchmark. This entails a welfare cost in terms of providing less insurance to “impatient” consumers; but this cost is now justified by the reduction in the probability and the size of a bank run.

The broader lesson emanating from this paper, as well as from Morris and Shin (1998), is the following: just as the optimal contract characterized in Diamond and Dybvig is dependent on the complete-information assumption, many of the policy recommendations found in the literature are dependent on shutting down frictions in coordination and are often driven by arbitrary equilibrium selections. By contrast, the global-games methodology offers a way to study policy without these caveats.

5.2 Big Players and Lenders of Last Resort

Corsetti, Dasgupta, Morris, and Shin (2004) sheds new light on the role of large players in speculative markets. The model is similar to the one in Morris and Shin (1998), except for one key difference: a large speculator (“Soros”) coexists with a continuum of small investors.

When a small speculator chooses whether to attack, she takes the probability of devaluation as exogenous to her own choice. When, instead, the large player chooses whether to attack, she takes into account that she has a non-trivial impact on the probability of devaluation. Other things equal, this effect makes the large speculator more aggressive than any small speculator.⁴² Perhaps more interestingly, the presence of a large player facilitates more coordination among the small players, even if her actions are not observable: because it is known that the large player is more aggressive, each small player finds it optimal to be more aggressive for any realization of her own signal, each small player expects the same from other small players, and so on. What is more, as the small players get more aggressive, the large player finds it optimal to become even more aggressive, and so on. It follows that the introduction of a large player in market can have a disproportionately strong effect on equilibrium outcomes, even if her actions are not observable. Allowing the small players to observe the actions of the large player further amplifies this property.

⁴¹The original paper had a mistake in the characterization of the comparative statics of the equilibrium with respect to the aforementioned tax. See Heinemann (2000) for the correction.

⁴²In effect, the large player is like a pool of small players that have managed to act in a coordinated fashion, thus overcoming the friction faced by other small players.

Davila (2012) studies the role of a different type of large players in a different context: the role of large banks in the context of the subprime crisis. The starting point of the analysis is Farhi and Tirole (2012), who identified the following collective moral hazard problem: because the government is more likely ex post to bailout an individual bank when the entire banking system is in trouble, each bank is ex ante more willing to expose itself to an aggregate risk when it expects other banks to do the same. In Farhi and Tirole (2012), this collective moral hazard problem is modeled as a complete information coordination game that ultimately features multiple equilibria: one with low “systemic risk” and another with high.⁴³ By contrast, Davila (2012) uses a global-game adaptation to obtain a unique equilibrium. The presence of large banks is shown to intensify the severity of the collective moral hazard problem, increase economy-wide leverage, and make the crisis occur for a large set of fundamentals. The intriguing policy implication is then that systemic risk can be reduced by breaking up the large banks into multiple small ones.

Rochet and Vives (2004) considers a bank-run model in which investors may refuse to renew their credit on the interbank market during the liquidity crisis. In their “globalized” model the equilibrium is unique and the probability of a crisis is linked to the fundamentals. In this unique equilibrium, there is an intermediate range for bank’s fundamentals in which in the absence of the government intervention the bank is solvent but may still fail if too large a proportion of investors withdraw their money. In other words, there exists a potential for coordination failure. The authors proceed to study the interaction between ex ante regulation of solvency and liquidity ratios and ex post provision of emergency liquidity assistance. They show that liquidity and solvency regulation can solve the coordination problem, but typically the cost is too high. These prudential measures must be complemented with emergency discount-window loans.

Corsetti, Guimaraes, and Roubini (2006) build a similar model as the one in Rochet and Vives (2004), featuring liquidity crises caused by the interaction of bad fundamentals and self-fulfilling runs. The authors focus on the “catalytic” effects of the liquidity provision from the official lender. Drawing on the results of Corsetti et al. (2004) about the role of large players, the paper models the official creditor (IMF) as a large player in the world economy. Even if IMF by its own does not have enough resources to close the large financing gaps generated by speculative runs, the range of economic fundamentals over which liquidity crises do not happen is enlarged through the “catalytic” effect. As a result, liquidity provision during crisis is justified even if it exacerbates debtor’s moral hazard problem. A similar point is made in an independent, contemporaneous contribution by Morris and Shin (2006).

The policy lessons that emerge from these papers rely on the uniqueness of the equilibrium. They therefore could not have been obtained in the complete-information versions of the relevant models. However, these lessons also hinge on abstracting from the possibility that policy actions, when made intentionally, convey valuable information about the willingness and ability of the policy maker to defend the status quo (a currency peg, a particular bank, or the entire banking system). As we discuss later on, this possibility is explored in Angeletos, Hellwig and Pavan (2006) and is shown to lead to a potentially very different lessons.

5.3 Debt runs, Default Premia, and Credit Ratings

Morris and Shin (2004a) globalize the coordination problem among creditors of a distressed borrower in an otherwise conventional debt-run model a la Calvo (1988). There are three periods. The borrower—who is interpreted as a firm but could also be a country—owns a long-term investment project and owes a certain

⁴³Note that this rests on the government lacking commitment: if the government could commit ex ante not to bailout the banking system ex post, it could alleviate the problem. See Ennis and Keister (2009) for a related point on how the ex post optimal policy response to a bank run distorts ex ante incentives, raising the occurrence of self-fulfilling runs.

level of short-term debt. The return of the project materializes in period 3. The debt expires in period 2 and must be rolled over to period 3, or else the borrower “dies” and the project is destroyed. There is a large pool of lenders, each of whom is too small relative to the size of the borrower. It follows that the borrower can survive if and only if enough lenders coordinate on rolling over their debts in period 2. By the same token, a coordination failure is possible: even when the borrower’s fundamentals are sound, a creditor’s fear of other creditors’ premature foreclosure may lead to pre-emptive action, thus undermining the project.

Incomplete information guarantees that this coordination failure materializes when the profitability of the project is positive but not high enough: there exists a unique equilibrium in which default takes place in period 2, not only when the project is unprofitable, but also when profitability is positive but not high enough to preclude a self-fulfilling debt run. In period 1, market price debt anticipating that outcomes in period 2 will be determined in the aforementioned fashion. It follows that thanks to the uniqueness of the equilibrium default outcome, there is now also a unique equilibrium price—a property that is not shared by the complete-information version of the model. Furthermore, the equilibrium default risk can be now decomposed in two components: one that measures insolvency (the probability that the project is unprofitable), and another that measures roll-over risk (the probability that a coordination failure materializes).

In a follow-up paper, Zabai (2014) endogenizes the behavior of borrowers and shows how optimal borrowing internalizes the effect that the size of debt has on the probability of coordination failure and thereby on default premia. Together, these papers therefore answer two important questions: how the risk of coordination failure is priced; and how policy makers can influence it.

In another related paper, Holden, Natvik, and Vigier (2014) study the role of credit-rating agencies. By influencing the behavior of creditors, the ratings published by such agencies affect the probability of default, which in turn affects the ratings of other agencies. Through this feedback effect, the presence of credit-rating agencies can exacerbate volatility and reduce welfare.

Finally, He and Xiong (2012) study the dynamics of debt runs in a model that obtains uniqueness with the help of the techniques found in Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2002). More specifically, the paper characterizes how the occurrence and the dynamics of debt runs depend on the level fundamental volatility, the presence of credit lines, and the debt maturity structure. Interestingly, they find that commonly used measures that are meant to ease runs, such as temporarily keeping the firm alive and increasing debt maturity, can actually backfire and exacerbate runs—yet another example of how policy conclusions drawn from complete-information models can go wrong.

5.4 “Liquidity Black Holes” and “Market Freezes”

Morris and Shin (2004b) study how the incompleteness of information can help explain sudden drops in asset prices and liquidity in otherwise smoothly functioning markets. A set of sophisticated, risk-neutral traders (who can be thought as “hedge funds”) interact with a representative unsophisticated, risk-averse, trader. This means that the value of the asset, and hence also its price, is higher when it is held by the former rather than by the latter. Importantly, the sophisticated traders have short horizons and face privately known loss limits. When the price of the asset is sufficiently above the loss limits of the sophisticated traders, their trades are strategic substitutes: the less the other traders buy of the asset, the cheaper (and more attractive) it becomes for an individual trader to buy. However, once the price of the asset falls sufficiently close to, or below, the loss limits, sales of the risky asset become strategic complements: the more the others sell, the more one has to sell “in distress”. This opens the door to a “liquidity black hole” analogous to the run equilibrium in Diamond and Dybvig (1983). A global-game specification is then used to select a unique

equilibrium and to obtain a unique trigger point in the fundamentals, below which the liquidity black hole comes into existence. The paper proceeds to show how this helps generate a sharp V-shaped pattern in prices around the time of the liquidity black hole and connect to historical experiences.

Bebchuk and Goldstein (2011) study a debt crisis model in which a credit-market freeze—a situation in which banks abstain from lending to good firms—may arise because of the banks’ self-fulfilling expectations that other banks will not lend. The authors “globalize” the model to study the effectiveness of alternative government interventions, including interest rate cuts, infusions of capital into banks, direct lending to operating firms, and provisions of government capital or guarantees to encourage privately managed lending.

5.5 Safe Assets

He et al. (2016) study the pricing of sovereign debt when debt valuations depend on both insolvency and roll-over risk. Their model features two countries that issue sovereign bonds on the international capital markets. An investor’s valuation of a sovereign bond depends not only on the country’s fundamental, but also on the number of other investors who purchase that bond. For a country’s bonds to be safe, the number of investors who invest in the bond must exceed a threshold. As a result, investor actions feature a strategic complementarity similar to the one found in Section 4. What is new is that the model also features strategic substitutability when the number of investors who invest in the bonds exceeds the threshold required to roll over debts: above the threshold, more demand for the bond drives up the bond price, leading to lower returns. Recall that a similar feature was present in Goldstein and Pauzner (2005).

He et al. (2016) “globalize” the environment by removing common knowledge of the relative fundamental of the two countries. Under certain conditions, they prove the existence of a unique, monotone equilibrium. In this equilibrium, a sovereign’s debt is more likely to be safe if its fundamentals are strong relative to the other country, but not necessarily strong on an absolute basis. Investors coordinate to invest in the country with relatively better fundamentals, and thus relative valuation determines which country’s bonds have less rollover risk and, in this sense, more safety. This prediction may help explain why the valuation of US debt may increase despite a deterioration in its fiscal conditions: the fact that other countries are in even worse shape means that US roll-over risk has gotten smaller.

An additional insight is that a higher level of debt may, perhaps paradoxically, carry a higher price. This happens when the global demand for assets is strong. When this is true, investors are attracted to the country with the highest level of debt, because this helps to satisfy their high demand for assets. But this means that the country with the higher level of debt faces lower roll-over risk, which in turn means that the bond of this country is safer and commands a higher price. By contrast, when the global demand for safe assets is low, investors coordinate on the country with the smallest debt size.

Finally, the authors use their results to evaluate the recent proposal for “Eurobonds”, that is, of a common bond for a club of countries. When a significant amount of such bonds is issued, all countries within the club can benefit by reducing the overall roll-over risk. It follows that Eurobonds can be beneficial for, say, both Germany and Greece. For the same reason, however, they can also upset the dominance of US debt as the international benchmark for safety.

5.6 Business Cycles and the Great Recession

A few papers sought to use the global-games methodology to shed light on business-cycle phenomena. The first attempt was Chamley (1999), which studies the dynamics of regime switches in a relatively abstract

investment game. A related attempt was made by Frankel and Burdzy (2005), using the technique reviewed in Subsection 6.1.

More recently, Schaal and Taschereau-Dumouchel (2015) develop a global-games adaptation of the RBC framework and use it to study the Great Recession. We find this paper to be particularly interesting because it succeeds to merge the more abstract global-games literature with the more canonical business-cycle paradigm. What is more, it takes up the challenge of confronting the theory with the data. Below, we briefly explain the key ingredients of the paper.

Firms have the option to pay a fixed cost in order to increase their productivities and reduce their variable costs. This option introduces a non-convexity in technology. The combination of this non-convexity with monopolistic competition and aggregate demand externalities introduces a coordination problem that is akin to the one we have formalized in our binary-action regime-change setting: a firm is willing to pay the aforementioned fixed cost if and only if it expects aggregate demand to be sufficiently high, which in turn is true if and only if enough other firms also pay the aforementioned fixed cost.

Under complete information, the model admits multiple equilibria. Schaal and Taschereau-Dumouchel consider a perturbation that maintains common knowledge of past outcomes (including the capital stock) but removes common knowledge of the current TFP level. This permits them to obtain, not only a unique equilibrium, but also tractable dynamics. Although the equilibrium is now unique, there are two locally stable steady states, leading to very different dynamics from those in the standard RBC model. In particular, a large transitory shock may push the economy into a quasi-permanent recession, helping explain the slow recovery and other salient features of the Great Recession. Importantly, these outcomes are the product of a coordination failure. But since this coordination failure is not the symptom of arbitrary equilibrium selection, a meaningful policy analysis is also possible. In particular, the framework justifies certain types of fiscal stimuli when the economy is transitioning between steady states.

A complementary policy point is made in Guimaraes, Machado, and Ribeiro (2015). Similarly to the above paper, this paper studies a micro-founded model with a non-convexity in production; but it abstracts from capital and productivity shocks. This precludes the type of quantitative assessment that Schaal and Taschereau-Dumouchel (2015) are after; but it also facilitates a sharper analysis of the equilibrium effects of fiscal policy. The key contribution of the paper is then to formalize and characterize the expectational/coordination channel of fiscal policy—a channel that is customarily invoked in real-world policy debates but does not have a sufficiently meaningful counterpart in standard macroeconomic models.

5.7 Prices and Other Endogenous Public Signals

The preceding applications have treated the information structure as exogenous and have used it primarily as a device to select a unique equilibrium. In fact, the literature on applied global games has often focused attention on the limit in which public information is minimal and strategic uncertainty is maximal, in the sense defined earlier on. During times of crises, however, it is unlikely that agents are in the dark. On the contrary, they are closely monitoring the activity of others.

More broadly, agents may have access to various public signals of the underlying fundamentals and/or the actions of others. Think of the price of a peso forward in the context of currency crises, or of media reports about long ATM queues in the context of bank runs. Alternatively, think of speculators observing the actions of a policy maker who is anxious to preempt a speculative attack. How does this kind of information affect the determinacy of the equilibrium and the observables of the model? In the remainder of this section, we briefly review a few papers that have sought to answer this question.

We start with Angeletos and Werning (2006). This paper augments the basic coordination game introduced in Subsection 4.1 with two kinds of endogenous public signals. The first public signal is a price of an asset whose dividend is correlated either with the fundamental, θ , or the size of attack, K . The second signal is a direct signal of the size of the attack. One can think of a country's stock market or the forward foreign exchange market as examples of the first type, and the ATM queues or the size of protests as examples of the second. In what follows, we focus on the second type, because it can easily be accommodated by our framework; the first type leads to similar lessons.

The payoff structure is the same as in Subsection 4.1, and so is the specification of private information. In particular, agent i 's private signal is once again given by $x_i = \theta + \sigma_\epsilon \epsilon_i$, where the noise $\epsilon_i \sim N(0, 1)$ is i.i.d. across agents and $\sigma_\epsilon > 0$. The only change is in the specification of the public signal. The latter is now given by a noisy signal of the size of the attack:

$$z = S(K, \zeta),$$

where S is a monotone function, K is the size of attack, and $\zeta \sim N(0, 1)$ is noise independent of θ and $\{\epsilon_i\}$.

Equilibrium can now be defined as follows.

Definition 14. *An equilibrium consists of an endogenous signal, $z = Z(\theta, \zeta)$, an individual attack strategy, $k(x, z)$, and an aggregate attack, $K(\theta, \zeta)$, such that*

$$k(x, z) \in \arg \max_{k \in \{0, 1\}} \mathbb{E}[U(k, K(\theta, z), \theta) \mid x, z] \quad \forall(x, z) \quad (11)$$

$$K(\theta, z) = \mathbb{E}[k(x, z) \mid \theta, z] \quad \forall(\theta, z) \quad (12)$$

$$z = S(K(\theta, z), \zeta) \quad \forall(\theta, \zeta, z) \quad (13)$$

Equation (11) requires individual choices to be optimal given all available information, including the one contained in the realized signal z . Equation (12) aggregates. Equation (13) imposes that the signal is generated by the joint equilibrium behavior of the agents. The only novelty in the above definition relative to Definition 3 is the fixed-point relation between the way that agents react to their available information and the way some of that information (namely the signal z) is generated in the first place. This fixed-point relation is standard in noisy rational-expectations models, including those used in the context of financial markets (e.g., Grossman and Stiglitz, 1980, Hassan and Mertens, 2014a).

Note that, in equilibrium, K is correlated with θ . It follows that, *in equilibrium*, z becomes also a signal of θ . This fact permits to map the equilibrium analysis of the present model to that of Section 4. What is different from before is that the precision of the information contained in z about θ is now endogenous to the actions of the agents.

To preserve Normality of the endogenous information structure, Angeletos and Werning (2006) impose the following functional form for the S function:

$$S(K, \zeta) = \Phi^{-1}(K) + \sigma_\zeta \zeta,$$

where Φ^{-1} is the inverse of the c.d.f. of the standardized Normal distribution and σ_ζ is a scalar parameterizing the *exogenous* noise in the signal. This functional form, which was first proposed by Dasgupta (2007), guarantees the existence of equilibria in which the observation of z is equivalent to the observation of a

Gaussian signal of the following form:

$$\tilde{z} = \theta + \tilde{\sigma}_\zeta \zeta,$$

for some scalar $\tilde{\sigma}_\zeta$, which is itself determined as part of the equilibrium. Taking $\tilde{\sigma}_\zeta$ as given, the equilibrium strategies and the regime outcome can be characterized in exactly the same fashion as in the case where information was exogenous (modulo replacing the scalar σ_ζ that appears in the analysis of Section 4 with the scalar $\tilde{\sigma}_\zeta$ introduced herein). What remains then is to characterize the equilibrium value of $\tilde{\sigma}_\zeta$.

To complete this task, we must solve the fixed-point relation between the strategies and the information. Intuitively, the sensitivity of $k(x, z)$ to x determines the sensitivity of $K(\theta, z)$ to θ , which in turn determines the precision of the signal \tilde{z} with respect to θ , which in turn determines the sensitivity of $k(x, z)$ to x , and so on. Solving this fixed-point problem ultimately yields the following relation between the endogenous $\tilde{\sigma}_\zeta$ and the exogenous parameters of the model:

$$\tilde{\sigma}_\zeta = \sigma_\epsilon \sigma_\zeta.$$

In other words, the precision of the endogenous public signal is proportional to the precision of the exogenous private signal, and inversely proportional to the exogenous noise in the observation of the size of attack.

The intuition is simple. When private signals are more precise, individual decisions are more sensitive to private information. As a result, the equilibrium aggregate attack reacts relatively more to the underlying fundamental θ than to the common noise ζ , thus also conveying more precise information about θ .

This elementary observation has important implications for the determinacy of equilibria and the volatility of outcomes. As the precision of private information increases, the *endogenous* increase in the precision of the available public information permits agents to better forecast one another's actions and thereby makes it easier to coordinate. Consequently, uniqueness need not obtain as a perturbation away from the perfect-information benchmark. Indeed, in the present model, multiplicity obtains when noise is small.

Proposition 10. *In the model described above, an equilibrium always exists. There are multiple equilibria if either of the exogenous noises is small, namely if $\sigma_\zeta^2 \sigma_\epsilon < 1/\sqrt{2\pi}$.*

This result does not upset the theoretical insights developed earlier: it remains true that uniqueness obtains if and only if strategic uncertainty is sufficiently high, which in turn is true if and only if private information overwhelms public information. It nevertheless weakens the empirical relevance of the global-games uniqueness result: in many applications of interest, there may be good reasons to expect that the information structure may be such that multiplicity obtains or, more generally, that outcomes remain highly sensitive to shocks that are largely or totally unrelated to the underlying fundamentals.

Angeletos and Werning (2006) reinforce this message by showing the following result: even if we restrict attention to the region of the parameter space in which the equilibrium is unique (namely, the region of σ_ζ and σ_ϵ such that $\sigma_\zeta^2 \sigma_\epsilon > 1/\sqrt{2\pi}$), a local reduction in either σ_ϵ or σ_ζ leads to an increase in the sensitivity of the equilibrium regime outcome to ζ relative to θ . In short:

Corollary. *Less exogenous noise can contribute to more non-fundamental volatility, not only by opening the door to multiple equilibria and sunspot volatility, but also by raising the sensitivity of the equilibrium outcome to higher-order uncertainty even when the equilibrium outcome is unique.*

This result, a variant of which we will encounter in Section 7, contrasts with results in single-agent decision-theoretic settings: in such settings, less noise, whether in the form of more information or in

the form of lower cognitive friction, contributes to fewer “mistakes” and less non-fundamental volatility in outcomes. A similar property holds in noisy rational-expectations models of financial markets as those studied in Grossman and Stiglitz (1980) and Hassan and Mertens (2014b): in these settings, less exogenous noise, either in the form of a lower volatility in the demand of noise traders or in the form of more precise information in the hands of sophisticated traders, typically leads to more informative asset prices and less non-fundamental volatility in investment. What differentiates these settings from the models we study in this chapter is the absence of strategic complementarity. Once such complementarity is present, the “paradoxical” result reviewed above emerges, whether we consider models that are prone to multiple equilibria (as those studied presently) or models that are prone to equilibrium uniqueness (as those studied in Section 7).

We conclude by mentioning a few additional papers that also concern the informational role of markets in global-games settings: Hellwig, Mukherji, and Tsyvinski (2006), Tarashev (2007), and Ozdenoren and Yuan (2009). Each of these papers focuses on additional institutional details and additional different specific insights. They nevertheless share the same core message with Angeletos and Werning (2006): there are good reasons to think that the price mechanism, or other endogenous signals of the actions of others, may “tame” the level of strategic uncertainty relative to what is assumed in Morris and Shin (1998, 2001), thus contributing to non-fundamental volatility. Somewhat complementary is also the paper by Iachan and Nenov (2015), which studies how the quality of private information can contribute to instability, even if one holds constant the precision of public information.

5.8 Policy Interventions

All the applications reviewed earlier on contain important policy insights. They nevertheless abstract from the possibility that active policy interventions may reveal valuable information about the state of the economy and/or the objectives and intentions of the policy maker. Importantly, because such policy interventions are highly visible, they could affect the level of strategic uncertainty, opening the door to effects that were assumed away in the aforementioned applications.

Motivated by this elementary observation, Angeletos, Hellwig and Pavan (2006) and Angeletos and Pavan (2013) study global-game settings that add a policy maker, who is informed about the underlying fundamentals and can take an action that may influence the size of attack. Think of a central bank trying to preempt a speculative attack by raising domestic interest rates or imposing a tax on capital outflows; or think of a dictator deploying the police or the army to quell a protest. Such actions are likely to reveal that the status quo is neither too strong nor too weak: if the status quo were sufficiently strong, the policy maker would not bother to intervene, for there would be little risk of an attack; and if the status quo were sufficiently weak, it would be pointless to pay the cost of such interventions, for the status quo would be doomed in any case.

Angeletos, Hellwig and Pavan (2006) thus show that such policy interventions may induce high common belief that the fundamental θ is within the critical region, even if they do not reveal precise information about the exact value of θ . They then proceed to show how this can open the door to multiple equilibria, not only in the regime outcome, but also in the policy action itself. In a certain sense, the policy maker therefore may find himself trapped in a situation in which he has to conform to market expectations that are beyond his control, as opposed to being able to shape economic outcomes with seemingly powerful tools such as taxes and interest rates.

While this result raises an important caveat to policy applications of global-games results, it does not necessarily eliminate the predictive power of the global-games approach. Angeletos and Pavan (2013) show

that, even though the signaling role of policy interventions brings back multiple equilibria, the resulting multiplicity is “much smaller” than the one that is present in the complete-information benchmark. We refer the reader to that paper for the explanation of what this “much smaller” means. The bottom line, however, is that a number of concrete policy predictions emerge that are robust across all equilibria and that could not have been reached without the introduction of incomplete information.

Edmond (2013) studies a different kind of policy interventions: interventions that can manipulate the information that is available to the agents, without however directly signaling information to them. The particular application concerns a dictator manipulating the media or the internet in an attempt to preempt a revolution. The action of the dictator is not directly observable. Instead, it is confounded with the signal the agents observe about the underlying fundamentals. This helps preserve equilibrium uniqueness. Yet, the nature of the optimal manipulation hinges on the effect it has on strategic uncertainty, as opposed to merely the information the agents extract about the fundamentals.

Combined, these papers indicate more generally how policy can interact with frictions in coordination, in manners that are absent from workhorse macroeconomic models. Exploring this possibility within a business-cycle context seems an interesting direction for future research.

5.9 Dynamics

The framework studied in Section 4 is static, and so are the applications we have studied so far. The preceding analysis has thus abstracted from the possibility that agents take multiple shots against the status quo and that their beliefs about their abilities to induce regime change vary over time. Yet, these two possibilities are important from both an applied and a theoretical perspective. First, crises are intrinsically dynamic phenomena. In the context of currency crises, for example, speculators can attack a currency again and again until they force devaluation; and their expectations about the ability of the central bank to defend the currency in the present may naturally depend on whether the bank has successfully defended it in the past. Second, learning in a dynamic setting may critically affect the level of strategic uncertainty (i.e., uncertainty about one another’s actions) and thereby the dynamics of coordination and the determinacy of equilibria.

Motivated by these considerations, Angeletos, Hellwig and Pavan (2007) study a repeated version of the regime-change game considered in the previous section. Whenever the regime survives an attack, it leads to common knowledge (or at least a strongly correlated belief) that the regime is “not too weak”, or else it would not have survived the attack. This kind of endogenous shift in beliefs opens again the door to multiple equilibria, but it also leads to a number of distinct predictions. First, fundamentals may predict the eventual regime outcome but not the timing or the number of attacks. Second, equilibrium dynamics tend to alternate between phases of *tranquility*, where no attack is possible, and phases of *distress*, where a large attack can occur. Finally, attacks take the form of relatively infrequent and acute changes in the level of economic activity.

Constain (2007) provides a complementary result, showing how the public observation of past actions can lead to herding as well as to multiplicity. By contrast, Heidhues and Melissas (2006), Dasgupta (2007), Dasgupta, Steiner, and Stewart (2012), and Steiner and Kovac (2013) study settings where learning is private, helping preserve the uniqueness of the equilibrium; these papers then proceed to study questions that have to do with the structure of the dynamics, such as the role of irreversible actions, the option value of waiting-to-see, and the synchronization of actions.

Other dynamic global-games applications that maintain the uniqueness of the equilibrium include Chamley (1999), Giannitsarou and Toxvaerd (2006), Guimaraes (2006), He and Xiong (2012), Huang (2013),

Mathevet and Steiner (2013), Steiner (2008), and Toxvaerd (2008). Finally, Chassang (2010) studies a dynamic exit game in which multiplicity survives but is much “smaller” than the one that obtains under complete information, thus helping deliver sharper predictions—a message echoed in Angeletos, Hellwig and Pavan (2007) and Angeletos and Pavan (2013).

Finally, He and Manela (2014) consider a variant of Abreu and Brunnermeier (2003) that derives interesting dynamics from the interaction of the same two features as in Angeletos, Hellwig and Pavan (2007): the public signal generated by the fact that the regime has survived past attacks, and the arrival of new private information over time. An important difference, however, is the arrival of private information is endogenous: the paper studies the incentives of acquiring information and show how this endogenously leads to the possibility of a new run after an unsuccessful one.

5.10 Other Related Research

We conclude our review of the global-games literature by briefly mentioning three additional lines of research, one theoretical and two empirical.

The first line endogenizes the acquisition of information in global games. Important contributions include Hellwig and Veldkamp (2009) and Yang (2015). The first paper restricts the agents’ information choice in a fixed set of private and public Gaussian signals. The second paper considers an attention-allocation problem as in Sims (2003): the agents are free to choose any signal they wish, including a non-Gaussian signal, at a cost that is proportional to the entropy reduction attained by the chosen signal structure. Both papers reach a similar conclusion: because strategic complementarity raises the incentive to observe the *same* information, there is yet another reason why multiplicity may survive. See, however, Denti (2016) and Morris and Yang (2016) for important qualifications to this conclusion and for even more flexible specifications of the information-acquisition technology.

The second line focuses on testing the empirical implications of global games. Prati and Sbracia (2002) use data on consensus forecasts for six Asian countries to measure the cross-sectional heterogeneity of beliefs in the context of speculative attack; they then proceed to document that both the level effect of this heterogeneity and its interaction with measures of fundamentals is consistent with the predictions of global games. Bannier (2006) and Tillmann (2004) provide complementary reduced-form evidence on the role of informational disparities and belief heterogeneity. Danielsson and Peñaranda (2011) offer a structural estimation of a global game within the context of carry trades in the yen-dollar market. Chen, Goldstein, and Jiang (2010) use mutual fund data to provide evidence that strategic complementarities among investors contribute to financial-market fragility, and interpret this evidence under the lenses of a global game. Nagar and Yu (2014) provide evidence on the coordinating role of public information, interpreting accounting data as a form of public signal. Finally, Nimark and Pitschner (2015) provide complementary evidence on the endogenous correlation of the information disseminated by newspapers.

The third line is also interested in testing the predictions of global games, but does so within the context of laboratory experiments. See Cabrales, Nagel, and Armenter (2007), Cornand (2006), Duffy and Ochs (2012), Heinemann, Nagel, and Ockenfels (2004, 2009), and Shurchkov (2013).

6 Coordination and Synchronization

So far we have focused on settings where incomplete information impedes coordination within any given period. We now shift attention to another aspect: the agents’ ability to synchronize their choices.

More specifically, we review two important contributions. The first one, which is by Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2001), studies the role of adding a Calvo-like friction in a dynamic game of regime change. By preventing synchronous choice, this friction is shown to help select a unique equilibrium. Importantly, this is true even when the friction is vanishingly small. The second contribution, which is by Abreu and Brunnermeier (2003), is also concerned with asynchronous choice. However, the primitive friction is now the asynchronous awareness of a change in the environment, as opposed to a Calvo-like friction in the agents' ability to act. Such asynchronous awareness is shown to cause a significant delay in the response of the economy, even if all agents become aware pretty fast. We discuss how these contributions shed light on the subtle relation between synchronization and coordination; how they are connected to the global-games literature as well as to one another; and how they illustrate, once again, importance of strategic uncertainty and the potentially fragility of standard macroeconomic models.

6.1 The Calvo friction, Asynchronous Choice, and Coordination

In this subsection we discuss the contribution of Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2001). Unlike the “canonical” example in the global-games literature, these papers study dynamic games of regime change in which the fundamental is perfectly observed in every period. They then proceed to show how equilibrium uniqueness can be induced by the combination of persistent shocks to the fundamentals together with a Calvo-like friction in the ability of the agents to reset their actions.

At first glance, this result appears to be distinct from the one in the global-games literature, and possibly of higher relevance. A closer look, however, reveals a tight connection. We elaborate by reviewing Frankel and Pauzner (2000). Burdzy, Frankel and Pauzner (2001) contains an extension that may be useful for applications but is not needed for our purposes.

The model in Frankel and Pauzner (2000) is a stochastic version of the two-sector model in Matsuyama (1991), which itself a dynamic variant of the “big-push” model of Murphy, Shleifer, and Vishny (1989). Time is continuous, indexed by $t \in [0, \infty)$, and agents are infinitely lived. At each point of time, each agent can be in one of two idiosyncratic states, state 0 or state 1. Let $k_{i,t} \in \{0, 1\}$ denote the state at which agent i is at time t . The interpretation is analogous to the interpretation of the actions in our static framework. In the aforementioned papers, for example, $k_{i,t} = 0$ corresponds to living in the village and working in agriculture, whereas $k_{i,t} = 1$ corresponds to living in the city and working in manufacturing. Furthermore, there is again strategic complementarity: it is more profitable to “attack”, that is, to live in the city and work in manufacturing, when enough other agents do the same. What is different is that each agent can not instantaneously switch between “attacking” and “not attacking”, which in turn explains why $k_{i,t}$ is a state variable rather than a control variable.

Let $K_t \in [0, 1]$ denote the fraction of the population that rests at state 1 (city/manufacture) at time t . The life-time utility of agent i is given by

$$\int_0^{\infty} e^{-\rho t} U(k_{i,t}, K_t, \theta_t) dt,$$

where $\rho > 0$ is the discount rate, θ_t is an exogenous fundamental that affects the relative return to being in the village/agriculture, and $U(k_{i,t}, K_t, \theta_t)$ is the flow utility. The latter is given as in Assumption 3: the flow payoff from $k_{i,t} = 0$ is normalized to zero, whereas the flow payoff from $k_{i,t} = 1$ is given by $b - c > 0$ whenever $K_t \geq \theta_t$ and by $-c < 0$ otherwise.⁴⁴

⁴⁴Frankel and Pauzner (2000) allow the net payoff from $k_{i,t} = 1$ to be a more general function $A(\theta_t, K_t)$ that is increasing in

We introduce stochasticity in the underlying fundamental by assuming that θ_t follows a Brownian motion with zero drift and volatility parameter $\sigma > 0$: $d\theta_t = \sigma dv_t$. For future reference, we note that the case of no aggregate shocks can be approximated by taking the limit as $\sigma \rightarrow 0$. We also let $\theta_L \equiv 0$ and $\theta_H \equiv 1$; these points identify the boundaries of dominance regions in the static benchmark (the one-shot game we studied in the preceding subsections).

Unlike the rest of this chapter, we rule out private information: the realized values of both the exogenous aggregate state θ_t and the endogenous aggregate state K_t are public information at t . Instead, the key friction is idiosyncratic inertia of the same type as in Calvo (1983): each agent's option to switch follows an idiosyncratic Poisson process with arrival rate $1/\lambda$. The scalar λ therefore parameterizes the level of inertia in individual adjustment: the higher λ is, the more likely that an agent is "stuck" in her current location for a while. For future reference, note that instantaneous adjustment is approximated by the limit $\lambda \rightarrow 0$.

For comparison purposes, let us momentarily consider the knife-edge case in which $\sigma = 0$ and $\lambda = 0$. By this we do not mean the double limit of the aforementioned model as $\sigma \rightarrow 0$ and $\lambda \rightarrow 0$. Rather, we mean a variant that *literally* shuts down both the shocks to fundamentals and the Calvo-like friction. The following is then trivially true.⁴⁵

Proposition 11. *Suppose $\lambda = \sigma = 0$ and let $\theta \in (\theta_L, \theta_H]$. For any t , and regardless of the history of past play, there is an equilibrium in which all agents "attack" (locate in the city/work in manufacture), as well as as well as an equilibrium in which all such agents choose the opposite.*

This is essentially a multi-period version of the multiplicity result we encountered in Subsection 4.3. The only novelty, in terms of observable implications, is that the multiplicity can take the form of rich fluctuations over time: there can be equilibria in which K_t is constant at either 0 or 1 forever, as well as equilibria in which K_t jumps up and down at arbitrary points of time. In short, sunspot volatility manifests in the time series of aggregate economic activity.

Consider next the case in which $\sigma > 0$ but $\lambda = 0$. That is, allow for shocks to fundamentals, but continue to rule out idiosyncratic inertia. Clearly, the multiplicity result survives. The only minor difference is that the possibility of sunspot volatility may now vary over time, depending on whether the current value of the fundamental happens to lie within the critical region $(\theta_L, \theta_H]$.

Consider next the diametrically opposite case in which $\lambda > 0$ but $\sigma = 0$. That is, introduce the Calvo friction, but rule out shocks to fundamentals. Now, K_t becomes a state variable, which moves slowly over time. This expands the regions over which it is dominant for an agent to chose either location. In particular, if K_t is close enough to 1, then it is dominant for an agent to choose city/manufacture in a neighborhood of θ_t above 0, simply because that agent knows that many other agents will be "stuck" in the same location for a while; and symmetrically, when K_t is close enough to 0, it is dominant to choose village/agriculture in a neighborhood of θ_t below 1. Nonetheless, at least insofar as λ is not too high, multiplicity survives.

This case is illustrated in Figure 2. One can readily show that there exist increasing mappings $\kappa_L : [\theta_L, \theta_H] \rightarrow [0, 1]$ and $\kappa_H : [\theta_L, \theta_H] \rightarrow [0, 1]$, illustrated by the solid lines in the figure, such that the following properties are true. Consider the agents who have the option to act at time t . If $K_t < \kappa_L(\theta_t)$, it is dominant to attack (i.e., choose village/agriculture). If $K_t > \kappa_H(\theta_t)$, it is dominant not to attack (i.e., choose city/manufacture). Finally, if $\kappa_L(\theta_t) < K_t < \kappa_H(\theta_t)$, there is an equilibrium in which all current agents

K_t and decreasing in θ_t . The restriction we adopt is only for expositional simplicity. Also, we have changed the notation to make it consistent with the one we use in the rest of the chapter. In particular, note that θ_t equals $-z_t$ in their notation: our fundamental is the opposite of theirs. This explains why the functions κ_L, κ_H , and κ^* below have the opposite monotonicity that the corresponding functions in the original paper.

⁴⁵The solution concept is Perfect Bayesian Equilibrium.




Figure 2: Multiple Equilibria with $\lambda > 0$ and $\sigma = 0$.

choose village/agriculture, as well as an equilibrium in which all current agents choose city/manufacture. These equilibria are sustained by the expectation that all *future* agents will do the same as the current agents. In short, the mappings κ_L and κ_H identify the boundaries of, respectively, the lower and the upper dominance regions. For any $\lambda > 0$, we have that $0 < \kappa_L(\theta)$ and $\kappa_H(\theta) < 1$ for all $\theta \in (\theta_L, \theta_H)$. This fact reflects the expansion of the dominance regions mentioned above. But as λ converges to 0, these dominance regions converge uniformly to the dominance regions of the aforementioned variant model in which $\lambda = 0$. The latter are identified by the dotted vertical lines in the figure.

Let us now consider the case of interest, which is the case in which both λ and σ are strictly positive (although potentially arbitrarily small). The following is then true.

Proposition 12 (Frankel-Pauzner). *Suppose $\lambda > 0$ and $\sigma > 0$. There exists a unique, and increasing, mapping $\kappa^* : \mathbb{R} \rightarrow [0, 1]$ such that the following is true in equilibrium: all the agents who have the option to act at time t choose not to attack (i.e., locate in the city/work in manufacture) whenever the current aggregate state satisfies $K_t > \kappa^*(\theta_t)$, whereas they make the opposite choice whenever $K_t < \kappa^*(\theta_t)$.*

This result is illustrated in Figure 3. The two dashed lines give the boundaries of the dominance regions; the solid line gives the mapping κ^* , which of course lies strictly in between the two dominance regions. Different realizations of the Brownian motion may induce the aggregate state (θ_t, K_t) to travel anywhere in the space $\mathbb{R} \times [0, 1]$. However, the path of K_t is uniquely pinned down by the path of θ_t . In particular, K_t has a positive drift whenever (θ_t, K_t) is on the left of κ^* , and a negative drift on the right of κ^* . In the limit as shocks vanish, K_t therefore converges either to 0 or to 1, depending on initial conditions.

The above result applies even when λ and σ are vanishingly small. This means that the multiplicity result in Proposition 11 is fragile: a small perturbation selects a unique equilibrium.

Furthermore, if we take the limit as the level of idiosyncratic inertia becomes arbitrarily small, we obtain the following sharp characterization of the unique equilibrium.

Proposition 13. *Let $\theta^* \equiv 1 - \frac{c}{b}$, fix any $\sigma > 0$, and let $\lambda \rightarrow 0$. The following properties hold in this limit:*

- (i) *The mapping κ^* becomes vertical at θ^* : $\kappa^*(\theta) \rightarrow 0 \forall \theta < \theta^*$ and $\kappa^*(\theta) \rightarrow 1 \forall \theta > \theta^*$.*
- (ii) *The distribution of K_t conditional on θ_t converges to a Dirac at 1 for all $\theta_t < \theta^*$ and to a Dirac at 0 for all $\theta_t > \theta^*$.*

This result is qualitatively the same as the Morris-Shin limit result in Proposition 8: in essence, every agent is “attacking” whenever the current fundamental is below θ^* , and nobody is attacking whenever the




Figure 3: Unique Equilibrium with $\lambda > 0$ and $\sigma > 0$.

current fundamental is above θ^* , regardless of the history of either past fundamentals or past play. What is more, the threshold θ^* that shows up here is exactly the same as the one that shows up in Proposition 6.

At first glance, this coincidence is surprising, even mysterious. The perturbation considered here is of a different nature than the one considered before: instead of removing common knowledge of fundamentals and introducing private information and strategic uncertainty, we only had to assume (a bit of) idiosyncratic inertia at the individual level along with aggregate shocks to fundamentals. How could it be that two very different perturbations lead to essentially the same selection?

The answer is that there is a deep connection: the combination of idiosyncratic inertia and aggregate shocks transforms the dynamic model under consideration to a game that has a similar mathematical structure as the static global game we have studied in the preceding sections. This fact is most evident if one compares the proof of the uniqueness result in Frankel and Pauzner (2000) with the proof of the uniqueness result that Frankel, Morris and Pauzner (2005) provide for a class of static global games with multiple actions: the proofs in the two papers are essentially the same!

To economize space on space and effort, we do not dwell into the formal details. Instead, we provide a sketch of the proof of the uniqueness result in Proposition 12 and relate it to the proof of Proposition 6.

Let us start with the following preliminary observations. At any given point of time, both the exogenous fundamental θ_t and the endogenous state variable K_t are publicly known to all agents. Yet, the mass of agents that have the option to act at any point of time is zero. We can therefore reinterpret the model as one in which a single agent acts at each moment. Because this agent will be “stuck” at her currently chosen action for a while, she must forecast the actions of the agents that will act in the near future. But these agents must themselves forecast the actions of other agents that will move further away in the future. It follows that we can understand the behavior of an agent as a function of her hierarchy of beliefs about the actions of other agents that will have the option to act in either the near or the far future.

These forward-looking higher-order beliefs would not have been relevant if λ were identically zero: it is essential that choices are asynchronous (due to the Calvo-like friction). This, however, does not necessarily mean that the role of higher-order belief vanishes as λ converges to zero from above. As long as both λ and σ are positive, the agents moving at any given point face higher-order uncertainty; and because we are in an environment with strong complementarities, higher-order uncertainty can leave its mark on equilibrium outcomes even if both λ and σ become vanishingly small.

To elaborate, let us first make the following elementary observation. No matter how small σ is, as long as σ is strictly positive, the fundamentals can drift away over time to either very low or very high levels. As

a result, the aggregate state variable (θ_t, K_t) can eventually enter either one of the two dominance regions, even if it is currently far away from both of them.

Consider a moment in time such that the aggregate state is outside the dominance regions but arbitrarily close to one of them, say, the upper one. An agent who acts at that moment does not have a dominant action. However, this agent expects the aggregate state to drift into the upper dominance region with probability close to $1/2$, at which point it becomes dominant for future agents to attack (i.e., choose village/agriculture). Given this, it becomes optimal—indeed, iteratively dominant—for the current agent to attack as well.

Repeating this argument gives a process of iterated deletion of dominated strategies from “above”, which is similar to the one we encountered in the proof of Proposition 6. The only difference is that the process is now over the space of (θ, K) , whereas before it was over the space of the private signal x_i . This underscores that the current value of the aggregate state plays a similar role as private information plays in global games.

One can construct a symmetric process of iterated deletion from “below” and can then attempt to prove the uniqueness result by showing that both processes converge to the same point, namely to the strategy described by the mapping κ^* . As it turns out, it is more convenient to consider a modification of the one of the two processes; see Frankel and Pauzner (2000). The substance, however, remains the same: as with the global games, the key mechanism is the impact of higher-order uncertainty, manifested as a contagion effect from the dominance regions.

A natural question is then the following. In global games, uniqueness obtains when private information is sufficiently precise, whereas multiplicity survives if public information is sufficiently precise. Is there an analogue in the present context?

Frankel and Pauzner do not provide an answer to this question. We nevertheless conjecture that multiplicity survives if the fundamental is mean-reverting towards a point in the critical region and, importantly, the mean reversion is sufficiently strong relative to λ and σ . Our intuition is that strong mean reversion in the fundamental has a similar coordinating effect as public information: it helps the current and future agents reach a high common belief about the state.⁴⁶

To recap, the results of Frankel and Pauzner (2000) underscore that, once combined with strategic complementarity and changing fundamentals, asynchronous choice in the form suggested by Calvo (1983) can be interpreted as a friction in coordination.

At some level, this idea is not entirely surprising. In the context of New-Keynesian models, the Calvo friction—and staggered pricing more generally—captures, not only price-stickiness at the individual level, but also the asynchronous price adjustment. It is this asynchronicity that one can reasonably interpret as a friction in coordination, even if one knows nothing about either the global-games literature or the related work of Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2001).

Yet, the findings of these papers indicate that the workings of the Calvo friction can be more subtle than those understood so far. They also raise the following questions. If the Calvo friction is meant to capture frictions in coordination, does it make sense to impose it *only* on the price-setting behavior of firms, or should we impose it on real choices as well? Also, if the Calvo friction is a proxy for a more primitive friction, what is the latter? Is it menu costs? Or is it an informational friction?

These questions add to the motivation for a later topic of our chapter, namely the study of business-cycle models in which nominal rigidity originates from incomplete information as opposed to Calvo-like asynchronicity. For now, we conclude with few remarks.

⁴⁶Burdzy, Frankel and Pauzner (2001) allow for mean-reversion, but take the limit as $\lambda \rightarrow 0$. This seems akin to allowing for public information but taking the limit as private information becomes infinitely precise.

Remark 1. In the model described above, equilibrium outcomes are discontinuous in λ at $\lambda = 0$, reflecting the non-vanishing impact of beliefs of arbitrarily high order. This kind of discontinuity, which is similar to the one encountered in global games, rests on the underlying strategic complementarity being of the strong form. If, instead, complementarity is of the weak form, as in the beauty contests and the business-cycle applications we study in Sections 7 and 8, the impact of beliefs of order h vanishes as $h \rightarrow \infty$ at a rate high enough that the aforementioned discontinuity does not emerge.⁴⁷

Remark 2. The methods developed in Frankel and Pauzner (2000) and Burdzy, Frankel and Pauzner (2001) should appeal to macroeconomists, because equilibrium uniqueness is obtained with the help of two familiar modeling ingredients: aggregate shocks to fundamentals and a Calvo-like friction. It is then somewhat surprising that this approach has not attracted more attention in applied research. Exemptions include Burdzy and Frankel (2005) on business cycles, Guimaraes (2006) on currency attacks, and He and Xiong (2012) on debt runs.

Remark 3. The Calvo friction helps captures asynchronous choice in a brute way. A possible alternative is that asynchronous choice is the byproduct of an information friction. Dasgupta, Steiner, and Stewart (2012) offer an analysis that has such a flavor: they study a dynamic global game in which there is value to synchronize actions and show how lack of common knowledge can impede synchronization. An important earlier contribution in the same vein is Abreu and Brunnermeier (2003), which we review next: in that paper, asynchronous choice is the byproduct of asynchronous awareness.

6.2 Asynchronous Awareness and Coordination (with Application to Bubbles)

We now turn attention to Abreu and Brunnermeier (2003), an important contribution that highlights how asynchronous awareness breaks common knowledge and how this in turn may cause significant delay in the response of equilibrium outcomes to changes in the environment. The particular application considered in Abreu and Brunnermeier (2003) regards asset bubbles: asynchronous awareness is shown to delay the burst of a bubble. The lessons that emerge, however, extend well beyond the particular application.

The model is unlike the ones we have studied so far, because it admits a unique equilibrium regardless of whether information is complete or incomplete. It nevertheless shares a similar contagion mechanism working through higher-order beliefs. The paper’s key contribution is to show how this kind of mechanism can help explain why a bubble may persist long after every trader has recognized that the price is unsustainable.

More specifically, the paper studies a continuous-time model in which a unit mass of sophisticated traders, each with relatively shallow pockets, decide whether to ride or attack an asset bubble. If enough of them attack the bubble, then the bubble will burst immediately. If, instead, enough of them choose to ride the bubble, then the bubble will survive for a while—but not forever. In particular, the bubble will burst at some random date $t = T$, even if no trader ever attacks it.

Before we proceed, we should clarify what a bubble means. In Abreu and Brunnermeier (2003), a bubble is defined in terms of two possible price paths. The first path, which is interpreted as the “bubble path”, features both a higher level and a higher growth rate than the second, the “fundamental path”. Both of these paths are exogenous. What is endogenous is only the crash of the bubble, that is, the transition from the first to the second path. The contribution of the paper rests in the characterization of this transition.

This transition is characterized under two scenarios. In the first scenario, Nature publicly announces at $t = 0$ the aforementioned date T to all traders at once. In the second scenario, Nature gradually reveals this date to the traders over a small time interval of length δ ; think of this as each trader becoming aware of,

⁴⁷See Weinstein and Yildiz (2007b) for a formalization of this kind of intuition, albeit within the context of a static game.

or “wakening up” to, the fact that there is a bubble, without however knowing right away how many other traders have also waken up. The two scenarios lead to different predictions, even when δ is arbitrarily small. This is because in the first scenario the aforementioned fact becomes common knowledge instantly, whereas in the second scenario asynchronous awareness implies that common knowledge of the aforementioned fact is not achieved even long time after every single agent has herself become aware of the fact.

Let us elaborate. To start, consider the first scenario, in which all agents wake up simultaneously. In this case, the bubble bursts immediately (at $t = 0$). The proof is as follows. Suppose that a trader believes that no other trader will ever attack. Then, this trader knows that the asset will continue to have abnormal returns for any $t < T$ and that its price will jump down exactly at T . It is then optimal for the trader to ride the bubble up to $t = T - \epsilon$ and sell the asset just at that moment.⁴⁸ But if all traders do the same, then the bubble will burst at $T - \epsilon$ rather than T , in which case it becomes optimal for the individual trader to sell the asset at $t = T - 2\epsilon$ rather than at $t = T - \epsilon$. Repeating this argument proves that the bubble bursts immediately, no matter how far in the future T happens to be.

Consider next the second scenario, in which agents wake up in an asynchronous manner. In this case, the bubble can persist for long time after *all* traders have waken up. To explain why, suppose that it takes half of the traders to attack in order for the bubble to burst before T . It follows that the bubble will survive at least till $t = \delta/2$, because not enough traders are awake prior to this point for a coordinated attack to trigger a crash. (Recall that δ denotes the length of time it takes for all traders to wake up; think of δ as small relative to T .)

Consider a trader who just woke up. At this moment, the trader does not know whether she is the first in line to wake up or the last one. As long as she is strictly below the middle of the line, she is better off riding the bubble, because it will take more time for half of the traders to be awake and therefore for a crash to be possible. But if this is true, it is optimal for the trader to ride the bubble even if she happens to be on or just above the middle of the line. Iterating, it is possible to prove that a trader who just woke up finds it optimal—indeed dominant—not to attack the bubble right away.

We can now use this fact to initiate a process of iterated deletion of dominated strategies from “below”: given that it is dominant for a trader not to attack at the moment she wakes up, it becomes iteratively dominant for a trader to delay attacking for a certain time interval after she has woken up, which in turn makes it iteratively dominant not to attack for an even bigger time interval, and so on. The opposite process, the one from “above”, gets initiated from the fact that it is dominant for a trader to attack once she reaches date $t = T - \epsilon$, for arbitrarily small ϵ .

Abreu and Brunnermeier (2003) prove that the aforementioned two processes converge to the same point, which means that the equilibrium is unique and is associated a critical threshold τ^* such that the following is true: the trader keeps riding the bubble until τ^* units of time have passed since she woke up, and starts attacking the bubble thereafter. It follows that the bubble bursts at date $t = t^* \equiv \delta/2 + \tau^*$, which is how long it takes for exactly 1/2 of the traders to attack. Importantly, t^* can be much larger than δ , although, of course, it has to be smaller than T . It follows that the bubble can survive for long time after it has become known—but not common knowledge—to all traders that the bubble will burst.

To recap, the unravelling that triggers the immediate burst of the bubble in the first scenario hinges on the assumption that traders can instantaneously reach common knowledge of the fact that there is a bubble that can be burst by a coordinated run. Once this assumption is relaxed by allowing for asynchronous

⁴⁸Although the model is in continuous time, it helps to recast it in discrete time, letting ϵ denote the length between any two consecutive points of time and thinking of ϵ as arbitrarily small.

awareness of the aforementioned fact, there can be significant delay in the kind of coordinated action that is needed for the bubble to burst and the market price to get in line with the asset’s fundamental valuation.

Remark 1. We will encounter similar delay effects within the class of “beauty contests” and the related business-cycle applications that we study in Sections 7 and 8. In that context, the relevant delay effects will manifest as rigidity, or inertia, in the response of certain outcomes to innovations in the underlying fundamentals.

Remark 2. As anticipated, the effects of asynchronous awareness that we reviewed here have a similar flavor as the effects of asynchronous choice—or the Calvo friction—that we discussed in the previous subsection. This supports our earlier claim, namely that there is a tight relation between asynchronous choice and informational frictions. Understanding the exact nature of this relation is an open question; applying the insights to business-cycle settings is another fruitful direction for future research.

7 Beauty Contests: Theory

In this section we turn attention to a class of games that feature weak complementarity and linear best responses. The incomplete-information versions of these games are often referred to as “beauty contests”, due to a resemblance with Keynes’ parable for financial markets. For applied purposes, one can think of this class of games as a stylization of unique-equilibrium macroeconomic models whose equilibrium conditions can be represented as, or at least be approximated by, a system of (log)linear equations. The lessons we develop in the present section are therefore relevant for canonical business-cycle models of either the RBC or the New-Keynesian type, as well as for a class of asset-pricing models. Such applications are discussed in the next section; here, we develop some basic insights, building in part on the analyses of Morris and Shin (2002), Angeletos and Pavan (2007, 2009), and Bergemann and Morris (2013).

Remark. The settings studied in this and the next section differ from those studied in Sections 4–6 in two related respects. First, they admit a unique equilibrium—in fact, a unique rationalizable outcome—irrespective of the structure of information. Second, they let higher-order beliefs matter in a more modest manner than in global games: beliefs of order h have a vanishingly small effect as $h \rightarrow \infty$.

7.1 Set Up

In this section we use the framework introduced in Section 2 and impose the following restrictions.

Assumption 4 (Beauty Contest). $D_k = D_\theta = \mathbb{R}$ and U is quadratic in (k, K, θ) . Furthermore, U satisfies the following restrictions: $U_{kk} < 0$, $0 \leq -U_{kK}/U_{kk} < 1$, and $U_{k\theta} > 0$.

The first restriction, $U_{kk} < 0$, imposes concavity at the individual level, ensuring that best responses are well defined. The second restriction, $0 \leq -U_{kK}/U_{kk} < 1$, is equivalent to imposing that the slope of best responses with respect to aggregate activity is positive but less than one. This means that strategic complementarity is of the weak form and guarantees that the equilibrium is unique. The last restriction, $U_{k\theta} > 0$, is innocuous: it means that, other things equal, a higher fundamental causes the agent to take a higher action. (If the converse had been true, we could have simply redefined θ as $-\theta$).

Remark. Although in this section we restrict attention to the case of weak strategic complementarity (namely $0 \leq -U_{kK}/U_{kk} < 1$), many of the key lessons, including those about the role of incomplete information in sustaining inertia and animal spirits, extend also to the case of strategic substitutability (namely $-U_{kK}/U_{kk} < 0$).

7.2 Complete-Information Benchmark

As in the case of global games, we start by studying the benchmark of complete information, which in turn nests the case of perfect information.

Proposition 14. *[Complete Info] There exist coefficients $(\kappa_0, \kappa_1, \kappa_2)$, pinned down by the payoff function U , such that the following is true: whenever information is complete, the equilibrium action is given by*

$$k_i = \kappa (\mathbb{E}_i \theta_i, \mathbb{E}_i \bar{\theta}) \equiv \kappa_0 + \kappa_1 \mathbb{E}_i \theta_i + \kappa_2 \mathbb{E}_i \bar{\theta}, \quad (14)$$

where $\bar{\theta} \equiv \int \theta d\Theta(\theta)$ is the average fundamental and \mathbb{E}_i is the rational expectation conditional on ω_i .

The coefficients $(\kappa_0, \kappa_1, \kappa_2)$, which are characterized in the Appendix, depend on U and thereby on the micro-foundations of the particular application that lies behind U . Understanding what determines these coefficients and how they can be identified in the data is important within any individual application.⁴⁹

Recall that complete information means common knowledge of the distribution of information, Ω , but allows for the possibility that agents face uncertainty either about their own fundamental (θ_i) or about the aggregate fundamentals (Θ , or $\bar{\theta}$). This explains why the forecast of the fundamentals, $\mathbb{E}_i \theta_i$ and $\mathbb{E}_i \bar{\theta}$, show up in the above characterization. Trivially, the case of perfect information is then nested simply by letting $\mathbb{E}_i \theta_i = \theta_i$ and $\mathbb{E}_i \bar{\theta} = \bar{\theta}$. That is, as we move from perfect information to imperfect but complete information, all we have to do is to replace the actual fundamentals with the forecasts of them. This offers a sharp illustration of the more general point we made in Section 3: insofar as strategic uncertainty is ruled out, it makes little difference whether the agents face uncertainty about the fundamentals or know them perfectly.

7.3 Equilibrium with Incomplete Information

Suppose now that information is incomplete. In this case, equilibrium actions need not be pinned down by fundamentals (or forecasts of fundamentals), because the latter are generally not sufficient for forecasting the actions of others. This makes equilibrium outcomes vary away from, and around, their complete-information counterparts, in a manner that is formalized in the next two propositions.

Proposition 15 (Equilibrium and Coordination). *The equilibrium satisfies the following fixed-point relation*

$$k_i = \mathbb{E}_i [\kappa (\theta_i, \bar{\theta})] + \alpha \cdot \mathbb{E}_i [K - \kappa (\bar{\theta}, \bar{\theta})], \quad (15)$$

where $\alpha \equiv -U_{kK}/U_{kk} \in [0, 1)$.

Proposition 16 (Equilibrium and Higher-order Beliefs). *Suppose that each agent knows her own fundamental. Then, the equilibrium action of agent i is given by*

$$k_i = \kappa_0 + \kappa_1 \theta_i + \kappa_2 \mathbb{E}_i \left\{ \sum_{h=0}^{\infty} (1 - \alpha)^h \bar{\mathbb{E}}^h [\bar{\theta}] \right\}, \quad (16)$$

⁴⁹For instance, in the context of the neoclassical economy introduced in Section 2, κ_1 measures the “micro” elasticity of the response of a farmer’s output to an idiosyncratic shock to her own productivity, whereas $\kappa_1 + \kappa_2$ measures the “macro” elasticity of the response of aggregate output to an aggregate productivity shock. The former is representative of ZIP- or state-level elasticities that are estimated in Mian et al. (2013), Mian and Sufi (2014) and Nakamura and Steinsson (2014) on the basis of appropriate cross-sectional variation in the data; the latter is what macroeconomists are most often concerned with; the two differ because of general-equilibrium effects, which are captured here by the dependence of U (and of U_K in particular) on K .

where $\bar{\mathbb{E}}^h[\bar{\theta}]$ denotes the h -th order average forecast of the mean fundamental.⁵⁰

Proposition 15 highlights the dependence of equilibrium allocations on beliefs regarding aggregate activity (forecasts of the actions of others). To understand condition (15), recall that $\mathbb{E}_i[\kappa(\theta_i, \bar{\theta})] = \kappa(\mathbb{E}_i\theta_i, \mathbb{E}_i\bar{\theta})$ is the action agent i would have taken in equilibrium had information been complete. How much an agent deviates from this benchmark when information is incomplete depends on $\mathbb{E}_i[K - \kappa(\bar{\theta}, \bar{\theta})]$, which is her forecast of the deviation of the average action of the rest of the population from this benchmark, weighted by the coefficient α . In this sense, the coefficient α measures how much each *individual* cares about aligning her action with that of others, or equivalently the private motive to coordinate: it identifies the degree of strategic complementarity.

Proposition 16 then restates this result in terms of the hierarchy of beliefs (forecasts of the forecasts of others). For this part we have added the restriction that each agent knows her own fundamental. This restriction allows us to isolate the uncertainty that agents face about one another's actions from the uncertainty that each agent may face about her own preferences and technologies: we shut down the latter and concentrate on the former. Comparing the above result to Proposition 14, we then see that the key difference as we move from complete to incomplete information is that $\mathbb{E}_i\bar{\theta}$ has been replaced by $\mathbb{E}_i\{\sum_{n=0}^{\infty}(1-\alpha)\alpha^n\bar{\mathbb{E}}^n[\bar{\theta}]\}$, a weighted average of the entire hierarchy of beliefs about the underlying aggregate fundamental. This is because an agent's first-order belief of the aggregate shock is no longer sufficient for forecasting the equilibrium level of aggregate activity; the agent needs to forecast the forecasts of others. The coefficient α then determines the sensitivity of the equilibrium action to higher-order beliefs: the stronger the degree of complementarity, the stronger the impact of higher-order beliefs relative to first-order beliefs.

That said, note that the specifics of higher-order beliefs do not matter per se: if we change, say, the 3rd and the 7th order belief of an agent while keeping $\mathbb{E}_i\{\sum_{n=0}^{\infty}(1-\alpha)\alpha^n\bar{\mathbb{E}}^n[\bar{\theta}]\}$ constant, then her optimal action does not change. This is because the agent's best response depends merely on her first-order belief of K , not the details of either the belief hierarchy about the fundamentals or the underlying information structure. This underscores, once again, that incomplete information and higher-order uncertainty are modeling devices that allow the researcher to accommodate interesting, and potentially testable, variation in the agents' expectations (first-order beliefs) of economic outcomes. We will discuss how this can be understood in a manner that is robust to the underlying micro-foundations.

7.4 Simplifying Assumptions

To simplify exposition, we make a few additional assumptions. First, we rescale the fundamental so that the following is true.

Assumption 5. $\kappa_0 = 0$ and $\kappa_1 + \kappa_2 = 1$.

Using the characterization of the coefficients κ_1 and κ_2 (equation (51) in the Appendix), one can then show that $\kappa_1 = 1 - \alpha$ and $\kappa_2 = \alpha$. This means that, under complete information, $1 - \alpha$ measures the *micro* elasticity of the individual's action to her own fundamental, α measures the general-equilibrium feedback in the case of aggregate shocks; and $1 = (1 - \alpha) + \alpha$ measures the *macro* elasticity of the aggregate action to an aggregate shock. It also implies that the agent's best response under incomplete information, or equation (15), reduces to the following:

$$k_i = (1 - \alpha) \mathbb{E}_i[\theta] + \alpha \cdot \mathbb{E}_i[K]. \quad (17)$$

⁵⁰The operator $\bar{\mathbb{E}}^h$ is defined recursively by letting, for any variable X , $\bar{\mathbb{E}}^0[X] = X$, $\bar{\mathbb{E}}^1[X] \equiv \bar{\mathbb{E}}[\bar{\mathbb{E}}^0[X]] \equiv \int \mathbb{E}_i[X] di$ and $\bar{\mathbb{E}}^h[X] \equiv \bar{\mathbb{E}}[\bar{\mathbb{E}}^{h-1}[X]] = \int \mathbb{E}_i[\bar{\mathbb{E}}^{h-1}[\bar{\theta}]] di \forall h \geq 2$.

To simplify the exposition and also to stay close to the applied literature, we rule out payoff heterogeneity and assume a Gaussian specification for the underlying common fundamental and for the information the agents receive about it.

Assumption 6. *There is a common fundamental: $\theta_i = \bar{\theta} = \theta$, for all i and all states of nature.*

Assumption 7. *The fundamental θ is drawn from a Normal distribution with mean 0 and variance $\sigma_\theta^2 > 0$.⁵¹ Each agent observes two signals, a private one and a public one. The private signal takes the form*

$$x_i = \theta + u + \epsilon_i,$$

where u and ϵ_i are common and idiosyncratic noises that are independent of one another as well as of θ , and are drawn from Normal distributions with mean 0 and variances, respectively, $\sigma_u^2 \geq 0$ and $\sigma_\epsilon^2 \geq 0$. The public signal takes the form

$$z = \theta + \zeta,$$

where ζ is noise, independent of θ, u and all ϵ_i , and drawn for a Normal distribution with mean 0 and variance $\sigma_\zeta^2 \geq 0$.

Incomplete information requires $\sigma_\epsilon > 0$ and $\sigma_\zeta > 0$, that is, non-trivial private information about the aggregate fundamental, but does not necessarily require $\sigma_u > 0$. In fact, much of the related literature (e.g., Morris and Shin 2002, Angeletos and Pavan 2007) is nested here by restricting $\sigma_u = 0$, that is, by shutting down the correlated noise, u , in the private signal. The reason why we allow for such correlation is twofold. First, it helps mimic the correlation in beliefs that emerges as agents talk to their neighbors or participate in localized markets. Second, the noise u in combination with the noise ζ allow us to disentangle the aggregate variation in beliefs of aggregate activity from the variation in either the fundamental or the beliefs of it.

To formalize this point, consider the residuals of the projection of $\bar{\mathbb{E}}\theta$ and $\bar{\mathbb{E}}K$ on θ .⁵² From Lemma 17 and Proposition 18 (see below), we have that the two residuals are linear transformations of u and ζ . When u is shut down, as is done in much of the existing literature, the two residuals become proportional to one another, implying that the noise in beliefs of endogenous economic outcomes is perfectly correlated, and thus indistinguishable from the noise in the beliefs of the fundamentals. When instead u is switched on, as we do here, the two residuals cease to be perfectly correlated, thus permitting us to disentangle the two types of noise. By the same token, the introduction of u helps accommodate a strong form of “animal spirits”: as emphasized in Proposition 19 below, expectations of economic activity (and actual activity) can fluctuate holding constant *both* the true fundamentals and the agents’ beliefs of them.

It is also worth noting that the assumed information structure is less restrictive than it seems. Let A be the set of joint distributions for $(\theta, \bar{\mathbb{E}}\theta, K, \bar{\mathbb{E}}K)$ that can be obtained in equilibrium under the assumed information structure, for some $(\sigma_u^2, \sigma_\epsilon^2, \sigma_\zeta^2)$. Next, let B be the set of joint distributions for $(\theta, \bar{\mathbb{E}}\theta, K, \bar{\mathbb{E}}K)$ that can be obtained in equilibrium under *arbitrary* Gaussian information structures. The latter nests cases in which the agents observe an arbitrary collection of Gaussian signals, not only about the underlying fundamental (θ), but also about the endogenous outcome (K) or even about one another’s signals; it can also nest situations where the agents talk to one another, either directly, or through some market mechanism. Using a similar argument as in Bergemann and Morris (2013), one can show that $A = B$. One can thus think

⁵¹ Allowing for a nonzero prior mean for θ only adds a constant in the equilibrium action, without affecting any of the positive properties we document below.

⁵² To ease notation, we henceforth let $\bar{\mathbb{E}}X$ be a short-cut for $\bar{\mathbb{E}}[X]$, for any variable X .

of the assumed signal structure merely as a convenient parameterization of the type of outcomes that can be obtained under arbitrary information structures, as opposed to seeking a literal interpretation of it.

Unless stated otherwise, we impose $\sigma_\epsilon, \sigma_\zeta, \sigma_u > 0$.

7.5 The Structure of Higher-order Beliefs

As a stepping stone towards the characterization of the equilibrium, we next study how the hierarchy of beliefs varies with the underlying aggregate shocks, namely the fundamental θ and the noises ζ and u .

Proposition 17 (Higher-Order Beliefs). *For every $h \geq 1$, there exist positive scalars $(\omega_\theta^h, \omega_u^h, \omega_\zeta^h)$ such that the h -th order average expectation of the fundamental is given by*

$$\bar{\mathbb{E}}^h[\theta] = \omega_\theta^h \theta + \omega_\zeta^h \zeta + \omega_u^h u,$$

for every realization of (θ, ζ, u) . Furthermore, for every $h \geq 1$,

$$1 > \omega_\theta^h > \omega_\theta^{h+1} > 0 \quad \text{and} \quad 0 < \frac{\text{Var}(\bar{\mathbb{E}}^h[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^h[\theta])} < \frac{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta])} < 1.$$

This result means that beliefs of any given order h vary with the fundamental, θ , and the two sources of noise, u and ζ ; but the higher the order h is, the smaller the absolute response of the h -th order belief to the fundamental; and the higher the contribution of the noise relative to that of the fundamental in driving the volatility in h -th order beliefs. The first property, namely the lower response to the fundamental, is the mirror image of the fact that higher-order beliefs are more anchored to the prior (which of course does not move with the realized fundamental) than lower-order beliefs. The second property, on the other hand, means that higher-order beliefs are more “noisy” than lower-order beliefs, in the sense that the R-square of the projection of the former on the fundamental is lower than that of the projection of the latter on the fundamental.

These properties have been established here only for a very specific information structure. However, using methods similar to those of Bergemann and Morris (2013), it is possible to show that these properties are shared by essentially any Gaussian information structure. By the same token, the positive implications we document next are not driven by any ad-hoc specifications of the information structure; rather, they are *generic* to the incompleteness of information.

7.6 Positive Implications

We now proceed to translate the preceding properties of higher-order beliefs into those of the observables of the model. To this goal, we must first take a stand on what the model’s observables are. For the purposes of the present exercise, we envision that the researcher collects data on the exogenous fundamental θ , the endogenous aggregate outcome K , and possibly the average beliefs of these objects, namely $\bar{\mathbb{E}}\theta$ and $\bar{\mathbb{E}}K$. For instance, in the context of the neoclassical economy introduced in Section 2 and taken up again in the next section, K corresponds to aggregate output and θ to labor productivity. The researcher can observe these variables in standard macroeconomic data sets, and can extract information about agents’ beliefs of these objects from surveys of expectations.

Definition 15. *The model’s observables are $(\theta, \bar{\mathbb{E}}\theta, K, \bar{\mathbb{E}}K)$. The model’s predictions are the restrictions the model imposes on the joint distribution of $(\theta, \bar{\mathbb{E}}\theta, K, \bar{\mathbb{E}}K)$.*

Note that the stochastic properties of θ and $\bar{\mathbb{E}}\theta$ are dictated directly by the assumptions the researcher makes about shocks and information. A model's essence is in the “cross-equation restrictions” it imposes on the stochastic properties of the endogenous objects K and $\bar{\mathbb{E}}K$. We now characterize these restrictions.

Proposition 18 (Positive Properties). *For any given information structure, there exist positive scalars $(\phi_\theta, \phi_u, \phi_\zeta)$ and $(\psi_\theta, \psi_u, \psi_\zeta)$, such that the aggregate outcome and the average forecast of it satisfy the following:*

$$K = \phi_\theta\theta + \phi_\zeta\zeta + \phi_uu \quad \text{and} \quad \bar{\mathbb{E}}[K] = \psi_\theta\theta + \psi_\zeta\zeta + \psi_uu.$$

Furthermore,

$$1 > \phi_\theta > \psi_\theta > 0, \tag{18}$$

and

$$0 < \frac{\text{Var}(K|\theta)}{\text{Var}(K)} < \frac{\text{Var}(\bar{\mathbb{E}}[K]|\theta)}{\text{Var}(\bar{\mathbb{E}}[K])} < 1. \tag{19}$$

This result represents the equilibrium counterpart of Lemma 1. To interpret it, note that ϕ_θ identifies the slope of the aggregate outcome to the fundamental, ψ_θ identifies the corresponding slope of the average forecast of the aggregate outcome, and finally the (normalized) slope of either object under complete information is 1. Condition (18) therefore means the average forecast under-reacts relatively to the actual outcome, and that both of them under-react relative to the complete-information benchmark. Condition (19), on the other hand, means that the average forecast of K is relatively more noisy than the actual K , in the sense that the relative contribution of the noises ζ and u to the volatility of $\bar{\mathbb{E}}K$ is larger than those to the volatility of K . In this sense, forecasts under-react to fundamentals but overreact to noise.

This result suggests a unifying explanation for two seemingly contradictory sets of facts: the *under-reaction* of forecasts documented in Coibion and Gorodnichenko (2012); and the *overreaction* of forecasts documented in Greenwood and Shleifer (2014) and Gennaioli, Ma and Shleifer (2015). We expand on the relation to Coibion and Gorodnichenko (2012) in the next section; we leave the potential connection to Greenwood and Shleifer (2014) and Gennaioli, Ma and Shleifer (2015) to future work.

We also note that the actual level of aggregate activity and the average belief of it are positively but only imperfectly correlated:

$$1 > \text{Corr}(K, \bar{\mathbb{E}}[K]) > 0.$$

By contrast, when information is complete, the above correlation becomes perfect.

Pushing this observation a bit further, we have the following result.

Proposition 19 (Animal Spirits). *The economy features “animal spirits”, not only in the sense of Definition 13, but also in the sense that*

$$\text{Var}(K|\theta, \bar{\mathbb{E}}\theta) > 0.$$

Furthermore, the following is necessarily true:

$$\text{Var}(\bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta) > 0 \quad \text{and} \quad \text{Cov}(K, \bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta) > 0.$$

Recall that Definition 13 identified animal spirits with situations in which outcomes vary without commensurate variation in the agents beliefs of their *own* fundamentals. In the present context, own and aggregate fundamentals coincide ($\theta_i = \bar{\theta}$ for all i and all states of nature). It follows that animal spirits obtain in the sense of Definition 13 if and only if $\text{Var}(K|\bar{\mathbb{E}}\theta) > 0$.

The first part of Proposition 19 reinforces this possibility by showing that variation in outcomes can obtain holding constant, not only the beliefs of the fundamental, but also the realized fundamental. This helps disentangle “animal spirits” from the noise in beliefs of the fundamental: aggregate activity varies while holding constant both the true fundamental and the aggregate error in beliefs of the fundamental. It also underscores that the notion of “animal spirits” used here, and the related notion found in Angeletos and La’O (2013), is both conceptually and empirically distinct from that used in Lorenzoni (2009) and Barsky and Sims (2012): in these papers, “animal spirits” is defined as noise in beliefs of fundamentals. Instead, the notion used here is closely connected to the one used in the literature on multiple equilibria and sunspot fluctuations; as noted before, the only difference is that we engineer animal spirits with the help of rich higher-order beliefs as oppose to multiple equilibria and correlation devices.

The second part of Proposition 19 states that the volatility attributed to “animal spirits” is features positive comovement between actual and expected outcomes, underscoring once more its self-fulfilling nature.

We conclude that incomplete information allows the researcher to rationalize empirically relevant imperfections in the co-movement of fundamentals, economic outcomes, and beliefs thereof—thereby accommodating a number of facts that may be prima-facie inconsistent with workhorse macroeconomic models. We elaborate on this point in the next section, within the context of specific applications.

7.7 Complementarity, Information, and Volatility

Before moving to the applications, we study how the degree of strategic complementarity α affects the different kinds of volatility in the model’s observables. Clearly, the stochastic properties of θ and $\bar{\mathbb{E}}\theta$ are invariant to α ; we thus need to consider only the effect of α on the stochastic properties of K and $\bar{\mathbb{E}}K$, which is what the next proposition does.

Proposition 20 (The Effect of Complementarity). *Stronger strategic complementarity (higher α) results in all of the following:*

(i) *Less covariation between the fundamental and either the aggregate outcome or the average belief of it:*

$$\frac{\partial Cov(K, \theta)}{\partial \alpha} < 0 \quad \text{and} \quad \frac{\partial Cov(\bar{\mathbb{E}}K, \theta)}{\partial \alpha} < 0$$

(ii) *Higher portion of volatility driven by noise:*

$$\frac{\partial}{\partial \alpha} \left(\frac{Var(K|\theta)}{Var(K)} \right) > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} \left(\frac{Var(\bar{\mathbb{E}}K|\theta)}{Var(\bar{\mathbb{E}}K)} \right) > 0$$

(iii) *More room for animal spirits, not only in the sense of Definition 13, but also in the stronger sense defined in Proposition 19:*

$$\frac{\partial Var(K|\theta, \bar{\mathbb{E}}\theta)}{\partial \alpha} > 0 \quad \frac{\partial Var(K|\theta, \bar{\mathbb{E}}\theta)}{\partial \alpha} > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} (Cov(K, \bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta)) > 0$$

To appreciate this result, note that under the maintained normalization $\kappa_1 + \kappa_2 = 1$, the value of α is completely immaterial for aggregate outcomes when information is complete. The above result therefore reveals how incomplete information interacts with strategic complementarity (or general-equilibrium effects) to shape the covariation of actual economic outcomes and their forecasts. To put it differently, this result hinges entirely on strategic uncertainty.

We conclude by noting that while incomplete information helps accommodate animal spirits and overreaction to certain forms of noise, it can do so only at the expense of dampening the overall volatility.

Proposition 21 (Volatility Bound). *The unconditional variance of K under incomplete information is lower than its perfect-information counterpart.*

This result follows from the basic fact that the variance of the forecast of any variable is always lower than the variance of the variable itself. Applying this fact recursively gives that the variance of higher-order forecasts is lower than the variance of lower-order forecasts, which in turn is lower than the variance of the fundamental. Because the equilibrium K is a weighted average of the hierarchy of forecasts, it then follows that the volatility of K is maximal when information is perfect.

That said, we wish to emphasize that this result concerns the *unconditional* variance of K , whereas researchers are often interested in certain *conditional* moments, such as those conditional on certain shocks or, in dynamic contexts, on certain past outcomes. When it comes to this kind of conditional moments, incomplete information may actually contribute to more volatility.⁵³ Finally, although the volatility of the average fundamental places an upper bound on the volatility of the aggregate outcome in the current setting, this bound can be relaxed in settings that allow for local interactions and/or idiosyncratic shocks; see Angeletos and La'O (2013) and Bergemann, Heumann, and Morris (2015).

7.8 Dynamics and Learning

So far the analysis has been confined to a static framework in order to deliver key insights in a sharp and transparent manner. In applications, however, it is often central to allow for two different kinds of dynamics. The first has to do with learning over time (the dynamics of beliefs). The second has to do with intertemporal payoff interdependences, which make the decisions depend on past and/or future actions (backward- and/or forward-looking effects).

Slow learning and inertia

To illustrate the role of the dynamics that obtain from slow learning, we now consider a repeated variant of our static beauty contest game. In this variant, the best responses remain static, in the sense that they depend only on the contemporaneous actions of others. Yet, interesting dynamics obtain because the fundamental is persistent and there is slow private learning about it.

Example 1. *The fundamental of each agent i follows a Gaussian random walk:*

$$\theta_{it} = \theta_{it-1} + v_t + \xi_{it},$$

where $v_t \sim N(0, \sigma_v^2)$ is an aggregate innovation and $\xi_{it} \sim N(0, \sigma_\xi^2)$ is an idiosyncratic one. Furthermore, in every period t , player i observes a private signal of her own fundamental:

$$x_{it} = \theta_{it} + \epsilon_{it} \tag{20}$$

where $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ is an idiosyncratic noise.

⁵³As a trivial example, recall that $Var(K | \mathbb{E}\theta) > 0$ and $Cov(K, \mathbb{E}K) < 1$ only when information is incomplete: with complete information, $Var(K | \mathbb{E}\theta) = 0$ and $Cov(K, \mathbb{E}K) = 1$.

The optimal action of agent i is still given by

$$k_{it} = \mathbb{E}_{it} [(1 - \alpha)\theta_{it} + \alpha K_t].$$

What is novel relative to our static setting is only the dynamic evolution of beliefs implied by the combination of persistent fundamentals and noisy learning.

To see how these features matter for the observables of the model, let us first consider the special case in which there is no complementarity ($\alpha = 0$). In this case, equilibrium outcomes depend merely on first-order beliefs of fundamentals. Furthermore, thanks to the Gaussian specification, the dynamics of first-order beliefs can be shown to obey the following rule:

$$\mathbb{E}_{it}\theta_t = \mathbb{E}_{it-1}\theta_{it} + \lambda [x_{it} - \mathbb{E}_{it-1}x_{it}],$$

for some $\lambda \in (0, 1)$ that depends on the standard deviations of the innovations in the fundamental and the noise in the signal. This result is an application of the Kalman filter, with λ being the Kalman gain.⁵⁴ Using $\mathbb{E}_{it-1}x_{it} = \mathbb{E}_{it-1}\theta_{it}$ and rearranging gives $\mathbb{E}_{it}\theta_{it} = (1 - \lambda)\mathbb{E}_{it-1}\theta_{it-1} + \lambda x_{it}$. Aggregating implies that

$$\bar{\mathbb{E}}_t\theta_t = (1 - \lambda)\bar{\mathbb{E}}_{t-1}\theta_{t-1} + \lambda\theta_t.$$

Finally, when $\alpha = 0$, we have that $K_t = \bar{\mathbb{E}}_t\theta_t$. It follows that we can express the dynamics of K_t as follows:

$$K_t = (1 - \lambda)K_{t-1} + \lambda\theta_t,$$

When information is perfect ($\sigma_\epsilon = 0$), λ becomes 1 and the expression above reduces to $K_t = \theta_t$, implying that K_t follows a random walk, just as the underlying fundamental. When instead information is imperfect ($\sigma_\epsilon > 0$), λ is strictly lower than 1, implying that K_t features sluggish response to the innovation in the fundamental. In particular, if we let IRF_j denote the cumulative effect of an aggregate innovation on the level of aggregate activity j periods after the innovation occurs,⁵⁵ we have the following characterization:

$$IRF_0 = \lambda, \quad IRF_1 = \lambda[1 + (1 - \lambda)], \quad IRF_2 = \lambda[1 + (1 - \lambda) + (1 - \lambda)^2], \text{ etc.}$$

Note then that IRF_j is an increasing sequence that starts from λ and converges to 1, with a higher speed of convergence when λ is higher. Combining this observation with the fact that λ is decreasing in σ_x , we infer that more noise (higher σ_x) induces more inertia in the response of the aggregate outcome K to any innovation in the underlying fundamental.

It is important to recognize that, when $\alpha = 0$, the information that matters is only the one that each agent has about her *own* fundamental, not the one about the aggregate fundamental. What explains the inertia documented above is therefore the lack of (first-order) knowledge of own fundamentals, not the lack of common knowledge of aggregate fundamentals. By the same token, this kind of inertia can be quantitatively important only insofar as the agents are sufficiently uninformed about their own fundamentals.

We can summarize all these points as follows.

⁵⁴Strictly speaking, λ should be allowed to be a deterministic function of t , reflecting the dependence on the initial prior. But as time passes, the impact of the prior vanishes, and λ_t converges to a constant, which identifies the steady-state solution of the Kalman filter. Here, as in most of the applied literature, we focus on this solution.

⁵⁵Formally, for any $j \geq 0$, $IRF_j = \partial\mathbb{E}[K_{t+j}|K_{t-1}, v_t]/\partial v_t$.

Proposition 22. *Slow learning can induce inertia by itself, even if $\alpha = 0$ (i.e., in the absence of strategic, or general-equilibrium, interdependence). However, the following properties hold whenever $\alpha = 0$:*

(i) *The inertia that is present at the aggregate level (in the response of aggregate outcomes to aggregate shocks) coincides with the inertia that is present at the individual level (in the response of idiosyncratic outcomes to idiosyncratic shocks).*

(ii) *Both types of inertia vanish as the agents learn their own fundamentals, regardless of the information they may or may not have about aggregate economic conditions.*

By contrast, when $\alpha > 0$, inertia can result from lack of common knowledge of the aggregate fundamental and can be quantitatively important, even if each agent is arbitrarily well informed about both her own or the aggregate fundamentals. This is because higher-order beliefs are less sensitive to innovations in fundamentals than lower-order beliefs, not only contemporaneously (which is the effect already documented in the static setting), but also in terms of their dynamic adjustment.

We illustrate this property with the following modification of the preceding example.

Example 2. *The fundamental of each agent i is given by $\theta_{it} = \bar{\theta}_t + \xi_{it}$, where $\xi_{it} \sim N(0, \sigma_\xi^2)$ is an idiosyncratic shock and $\bar{\theta}_t$ is the aggregate fundamental. The latter follows a random walk:*

$$\bar{\theta}_t = \bar{\theta}_{t-1} + v_t,$$

where $v_t \sim N(0, \sigma_v^2)$ is an aggregate innovation. Finally, in every period t , agent i observes her own fundamental, θ_{it} , along with the following additional private signal about the aggregate fundamental:

$$x_{it} = \bar{\theta}_t + \epsilon_{it},$$

where $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ is idiosyncratic noise.

We can then show the following.

Proposition 23. *[Complementarity and Inertia] When $\alpha = 0$, $K_t = \bar{\theta}_t$. When instead $\alpha \in (0, 1)$,*

$$K_t = \gamma_K K_{t-1} + \gamma_\theta \bar{\theta}_{t-1} + \gamma_v v_t,$$

for some scalars $\gamma_K, \gamma_\theta, \gamma_v$ that depend on $(\alpha, \sigma_\epsilon, \sigma_\xi, \sigma_v)$.

This result follows from Angeletos and La'O (2010)'s adaptation of Woodford (2003). Woodford (2003) studies an environment where there is no payoff heterogeneity, thus combining two mechanisms: the inertia induced by mere lack of knowledge of own fundamentals, with the inertia induced by strategic complementarity and lack of common knowledge of the aggregate fundamentals. By contrast, Angeletos and La'O (2010) and the closely related example we study here isolate the second mechanism by imposing perfect information about own fundamentals. The dependence of K_t on K_{t-1} and $\bar{\theta}_{t-1}$ documented in the above result therefore reflects the endogenous persistence introduced by, and only by, the friction in coordination.⁵⁶

To gain further insight, we consider a numerical example. Figure 4, which is borrowed from Woodford (2003), illustrates the dynamic response of the average first- and higher-order forecasts of $\bar{\theta}_t$ to a positive innovation in it. For any order h , the response of $\bar{\mathbb{E}}_t \bar{\theta}_t$, the h -th order forecast, eventually converges to 1 as

⁵⁶The exact characterization of this dependence was elusive to both Woodford (2003) and Angeletos and La'O (2010), but becomes possible with the method developed in Huo and Takayama (2015a,b).




Figure 4: Dynamics of First- and Higher-Order Beliefs

time passes, reflecting the accumulation of more and more information over time. Yet, in any finite period after the innovation has occurred, higher-order beliefs remain further away from 1 than lower-order beliefs, mirroring the property documented in Lemma 17 on the static framework. What is more, higher-order beliefs converge to 1 more slowly than lower-order beliefs. All in all, Figure 4 therefore reveals that higher-order beliefs exhibit, not only lower sensitivity on impact, but also more sluggishness over time.

The translation of these properties to the observables of the model is straightforward. Importantly, if we fix the information structure (and hence also fix the dynamics of first- and higher-order beliefs) but increase the degree of strategic complementarity, we obtain both a smaller immediate response and a more sluggish adjustment of the equilibrium level of activity.

Last but not least, note that the mechanism we have documented above operates at the aggregate level but not at the idiosyncratic level (or in the cross section): because the agents are perfectly informed about their own fundamentals, they exhibit no inertia at all in their response to idiosyncratic shocks.

In the next section, we present two applications of this mechanism: one in the context of an RBC economy hit by technology shocks, another in the context of New-Keynesian economy hit by a monetary shock. Moving beyond these examples, the broader lesson is the following.

Fact. *Frictions in coordination, modeled as incomplete information coupled with strategic complementarity, can help generate significant inertia in the response of aggregate outcomes to aggregate shocks to fundamentals, even when the available information about the fundamentals is precise and even when there is no such inertia in the response of idiosyncratic outcomes to idiosyncratic shocks.*

Remark. As anticipated in Subsection 6.2, the inertia we have documented here is of a similar nature as the delay effects documented in the asynchronous-awareness setting of Abreu and Brunnermeier (2003). In both cases, the key is that higher-order beliefs adjust more slowly than first-order beliefs—or, equivalently, that agents continue to lack confidence on whether others will adjust to a change in the environment long after they have themselves become confident that this change took place.

Intertemporal payoff interdependencies

We now turn attention to the role of intertemporal payoff interdependencies. In particular, we have in mind settings in which an agent's actions in any given period depend on her expectations of the actions of other agents, not only in the same period, but also in the future. Such forward-looking features are standard in dynamic macroeconomic models. Importantly, these features expand the room for higher-order beliefs.

Consider, for example, the textbook RBC model. In this model, the labor-supply and consumption-saving choices of a household in any given period depend, not only on the current wage and interest rate, but also on her expectations of wages and interest rates in all subsequent periods. Furthermore, the wage and the interest rate that clear the relevant markets in any given period depend on the joint behavior of all consumers and firms, which in turn depend on their own expectations of future wages and interest rates, and so on. It follows that the kind of high-order beliefs that are relevant can be vastly richer than the one we have allowed so far. In particular, even if the past and current fundamentals are common knowledge, one can sustain frictions in coordination by introducing private information about future fundamentals.

To illustrate these points, consider a dynamic game in which the best-response function of player i happens to take the following form:

$$k_{it} = \mathbb{E}_{it} [g(\theta_{it}, K_t; \theta_{it+1}, K_{t+1})], \quad (21)$$

where g is a linear function. The dependence of g on θ_{it} and K_t mirrors the one accommodated in our earlier static analysis; the novelty here is the inclusion of the next-period variables θ_{it+1} and K_{t+1} . This inclusion captures the forward-looking aspects of the environment.

To give an example, let us consider a monetary model that merges Woodford (2003) with Taylor (1979, 1980). In particular, firms face two frictions in their price-setting behavior: first, they have incomplete information about aggregate economic conditions (as in Woodford); and second, they can reset their prices only every two periods (as in Taylor). In the absence of these two frictions, the optimal price in period t would have been given (after an appropriate log-linearization) by

$$p_{it}^* \equiv (1 - \alpha)\theta_{it} + \alpha P_t$$

where θ_{it} summarizes exogenous demand and supply shocks hitting the firm in period t , P_t is the endogenous price level in period t , and α is the degree of strategic complementarity in pricing decisions. When only the first friction (incomplete information) is present, the optimal price set by firm i is

$$p_{it} = \mathbb{E}_{it} p_{it}^* = \mathbb{E}_{it} [(1 - \alpha)\theta_{it} + \alpha P_t]. \quad (22)$$

When instead both of the aforementioned frictions are present, the optimal reset price is given by

$$p_{it} = \mathbb{E}_{it} \left[\frac{1}{1+\beta} p_{it}^* + \frac{\beta}{1+\beta} p_{it+1}^* \right] = \mathbb{E}_{it} \left[\frac{1}{1+\beta} ((1 - \alpha)\theta_{it} + \alpha P_t) + \frac{\beta}{1+\beta} ((1 - \alpha)\theta_{it+1} + \alpha P_{t+1}) \right] \quad (23)$$

where β is the firm's discount factor. Clearly, while equation (22) is nested in our static framework, (23) is nested in the dynamic extension we have just introduced.⁵⁷

⁵⁷As already noted, the example considered here is a hybrid of Woodford (2003) and Taylor (1979, 1980). If one considers hybrids of Woodford (2003) and Calvo (1983), like those studied in Angeletos and La'O (2009) and Nimark (2011), the key difference is that p_{it} then depends on the expectation of $(\theta_{i\tau}, P_\tau)$, not only for $\tau \in \{0, 1\}$ as in the present example, but also for $\tau > 1$. This enriches the type of strategic uncertainty that is relevant, which can be important for quantitative purposes; but it only reinforces the message we wish to convey here.

Now we go back to the general set up, equation (21). We will show that strategic uncertainty can originate, not only from differential *current* information about the current and future fundamentals, but also from uncertainty of the information that others may receive in the *future*. To isolate the second channel, we shut down the former by restricting $g_2 = 0$. In other words, only future actions of other agents, not current actions of others, matter for an agent's best response. For simplicity, we also assume that there are no idiosyncratic shocks, so that $\theta_{it} = \theta_t$. Then, using the linearity of g and aggregating, we get

$$K_t = g_1 \bar{\mathbb{E}}_t \theta_t + g_3 \bar{\mathbb{E}}_t \theta_{t+1} + g_4 \bar{\mathbb{E}}_t K_{t+1}. \quad (24)$$

Next, let

$$\tilde{\theta}_t \equiv \theta_t + \frac{g_3}{g_1} \theta_{t+1},$$

Clearly, we can then rewrite the above as

$$K_t = g_1 \bar{\mathbb{E}}_t \tilde{\theta}_t + g_4 \bar{\mathbb{E}}_t K_{t+1}. \quad (25)$$

With some abuse of notation, we henceforth redefine θ_t as $\tilde{\theta}_t$. This indicates that the appropriate interpretation of what the “period- t fundamentals” mean in the theory is neither the exogenous payoff relevant variables that happen to get realized in period t nor the exogenous shocks that enter the period- t flow payoff of an agent; rather, it is the collection of all payoff-relevant variables that are relevant for period- t decisions, regardless of the time the values of these variables materialize or appear in the player's flow payoffs.

Notwithstanding these observations, let us now solve for the equilibrium K_t . To this goal, let $z_t^0 \equiv \bar{\mathbb{E}}_t \theta_t$ and, for all $j \geq 1$, let

$$z_t^j \equiv \bar{\mathbb{E}}_t z_{t+1}^{j-1} = \bar{\mathbb{E}}_t \left\{ \bar{\mathbb{E}}_{t+1} \left\{ \dots \left\{ \bar{\mathbb{E}}_{t+j} \left\{ \theta_{t+j} \right\} \dots \right\} \right\} \right\}.$$

Iterating (25) yields

$$K_t = g_1 \sum_{j=0}^{+\infty} (g_4)^j z_t^j. \quad (26)$$

It follows that K_t depends not only on today's forecasts of fundamentals (the forecasts captured in z_t^0) but also on today's forecasts of *tomorrow's* forecasts of fundamentals (the forecasts captured in z_t^1), on today's forecasts of tomorrow's forecasts of forecasts two period ahead (z_t^2), and so on. Uncertainty about the information that others may receive in the *future* can thus serve as an independent source of higher-order uncertainty in the *present*. Nevertheless, there is no need to get lost in the wilderness of higher-order beliefs. If we put aside condition (26) and go back to (24), we can readily see that, conditioning on the equilibrium first-order beliefs of K_t and K_{t+1} , the actual outcome K_t is invariant to the underlying hierarchy of beliefs about the fundamentals. As with our static framework, it therefore remains true that the practical value of incomplete information is to enrich the testable restrictions on the observables of the model.

7.9 Endogenous Information Acquisition

We now discuss a strand of the literature that endogenizes the acquisition of information, or the allocation of attention, within beauty contests. We follow Myatt and Wallace (2012) and Pavan (2015) because their framework is essentially the same as ours. Other notable contributions include Hellwig and Veldkamp (2009), Vives (2013), Chahrour (2014), Colombo, Femminis and Pavan (2014), Llosa and Venkateswaran (2015), Yang (2015), and Denti (2015).

This line of work is closely related to the one on rational inattention, such as Sims (2003) and Mackowiak and Wiederholt (2009). It nevertheless departs from it in that it pays attention to the following fact: economic agents may wish to collect information that is useful for predicting and tracking the actions of others, even if that information is not particularly useful for predicting and tracking the underlying fundamentals.

We use the payoff structure introduced in Section 7, along with the simplifying assumptions introduced in Subsection 7.1 and 7.4. The only novelty is the introduction of an “information technology”, by which we mean the following.

Each agent i has access to $N \in \mathbb{N}$ potential sources of information about θ , the underlying common fundamental, which is drawn from $N(0, \sigma_\theta^2)$. We let $\pi_\theta \equiv \sigma_\theta^{-2}$ denote the precision of the common prior about θ . The information contained in each source $n \in \{1, \dots, N\}$ is given by

$$y_n = \theta + \epsilon_n,$$

where $\epsilon_n \sim N(0, \eta_n^{-1})$ is i.i.d. Normally distributed noise, independent of θ . By paying attention $z^i = (z_n^i)_{n=1}^N \in \mathbb{R}_+^N$ to the available sources, agent $i \in [0, 1]$ then receives private signals $x^i \equiv (x_n^i)_{n=1}^N \in \mathbb{R}^N$ as

$$x_n^i = y_n + \xi_n^i,$$

where $\xi_n^i \sim N(0, (t_n z_n^i)^{-1})$ is i.i.d. Normally distributed noise, independent of θ and $\epsilon \equiv (\epsilon_n)_{n=1}^N$. Myatt and Wallace (2012) and Pavan (2015) interpret the parameter η_n as the *accuracy* of the source (the “sender’s noise”), and the parameter t_n as the *transparency* of the source (how a marginal increase in the attention increases the precision of private signal that each agent ultimately obtains from “listening” to the source).

Agents incur a cost when paying attention to any given source. In particular, payoff are now given by

$$U(k_i, K, \theta) - C(z^i) \tag{27}$$

where U is as in Section 7 and C is an increasing and continuously differentiable function.

To simplify exposition, we normalize $U_{kk} = -1$ and restrict attention to the case in which attention is perfectly substitutable across different sources in the following sense.⁵⁸

Assumption 8. $C(z^i) = c\left(\sum_{n=1}^N z_n^i\right)$, where $c(\cdot)$ is strictly increasing, strictly convex, and differentiable.

One can then establish the following result regarding the equilibrium allocation of attention.

Proposition 24. *There exists a threshold $\Delta > 0$ such that, in the unique symmetric equilibrium,*

$$z_n = \frac{\eta_n \max\left\{\left(\Delta - \frac{1}{\sqrt{t_n}}\right), 0\right\}}{(1 - \alpha) \sqrt{t_n}}.$$

Proposition 24 follows from Proposition 2 in Myatt and Wallace (2012), as well as from Proposition 1 and Corollary 1 in Pavan (2015). Transparency (t_n) determines whether a source of information will receive any attention and also how much attention it receives. Accuracy (η_n) instead only influences how much attention each information source receives, conditional on that source being used. The intuition why accuracy does not play a role in determining whether an information source receives positive attention is as follows. When z_n is small, the total amount of noise from information source n is dominated by the receiver noise, ξ_n^i .

⁵⁸See Myatt and Wallace (2012) and Pavan (2015) for the more general case.

As a result, when thinking about which information source to pay attention, a player starts with the most transparent source.

Let us now define the *publicity* of an information source as the correlation of the noises in signals that two different agents receive from this source.⁵⁹

$$\rho_n \equiv \text{corr}(x_n^i, x_n^j | \theta) = \frac{z_n t_n}{z_n t_n + \eta_n}$$

We can now state the following result, which is due to Myatt and Wallace (2012) and regards the interaction of the coordination motive with the acquisition of information.

Proposition 25. *As the degree of strategic complementarity α rises, attention moves away from more private signals and towards more public signals: there is a $\hat{\rho}$ such that the attention paid to source n is locally increasing in α if and only if $\rho_n > \hat{\rho}$.*

Intuitively, public signals act as effective focal points for players' coordination. As the desire for coordination strengthens, agents pay more attention to such signals.

Pavan (2015) investigates the question of whether the equilibrium allocation of attention is efficient from the perspective of a social planner that is interested in maximizing ex ante welfare. This complements the earlier work by Angeletos and Pavan (2007), which studied the efficiency of the equilibrium use of information, taking the latter as exogenously given. We refer the reader to these works for an analysis of these normative questions; and to Angeletos and La'O (2012), Llosa and Venkateswaran (2015), and Paciello and Wiederholt (2014) for applications related to business cycles and macroeconomic policy.

We close this subsection by circling back to the connection between the present framework and the rational inattention literature. The applications of this literature that relate to monetary policy are briefly noted in Subsection 8.5; for a more extensive review, see Sims (2011). Here, we instead focus on two conceptual issues: the sense in which the framework we study can nest the works of Sims (2003) and Mackowiak and Wiederholt (2009); and the sense in which it allows for an important departure from these works.

Let $x^i = (x_1^i, \dots, x_N^i)$ and $y = (y_1, \dots, y_N)$. Next, assume that instead of facing a cost for paying attention to different sources of information, the agents face a "capacity constraint" of the following form:

$$\Gamma(x^i, y) \leq \bar{\Gamma}, \tag{28}$$

for some function Γ . The agent's attention-allocation problem is then to choose z^i so as to maximize

$$\mathbb{E}[U(k_i, K, \theta) - \lambda \Gamma(x^i, y)]$$

where $\lambda \geq 0$ is the Lagrange multiplier associated with the capacity constraint. This is obviously nested in our present framework by letting $C(z^i) = \lambda \Gamma(x^i, y)$, no matter what the function Γ is.

To nest the information-acquisition problems studied by Sims (2003) and much of the related rational-inattention literature, we need to make two assumptions. The first is that Γ is the function that measures the mutual information between two random variables; in a Gaussian context like the one we are studying here, this means letting $\Gamma(x^i, y) = \frac{1}{2} \log [\det[V(x^i)] / \det[V(x^i|y)]]$, where $\det[\cdot]$ denotes the determinant of a matrix, $V(x^i)$ denotes the covariance matrix of the random variable x^i , and $\text{Var}(x^i|y)$ denotes the covariance matrix of x^i conditional on y .

⁵⁹This is related to the notion of the *commonality* of information used in Angeletos and Pavan (2007). See that paper for how this notion helps understand both the positive and the normative properties of the class of games we have been studying.

The second assumption is that there is a single and perfect source of information about the fundamental: $y = \theta$. This means that the attention-allocation problem reduces to one of trying to track *only* the fundamental, as opposed to various sources of information. This is not a serious limitation in a single-agent, decision-theoretic context. In a context with strategic, or general-equilibrium interactions, however, this restriction means there is no more a useful distinction between tracking the fundamental and tracking the actions of others: under this restriction, all the noise in the private signals of the agents becomes *purely* idiosyncratic, implying that the equilibrium value of \mathbf{K} is pinned down by the fundamental θ . In equilibrium, predicting \mathbf{K} is therefore the same as predicting θ .

It is this last property that is more relevant for our purposes. If none of the available sources of information is perfect, the underlying noise in the sources becomes correlated noise in the private signals received by different agents. As noted before, such correlated noise helps disentangle the expectations of endogenous outcomes from the expectations of exogenous fundamentals, for it accommodates aggregate shocks in the gap between first- and higher-order beliefs. The more flexible forms of information acquisition proposed by Myatt and Wallace (2012) and Pavan (2015) are therefore best suited for understanding how strategic, or general-equilibrium interaction may *endogenously* lead to more correlated noise.⁶⁰

On the other hand, an important feature of Mackowiak and Wiederholt (2009) that is missing from both Myatt and Wallace (2012) and Pavan (2015) is the trade off in allocating attention to idiosyncratic versus aggregate shocks. Merging the two frameworks is an interesting venue for future research.⁶¹

7.10 A Note on Solution Methods

Dynamic models with time-varying fundamentals and dispersed information can be hard to solve. Townsend (1983) first suggested that a finite state space solution for the equilibrium dynamics need not exist, reflecting an infinite-regress problem in forecasting the forecast of others.

To circumvent the problem, Townsend (1983) proposed the following short-cut: let the exogenous state of nature become common knowledge after T periods, where $T \geq 1$ is potentially large but finite. Then, for a class of linear models, one can guess and verify the existence of an equilibrium in which aggregate outcomes can be expressed as linear functions of the history of shocks over the last T periods, along with the relevant commonly-known state variables prior to that period.⁶²

It is worth noting here that Townsend (1983) was concerned with settings in which there were no strategic complementarity in actions and higher-order beliefs entered decisions only through a signal-extraction problem.⁶³ The method, however, readily extends to the class of environments we are interested in. For recent

⁶⁰We refer the reader also to Denti (2015) for an analysis that covers games with either a finite number or an infinity of players. See in particular the discussion in Section 9.4 of that paper regarding the significance of allowing for noise in the underlying sources of information if one wishes to obtain non-fundamental volatility in games with an infinity of players, like those studied in the present chapter.

⁶¹Myatt and Wallace (2012) and Pavan (2015) restrict attention to settings where, as in the analysis above, there is a single fundamental that is common to all agents. Mackowiak and Wiederholt (2009) study a setting in which there are two types of fundamentals: an aggregate monetary shock, and a purely idiosyncratic TFP shock. They study the optimal allocation of attention across these two types of fundamentals, under the restriction that each agent observes a pair of independent noisy private signals for the two fundamentals (as opposed to observing a joint signal of both fundamentals). Denti (2015) considers a more general structure that allows for arbitrary correlation among the fundamentals of different players, but does not focus on the aforementioned trade off and, unlike Myatt and Wallace (2012) and Pavan (2015), rules out the noise in the underlying sources of information.

⁶²For instance, if we consider Example 1 from Subsection 7.8, this would mean expressing K_t as a linear function of θ_{t-T-1} and of $(v_{t-T}, \dots, v_{t-1}, v_t)$. More generally, one would have to include the values of any endogenous state variable (such as capital), as well as any aggregate shock to information (such as the noise in public signals), not just fundamentals.

⁶³More specifically, Townsend (1983) used two examples. In the first example, there are two industries, each trying to extract a signal about the underlying fundamental from the price of the other industry. In the second, there is a continuum of industries, each trying to extract a signal of the underlying fundamental from aggregate outcomes. Townsend suggested that an infinite-

applications, see Bacchetta and Van Wincoop (2006), Hellwig (2005), Hellwig and Venkateswaran (2009), and Lorenzoni (2009).

It is tempting to interpret Townsend’s short cut as an approximate solution for the situation in which the past shocks never become common knowledge. However, the conditions under which this is true are not completely understood. On the one hand, Sargent (1991) and Kassa (2000) verify that this is indeed the case for the leading example in Townsend (1983).⁶⁴ On the other hand, Rondina and Walker (2014) present a version of Singleton’s (1987) asset-pricing model in which the equilibrium is perfectly revealing when the lag T is finite (no matter how large), but not when T is infinite.

An alternative approximation method involves a truncation of the hierarchy of beliefs, rather than a truncation of the history of shocks. Nimark (2011) pursues this method. He also shows that the approximation error vanishes as the truncation gets larger for a class of dynamic linear models that exhibit a similar “stability” property as the static games studied in Weinstein and Yildiz (2007b), namely the property that the impact of higher-order beliefs on actions vanishes at an exponential rate as the order increases. Notwithstanding the theoretical appeal of this method, its computational efficiency and its comparison to Townsend’s approach remain unclear.

Barring the above kinds of shortcuts and approximations, the literature has long struggled to obtain *exact* finite-state solutions for models with infinite horizons and perpetually dispersed information. A few papers have considered special examples in which it is possible to guess and verify the existence of such a solution using the Kalman filter; see Pearlman and Sargent (2005), Singleton (1987), and Woodford (2003). Another line of work has obtained analytic characterizations of the equilibrium by transforming the problem into the frequency domain; see Futia (1981), Kasa (2000), Kasa, Walker and Whiteman (2007), Acharya (2013), and Rondina and Walker (2014). These papers have therefore been able to provide important insights into the dynamics of higher-order beliefs, often in settings with endogenous signals. Yet, the results of these papers appear to rest on special settings, which has limited their applicability and has not helped advance quantitative work.

A significant breakthrough was recently made by Huo and Takayama (2015a,b). The authors show that a finite-state space solution is possible for a large class of linear models insofar as the observed signals are exogenous and the underlying shocks (fundamentals and noises) follow ARMA processes. In particular, the solution itself has an ARMA structure, although typically of a higher order than those of the underlying shocks. The method may still be computationally intensive if there are multiple endogenous state variables and multiple shocks, but it is flexible and powerful.⁶⁵

A more radical approach is proposed in Angeletos, Collard, and Dellas (2015). The key idea is to use a heterogeneous-prior specification that helps proxy the dynamics of higher-order beliefs that are induced by incomplete information, while bypassing all the computational difficulties. Each agent is assumed to believe that there is a bias in the signals of others. Variation in this perceived bias moves higher-order beliefs without moving first-order beliefs. The noise is taken to zero, shutting down noisy learning and forcing a

regress problem exists in both examples, because an agent’s signal extraction depends on the actions of others, which in turn depend on their signal extraction, and so on. However, Sargent (1991) points out that the apparent infinite-regress problem vanishes in the first example: the solution turns out to be perfectly revealing due the limited dimensionality of the underlying uncertainty. The problem remains in the second example, which is the one that Townsend (1983) actually concentrates on, thanks to the fact that available signals are contaminated by measurement error.

⁶⁴That is, the example with a continuum of industries studied in Section VIII of Townsend (1983).

⁶⁵The method also rules out signals of endogenous outcomes. This need not be a serious limitation for certain quantitative purposes, because the informational content of endogenous signals can be mimicked by appropriately chosen exogenous signals. Furthermore, an extension of the method can be used to obtain an approximate solution to models with endogenous signals; see Huo and Takayama (2015b) for details.

low-dimensional representation of the stochastic process of the entire belief hierarchy. All in all, Angeletos, Collard, and Dellas (2015) are able to augment a large class of DSGE models with rich yet highly tractable higher-order belief dynamics: the belief-augmented models can be solved and structurally estimated with the same ease and computational efficiency as standard, representative-agent models.⁶⁶

We review two related applications of the last two methods in Subsection 8.7 and use them to argue that the waves of optimism and pessimism that are rationalized by higher-order belief dynamics can offer a quantitatively potent explanation of business-cycle phenomena. Additional works that seek to push the quantitative frontier of the literature include Melosi (2014), who estimates a version of Woodford (2003); Mankiw and Reis (2006), Kiley (2007), and Reis (2009), who estimate DSGE models with sticky information; Mackowiak and Wiederholt (2015), who study the quantitative performance of a DSGE model in which both firms and consumers are rationally inattentive; and David, Hopenhayn, and Venkateswaran (2014), who use firm-level data to gauge the cross-sectional misallocation of inputs caused by informational frictions.

8 Beauty Contests: Applications

The analysis of the previous section has indicated that incomplete information has two key positive implications. First, it dampens the response of equilibrium actions to changes in the underlying fundamentals. Second, it accommodates forces akin to animal spirits along the unique equilibrium. In this section we explore what these properties mean within the context of specific applications; how they can help accommodate important facts that are inconsistent with certain workhorse macroeconomic models; and how they can inform policy. We also make a digression to discuss the connection between the mechanisms we study here and those in the complementary literature on sticky information (Mankiw and Reis, 2002, Reis, 2006) and rational inattention (Sims, 2003, Mackowiak and Wiederholt, 2009).

8.1 Real Rigidity and Technology Shocks

In this subsection we discuss how frictions in coordination can be the source of real rigidity at the macro level⁶⁷ and how this in turn can help reconcile the RBC paradigm with Gali (1999).

Gali (1999) argues that a key implication of the RBC model regarding the response of employment to technology shocks is grossly inconsistent with the data. More specifically, Gali uses a structural VAR method to estimate the impulse response functions of macroeconomic aggregates to an identified technology shock in US data: he run a VAR on key macroeconomic variables and identified the technology shock as the only shock that drives labor productivity in the long run.⁶⁸ He then compares the impulse response functions obtained in the data with those predicted by the model.

Figure 5, which is a replica of Figure 4 in Gali (1999), depicts the impulse responses of labor productivity, output, and employment obtained in the data. As it is evident in the first panel, the identified shock in the

⁶⁶The potential downside is that this method gives the researcher the freedom to specify higher-order beliefs “at will” (i.e., without the restrictions imposed by the combination of the common-prior assumption and specific information structures). How this freedom is used in practice is a delicate balancing act. Angeletos, Collard, and Dellas (2015) propose that data on forecasts can offer much of the needed discipline, but do not pursue this idea in detail.

⁶⁷By “real rigidity” we refer to inertia in the response of real quantities, such as employment and output, to shocks in fundamentals, such as preferences and technology. See Angeletos, La’O and Iovino (2016b) for a complementary discussion of the sense in which incomplete information can be the source of either real or nominal rigidity (or of both) in the related applied literature.

⁶⁸Gali (1999) run a baseline bivariate VAR on hours and labor productivity; the same finding obtains in a five-variable VAR on hours, labor productivity, money supply, interest rates and price level. Figure 5 below is based on the five-variable specification.

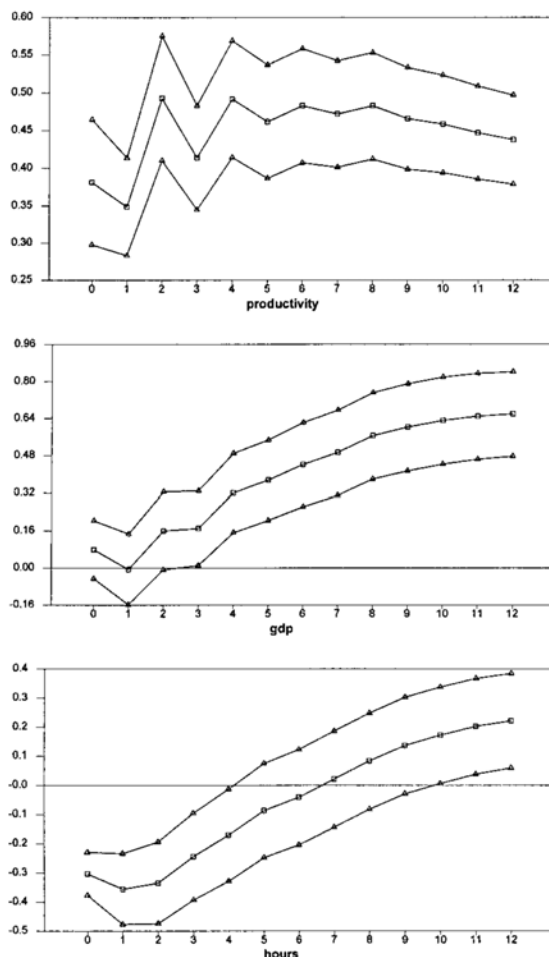


Figure 5: Estimated Impulse Responses to a Technology Shock, from Gali (1999). Point estimates and ± 2 standard error confidence intervals; variables given in percentages.

data triggers an immediate and persistent increase in productivity. This makes the identified shock in the data comparable to the theoretical technology shock in the RBC model. And yet, as seen in the last panel, the identified technology shock leads to a decline in hours/employment in the short run, which is exactly the opposite of what the RBC model predicts.

Gali (1999) proceeds to argue that this fact is consistent with the New-Keynesian framework. Suppose that prices are sticky. Suppose further that monetary policy fails to “accommodate” the technology shock, that is, it responds to it by contracting aggregate demand *relative* to the level that would have obtained under flexible prices. Finally, suppose that this contractionary effect is strong enough. Then, aggregate output may increase less than productivity. By the same token, equilibrium employment may fall under sticky prices, even though it would have increased under flexible prices.

It remains debatable whether this is a satisfactory theoretical explanation, as well as whether the fact itself is robust to alternative empirical methodologies: see Gali and Rabanal (2005) for complementary evidence, Christiano, Eichenbaum and Vigfusson (2004) and McGrattan (2004) for criticisms. Here, we take the fact for granted and proceed to make two points: one regarding its structural interpretation and one regarding its policy implications.

The first is that, while the fact is inconsistent with the standard, complete-information, RBC model, it is not necessarily inconsistent with an incomplete-information extension of it. The reason is the dampening effect documented in Proposition 18: lack of common knowledge may cause significant inertia in the response of aggregate output, perhaps even a negative initial response in employment.

The second is that this dampening can be a reflection of the efficient use of the information that is dispersed in the economy.⁶⁹ This contrasts the New-Keynesian explanation, which prescribes the observed fact to a failure of the monetary authority to replicate flexible-price allocations. To put it differently, the success of the New-Keynesian model in matching the fact in Gali (1999) is, in the context of that model, a manifestation of the failure of the policy maker to do the right thing. By contrast, augmenting the RBC model with a friction in coordination can help rationalize the fact without any policy failure.

These points are formalized in Angeletos and La’O (2010) by studying an incomplete-information version of the neoclassical economy we introduced in Subsection 2.2. This economy is essentially the same as the textbook RBC model, except for three modifications. First, capital is assumed away in order to simplify the analysis. Second, product differentiation is introduced, giving rise to strategic complementarity. Lastly, information is incomplete, giving rise to imperfect coordination.

More specifically, the economy is split into a large number of islands, indexed by $i \in [0, 1]$. Each island is populated by a competitive firm and a competitive worker, who produce a differentiated, island-specific, intermediate good. The goods produced by different islands enter the production of a single, final, consumption good. This introduces strategic complementarity across islands: the optimal production on each island depends on other islands’ production. Incomplete information is introduced by assuming the following: when the firm and the worker of any given island decide on the local employment and output, they are uncertain about the productivity, employment and output of other islands. Despite the incompleteness of information, incomplete risk-sharing is assumed away by letting workers pool their income at the end of each period: everybody belongs to the same “big family”, whose income and consumption is fully diversified against the idiosyncratic information and choices of the various family-members. Finally, it is assumed that the local productivity of each island is perfectly observed by the firm and the worker of the island,⁷⁰ which guarantees that any other information is relevant only insofar as it helps predict the production levels of other islands.

To facilitate an exact mapping to the class of beauty-contest games we studied in the previous section, information is assumed to be Gaussian and preferences and technologies are assumed to take a power-form specification. Without further loss of generality, let us impose a linear technology: $y_{it} = a_{it} + n_{it}$, where y_{it} is the (log)level of the output produced in island i , a_{it} is the (log)level of the exogenous local productivity, and n_{it} is the (log)level of local employment. Angeletos and La’O (2010) show that the general equilibrium of the model reduces to the solution of the following fixed-point relation:

$$y_{it} = (1 - \alpha)\chi a_{i,t} + \alpha \mathbb{E}_{it} Y_t \tag{29}$$

where Y_i is the (log)level of aggregate output, and $\chi > 0$ and $\alpha < 1$ are constants that depend on the underlying preference and technology parameters. With the output levels obtained from the solution of (29), local and aggregate employment are given by, respectively, $n_{it} = y_{it} - a_{it}$ and $N_t = Y_t - A_t$.

Clearly, the fixed-point relation in (29) is mathematically equivalent to the best-response condition encountered in the previous section. It follows that the insights developed there are readily applicable to the present context. We now elaborate what this means for the empirical properties of the RBC framework.

⁶⁹By “efficient” we mean the notion of constrained efficiency discussed in Section 9.

⁷⁰This is the analogue of the assumption we made earlier on in Proposition 16.

Assume that local productivity is given by the sum of an aggregate and an idiosyncratic component: $a_{it} = A_t + \xi_{it}$, where A_t is the aggregate technology shock and ξ_{it} is an idiosyncratic one. The latter is i.i.d. across both i and t (for simplicity), whereas the former follows a random walk:

$$A_t = A_{t-1} + u_t,$$

where u_t is the period- t innovation.

When information is complete (equivalently, coordination is perfect), we have that $\mathbb{E}_{it} Y_t = Y_t$ for all i . Substituting this into the above fixed-point relation and aggregating across i yields the following:

Proposition 26. *Under complete information, aggregate output and employment are given by, respectively,*

$$Y_t = \chi A_t \quad \text{and} \quad N_t = (\chi - 1) A_t.$$

It follows that equilibrium outcomes are pinned down by the technology shock; that χ and $\chi - 1$ identify the general-equilibrium elasticities of, respectively, aggregate output and aggregate employment to technology shocks, as predicted by the complete-information RBC model; and that, holding χ constant, the scalar α is irrelevant for macroeconomic outcomes.

When instead information is incomplete (equivalently, coordination is imperfect), equilibrium outcomes depend, in effect, on the entire hierarchy of beliefs about A_t . The scalar α then starts playing a crucial role for the observables of the model, because it determines the relative importance of higher-order beliefs.

Proposition 27. *Under incomplete information, aggregate output is given by*

$$Y_t = \chi \sum_{h=0}^{\infty} (1 - \alpha) \alpha^h \bar{\mathbb{E}}_t^h [A_t] \tag{30}$$

How α depends on the underlying primitives is discussed in detail in Angeletos and La'O (2010). There are two opposing mechanisms: a “demand-side” effect that contributes towards strategic complementarity ($\alpha > 0$) and a “supply-side” effect contributing towards strategic substitutability ($\alpha < 0$).

Let us explain. When an island expects aggregate output to go up, the demand for the local good, its relative price, and the local real wage are all expected to go up as well. This effect motivates the local worker to work more and the local firm to produce more. This is the “demand-side” effect, which induces local output to increase with expectations of aggregate output. The opposing “supply-side” effect originates from an income effect: when an island expects aggregate output to go up, income is also expected to go up, which tends to discourage labor supply and production. Whether α is positive or negative depends on which of the aforementioned two effects dominates. For the remainder of our analysis, we impose $\alpha > 0$.

If we vary α holding χ constant,⁷¹ we vary the incomplete-information outcomes without varying their complete-information counterparts. We can thus think of α as a “sufficient statistic” of the primitives of the environment that regulate the macroeconomic effects of strategic uncertainty, and thereby the observable aggregate implications of the incomplete-information model, holding constant the observable aggregate implications of the standard, complete-information macroeconomic model.⁷²

⁷¹By varying the underlying preference and technology parameters one can match any pair of values (χ, α) , which explains why one can indeed vary α holding χ constant.

⁷²Note, however, that the combination of α and χ also matters for the observable implications of both models in the cross section: in the economy under consideration, the product $(1 - \alpha)\chi$ is identified by the micro-elasticity of the response of local output to island-specific productivity shocks. This is indicative of how the combination of micro and macro data could help discipline the mechanisms we study in this section, in a manner that complements the approach reviewed in Subsection 8.6.

This last observation seems particularly useful if we adopt a “relaxed” interpretation of the model at hand as a representative of a broader class of RBC models in which complementarity may originate, not only in the aforementioned demand-side effect, but also in “financial multipliers” or other kinds of market interactions that make the fate of any given firm and consumer sensitive to aggregate economic activity. For the purposes of the subsequent discussion, we invite the reader to entertain such a flexible interpretation of the strategic complementarity and the associated coefficient α .

Let us now go back to the original motivation of this subsection, namely the negative employment response to technology shocks. As already noted, any empirically plausible calibration of the standard, complete-information, RBC model makes the opposite prediction. In the context of the present framework (which assumes away capital), this prediction is equivalent to imposing $\chi > 1$, but leaves α unrestricted. But once information is incomplete, the predicted employment response crucially depends on α and can turn negative irrespective of the value of χ . That is, we can reconcile the RBC framework with Gali’s findings by removing common knowledge of the technology shock and letting α be large enough.

Angeletos and La’O (2010) demonstrate the above possibility with numerical examples that allow learning to be slow and the incompleteness of information to persist indefinitely. Here, we illustrate the key idea with an example that forces the incompleteness of information to be transient, lasting for only one period.

Assumption 9. *In each period t , the previous-period aggregate technology shock, A_{t-1} , becomes commonly known. Furthermore, for every i , any information that island i has about the current-period shock, A_t , is summarized in a private signal of the form*

$$x_{it} = A_t + \epsilon_{it},$$

where $\epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is i.i.d. across (i, t) and independent of A_t for all t .⁷³

The first part of this assumption rules out noisy dynamic learning and guarantees that the economy reduces, in effect, to a repeated version of the static beauty-contest game studied in Section 7: in each period, it is as if the islands play a beauty-contest game in which the fundamental $\theta = A_t$ and the common prior about it is a Normal distribution with mean $\mu = A_{t-1}$. The second part of the assumption abstracts from public signals and any other form of correlated noise, thus isolating the mechanism we care about in this section, namely the role of strategic uncertainty in dampening the response to fundamentals. The alternative function of accommodating forces akin to animal spirits is the subject of Subsection 8.7.

Following an approach similar to that used in the proof of Proposition 18, we can establish the following.

Proposition 28. *Under Assumption 9, equilibrium output is given by*

$$Y_t = \chi A_{t-1} + \phi u_t,$$

for some positive scalars ϕ that satisfy $0 < \phi < \chi$. Furthermore, for any given $\chi > 0$ and $\sigma_\epsilon > 0$, ϕ converges to 0 from above as α converges to 1 from below.

This result illustrates two properties that go beyond the specific information structure assumed above. The first property is that the *long-run* effect of a technology shock under incomplete information remains the same as the one under complete information (which is χ). The second property is that the short-run effect

⁷³Because local productivity is itself a private (local) signal of aggregate productivity, x_{it} is meant to be a sufficient statistic for both the information contained in local productivity and any other private information that island i has about A_t .

(given by ϕ) is smaller and can even be arbitrarily close to zero if α is large enough. The first property holds because the shock becomes common knowledge in the long run (where “long run” means “next periods” in the above example and “asymptotically” in the example studied below). The second property holds because, as long as the fundamental is not common knowledge, the following is true: as h goes to ∞ , the h -th order belief of A_t converges to the common prior, no matter what the current innovation u_t is; letting α be high enough therefore guarantees that the equilibrium belief of Y_t does not move much with u_t , and that actual output does not move much either.

To derive the aggregate employment response, recall that $N_t = Y_t - A_t$. It follows that the effect of the technology shock on aggregate employment is given by $\phi - 1$, which is negative for high enough α .

Corollary. *Suppose $\chi > 1$ and take for granted Gali’s finding that the short-run response of employment to technology shocks is negative. This finding rejects the frictionless version of the RBC model, but does not reject its incomplete-information extension insofar as α is large enough.*

The above result has so far been established only under a narrow interpretation of what the “short run” is: the response of employment turns positive after a one-period lag. However, this last property is an artifact of the simplifying assumption that the technology shock becomes common knowledge after a lag of only one period: if we allow the lack of common knowledge to persist for more periods, then we can accommodate a negative employment response for more periods. See Angeletos and La’O (2010) for details.

Angeletos and La’O (2010) further show that the equilibrium response of the economy to technology shocks is constrained efficient in the sense that there is no allocation that can improve upon the equilibrium without violating either resource feasibility or the informational constraints of the economy.⁷⁴ It follows that, unlike the New-Keynesian interpretation, the Gali fact is not more a symptom of suboptimality (or constraints) in monetary policy.

To recap, insofar as there is sufficient lack of common knowledge and sufficient strategic complementarity, the RBC framework can be reconciled with the Gali fact. This, however, pegs the question of how strong the relevant strategic complementarity is. In Angeletos and La’O (2010), the only source of strategic complementarity is that induced by the Dixit-Stiglitz preference specification. For conventional parameterizations, this implies a rather low value for α , and therefore also a rather low role for higher-order uncertainty. Financial frictions and feedback effects as those in Kiyotaki and Moore (1997) can “boost” the degree of complementarity that is present in the RBC framework, while at the same time deliver different normative conclusions. A promising direction for future research is therefore to introduce incomplete information in the models of the growing literature on financial frictions that has been spurred by the recent crisis.

We conclude by discussing a few complementary papers which have used informational frictions to generate real rigidity in different contexts, most notably in the canonical consumption-saving problem. Sims (2003, 2006), Luo (2008) and Tutino (2013) show how rational inattention can induce inertia in the response of consumption to income shocks. Luo, Nie and Young (2015) show how slow learning can help generate excess smoothness in durable and nondurable consumption, bringing the model closer to the data. Alvarez, Guiso and Lippi (2011) study the interaction of transaction and observation costs for liquidity and consumption. The results in all these papers have a similar flavor to those derived here. But there is a key difference. These papers feature no strategic, or general-equilibrium, interaction. As a result, the rigidity they document is a rigidity at the *micro* level: it dampens the response of individual outcomes to individual fundamentals. It also requires enough noise in the observation of such fundamentals; that is, first-order beliefs of fundamentals

⁷⁴See Section 9 for a definition and a discussion of this kind of efficiency concept.

must themselves be rigid. By contrast, the rigidity we have documented obtains at the *macro* level; it rests on lack of common knowledge (i.e., on higher-order beliefs being rigid) as opposed to individual uncertainty; and it can thus be consistent with considerable flexibility at the micro level.

8.2 General-Equilibrium Dampening, and Rigidity vs Overshooting

Angeletos and Lian (2016b) push the aforementioned insights further, showing (i) that incomplete information is equivalent to a certain relaxation of the solution concept and (ii) that it can dampen the general-equilibrium effects of macroeconomic models, while holding constant the underlying micro elasticities. This helps reduce the disconnect between recent empirical works such as Mian et al. (2013) and Mian and Sufi (2014), which identify the cross-sectional effects of regional shocks, and the key questions of interest, namely the macroeconomic effects of aggregate shocks.

The results in Angeletos and Lian (2016b) also clarify that the same mechanism—the dampening of general-equilibrium effects—could mean either rigidity or overshooting. If the general-equilibrium effect works in the same direction as the corresponding partial-equilibrium effect, meaning that the macro elasticity is higher than the micro one under complete information, then the introduction of incomplete information contributes to rigidity. If, instead, the general-equilibrium works in the opposite direction, meaning that the macro elasticity is smaller than the micro one under complete information, then the introduction of incomplete information contributes to overshooting.

The first scenario corresponds to settings in which the general-equilibrium effect is akin to strategic complementarity. The second scenario corresponds to settings in which the general-equilibrium effect is akin to strategic substitutability. The application by Angeletos and La’O (2010) that we discussed above and the monetary applications by Woodford (2003) and others that we review in the sequel are, in effect, examples of the first case. Angeletos and Lian (2016b) contain examples of the second case.

An interesting example of the second case is also Venkateswaran (2014). This paper considers an incomplete-information version of the Diamond-Mortensen-Pissarides model and shows that the incompleteness of information increases the volatility in aggregate unemployment, helping reconcile the model with the business-cycle data.

8.3 Nominal Rigidity and Monetary Shocks

In this subsection, we review Woodford (2003), an influential contribution that illustrated how incomplete information offers a potent substitute to more conventional formalizations of nominal rigidity.

Before reviewing Woodford’s contribution, it is worth to recall some empirical evidence on the macroeconomic effects of monetary shocks. Christiano et al. (1999) and Christiano et al. (2005) use a structural VAR method to identify the impulse responses of macroeconomic variables to monetary shocks. Two of their key findings are the following: (i) Output, consumption, and investment respond in a hump-shaped fashion, peaking after about one to one-and-a-half years and returning to pre-shock levels after about three years; and (ii) inflation responds also in a hump-shaped fashion, peaking after about two years.

These findings are illustrated in Figure 6, which is borrowed from Figure 1 of Christiano et al. (2005). The figure shows the estimated responses of output and inflation to an identified monetary shock.⁷⁵ These

⁷⁵The estimated responses of consumption and investment are omitted because we focus on a simple model without capital and consider only the response of output and inflation. Units on the horizontal axis are quarters; on the vertical, it is deviation from the unshocked path (annualized percentage points for inflation, percentages for output). Grey areas indicate 95 percent confidence intervals.




Figure 6: Impulse responses of inflation and output to monetary shocks. Solid lines with plus signs are VAR-based impulse responses. Grey areas are 95 percent confidence intervals around the VAR-based estimates. Units on the horizontal axis are quarters. An asterisk indicates the period of the policy shock. The vertical axis units are deviations from the un-shocked path. In the case of inflation, the unit is in annualized percentage points (APR). In the case of output, the unit is in percentage. Solid lines are the impulse responses generated by medium-scale new-Keynesian DSGE model (Christiano et al. , 2005).

responses provide support for the New-Keynesian framework insofar as they indicate that monetary shocks have large and persistent real effects. However, they also raise a challenge, which we turn to next.

The baseline version of New-Keynesian Philips Curve, which is at the core of the New-Keynesian model, takes the following form:

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t \pi_{t+1},$$

where π_t denotes inflation and y_t the output gap, κ is a function of the Calvo probability of resetting prices (along with preference and technology parameters), and β is the discount factor.⁷⁶ Iterating the above condition gives current inflation as the best forecast of the present value of output gaps:

$$\pi_t = \kappa \mathbb{E}_t \left[\sum_{h=0}^{\infty} \beta^h y_{t+h} \right].$$

This implies that inflation must lead (or predict) output in response to monetary shocks: if the response of output to a monetary shock is hump-shaped as in the right panel of Figure 6, then the peak in inflation has to come *before* the peak in output. But this is the opposite of what seen in the left panel of Figure 6.

To overcome this failure, quantitative DSGE models have augmented the baseline New-Keynesian model with a number of ad hoc features, such as price indexation and certain kinds of adjustment costs, whose micro-foundations and immunity to the Lucas critique remain debatable. By contrast, Woodford (2003) shows that incomplete information, couple with strategic complementarity in price-setting decisions, can, not only substitute for the Calvo friction as a source of nominal rigidity, but also naturally produce the empirical pattern seen in Figure 6.

Let us elaborate. Woodford (2003) considers a monetary economy similar to the one we sketched in Subsection 2.2. Unlike the New-Keynesian model, there is no Calvo friction: firms are free to reset their prices period by period. Instead, nominal rigidity originates from the incompleteness of information about the underlying monetary shock. In particular, the (log) price set by the typical firm in period t is given by

$$p_{it} = \mathbb{E}_{it}[(1 - \alpha)\theta_t + \alpha P_t], \quad (31)$$

⁷⁶Strictly speaking, y_t measures the log deviation of the price-to-cost markup that obtains under sticky prices from the one that would have obtained if prices had been flexible. It is an open question what is the best empirical counterpart of this theoretical object. For simplicity, in the present discussion we interpret y_t as the output gap.

where θ_t denotes (log) Nominal GDP, P_t denotes the (log) price level, and α is the degree of strategic complementarity in the firms’ price-setting decisions. Nominal GDP is treated as an exogenous process. In line with the data, its growth rate is assumed to follow an AR(1) process:

$$\Delta\theta_t = \rho\Delta\theta_{t-1} + v_t,$$

where $\Delta\theta_t \equiv \theta_t - \theta_{t-1}$ is the growth of Nominal GDP, $\rho \in [0, 1)$ is the persistence, and v_t is the innovation. Finally, the information received by the firm in any given period t is assumed to be a private signal x_{it} of the form

$$x_{it} = \theta_t + \epsilon_{it},$$

where $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ is noise, i.i.d. across (i, t) , and independent of θ_τ for all τ .

Woodford (2003) motivates the above information structure on the basis of rational inattention as in Sims (2003): the noise ϵ_{it} is interpreted as the byproduct of a cognitive friction. This helps bypass one of the critiques of the first generation of models that sought to attribute monetary non-neutrality to informational frictions (Lucas, 1972, 1973; Barro, 1976, 1977, 1978). That earlier literature was based on the assumption that there was no readily available information about the current monetary policy and the current aggregate price level. At face value, this assumption seems unrealistic. However, if agents face cognitive constraints in their ability to attend to and digest the available information, they may well act *as if* that information were unavailable: rational inattention is akin to adding noise in the observation of the underlying shocks.

In the absence of strategic complementarity, the above point would have provided merely a new rationale for the type of informational frictions that were assumed by Lucas and Barro. The hallmark of Woodford’s contribution, however, is the interaction between the informational friction and the strategic complementarity in pricing decisions, and the associated role of higher-order beliefs.

As we show in the next subsection, the setting studied in Lucas (1972) is akin to imposing the restriction $\alpha = 0$ (no strategic complementarity) and thereby shutting down the role of higher-order beliefs. In a nutshell, the theoretical mechanism in Lucas (1972) had to do only with the response of first-order beliefs—which is also why Barro went after measuring “unanticipated changes” in monetary policy.

By contrast, Woodford (2003) shifts the focus to strategic complementarity and higher-order beliefs. As explained in the previous section, the response of higher-order beliefs to innovations in fundamentals is both weaker and more sluggish than that of lower-order beliefs. It follows that, holding constant the precision of the available information and the speed of learning, a stronger complementarity translates into a more muted and more sluggish response of equilibrium prices to the underlying monetary shock—and hence also to larger and more persistent real effects.

Putting the above observations together, Woodford’s first contribution is to explain why there can be significant nominal rigidity at the aggregate level even if (i) there is a lot of readily available information about the underlying monetary shocks; and (ii) each firm alone is only modestly unaware of, or confused about, the underlying monetary shocks. The cognitive friction explains why each firm may be confused in the first place; the strategic complementarity explains why there can be significant inertia at the aggregate level even if the aforementioned confusion is modest.

This kind of amplification effect distinguishes Woodford’s contribution, not only from the earlier literature by Lucas and Barro, but also from some more recent literature that has proposed “observation costs” as a source of nominal rigidity but has abstracted from strategic interactions. Consider, in particular, Alvarez and Lippi (2014) and Alvarez, Lippi and Paciello (2015). These papers use models in which firms update

their information sets infrequently because they have to pay a fixed cost whenever they do so. Because the updating is asynchronized, this implies that at any given point of time, firms are differentially informed about the underlying shocks, opening the door to higher-order uncertainty. However, these papers abstract from strategic complementarity in pricing decisions, thus ultimately shutting down the effects of higher-order beliefs.⁷⁷ It follows that the macro-level nominal rigidity that is formalized and quantified in those papers is tied to micro-level rigidity. By contrast, Woodford’s work highlights that the nominal rigidity can be large at the aggregate, even if the underlying micro-level rigidity is modest. An important open question for future research is how this elementary insight matters for the mapping between micro-level data and macro-level responses developed in Alvarez and Lippi (2014) and Alvarez, Lippi and Paciello (2015), and the quantitative conclusions that are drawn with the help of those mappings.

Let us now return to the motivating evidence reviewed in the beginning of this subsection. How does incomplete information help the model match this evidence? And does it offer a superior structural interpretation of the evidence than that offered by the Calvo friction?

Woodford (2003) addresses these questions with the help of Figure 7 (which copies Figures 3 and 4 from the original paper). These figures illustrate the impulse response of inflation and real output to a positive innovation in θ_t , under four alternative parameterizations of ρ , the autocorrelation of the growth rate of Nominal GDP. The right column corresponds to the incomplete-information model studied here; the left column gives, for comparison purposes, the impulse response of a New-Keynesian variant, which replaces the informational friction with standard, Calvo-like, sticky prices.⁷⁸

For the reasons already explained, incomplete information generates inertia in the response of the price level, P_t , to the underlying shocks. However, the evidence requires that the model delivers two stronger properties: first, we need inertia in the response of *inflation*, not just in that of the price level; second, we need the inertia of inflation to be more pronounced than that of real output.

As is evident in the top panel of Figure 7, both the incomplete-information and the Calvo model fail to deliver the required empirical properties when $\rho = 0$, i.e., when θ_t follows a random walk: in this case, both inflation and output peak on impact (at the moment that the innovation occurs). What is more, Woodford (2003) shows that the two models are isomorphic with regard to the response of inflation and output:

Proposition 29. *Suppose $\rho = 0$. For any parameterization of the Calvo model, there is a parameterization of the incomplete-information model that delivers identical impulse responses for output and inflation, and vice versa.*

This result complements the intuitions we developed in Section 6. Not only can the Calvo friction be thought of as a friction in coordination, we now have a concrete example in which the Calvo friction is observationally equivalent to an information-driven friction in coordination.

That said, this equivalence is exact *only* when $\rho = 0$. When instead $\rho > 0$, this equivalence breaks, indeed in a manner that helps incomplete information outperform the Calvo friction vis-a-vis the data.

⁷⁷Another difference is that these papers impose that an agent learns *perfectly* the entire state of nature once she updates her information. This implies that, at a given date t , an agent who updated her information at $t - j$ ($j \geq 0$) does not face any uncertainty about the beliefs of the set of agents who updated their information prior to $t - j$; she only faces uncertainty about the beliefs of the set of agents who updated their information after she did (which is an empty set if $j = 0$). The same point applies to Mankiw and Reis (2002), which we discuss in the sequel.

⁷⁸See Woodford (2003) for how the two models are calibrated. The calibration of the Calvo model is based on standard criteria; the calibration of the incomplete-information model is more ad hoc. See, however, Melosi (2014) for an estimated version of Woodford’s model, which delivers similar lessons.

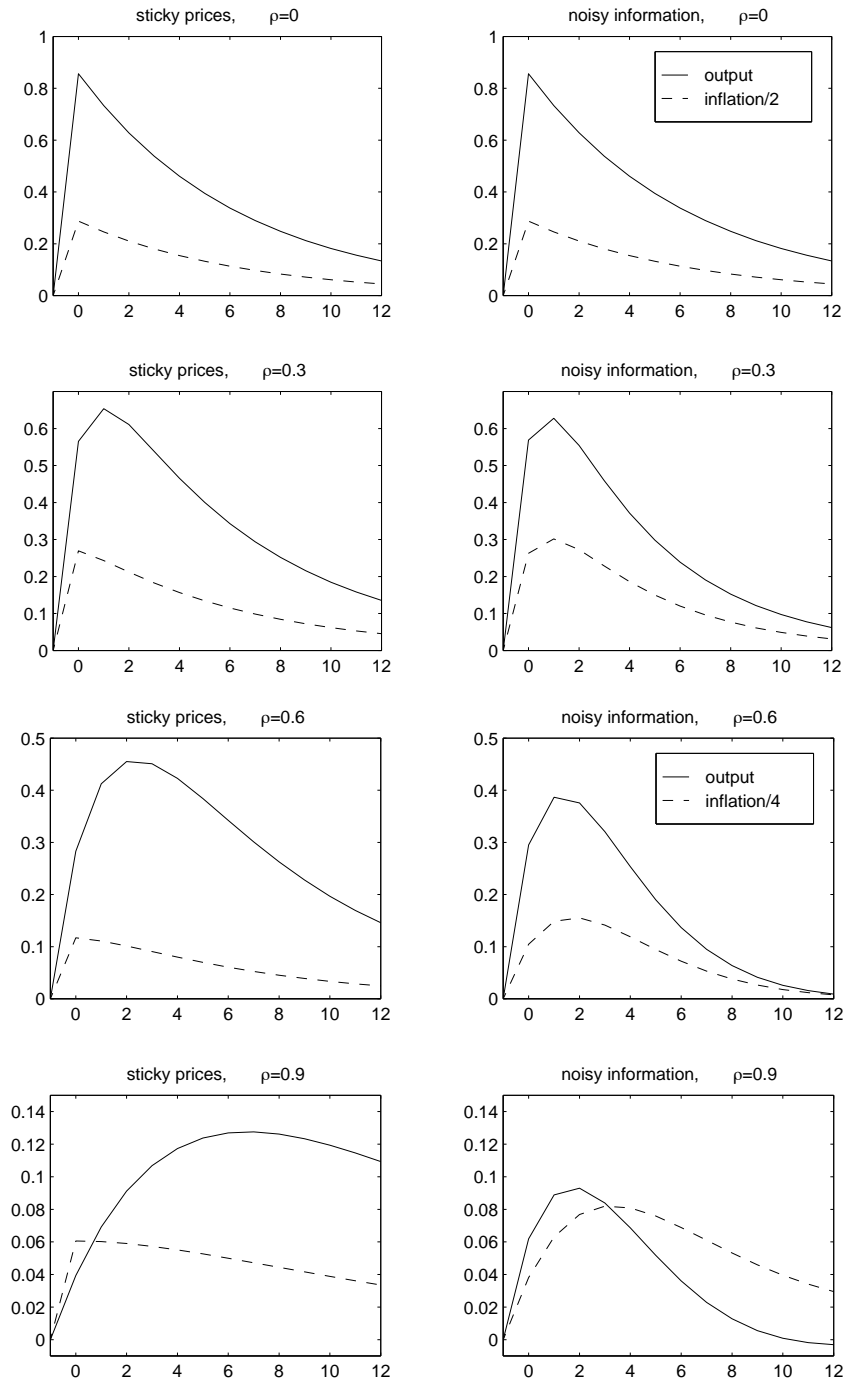


Figure 7: Comparison of impulse response functions predicted by the textbook New-Keynesian model and by the incomplete-information model of Woodford (2003), for different values of ρ .

Proposition 30. *Suppose that the output response is hump-shaped. For any parameterization, the Calvo model predicts that inflation peaks before output. By contrast, for some parameterizations, the incomplete-information model predicts that inflation peaks after output.*

The first part of the proposition follows directly from the property noted earlier that the New-Keynesian Philips Curve requires that inflation is the best predictor of the present value of future output gaps: insofar as a monetary (or any other) shock causes the output gap to follow a hump-shaped response like the one seen in the data, the NK model predicts that inflation has to peak before output.

The second part follows from the numerical examples presented in Woodford (2003) and repeated in the lower panels of Figure 7 here: once ρ is sufficiently high (such as $\rho = .6$ or $\rho = .9$ in the figure), the incomplete-information model delivers, not only a hump-shaped response for both inflation and output, but also a peak for inflation that comes later than the peak in output.

Let us explain why. A positive innovation v_t triggers a *gradual* increase in θ_t , from its initial level to a higher long-run level. This means that firms would find it optimal to increase their prices only gradually *even if* they become immediately aware of the shock and could perfectly coordinate their pricing decisions. The fact that each firm becomes only slowly aware about the shock causes each firm alone to delay its response. The fact that each firm expects other firms to do so adds further delay. This delay, however, is necessarily bounded. As time passes, not only first-order but also higher-order beliefs get closer to the true θ_t , which itself converges to its higher-long run level. It follows that the adjustment of prices accelerates after some point. Putting these observations together, we have that inflation is low early on and accelerates later on. By contrast, the growth rate of θ_t is, by assumption, high early on and slows down later on. It follows the difference between the two, which is the growth rate of output, can be high early on, exactly when inflation is still low. This explains why output can peak before inflation in the incomplete-information setting.

To recap, Woodford (2003) highlights how the inertia of higher-order beliefs can help rationalize nominal rigidity at the aggregate level in a manner that is empirically distinct from, and potentially superior to, the conventional Calvo-like formalization. Complementary subsequent work includes Nimark (2008) and Angeletos and La’O (2009), who study the interaction of the two frictions and find they tend to reinforce each other; Melosi (2014), who estimates Woodford’s model on the basis of output and inflation data and finds that the model can match the data with a modest level of informational friction; and Hellwig (2005), who elaborates on the micro-foundations and the welfare implications of the model. In what follows, we leave aside these extensions and instead relate Woodford’s contribution to two other approaches: “sticky information” as in Mankiw and Reis (2002); and “rational inattention” as in Sims (2003) and Mackowiak and Wiederholt (2009). But before doing this, we make an important parenthesis: we review a simplified version of Lucas (1972) and use this to clarify the novelty of the mechanism we have discussed in this subsection.

8.4 Parenthesis: Lucas (1972)

Consider the following simplified version of Lucas (1972). The economy consists of overlapping generations. Each agent lives for two periods, working when she is young and consuming when she is old. There is a continuum of islands, $i \in [0, 1]$. In each period t , each island has a continuum of young and old agents.

Fix a period t and an island i . Consider the young agents who are born in that period and work in that island. There is a measure one of such agents and they are all identical. Each of them supplies $N_{i,t}$ units of labor to produce

$$Y_{i,t} = N_{i,t} \tag{32}$$

units of the island's good. Let $P_{i,t}$ denote the nominal price of that good in period t .

In period $t + 1$, the aforementioned agents become old; they are randomly relocated to different islands, according to an assignment rule that will be described shortly; and they each receive a monetary transfer that is proportional to the nominal income they made when young. At the aggregate level, the monetary transfer is pinned down by the growth rate of money supply. The latter follows a random walk in logs:

$$M_{t+1} = M_t e^{v_{t+1}},$$

where $v_{t+1} \sim N(0, \sigma_v^2)$ is an aggregate shock. It follows that the cash-in-hand of an old agent, who was born in island i and is currently (i.e., in period $t + 1$) located in island j , is given by

$$M_{i,t+1} = P_{i,t} N_{i,t} e^{v_{t+1}}.$$

Her budget constraint is given by

$$P_{j,t+1} C_{i,j,t+1} = M_{i,t+1},$$

where $C_{i,j,t+1}$ denotes her consumption. Finally, the her realized utility is given by

$$U_{i,j,t} = C_{i,j,t+1} - \frac{1}{1+\kappa} N_{i,t}^{1+\kappa}.$$

We now elaborate on the assignment of old agents to different islands and the resulting demand for the local good of each island. In each period t , the old agents on any given island are a representative sample of all the agents who were born in the previous period and work on different islands. However, different islands receive samples of different sizes. In particular, the mass of old agents that island i receives in period t is given by $\bar{\xi} e^{\xi_{i,t}}$, where $\xi_{i,t} \sim N(0, \sigma_\xi^2)$ is i.i.d. across islands and across periods, as well as independent from v_t . We set $\bar{\xi}$ so that $\mathbb{E}[\bar{\xi} e^{\xi_{i,t}}] = 1$, that is, the average mass is one. It follows that the nominal demand on island i during period t is given by

$$D_{i,t} = \bar{\xi} e^{\xi_{i,t}} M_t,$$

where M_t is the aggregate quantity of money. Market clearing then imposes

$$P_{i,t} N_{i,t} = D_{i,t}. \tag{33}$$

The modeling role of $\xi_{i,t}$ is therefore to induce island-specific demand shocks—or, equivalently, variation in relative prices.

The above assumptions guarantee a simple characterization of the optimal labor supply of the young agents. From the budget constraint, we have that the consumption of a young agent who works on island i in period t and consumes on island j in period $t + 1$ is given by

$$C_{i,j,t+1} = \frac{M_{i,t+1}}{P_{j,t+1}} = \frac{P_{i,t} N_{i,t} e^{v_{t+1}}}{P_{j,t+1}} = \frac{P_{i,t}}{M_t} \cdot \frac{M_{t+1}}{P_{j,t+1}} \cdot N_{i,t}.$$

It follows that the optimal labor supply of that agent in period t is given by

$$N_{i,t} = \arg \max_N \mathbb{E}_{i,t} \left[\frac{P_{i,t}}{M_t} \cdot \frac{M_{t+1}}{P_{j,t+1}} \cdot N - \frac{1}{1+\kappa} N^{1+\kappa} \right] = \left(\mathbb{E}_{i,t} \left[\frac{P_{i,t}}{M_t} \cdot \frac{M_{t+1}}{P_{j,t+1}} \right] \right)^{\frac{1}{\kappa}}. \tag{34}$$

What remains then is to specify the information structure.

As in Lucas (1972), we assume that *previous-period* money supply, M_{t-1} , is public information, but *current* aggregate money supply, M_t , is unknown; equivalently, the current monetary shock v_t is unknown. We also assume that each agent observes the current nominal price in her island, but not those in other islands. It follows that the information set that enters the expectation in condition (34) is the pair $(M_{t-1}, P_{i,t})$.⁷⁹ Finally, we restrict attention to equilibria in which $P_{i,t}$ is log-normally distributed.

Let lower-case variables denote the logarithm of the corresponding upper-case variables, measured as deviations from steady state, and restrict attention to rational-expectations equilibria in which p_{it} is Normally distributed. There exists a unique equilibrium of this type and is characterized as follows.

Proposition 31 (Lucas 72). *Consider the version of Lucas (1972) described above. There exist scalars $\beta, \lambda \in (0, 1)$ such that the following properties hold:*

(i) *The nominal price level is given by*

$$p_t = (1 - \beta)m_t + \beta\bar{\mathbb{E}}[m_t]. \quad (35)$$

(ii) *Real output is given by*

$$y_t = m_t - p_t = \beta \{m_t - \bar{\mathbb{E}}[m_t]\}. \quad (36)$$

(iii) *The average forecast error of m_t (also known as the “unanticipated” change in the supply of money) is given by*

$$m_t - \bar{\mathbb{E}}[m_t] = \lambda v_t. \quad (37)$$

The combination of parts (ii) and (iii) reveal that monetary shocks have real effects. The intuition is simple. Young agents do not directly observe the shock in money supply. Instead, they only observe the movement in the local price, which confounds the aggregate monetary shock v_t with the island-specific demand shock ξ_{it} . As a result, they (rationally) confuse nominal price movements for relative price movements. This confusion then explains why it is optimal for young agents to exert more effort and produce more output in response to a monetary shock.

For our purposes, however, it is more useful to focus on part (i). Contrast condition (35) with condition (31), that is, the condition that pins down the price level in Woodford (2003).⁸⁰ It then becomes clear that the price level in Lucas’s model depends only on first-order beliefs of the monetary shock (the “fundamental”), whereas in Woodford’s model it also depends on higher-order beliefs. By the same token, the degree of monetary non-neutrality obtained in Lucas’s model can be large only insofar as agents are uninformed about the monetary shock, which in turn explains why Barro sought to test the theory by measuring unanticipated monetary shocks. By contrast, insofar as there is strong complementarity in price-setting decisions, the degree of monetary non-neutrality obtained by Woodford can be large even if agents are well informed about the monetary shock—for the key is now the lack of common knowledge, as opposed to the lack of individual knowledge.

Remark. There are two reasons why higher-order beliefs are absent in Lucas’s work. The one is the absence of strategic complementarity, which is emphasized above. The other is the assumption that the previous-period fundamental (m_{t-1}) is commonly known at the begin of each period. If the last assumption is appropriately relaxed, and in particular if different agents within any given island have differential infor-

⁷⁹We could also allow the agents to observe the entire history of money supply and nominal prices in all past periods. Given the rest of the assumptions we have made and the equilibrium we construct in the sequel, this history provides no additional information.

⁸⁰Note that θ_t in the previous subsection is the same as m_t presently.

mation about the past fundamentals, then higher-order beliefs become relevant through the signal-extraction problem: the interpretation of the local price signal by any young agent depends, in general, on her beliefs regarding the information and the beliefs of old agents, as well as that of other young agents. This kind of mechanism—the role of higher-order beliefs in the interpretation of signals of the activity of others—was the topic of Townsend (1983) and of the subsequent works by Sargent (1991) and Kassa (2000). But just as Lucas (1972) abstracted from strategic complementarity in actions, so did these works.

8.5 Sticky Information and Rational Inattention

Having clarified the key difference between Lucas (1972) and Woodford (2003), we now briefly comment on the connection of the latter with two other important recent contributions: that of Mankiw and Reis (2002), and that of Mackowiak and Wiederholt (2009, 2015).

In Mankiw and Reis (2002), firms are free to adjust their prices continuously over time, but may update their information sets only infrequently. In particular, it is assumed that, in each period, and regardless of its history up to that point, a firm gets to see the underlying state with probability $\lambda \in (0, 1)$, and it is otherwise stuck with its previous-period information set. By the same token, a fraction λ of firms are perfectly informed within any given period, whereas the rest must set prices on the basis of outdated information. This friction is taken as given in Mankiw and Reis (2002), but it is micro-founded on the basis of a fixed cost for observing the state of nature in Reis (2006), and Alvarez, Lippi and Paciello (2011, 2015).⁸¹

In their paper, Mankiw and Reis focus on the comparison of their model to the standard Calvo model. In particular, they show that their model imposes the following restriction on the joint dynamics of inflation and output, which can be interpreted as the Philips curve of the model:

$$\pi_t = \left(\frac{(1 - \alpha)\lambda}{1 - \lambda} \right) y_t + \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j \mathbb{E}_{t-1-j} (\pi_t + (1 - \alpha) \Delta y_t),$$

where α is the degree of strategic complementarity in pricing decisions and λ is the aforementioned probability of observing the state. The above result indicates a certain backward-looking aspect, unlike the forward-looking nature of the New-Keynesian Philips Curve: inflation today depends on past information, simply because that past information is relevant for the current pricing choices of firms that have not updated their information.

Notwithstanding the distinct applied contribution of Mankiw and Reis (2002), we now proceed to clarify the manner in which this approach is similar to that of Woodford (2003).

Because the best responses in the two models are the same, the equilibrium price level can be expressed as the *same* function of the hierarchy of beliefs in both models. It follows that the two models can deliver quantitatively different predictions for the dynamics of prices, inflation, and output only to the extent that they happen to feature sufficiently different dynamics in higher-order beliefs.

As with Woodford (2003), we assume that $\Delta\theta_t$ follows an AR(1) process with autocorrelation coefficient $\rho \in [0, 1)$. A clear benchmark emerges when $\rho = 0$, that is, when θ_t follows a random walk. For this case, we already studied the dynamics of higher-order beliefs implied by Woodford’s specification of the information structure. Turning to the dynamics implied by Mankiw and Reis’s specification, we prove the following.

⁸¹See also the review in Mankiw and Reis (2011).

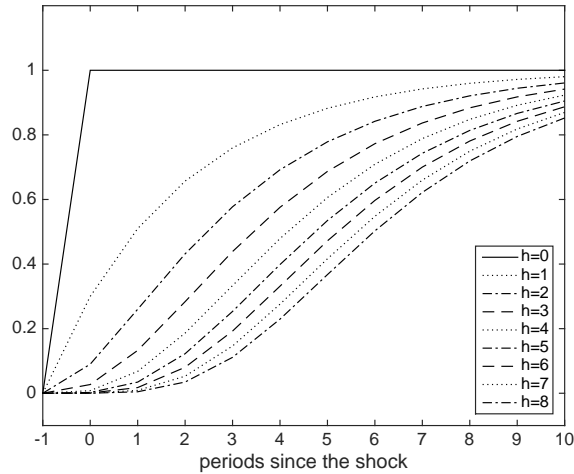


Figure 8: Impulse response function of higher-order beliefs in sticky information model

Proposition 32. *When $\rho = 0$, the first- and higher-order forecasts of θ_t are given by the following:*

$$\bar{\mathbb{E}}_t^h[\theta_t] = \sum_{j=0}^{+\infty} \left\{ \left(1 - (1 - \lambda)^{j+1}\right)^h v_{t-j} \right\} \quad \forall h \geq 1,$$

where $\lambda \in (0, 1)$ is the probability that a firm updates its information in any given period.

This result gives a closed-form solution for the IRFs of the entire hierarchy of forecasts: the effect of an innovation on the h -th order forecast after j periods is $\left(1 - (1 - \lambda)^{j+1}\right)^h$, which is clearly increasing in λ , increasing in j , and decreasing in h .

In Woodford's specification such a closed-form solution was not feasible (with the exception of first-order beliefs). Nevertheless, it should now be evident that the qualitative features of the two specifications are the same. In both specifications, higher-order beliefs exhibit a weaker and more sluggish response to fundamentals than lower-order beliefs. Furthermore, the effect that λ has in the above IRFs is essentially the same as the one that the reciprocal of σ_ϵ has in Woodford's model. This is because in both cases these scalars relate to the speed of learning.

We further illustrate these points in Figure 8. This figure depicts the impulse responses of first- and higher-order beliefs in the sticky-information model, under the maintained assumption that θ_t follows a random walk. Comparing this figure to Figure 4, we see that the qualitative dynamics of the belief hierarchy is nearly indistinguishable—and therefore so are the impulse response of inflation and output.

Moving beyond the random-walk case, Figure 9 revisits the last exercise we conducted for Woodford's contribution: it draws the impulse responses of inflation and output to a positive monetary shock under different values for ρ . A comparison of this figure to Figure 7 reveals that the predictions of the two models are closely related. The only noticeable difference is that, at least for the chosen parameterizations, the sticky-information model appears to have a relatively easier time in matching the fact that inflation peaks after output, even for low values of ρ .

To recap, Woodford (2003) and Mankiw and Reis (2002) appear to deliver similar dynamics for higher-order beliefs, and thereby also similar dynamics for inflation, output, and the average forecasts of any order.

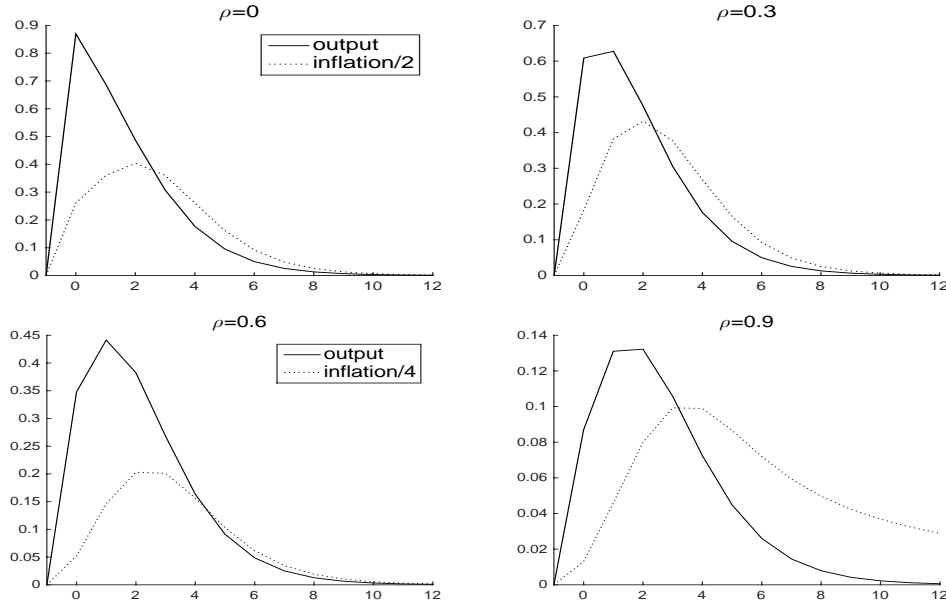


Figure 9: Impulse response functions of inflation and output in sticky-information model.

This, however, does not mean that the two models are observationally equivalent (or nearly equivalent) in *every* dimension or in every context. For instance, Woodford’s specification implies that the cross-sectional dispersion of forecasts—whether of the exogenous shock or of inflation and all the other endogenous variables—is constant over time, whereas Mankiw and Reis’s specification implies that the cross-sectional dispersion of forecasts *increases* following any innovation in the fundamentals.

Some authors have paid special attention to this last property.⁸² In our view, however, this property is not particularly interesting. We, economists, have little knowledge of the precise ways in which real-world people collect, exchange, and digest information. Accordingly, what we would like to retain as a lesson from *both* Woodford (2003) and Mankiw and Reis (2002) is their common predictions with regard to the inertia of higher-order beliefs and the consequent dynamics of inflation and output. These predictions are driven by the combination of strategic complementarity with private learning; they are therefore likely to be robust to other plausible specifications of the information structure.

Let us now consider yet another popular form of informational friction: the form of “rational inattention” proposed by Sims (2003). Because of space constraints, we will not review either the foundations of this approach or the extensive applied literature that followed Sims’s original contribution.⁸³ Instead, we limit our discussion to the relation of this approach to the central theme of this chapter.

Unlike the more ad hoc alternatives we have studied so far, Sims’s formalization of rational inattention is grounded on the idea that, even if arbitrarily precise information is readily available, people may have limited capacity in processing all that information and they may therefore act *as if* their information were noisy. In short, noise is present, not because information is lacking, but rather because of a cognitive friction.

⁸²Mankiw, Reis and Wolfers (2004) argue that this property helps explain certain time-varying patterns in the cross-sectional dispersion of inflation forecast. By contrast, Coibion and Gorodnichenko (2012) find no evidence of this property in the response of inflation forecasts to certain identified shocks (more on this in Subsection 8.6) and therefore conclude that the data is more in line with noisy information as in Woodford (2003) than with sticky information.

⁸³Important follow-up contributions include Luo (2008), Mackowiak and Wiederholt (2009, 2015), Paciello and Wiederholt (2014), Matějka (2015a,b), Matejka and Sims (2011), Matejka and McKay (2015), Matějka et al. (2015), Sims (2006), Stevens (2015), Tutino (2013), and Woodford (2009). See also the review in Sims (2011).

The basic idea is compelling. Applied to the context of monetary policy, it also allows us to avoid a critique of the earlier literature on imperfect information (Lucas, 1972), namely that information about monetary policy and the price level is readily available: even if this were true, rational inattention could explain why firms and consumers may act *as if* they did not have access to all the relevant data.

That said, we would also like to highlight that the foundations of rational inattention are decision theoretic. The issues that are central to this chapter—coordination, strategic or higher-order uncertainty, and solution concepts—were left completely out of the picture in Sims’s original contribution.

Subsequent works, most notably those by Mackowiak and Wiederholt (2009, 2015), applied Sims’s approach to general-equilibrium models. The works impose information structures that rule out correlated noise: in the equilibria of those models, aggregate economic outcomes (such as prices and quantities) are pinned down by the aggregate fundamentals (such as monetary and productivity shocks). This assumes away the type of “animal spirits” identified earlier. Nevertheless, because agents act as if they observe the fundamentals with *idiosyncratic* noise, the following properties hold: first, information is incomplete (in the sense of Definition 10), not just imperfect (in the sense of Definition 9); second, as long as there is strategic complementarity, the mechanism we studied before regarding the inertia of higher-order beliefs is active.

In this respect, the approach taken in Mackowiak and Wiederholt (2009, 2015) is closely related, and complementary, to Woodford (2003) and Mankiw and Reis (2002). Nevertheless, Mackowiak and Wiederholt’s approach makes two distinct predictions, which are of independent interest.

The first distinct prediction is that the de-facto noise is likely to be much larger for aggregate shocks than for idiosyncratic shocks. This is because of the following. Empirically, idiosyncratic shocks are an order of magnitude more volatile than aggregate shocks. It follows that each individual agent finds it more worthwhile to allocate her attention (capacity) to idiosyncratic shocks rather than to aggregate shocks. What is more, strategic complementarity reinforces this decision-theoretic effect: when other firms pay less attention to aggregate shocks, the individual firm has an even small incentive to pay attention to such shocks.

The second distinct prediction has to do with comparative statics. In a setting where the signals are exogenously specified, one may find it natural to vary elements of the payoff structure—e.g., the degree of strategic complementarity in pricing decision, or the specification of the policy rule followed by the monetary authority—keeping constant the precision of the available signals. In a rational-inattention setting, by contrast, the precision of the signals is tied to the underlying payoff characteristics. It follows that the two approaches may make different predictions about the effect of, say, a regime change in monetary policy, even if they deliver the same predictions about the effects of a monetary shock under a fixed policy regime.

Needless to say, the above point is not specific to rational inattention; it applies more generally to any model that endogenizes either the collection or the aggregation of information. Rational inattention imposes a particular structure on this kind of endogeneity. A plausible alternative is developed in Reis (2006) and Alvarez, Lippi and Paciello (2015): this approach replaces rational-inattention with fixed costs in the observation of the underlying shocks, but shares the prediction that information is endogenous to policy.⁸⁴

8.6 Survey Evidence on Informational Frictions

In Section 7, we highlighted that complete-information and incomplete-information models have distinct predictions regarding the joint distribution of the aggregate action, the average expectation of the aggregate

⁸⁴This discussion pegs the question of how the endogeneity of information influences the nature of optimal monetary policy. For recent advances into answering this question, see Paciello and Wiederholt (2014) and Angeletos, Iovino and La’O (2016b).

action, and the underlying fundamental. In the context of monetary models, these predictions regard, inter alia, the joint distribution of actual inflation, the average forecast of inflation, and the underlying shocks.

Consider the textbook New-Keynesian model, or any modern DSGE model that maintains the assumption of complete information. These models predict that firms face no uncertainty about the contemporaneous price level, or inflation, even if they face uncertainty about the underlying shock. It follows that, in these models, $\bar{\mathbb{E}}_t \pi_t$ coincides with π_t in all states of nature. Hence, the impulse response function of π_t to any shock coincides with that of $\bar{\mathbb{E}}_t \pi_t$. By contrast, incomplete-information models such as those in Woodford (2003), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009) predict that the two IRFs are distinct, and in particular that $\bar{\mathbb{E}}_t \pi_t$ responds more sluggishly than π_t to innovations in the underlying fundamentals.

In an important recent paper, Coibion and Gorodnichenko (2012) provide evidence in favor of the latter prediction. This paper estimates the impulse response functions of actual inflation and of inflation forecasts to three distinct shocks recovered from the data: a technology shock identified as in Galí (1999); an oil shock identified as in Hamilton (1996); and a news shock identified as in Barsky and Sims (2011).⁸⁵ Inflation forecasts are obtained from four sources: the Survey of Professional Forecasters (SPF); the University of Michigan Survey of Consumers; the Livingston Survey; and the FOMC blue book. For each of these sources, and for each of the aforementioned shocks, it is shown that the average inflation forecasts respond more sluggishly than actual inflation.

This finding is illustrated in Figure 10, which is borrowed from Coibion and Gorodnichenko (2012). This figure depicts the IRFs of actual inflation, π_t , and of the average forecast error, $\bar{\mathbb{E}}_t \pi_t - \pi_t$, to each of the aforementioned shocks, using the SPF-based measure for $\bar{\mathbb{E}}_t \pi_t$. Note that the technology shock and news shock are found to be disinflationary, whereas the oil shock is found to be inflationary. These properties are consistent with previous empirical work, as well as with the predictions of standard macroeconomic models. More importantly for our purposes, note that the response of the average forecast error is found to be negative in the first two cases, positive in the last case, and always lower in absolute value than the actual inflation response. This means that, for all the three shocks, the average forecast $\bar{\mathbb{E}}_t \pi_t$ moves in the same direction as actual inflation, π_t , but also in a more muted manner. Finally, note that the average forecast error vanishes with time.

Coibion and Gorodnichenko (2012) show that the same qualitative patterns hold if the SPF measure of inflation forecasts is replaced by forecasts measures from the University of Michigan Survey of Consumers, the Livingston Survey, and the FOMC blue book. To the extent that these measures are good proxies for the average inflation forecast $\bar{\mathbb{E}}_t \pi_t$ in the theory, the above evidence provides prima-facie support against the standard framework and in favor of belief inertia like the one implied by incomplete information.

That said, there is an important disconnect between the empirical exercise and the theoretical models of interest. When Coibion and Gorodnichenko (2012) seek to map their evidence to the theory, they treat the observed inflation dynamics as exogenous to the observed forecast dynamics. By contrast, the models we have reviewed impose that the dynamics of inflation are endogenous to the dynamics of forecasts, indeed in a manner that is central to the predictions of the theory. To put it differently, Mankiw and Reis (2002), Woodford (2003), and others have used informational frictions, not just to explain why the agents' forecasts of inflation may adjust slowly to exogenous shocks, but also to explain why this kind of belief inertia may itself feed into inertia in actual inflation, which in turn may feed into further belief inertia, and so on.

⁸⁵Given that Woodford (2003), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009) were primarily interested in the response of prices to monetary shocks, it may seem peculiar that Coibion and Gorodnichenko (2012) do not study such shocks. Coibion and Gorodnichenko justify this on the basis that identified monetary shocks only drive a relatively small fraction of the business-cycle variation in inflation and economic activity.

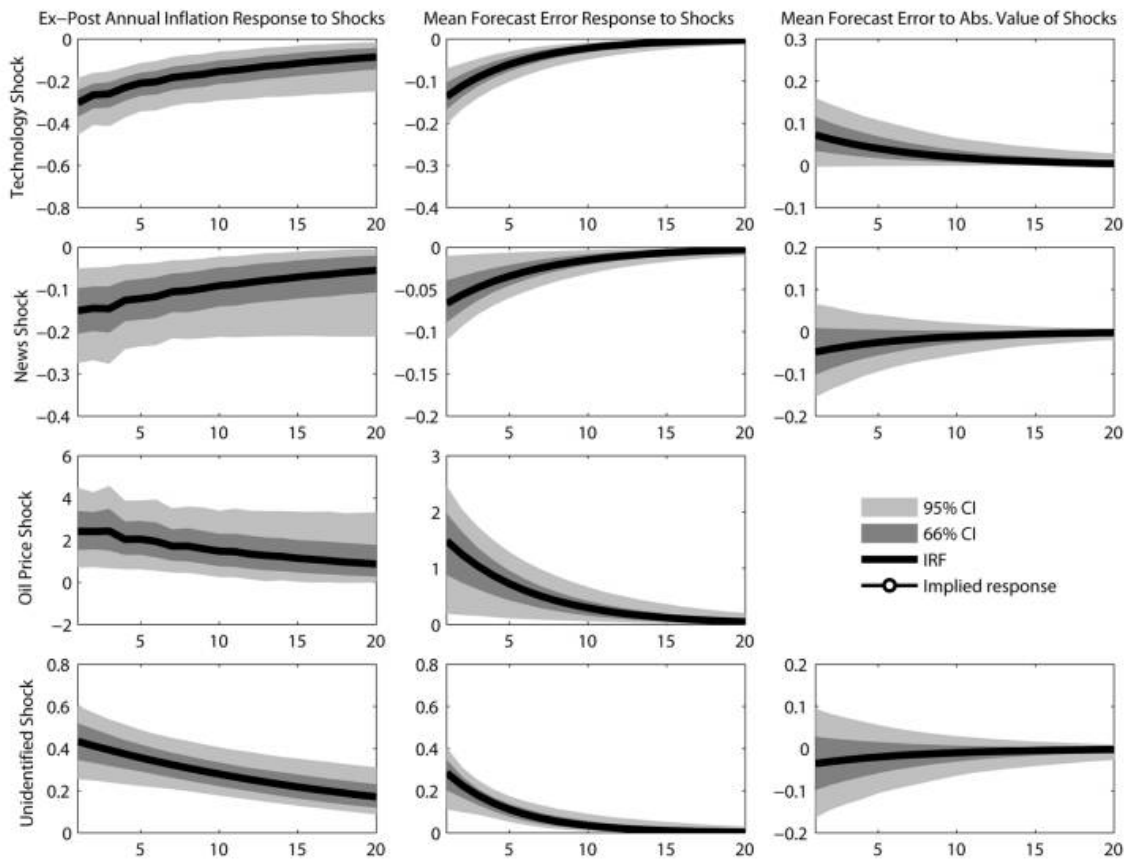


Figure 10: Response of forecast errors to structural shocks in SPF.

To recap, Coibion and Gorodnichenko (2012) make a significant contribution by documenting salient co-movement patterns in the joint responses of forecasts of inflation and of actual inflation to certain shocks. Nevertheless, the precise mapping from that evidence to the theory remains an open question: by treating inflation as an exogenous object, that paper has stopped short of quantifying the equilibrium mechanism that is at the core of the models we have reviewed.⁸⁶

We conclude by mentioning a few additional works that use survey data to information heterogeneity. Afrouzi et al. (2015) and Coibion et al. (2015) document widespread dispersion in firms' beliefs about macroeconomic conditions, especially inflation. Coibion and Gorodnichenko (2015) use the relationship between ex-post mean forecast errors and the ex-ante revisions in the average forecast to test the existence of information rigidities. Andrade and Bihan (2013), Branch (2007), Carvalho and Nechio (2014), Cavallo, Cruces and Perez-Truglia (2015), and Sarte (2014) provide additional evidence in favor of information frictions. Combined, these papers provide ample motivations for studying the role of information frictions in the context of business cycles and monetary policy.

8.7 Demand-Driven Fluctuations

In the preceding subsections we reviewed a literature that shows how informational frictions can offer a compelling micro-foundation for nominal rigidity. This literature complements the New-Keynesian paradigm, because nominal rigidity is central to this paradigm's ability to explain the observed business-cycles and to accommodate the notion of demand-driven fluctuations. We now turn attention to a different line of work, which goes against the New-Keynesian paradigm: we argue that informational frictions can achieve the aforementioned goals (i.e., explain the business-cycle data and accommodate the demand-driven fluctuations) even in the absence of nominal rigidity.

There is a long tradition that formalizes demand-driven fluctuations as the product of "animal spirits" within multiple-equilibrium models. See, e.g, Azariadis (1981), Benhabib and Farmer (1994), Cass and Shell (1983), Diamond (1982), Cooper and John (1988), and Guesnerie and Woodford (1993). This approach enjoys limited popularity in the modern business-cycle paradigm, in part because of debatable empirical foundations and in part because of the inconveniences that multiple-equilibrium models carry for estimation and policy evaluation. Nevertheless, what we find appealing with this tradition is that it disentangled the notion of demand-driven fluctuations from that of monetary non-neutrality. This contrasts with the New-Keynesian framework, which gives a central position to nominal rigidity and to monetary policies that fail to replicate flexible prices: in the absence of these features, the notion of demand-driven fluctuations evaporates, and so does the framework's ability to match salient features of the business-cycle data.

But now note that the modern business-cycle paradigm leaves no room for animal spirits, not only because it imposes a unique equilibrium, but also because it rules out imperfect coordination (in the sense we have defined in this chapter). This suggests that the introduction of incomplete information to otherwise canonical unique-equilibrium macroeconomic models may help accommodate a type of fluctuations in expectations and macroeconomic outcomes that resembles the sunspot fluctuations obtained in the older multiple-equilibrium literature, thus also providing a potent formalization of demand-driven fluctuations that does not rest on either nominal rigidity or "mistakes" in monetary policy.

⁸⁶Relatedly, the finding in Coibion and Gorodnichenko (2012) that there is little evidence for strategic complementarity should not be misinterpreted. This finding refers to a rejection of the hypothesis that professional forecasters distort their reported forecasts in an attempt to conform to the forecasts of others. It has nothing to say about the role of strategic complementarity in price-setting behavior or, more generally, in business-cycle phenomena.

This basic idea was pursued in Angeletos and La’O (2013). The authors consider a convex neoclassical economy in which agents are rational, markets are competitive, the equilibrium is unique, and there is no room for randomization devices. They also rule out aggregate shocks to preferences, technologies, or any other payoff-relevant fundamentals. This shuts down the type of higher-order uncertainty we have studied so far. And yet, the authors are able to obtain aggregate fluctuations, thanks to correlated higher-order beliefs about idiosyncratic trading opportunities.

Let us elaborate on the mechanics of that paper. The structure of the economy is similar to the one used in Subsection 8.1 to study real rigidity, except for two modifications: islands are randomly matched in pairs; and in any given period, islands can trade and communicate only with their trading partners.

As in the model introduced in Subsection 8.1, each island specializes in the production of a specific good and consumes also the good produced by at least one other island. This gives rise to trade and, thereby, to a certain kind of strategic complementarity. Unlike the model of Subsection 8.1, however, trade is decentralized and takes place through random matching: in each period, each island meets and trades with only one other, randomly selected, island. This implies that game-theoretic representation of the general equilibrium of that model takes the following form:

$$y_{it} = (1 - \alpha)\theta_i + \alpha\mathbb{E}_{it}[y_{m(i,t),t}], \quad (38)$$

where y_{it} is the output of island i in period t , θ_i is its exogenous and time-invariant productivity, $\alpha \in (0, 1)$ is the degree of strategic complementarity implied by the underlying preference and technology parameters, and $m(i, t)$ denotes the trading partner (or “match”) of island i in period t . It follows that the output of an island depends, not only on its own productivity and the expected productivity of its likely trading partner, but also on what the later expects from its own trading partner, and so on. The authors then proceed to engineer aggregate fluctuations from this kind of higher-order uncertainty.

Formally, the fluctuations obtained in Angeletos and La’O (2013) are aggregate manifestations of exogenous correlated shifts in higher-order beliefs. As discussed before, however, these higher-order belief shifts need to be taken too literally. Rather, they can be interpreted as a device for accommodating nearly self-fulfilling fluctuations in expectations of “demand”. This should be evident from equation (38): shifts in higher-order beliefs trigger shifts in actual output because, and only because, they rationalize shifts in the expectations (first-order beliefs) that each island forms about the demand for its product.⁸⁷ What is more, even though the equilibrium is unique, the resulting fluctuations have a similar flavor as those sustained in multiple-equilibrium models: when an island expects more demand for its product, it produces more, which in turn raises the demand for other islands’ products. Last but not least, these fluctuations are possible even if the aggregate fundamentals are constant, thus bypassing the limitation discussed in Proposition 21.

A related formalization of aggregate demand fluctuations appears in Benhabib, Wang, and Wen (2015). As in Angeletos and La’O (2013), there are no aggregate shocks to fundamentals (preferences or technologies) and the equilibrium is unique when information is perfect. But unlike that paper, multiple equilibria obtain once a certain informational friction is introduced. What opens the door to multiplicity is the endogeneity of the information. In the model, each firm observes an endogenous private signal about its demand, which in turn depends on the behavior of other firms. In one of the equilibria, which Benhabib, Wang, and Wen (2015) argue is the most plausible one, aggregate activity is shown to vary with a sunspot. The sunspot is not publicly observable. Instead, the signal that each firm receives about its demand acts partly as an imperfect signal of an idiosyncratic demand shock and partly as an imperfect signal of the aggregate sunspot. What sustains the equilibrium is then the signal-extraction problem that firms face with regard to figuring

⁸⁷The higher an island’s trading partner’s output, $y_{m(i,t),t}$, the higher demand for the good produced by the island.

out whether demand is driven by the one or the other shock.

Complementary are also the works of Gaballo (2015), Chahrour and Gaballo (2015), Benhabib, Liu and Wang (2015, 2016). Gaballo (2015) studies an economy in which final producers are informed about aggregate conditions only through the equilibrium prices of their local inputs; shows that multiple equilibria can arise when idiosyncratic shocks to intermediate production are small; and documents the existence of an interesting equilibrium in which prices are rigid with respect to aggregate shock. Chahrour and Gaballo (2015) show that non-trivial aggregate fluctuations may originate with vanishingly small common shocks to either information or fundamentals. Benhabib, Liu and Wang (2015) introduce endogenous information acquisition in a model where firms face both idiosyncratic and aggregate demand shocks and show that endogenous information acquisition makes economic volatility time-varying and countercyclical. Benhabib, Liu and Wang (2016) show that exuberant financial market sentiments can increase the price of capital, which signals strong fundamentals of the economy to the real side and consequently leads to an actual boom in real output and employment.

What all the aforementioned papers have in common with Angeletos and La'O (2013) is the central role played by incomplete information: if it were not for that feature, the models studied in all these papers would reduce to conventional, unique-equilibrium, neoclassical models, in which equilibrium outcomes would have been pinned down by fundamentals. What, however, distinguishes the aforementioned papers is the emphasis on the signal extraction problems that arise once information is incomplete and the additional volatility—including that in the form of multiple equilibria and sunspot fluctuations—which may obtain from such signal extraction problems.⁸⁸

Finally, Angeletos and Lian (2016a) study an environment that relates to the aforementioned papers in that it also rests on a signal-extraction problem between aggregate and idiosyncratic shocks, but shifts the emphasis away from multiple equilibria and sunspot fluctuations to a mechanism that helps formalize the notion of a Keynesian multiplier. In particular, a rational confusion between idiosyncratic and aggregate shocks—or at least the lack of common knowledge about aggregate shocks—is shown to explain why a negative aggregate shock to consumer spending⁸⁹ may lead firms to hire and produce less, which in turns leads consumers to spend less, and so on, ultimately leading to a recession. The same mechanism is then also shown to generate large fiscal multiplier: an exogenous increase in government spending may actually crowd in private consumption, leading to an increase in output that is higher than the exogenous increase in government spending, and helping undo the recession.

Combined, the papers we have discussed in this subsection indicate how incomplete information helps provide a set of complementary formalizations of the notion of demand-driven business cycles. Importantly, these formalizations do not require—but also do not preclude—either any form of nominal rigidity or any friction in monetary policy. Embracing these formalizations may thus affect, not only the structural interpretation of the available data, but also the policy implication one may wish to draw from the data. Indeed, if nominal rigidity is, by assumption, the only way one can make sense of demand-driven fluctuations, one is *forced* to think about frictions in monetary policy, such as the zero lower bound, as central to the observed phenomena. But if the key relevant mechanism is a friction in information and coordination, monetary policy could be a sideshow.⁹⁰

⁸⁸At some abstract level, all these signal-extraction problems are similar to the one first formalized in Lucas (1972). However, not only are the applications very different, but also higher-order beliefs come into play.

⁸⁹Such shocks may reflect shifts in preferences or, more plausible, be proxies for shocks to consumer credit and/or consumer expectations.

⁹⁰To be clear, we do not question the empirical relevance of nominal rigidity; we only question whether nominal rigidities and Philips curves are central to understanding either the notion of demand-driven fluctuations or the key regularities of the

| | Data | RBC plus HOB | NK with TFP shock plus... | | | |
|-----------------------------|------|--------------|---------------------------|----------------|--------------|----------------|
| | | | <i>I</i> shock | <i>C</i> shock | News shock | <i>M</i> shock |
| st.dev(<i>y</i>) | 1.42 | 1.42 | 1.24 | 1.15 | 1.29 | 1.37 |
| st.dev(<i>h</i>) | 1.56 | 1.52 | 1.18 | 0.97 | 1.02 | 1.44 |
| st.dev(<i>c</i>) | 0.76 | 0.76 | 0.86 | 0.95 | 0.84 | 0.77 |
| st.dev(<i>i</i>) | 5.43 | 5.66 | 7.03 | 7.04 | 7.24 | 6.20 |
| corr(<i>c</i> , <i>y</i>) | 0.85 | 0.77 | 0.42 | 0.37 | 0.43 | 0.73 |
| corr(<i>i</i> , <i>y</i>) | 0.94 | 0.92 | 0.82 | 0.75 | 0.84 | 0.90 |
| corr(<i>h</i> , <i>y</i>) | 0.88 | 0.85 | 0.80 | 0.77 | 0.86 | 0.84 |
| corr(<i>c</i> , <i>h</i>) | 0.84 | 0.34 | -0.19 | -0.29 | -0.07 | 0.24 |
| corr(<i>i</i> , <i>h</i>) | 0.82 | 0.99 | 1.00 | 1.00 | 1.00 | 0.99 |
| corr(<i>c</i> , <i>i</i>) | 0.74 | 0.47 | -0.17 | -0.33 | -0.13 | 0.35 |

Table 1: HOB shocks in the RBC model versus different kinds of demand shocks in the NK model. Boldface indicates significant difference between the model and the data.

Angeletos, Collard and Dellas (2015) and Huo and Takayama (2015a) push this research agenda further by seeking to quantify the aforementioned kind of belief fluctuations. These papers share a similar objective, but take rather different approaches. The first one relaxes the assumption of a common prior in order to accommodate rich, higher-order belief dynamics in an arbitrary linear DSGE model and proceeds to conduct a wide range of horseraces between higher order beliefs, nominal rigidity, and a variety of structural shocks commonly used in the literature. The second paper maintains the assumption of a common prior, thus disciplining the dynamics of higher-order beliefs, at the expense of a more limited range of quantitative explorations. Both papers nevertheless reach a similar conclusion: the kind of fluctuations in expectations and outcomes that are rationalized by shifts in higher-order beliefs can be quantitatively important and can offer a potent structural interpretation of salient features of the data.

An example of what this means is illustrated in Table 1, which is borrowed from Angeletos, Collard and Dellas (2015). This table compares the empirical performance of five alternative models. The first column in the table reports some key business-cycle moments of the US data; the other five columns report the corresponding moments of the five models. Let us explain what these models are.

The first model (column 2) is the authors’ baseline model. This is a variant of the textbook RBC model that features only two sources of volatility: a persistent technology shock, A_t , which moves the production possibilities of the economy; and a transitory belief shock, ξ_t , which moves higher-order beliefs of A_t for given A_t . In the equilibrium of this model, variation in ξ_t manifests as waves of optimism and pessimism about the short-term economic outlook. The authors refer to ξ_t as a “confidence shock”.

The remaining four models are versions of the New-Keynesian model. All four versions share the same RBC backbone as the baseline model of Angeletos, Collard and Dellas (2015). They differ from it only in two respects. First, they all add sticky prices a la Calvo and a Taylor rule for monetary policy, so as to accommodate the New-Keynesian transmission mechanism. Second, each one of them replaces the confidence shock with with one of the following alternative structural shocks, which have been used in the literature as proxies for “demand shocks”: an investment-specific shock; a consumption-specific, or discount-rate, shock; a news shock; and a monetary shock.

All five models are calibrated in a comparable manner: the technology and preference parameters, which

business-cycle data.




Figure 11: IRFs to HOB Shocks: a formalization of non-monetary demand shocks.

are common to all models, are set at conventional values; and the stochastic properties of the shocks, which differ across models, are chosen so as to minimize the distance of the model’s predicted volatilities for output, hours and investment from the corresponding volatilities in the US data. The models can then be judged in terms of how well they match these targeted moments and, most importantly, how well they match other salient features of the data.

As can be seen in the table, the belief-augmented RBC model does a good job in matching the US data. It also outperforms, in multiple fronts, the competing New-Keynesian models with either the investment shock, or the consumption shock, or the news shock. In this sense, the “confidence shock” is superior to these three conventional formalizations of “demand shocks.”

The only New-Keynesian model that does as well as the RBC model with confidence shock is the one featuring monetary shock. The problem with that particular model is that its empirical fit rests on allowing for a size of monetary shocks that is an order of magnitude higher than standard estimates of monetary shocks.⁹¹ Nevertheless, the virtual tie between the two models is a measure of the ability of higher-order uncertainty to provide a potent substitute for the quintessential formalization of “demand shocks”.

The reason for these findings can be found in Figure 11. This figure shows the impulse response functions of the model’s key endogenous outcomes to the exogenous shocks in higher-order beliefs. This shock triggers strong comovement in employment, output, consumption, and investment, without a strong comovement in either labor productivity (shown) or inflation (not shown). It is precisely these comovement patterns that are present in the data and that conventional structural shocks have difficulties in matching.

Angeletos, Collard and Dellas (2015) document that these lessons extend to an estimated medium-scale DSGE model that allows for richer propagation mechanisms (sticky prices, habit formation, and investment-adjustment costs) along with multiple structural shocks. In particular, the richer model includes both permanent and transitory TFP shocks, news shocks, investment shocks, consumption-specific shocks, and fiscal and monetary shocks. In the absence of the confidence shock, the estimated model delivers a similar picture as the one found in the extant DSGE literature. But once the confidence shock is included in the model, the picture changes dramatically: the confidence shock is estimated to account for more than half of the business-cycle volatility in output and other macroeconomic quantities.

Huo and Takayama (2015a) complement the above findings by studying a version of the RBC model that allows for a similar type of confidence shocks as in Angeletos and La’O (2013). Relative to Angeletos, Collard and Dellas (2015), the first key difference is the imposition of the common-prior assumption, which disciplines

⁹¹Relatedly, if we look at the model’s predicted moments for inflation and the nominal interest rate, these moments are far away from their empirical counterparts.

the magnitude and the persistence of the fluctuations in higher-order beliefs that can be entertained in the theory. The second key difference is the use of forecast data, which further discipline the aforementioned fluctuations. Despite these additional “constraints”, Huo and Takayama (2015a) reach a similar bottom line as Angeletos, Collard and Dellas (2015): confidence shocks are shown to generate realistic business-cycle patterns and to account for a sizable component of the volatility in the data.

Whether one takes these quantitative findings at face value or not, they offer a very different message from Ramey (2016) in this *Handbook*. That chapter concludes: “we are much closer to understanding the shocks that drive economic fluctuations than we were twenty years ago”. We contend that this conclusion hinges on structural interpretations of the data that rule out frictions in coordination and forces akin to market psychology and animal spirits. Once these elements are taken into account, both the interpretation of existing structural VAR evidence and the quantitative performance of existing DSGE models can be seriously upset. The state of our understanding is certainly *different* from what it was twenty years ago. But it is not necessarily closer to the “truth”, at least not insofar as the “truth” assigns a prominent role to the type of frictions we have studied in this chapter.⁹²

Last but not least, the works we have reviewed in this subsection suggest that aggregate demand can be “deficient” during or in the aftermath of certain recessions, not only because of sticky prices and constraints on monetary policy such as the zero-lower bound, but also—and perhaps primarily—because of difficulties in the coordination of the decentralized choices of firms and consumers. How these ideas apply to the Great Recession is an important open question.

8.8 Financial Markets

In this subsection we discuss how incomplete information can help explain certain asset-pricing puzzles, such as the deviation of asset prices (or exchange rates) from fundamentals, or momentum. To this goal, we use a simple, forward-looking, asset-pricing model with incomplete information, which is the backbone of Futia (1981), Singleton (1987), Allen, Morris and Shin (2006), Bacchetta and Van Wincoop (2006), Kasa et al. (2007), Rodina and Walker (2014), and others. All these papers share the same key structural equation, namely condition (39) below, but make different assumptions about the information structure and the underlying stochastic processes.

There is a continuum of agents, or traders, who participate in a competitive market for a risky asset. The market operates for $T + 1$ periods, where T can be either finite or infinite. In any period $t \leq T$, the demand for the asset of any given trader i is proportional to her expected excess return of the asset:

$$q_{it} = \mathbb{E}_{it}d_t + \beta\mathbb{E}_{it}p_{t+1} - p_t, \tag{39}$$

where q_{it} is the net position of trader i in period t , p_t is the price of the asset in period t , and d_t is the dividend the asset pays at the end of period t . In period $t = T + 1$ (which is valid only if T is finite), we instead have $q_{iT+1} = \mathbb{E}_{iT+1}d_{T+1} - p_{T+1}$, because there is no further trading after that period. In what follows, we focus on what happens in $t \leq T$.

Condition (39), or a slight variation of it, is present in all the aforementioned papers. The papers differ only in the specification of the information structure and of the stochastic properties of d_t and s_t .

⁹²Here it is important to clarify the following. Whenever macroeconomists talk about “shocks” and “propagation mechanisms”, they do *not* talk directly about the data (“the real world”). Shocks and propagation mechanisms are theoretical objects, which are defined in specific models, or in associated structural VARs, and which are used to *interpret* the data.

In what follows, we treat condition (39) as a primitive structural equation. It is worth noting, however, that this condition can be micro-founded in (at least) two ways.

First, suppose that the traders are risk neutral but need to pay a quadratic “holding” cost for any net position they hold in the risky asset. This means that the trader i 's per-period payoff is given by

$$U_{it} = -p_t q_{it} - \frac{1}{2} q_{it}^2 + \left(d_t + \frac{p_{t+1}}{1+r} \right) q_{it}.$$

where r is the risk-free rate. It follows that the optimal demand is given by condition (39), with $\beta = \frac{1}{1+r}$.

Alternatively, suppose that the traders are myopic and have CARA preferences, meaning that their per-period payoff is given by

$$U_{it} = -\frac{1}{\gamma} \exp \left\{ -\gamma \left(-p_t q_{it} + \left(d_t + \frac{p_{t+1}}{1+r} \right) q_{it} \right) \right\},$$

where $\gamma > 0$ is the coefficient of absolute risk aversion. Suppose furthermore that the information structure is Gaussian and let us focus on equilibria in which the equilibrium price is itself a Gaussian signal of the dividend. Then, the optimal demand for the asset i is given by

$$q_{it} = \frac{\mathbb{E}_{it} [d_{it} (1+r) + p_{t+1}] - (1+r) p_t}{\gamma \text{Var}_{it} (d_{it} (1+r) + p_{t+1})},$$

where $\text{Var}_{it}(X)$ denotes the variance of X conditional on the information of trader i in period t . Finally, suppose that all the underlying shocks have known time-invariant variances and let us focus on equilibria in which the conditional risk faced by the typical trader is also time-invariant. This means that $\text{Var}_{it} (d_{it} (1+r) + p_{t+1}) = V$, for some known constant V . The expression above then reduces to (39) if we let $\beta \equiv \frac{1}{1+r}$ and, without serious loss of generality, set $\gamma V = (1+r)$.

We now put aside these micro-foundations and focus on how the equilibrium price dynamics are affected by the incompleteness of information. To start with, let us make the following basic, but important, observation. In any $t \leq T$, a trader faces two kinds of uncertainty in the return to her investment choice: one about the end-of-period dividend, d_t , and another about the next-period price, p_{t+1} . The first kind of uncertainty regards an exogenous variable; it therefore corresponds to what we have called fundamental uncertainty. The second kind of uncertainty has to do with the demand of future traders; it is therefore an example of strategic uncertainty.

Let s_t denote the exogenous supply of the asset, or equivalently the (negative of the) exogenous demand of any “noise traders”. Market clearing requires $\int q_{it} di = s_t$. It follows that the equilibrium price satisfies

$$p_t = \beta \bar{\mathbb{E}}_t p_{t+1} + \bar{\mathbb{E}}_t d_t - s_t, \tag{40}$$

where $\bar{\mathbb{E}}_t$ denotes, as usual, the average expectation at t . Next, define the “fundamental” for this environment as $\theta_t \equiv \bar{\mathbb{E}}_t d_t - s_t$. This is somewhat at odds with the rest of our paper because θ_t now contains an average of first-order beliefs. If one finds this to be confusing, one can henceforth limit attention to the special case in which d_t is known, in which case one can also let $\theta_t \equiv d_t - s_t$. An additional advantage of this special case is that it isolates the role of strategic uncertainty: in the eyes of each trader, the only remaining uncertainty is the one about the behavior of future traders (as manifested in future prices).⁹³ Either way, condition (21)

⁹³The only minor caveat with this simplification is the following. If T is finite and the last-period dividend is known, the agents face no uncertainty in the last period of trading. Under the first of the two micro-foundations discussed above, the pricing

can be rewritten as

$$p_t = \beta \bar{\mathbb{E}}_t p_{t+1} + \theta_t, \quad (41)$$

which is essentially the same as the forward-looking condition we encountered in Subsection 7.8. This underscores that the type of strategic uncertainty that is relevant in the present context regards the *future* choices of other agents, as opposed to contemporaneous type of strategic interaction featured in either our abstract static framework or the business-cycles applications studied in Subsections 8.1–8.5.

As in Subsection 7.8, let

$$z_t^0 \equiv \theta_t, \quad z_t^1 \equiv \bar{\mathbb{E}}_t z_{t+1}^0 = \bar{\mathbb{E}}_t \theta_{t+1},$$

and, for all $j \geq 2$,

$$z_t^j \equiv \bar{\mathbb{E}}_t z_{t+1}^{j-1} = \bar{\mathbb{E}}_t \{ \bar{\mathbb{E}}_{t+1} \{ \dots \{ \theta_{t+j} \} \dots \} \},$$

Iterating (41) yields

$$p_t = \sum_{j=0}^{T-1} \beta^j z_t^j. \quad (42)$$

It follows that the equilibrium price in any given period depends, not only on today's fundamentals (z_t^0) and today's forecasts of tomorrow's fundamentals (z_t^1), but also on today's forecasts of tomorrow's forecasts of the fundamentals two periods ahead (z_t^2), and so on.

This result illustrates that higher-order beliefs of *future* fundamentals can be important determinants of asset prices. Importantly, this can be true even if the current fundamentals happen to be common knowledge: even if all traders are able to reach a common belief about θ_t in *every* period t , they may still fail to reach a common belief about the future demand-and-supply conditions, and can therefore disagree about future price movements. Finally, the relevant notion of the fundamentals now contains, not only the dividend, but also the supply of the asset or, equivalently, the residual demand of noise traders.

The related literature proceeds by making different special assumptions about the length of the horizon T , the stochastic process for the dividend and the supply, and the information that is available in each period. For example, Rondina and Walker (2014) assume that T is infinite, fix the dividend to a known constant (which can be normalized zero), and allow the supply shock to follow an autocorrelated process; they then focus on how higher-order uncertainty about the history of the underlying supply shock and therefore also about the future path of prices can persist even if the traders observe the entire history of past prices. By contrast, Allen, Morris and Shin (2006) assume that T is finite, that the supply shock is i.i.d. over time, and that the asset pays out a dividend only in the last period; they then focus on the higher-order uncertainty the traders face with regard to the final-period dividend. In the rest of this subsection, we review some of the key findings of the latter paper as well as that of Bacchetta and Van Wincoop (2006).

To nest Allen, Morris and Shin (2006) in our setting we let the dividend be zero in all but the last period, namely $d_{T+1} = \vartheta \sim \mathcal{N}(y, \sigma_\vartheta^2)$ and $d_t = 0 \forall 1 \leq t \leq T$; let the supply shock s_t be i.i.d. over time, drawn from $\mathcal{N}(0, \sigma_s^2)$; and finally assume that the traders receive no exogenous information about the supply shocks. The equilibrium price p_t always reveals some information about the current supply shock s_t . However, because the supply shock is i.i.d. over time, and because the traders receive no exogenous information about future supply shocks, we have that $\bar{\mathbb{E}}_t s_{t+j} = 0$ for all t and all $j > 1$. As a result, condition (42) reduces to the

condition $p_{T+1} = d_{T+1} - s_{T+1}$ is still valid. Under the second micro-foundation, however, the traders face no uncertainty and arbitrage now imposes $p_{T+1} = d_{T+1}$. For this case, we must therefore redefine the last-period fundamental as $\theta_{T+1} \equiv d_{T+1}$. (Note that we still have $\theta_t \equiv d_t - s_t$ for all $t \leq T$.)

following:⁹⁴

$$p_t = \beta^{T-t} \bar{\mathbb{E}}_t \bar{\mathbb{E}}_{t+1} \cdots \bar{\mathbb{E}}_{T+1}[\vartheta] - s_t.$$

In a nutshell, Allen, Morris and Shin (2006) is a special case in which the higher-order uncertainty has a relatively small dimension: it regards only the last-period dividend, as opposed to the entire path of future dividends and supply shocks. Without any loss, we henceforth set $\beta = 1$.

Consider, as a benchmark, the scenario in which the supply is either fixed or commonly known. In equilibrium, the price p_t would perfectly reveal the value of $\bar{\mathbb{E}}_t \bar{\mathbb{E}}_{t+1} \cdots \bar{\mathbb{E}}_{T+1}[\vartheta]$. Under a common prior, the average belief of a random variable can be commonly known only if all the agents share a common belief about that variable. It follows that, in this benchmark scenario, the observation of the equilibrium price induces all traders to form the same belief about the underlying ϑ at all t , and therefore also the same belief about p_{t+1} at all $t < T$. By the same token, there is no speculative trading in this benchmark: all traders would choose the same position and would expect to make zero profits from their trades.

Consider then the alternative case in which the supply shock is present and unknown. In general, the equilibrium price is then only a noisy signal of the average belief $\bar{\mathbb{E}}_t \bar{\mathbb{E}}_{t+1} \cdots \bar{\mathbb{E}}_{T+1}[\vartheta]$. It follows that different traders may maintain different beliefs about this object in equilibrium, which in turn means that they can also maintain different beliefs about either the capital gain $p_{t+1} - p_t$ they can make at any $t < T$, or the final-period yield $\vartheta - p_t$ they can make at $t = T$. In other words, speculation obtains in equilibrium.

The above insights are general. To obtain sharper predictions, however, Allen, Morris and Shin (2006) impose that all the information available to trader i in period t is given by the combination of the history of prices along with a private signal of the form $x_{it} = \vartheta + \epsilon_{it}$, with $\epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon)$ for some $\epsilon > 0$. This in turn is justified by assuming that the traders are short-lived and do not observe any of the private signals of the earlier generations of traders, although they do observe the public history of prices.⁹⁵

With the assumed information structure, one can guess and verify the existence of an equilibrium in which the price in each period t conveys the same information as a Gaussian public signal of the form

$$z_t = \vartheta - \chi_t s_t,$$

for some deterministic scalars $\chi = (\chi_t)_{t=0}^T$, which are themselves determined by the trading strategies of the agents. Together with the assumption that the private signal is also Gaussian, the above property guarantees a tractable dynamic structure for the posteriors of the traders—namely, a Kalman filter over the time-invariant variable ϑ . As it is standard in noisy rational-expectations settings, the scalars $\chi = (\chi_t)_{t=0}^T$ are then characterized by solving the fixed-point relation between (i) the forecasts that each trader forms on the basis of the available price signals and (ii) the price signal generated by the joint behavior of all the traders. The following two results emerge.

Proposition 33. *There exist weights $\{\lambda_t\}$ with $0 < \lambda_T < \lambda_{T-1} < \cdots < \lambda_2 < \lambda_1 < 1$ such that*

$$E_s(p_t) = \lambda_t y + (1 - \lambda_t) \vartheta$$

where $E_s(\cdot)$ denotes the average value over all the possible supply shocks (equivalently, the expectation con-

⁹⁴Strictly speaking, the equilibrium in Allen, Morris and Shin (2006) entails time-varying risk, which shows up in the demand for the asset (because of the assumption of CARA preferences) and thereby also in the equilibrium condition of the price. Here, we abstract from this nuisance.

⁹⁵For a discussion of this assumption and for the alternative of allowing long-term traders, see the original paper (Allen, Morris and Shin, 2006) and the references therein.

ditional on y and ϑ).

Proposition 34. *For all $t < T$*

$$E_s (|p_t - \vartheta|) > E_s (|\bar{E}_t(\theta) - \vartheta|).$$

Proposition 33 is reminiscent of the inertia effect we documented earlier: the equilibrium price is anchored to the initial prior y and converges to the fundamental ϑ only sluggishly over time. By the same token, the price path exhibits “momentum”: following certain innovations (namely, an innovation in ϑ), prices under-react in the short run relative to the long run; equivalently, they exhibit a drift reminiscent of the ones documented in empirical studies on momentum.⁹⁶

Proposition 33 complements the above lesson by establishing that the gap between the fundamental and the equilibrium price is more volatile than the gap between the fundamental and the average belief of it. This kind of “excess volatility” reflects the variation that obtains in higher-order beliefs for given first-order beliefs. Proposition 34 is therefore reminiscent of our earlier result regarding the role of incomplete information in sustaining animal spirits.

Combined, these results illustrate how incomplete information and higher-order uncertainty can help standard asset-pricing models accommodate momentum and other interesting volatility patterns in asset prices. In an insightful earlier contribution, Allen, Morris and Postlewaite (1993) show that higher-order uncertainty can also accommodate “bubbles”: there can exist events in which the equilibrium price is higher, not only than the average expectation of the fundamental, but also than the expectation of *every* trader. This particular possibility is ruled out in the present setting because the Gaussian specification implies that there always exists a trader whose private signal is arbitrarily high, and therefore whose forecast of ϑ is higher than the equilibrium price. Nevertheless, the excess volatility documented in Proposition 34 can be seen as a variant of the same possibility.

The paper by Allen, Morris and Postlewaite (1993) contains also an insightful discussion of the following issues: the appropriate definition of “fundamentals” in asset-pricing models; the close relation between incomplete-information common-prior settings like the one studied here and the kind of symmetric-information heterogeneous-prior settings studied in, inter alia, Harrison and Kreps (1978) and Scheinkman and Xiong (2003); and the potential richness of the relevant state space. The latter means that the revelation of information through prices can be quite limited in practice, or in sufficiently sophisticated models, even though it can be large in more simplistic models, therefore leaving significant room for higher-order uncertainty to drive asset prices. Complementary in this regard is also the message of Rondina and Walker (2014), although a quantitative evaluation of this insight is still missing.

Let us now shift attention to Bacchetta and Van Wincoop (2006). This paper uses a version of the framework we have introduced in this subsection to argue that incomplete information may explain the disconnect of exchange rates from macroeconomic fundamentals at short to medium horizons. In their model, the asset price is the nominal exchange rate between two countries, which we denote by e_t .⁹⁷ The key equilibrium condition is given by the following:⁹⁸


$$e_t = \beta \bar{E}_t e_{t+1} + f_t - b_t,$$

⁹⁶If we replace the private signals that the traders observe in each period t with a public signal of equal precision, then the weight λ_t remains positive but becomes smaller at each t . The dynamic effect documented above is therefore the product, not just of slow learning, but also of the heterogeneity of the information.


⁹⁷In Bacchetta and Van Wincoop (2006), the exchange rate is denoted by s_t ; we use the alternative notation, e_t , in order to avoid any confusion with our earlier notation, which lets s_t denote the supply of the asset.

⁹⁸As in our discussion of Allen, Morris and Shin (2006), we omit some potential time-varying coefficients on f_t and b_t .

Panel A Impulse Response Functions in Heterogeneous Information Model



Panel C Percent contribution b-shocks to $\text{var}(s_{t+k}-s_t)$



Panel D Connection between Exchange Rate and Observed Fundamentals: R^2 of regression of $s_{t+k}-s_t$ on observed fundamentals.*




Figure 12: Exchange Rate Disconnect in the Short and Long Run.

where f_t is the observable difference in money supply between two countries—which is what the authors interpret as the macroeconomic fundamentals—and b_t is an unobserved hedging demand. It follows that Bacchetta and Van Wincoop (2006) is nested in our setting simply by the following change in notation: $p_t = e_t$, $d_t = f_t$, and $s_t = b_t$.

The model is closed by assuming the following stochastic structure. The macroeconomic fundamental, f_t , follows a general MA process, $f_t = D(L)\epsilon_t^f$ where $D(L) = d_1 + d_2L + \dots$, L is the lag operator, and $\epsilon_t^f \sim N(0, \sigma_f^2)$. The hedging demand, b_t , follows an AR(1) process: $b_t = \rho_b b_{t-1} + \epsilon_t^b$, where $\epsilon_t^b \sim N(0, \sigma_b^2)$ and $\rho \in [0, 1)$. In each period, all traders observe the past and the current values of the fundamental and the exchange rate, but not of the hedging demand. In addition, each trader receives a private signal about the *future* fundamental. This signal is given by $x_t^i = f_{t+\Delta} + \epsilon_t^{xi}$, where $\epsilon_t^{xi} \sim \mathcal{N}(0, \sigma_x^2)$ is independent from f_τ for all τ , as well as of the noises in other agents' signals, and $\Delta \geq 1$.

The model's predictions are illustrated in Figure 12, which we borrow from Figure 2 in Bacchetta and Van Wincoop (2006).⁹⁹ Panel A shows that the exchange rate responds sluggishly to innovations in the fundamental, a recurring theme of this chapter. At the same time, the exchange rate responds non-trivially to shocks in hedging demand, and more so than in the complete-information variant of the model. Both effects originate in rational confusion at the individual level, but get amplified by our familiar beauty-contest mechanism.

The amplification effect of hedging demand is evident in Panel C, which reports the contribution of the hedging demands to the variance of the exchange-rate change $e_{t+k} - e_t$ at different horizons $k \geq 1$. Although hedging shocks contribute to short-run volatility in both the incomplete-information and common-knowledge

⁹⁹Panel B is omitted because it is not useful for our discussion.

versions of the model, their contribution is far greater in the former case. The mirror-image of this result is Panel D, which reports the R^2 of the regression of $e_{t+k} - e_t$ on the realization of the fundamentals up to $t+k$. In the short run, the R^2 is far lower in the incomplete-information version of model than in its common-knowledge counterpart. In the long run, R^2 converges to 1 in both models, reflecting the fact that both first- and higher-order beliefs converge to the realized fundamental as time passes and more and more information is accumulated. However, the convergence is slower in the incomplete-information version, due to the fact that higher-order beliefs converge more slowly than first-order beliefs. These properties are consistent with empirical findings that macroeconomic fundamentals have weak explanatory power for exchange rates in the short to medium run (Meese and Rogoff, 1983), but play a much more important role over longer horizons.

Complementary to the above works are Biais and Bossaerts (1998), Kasa et al. (2007), Makarov and Rytchkov (2012), and Rondina and Walker (2014). The first paper is an early contribution that also touches on the beauty-contest aspect of asset markets and explores on the distinct positive implications of incomplete information and heterogeneous priors in a finite-horizon model with a finite number of agents. The other three papers use frequency-domain techniques to solve infinite-horizon models, where there are recurring shocks to fundamentals and the traders' beliefs remain perpetually heterogeneous. Combined, these papers show how higher-order belief dynamics can explain apparent violations of variance bounds and other cross-equation restrictions that guide the empirical test of representative-agent models; how they can accommodate momentum; and how can also generate endogenous boom-and-bust cycles.

Two other important contributions are Cespa and Vives (2012, 2015). These papers revisit the question of whether asset prices are closer or further away from the underlying fundamentals than the consensus (average) forecast of fundamentals. Whereas the aforementioned works by Allen, Morris and Shin (2006), Bacchetta and Van Wincoop (2006), and others answer this question always in the affirmative, Cespa and Vives (2012, 2015) show that the answer depends critically on the persistence of the shocks to liquidity (noise) trades, which is typically assumed to be zero in the literature.

More specifically, Cespa and Vives (2012) studies a model similar to that in Allen, Morris and Shin (2006), except for two modifications: the informed, rational traders are long-lived; and the unobserved shocks to liquidity (noise) trades are allowed to be persistent. When this persistence is sufficiently low, equilibrium prices are systematically farther away from fundamentals than the average expectation, a pattern consistent with Allen, Morris and Shin (2006) and Bacchetta and Van Wincoop (2006). But when liquidity trades are sufficiently persistent, the opposite pattern emerges: rational traders chase long-term returns rather than short-term speculative gains, causing prices to be systematically closer to fundamentals than average expectations. Cespa and Vives (2015) provide a complementary result in a model where rational traders are short-lived. It nevertheless remains true that incomplete information can generate interesting time-series patterns, such as momentum.¹⁰⁰

All in all, the works we have reviewed in this subsection illustrate how higher-order uncertainty can offer a parsimonious explanation to various puzzles in asset pricing and international finance. These works are part of a broader literature that uses informational frictions—although not always higher-order uncertainty—in asset-pricing contexts. Grossman and Stiglitz (1980), Hellwig (1980), and Kyle (1985) are classics, which emphasize the aggregation of information through equilibrium prices. More recently, Albagli, Hellwig and Tsyvinski (2014, 2015) develop an extension that accommodates more general specifications of the joint distribution of the dividend and the private signals. This helps shed light on the differential effect that the

¹⁰⁰What is more, Cespa and Vives (2015) show that the combination of incomplete information with short-term horizons may open the door to multiple equilibria, which are ranked both in terms of their informational efficiency and in terms of the strength of the momentum effect.

heterogeneity of information can have on different claims on the same underlying asset, such as the bonds and the stocks of the same firm. Hassan and Mertens (2014a), on the other hand, develop a framework for augmenting DSGE models with a noisy rational-expectations asset market, and use that to study the macroeconomic effects of the news contained in asset prices about future fundamentals. Notwithstanding the contributions of these papers, it is useful to note that none of them study the beauty-contest effects that are the core of the papers we reviewed above, for they study setting in which either the fundamentals become common knowledge at the end of each period (Hassan and Mertens, 2014a) or trading takes place only once (all the other papers mentioned in this paragraph).¹⁰¹

More closely related to the papers we have reviewed above are, instead, Angeletos et al. (2010) and Hassan and Mertens (2011, 2014b). Angeletos et al. (2010) show that the information spillover between the real and the financial sector induces beauty-contest features in IPO activity and real investment, leading to excessive waves of optimism and pessimism in investment and asset prices. Hassan and Mertens (2011, 2014b) show how financial trading may amplify small correlated biases in traders' beliefs, while also reducing the individual incentive to correct one's beliefs from such biases, and how this mechanism can ultimately distort real investment. In these papers, strategic complementarity emerges endogenously and is key to the obtained results.

Finally, there is a growing strand of the asset-pricing literature that shifts the focus to the acquisition of information or the attention allocation (e.g., Veldkamp, 2006; Veldkamp and Nieuwerburgh, 2009, 2010; Abel, Eberly and Panageas, 2007, 2013), but is beyond the scope of this chapter. For more thorough discussions of the broader asset-pricing literature on informational frictions, we refer the reader to Vives (1993) and Veldkamp (2011).

We conclude by mentioning a line of work that touches on similar mechanisms, albeit with different modeling techniques. This is the asset-pricing literature on heterogeneous beliefs and speculation, such as Harrison and Kreps (1978) and Scheinkman and Xiong (2003). What is common between this literature and the works we reviewed above is that equilibrium trades and asset prices are driven by higher-order beliefs. What is different is that variation in the relevant higher-order beliefs is engineered with the combination of heterogeneous priors and the arrival of public information over time, as opposed to dispersed private information. This means a great deal of tractability, the absence of which could have impeded the discovery of important results. But it also comes with three potential costs. First, this approach permits the theorist to choose higher-order beliefs at will, a freedom that must be exercised with great moderation. Second, this approach shuts down the role of the price mechanism in aggregating information and coordinating beliefs, a mechanism that is of interest on its own right. Last but not least, this approach invites tricky normative questions, such as which prior must be used to evaluate the welfare implications of different allocations.¹⁰² Whether one should use this approach or the more "conservative" one employed by the papers we reviewed in this subsection is ultimately a function of one's research objectives.¹⁰³

¹⁰¹This statement should not be misinterpreted as a claim that that higher-order beliefs are entirely irrelevant. Because the agents extract a signal from prices, and because this signal is endogenous to the trades of other agents, higher-order beliefs matter in all the mentioned papers for essentially the same reason as the one suggested by Townsend (1983). But this mechanism, which has to do with the interpretation of endogenous signals, is different from the mechanism we discussed above, which has to do with speculation.

¹⁰²One possible answer to this last question is that the planner must evaluate the utility of each agent under the agent's *own* prior. This is effectively the assumption made in Debreu's classic proof of the two welfare theorems in economies with subjective probabilities. There are, however, other plausible, and potentially preferable, answers to the aforementioned question. See, for example, Brunnermeier, Simsek, and Xiong (2014) for a welfare criterion that instructs the planner to discard an allocation only if it is deemed inferior under multiple priors at once.

¹⁰³The discussion of this paragraph echoes our earlier discussion in Section 7.10 regarding solution methods and the heterogeneous-prior approach taken in Angeletos, Collard, and Dellas (2015). Similar points apply to the recent literature

9 Efficiency and Policy Implications

Although we often touched on policy implications in the preceding sections, we have not systematically studied either the normative content of the positive properties we have documented, or the implications of coordination frictions for the type of Ramsey policy problems that are commonly used in the study of optimal fiscal and monetary policy. In the RBC model considered in Subsection 8.1, we simply abstracted from policy altogether. In the monetary models studied in Subsection 8.3, nominal GDP was assumed to follow an exogenous stochastic process, bypassing the issue of what kind of monetary policies may or may not justify this assumption. Finally, in the applications of global games reviewed in Section 5, we restricted attention to specific policy instruments, such as a bailout from a lender of last resort in the context of bank runs and debt crises.

In this section, we take a different, and complementary, approach. We put aside any specific context and seek to identify a set of policy insights that may be robust across a variety of applications. In this regard, the approach taken in this section mirrors the one we have taken throughout this paper with regard to the study of the *positive* implications of frictions in information and coordination; the key difference is that we now shift attention to *normative* properties.

More specifically, we proceed as follows. First, we propose a notion of constrained efficiency, which suits our purposes, and discuss the ways in which the constrained efficient strategy may differ from either the first best or the equilibrium. Next, we compare the constrained efficient strategy to the equilibrium one within the class of beauty-contest games studied in Sections 7 and 8. This in turn paves the way to the last part of this section, where we consider a few applications. These applications regard the social value of information, the related argument in favor or against transparency in central bank communication, and the nature of optimal monetary policy in the presence of frictions in coordination.

Remark. Unlike the positive properties we have emphasized in this chapter, the normative properties of incomplete-information models are more sensitive to the underlying micro-foundations. This is naturally the case because the same strategic effects can be consistent with different externalities across the agents, and therefore the same positive properties can be consistent with different normative properties. A meaningful discussion of efficiency and policy therefore requires one to dig into the specifics of different applications, a task that is beyond the scope of this chapter with the exception of the few applications briefly considered in the last part of this section. What we nevertheless hope to accomplish in this section is to develop a useful way of posing the relevant normative questions.

9.1 Constrained Efficiency

The analysis follows Angeletos and Pavan (2007, 2009). The adopted efficiency concept corresponds to a benevolent planner who has the power to dictate to the agents how to act on the basis of their information, but can not collect information from one agent and send it to another. It therefore shares with Hayek (1945) and Radner (1962) the idea that information is dispersed and cannot be communicated to a “center.” Similar efficiency concepts have been used to study the welfare properties of large Cournot games (Vives, 1988) and of business-cycle models with informational frictions (Angeletos and La’O, 2010, Lorenzoni, 2010).

To be concrete, consider the abstract framework introduced in Section 2 and, without serious loss of generality, restrict attention to settings in which the payoff of an agent depends on the actions of others only

on heterogeneous beliefs and leverage, such as Geanakoplos (2010) and Simsek (2013).

through the average actions. Next, for any strategy $k(\cdot) : D_\omega \rightarrow D_k$, let¹⁰⁴

$$W(k(\cdot)) \equiv \mathbb{E} [U(k(\omega_i), K(\Omega), \theta_i)]$$

denote the ex-ante utility attained by this strategy, with the understanding (i) that $K(\Omega) \equiv \int k(\omega) d\Omega(\omega)$ is the average action induced by the particular strategy when the distribution of information is Ω and (ii) that the expectation is taken over the joint realizations of $s_i = (\omega_i, \theta_i)$ and \mathbf{S} . We can then define our efficiency benchmark as follows.¹⁰⁵

Definition 16. *A (constrained) efficient strategy is a strategy $k^*(\cdot)$ that maximizes $W(k(\cdot))$.*

This concept lets the planner internalize how the action of each agent affects the payoff of other agents, which of course does not happen in equilibrium. Similarly to equilibrium, however, the action of any given agent is restricted to depend only on the information that Nature reveals to that particular agent, which means that the planner is precluded from transferring information from one agent to another.

If the planner were able to transfer information from one agent to another, then he would also be able to condition the action k_i of any agent i , not only on the agent's own ω_i , but also on Ω . And if that were true, the planner could instruct the agents to play the first-best actions. But as long as the planner is restricted from transferring information across agents, the first-best actions are generally not attainable. This explains the precise sense in which the notion of efficiency defined above is “constrained”: the constraint is the measurability constraint that the action of each agent can not depend on the private information of other agents.

Although the planner is precluded from eliminating the heterogeneity in information, he can regulate the degree of strategic uncertainty by controlling the degree to which agents respond to their private information. For instance, the planner could eliminate strategic uncertainty altogether by instructing all agents to condition their actions only on publicly available information. Doing so, however, would mean that the agents completely disregard their private information, which is generally socially undesirable.

Studying the constrained efficient strategy therefore allows one to answer the following types of questions: What is the degree of coordination that is socially optimal given the underlying friction in communication? Should the planner use taxes, or other instruments that manipulate incentives, in an attempt to regulate the equilibrium use of information and the resulting strategic uncertainty? Do “animal spirits” justify policy intervention, and if yes of what kind?

A first step towards answering these questions is contained in the following proposition.

Proposition 35 (Efficiency). *Suppose that $k^*(\cdot)$ is an efficient strategy, that $k^*(\omega)$ is in the interior of D_k for all ω , and that U is differentiable. Then,*

$$\mathbb{E} [U_k(k^*(\omega), K^*(\Omega), \theta) + \Lambda^*(\Omega) \mid \omega] = 0 \quad \forall \omega, \tag{43}$$

with

$$\Lambda^*(\Omega) \equiv \int U_K(k^*(\omega), K^*(\Omega), \theta) d\Omega(\omega) \quad \text{and} \quad K^*(\Omega) = \int k^*(\omega) d\Omega(\omega) \quad \forall \Omega.$$

¹⁰⁴To avoid the confusion of strategies and realized actions, we henceforth denote strategies with $k(\cdot)$. That is, whereas k is an object in D_k , $k(\cdot)$ is an object in the space of functions that map D_ω to D_k .

¹⁰⁵Note that Definition 16 imposes symmetry: the planner must choose the same strategy for all agents. This is without serious loss given the symmetry of the environment and the purposes of the exercises we conduct in the rest of this section, but it could matter in general. Relatedly, we envision a planner that maximizes ex-ante utility “behind the veil of ignorance”, as opposed to a planner that favors any particular group of agents.

Condition (43) is simply the first-order condition of the planner’s problem. The equilibrium counterpart is the best-response condition of the individual agent, which (assuming again an interior solution) is simply given by the following:

$$\mathbb{E}[U_k(k(\omega), K(\Omega), \theta) | \omega] = 0 \quad \forall \omega.$$

The only essential difference between the two conditions is therefore the presence of $\Lambda^*(\Omega)$. This term captures the *average* externality that agents impose on one another when the cross-sectional distribution of information is Ω and when all the agents follow the strategy k^* . In this regard, condition (43) is similar to the condition that characterizes the first best. What is different is that agents need not share a common belief about $\Lambda^*(\Omega)$, due to the incompleteness of the information.¹⁰⁶ We conclude that the efficient strategy is similar to the first best in the sense that the agents are instructed to internalize the externalities they impose on one another (if any), but departs from the first best in the sense that the agents hold different beliefs about the value of these externalities (because, and only because, they are unable to communicate with one another).

A complementary interpretation of Proposition is the following. Note that we can represent the efficient strategy as the equilibrium of a fictitious game, in which the information structure remains the same as in the original game but the payoffs are appropriately modified, so that the best-response conditions of the fictitious game coincide with the first-order conditions of the planner’s problem in the original game.¹⁰⁷ Under this representation, condition (43) describes the socially optimal response of each agent to her beliefs of the actions of others. In this sense, the efficient strategy helps identify the degree of coordination that is socially optimal given the underlying friction in communication.

We develop below more concrete translations of these elementary insights by restricting attention to the class of beauty-contest games studied in Section 7 and to some of the related applications studied in Section 8. This focus is partly motivated by the fact that the efficiency concept defined above has not been sufficiently explored within the global-games literature.

A notable exemption to the last statement is Schaal and Taschereau-Dumouchel (2015), which characterizes, inter alia, the constrained efficient allocation within the global-games adaptation of an RBC model that we discussed in Subsection 5.6. See also Sakovics and Steiner (2012) and Frankel (2016), who study optimal subsidies within coordination games akin to the adoption of a network technology.

We conclude with two remarks regarding the applicability of the efficiency concept we have introduced.

First, it is straightforward to adapt this concept and the results developed in the next subsection to micro-founded macroeconomic models, at least insofar as one makes appropriate modeling choices. In particular, Angeletos and La’O (2010), Lorenzoni (2010), and Schaal and Taschereau-Dumouchel (2015) study business-cycle models in which the incompleteness of information introduces strategic uncertainty, without however introducing incomplete-risk sharing in consumption. This property is achieved by assuming that all relevant agents (firms, workers, etc) belong to a “big family” whose income and consumption is fully diversified against the idiosyncratic noise in the information of different agents. All in all, one can find a direct mapping between the abstract analysis of this section and the specific results of those papers, a point that we revisit below.

Second, it is also possible to extend the notion of efficiency to accommodate certain forms of endogeneity in the information structure, whether this means endogenous *aggregation* of information or endogenous *collection* of information. For examples of the first type, we refer the reader to Amador and Weill (2012),

¹⁰⁶Note that, as long as information is complete, $\Lambda(\Omega)$ remains commonly known even if information is *imperfect*: as with equilibrium, what matters is higher-order uncertainty, not first-order uncertainty.

¹⁰⁷That is, the payoffs of the fictitious game are given by $\tilde{U}(k, K, \theta, \Omega) \equiv \mathbb{E}[U(k, K, \theta) + \Lambda^*(\Omega)k]$.

Angeletos and Pavan (2009), Laffont (1985), Messner and Vives (2005), and Vives (2014); for example of the second type, see Angeletos, La’O and Iovino (2016a), Colombo, Femminis and Pavan (2014), Llosa and Venkateswaran (2015), and Pavan (2015).

9.2 Efficient Coordination in Beauty Contests

We now turn to the study of constrained efficiency within the class of beauty-contests settings introduced in Section 7.1. To guarantee the existence and uniqueness of the efficient strategy, we complement Assumption 4 with the following additional restriction on the payoff structure.

Assumption 10. $U_{kk} + 2U_{kK} + U_{KK} < 0$.

We reach the following characterization of the efficient strategy.

Proposition 36 (Efficient Coordination). *There exist scalars $(\kappa_0^*, \kappa_1^*, \kappa_2^*, \alpha^*)$, pinned down by U , such that the following is true:*

(i) *Whenever information is complete, the efficient action is given by*

$$k_i = \kappa^* (\mathbb{E}_i \theta_i, \mathbb{E}_i \bar{\theta}) \equiv \kappa_0^* + \kappa_1^* \mathbb{E}_i \theta_i + \kappa_2^* \mathbb{E}_i \bar{\theta}. \quad (44)$$

(ii) *Whenever information is incomplete, the efficient strategy solves the following fixed-point relation:*

$$k_i = \mathbb{E}_i [\kappa^* (\theta_i, \bar{\theta})] + \alpha^* \cdot \mathbb{E}_i [K - \kappa^* (\bar{\theta}, \bar{\theta})]. \quad (45)$$

Parts (i) and (ii) are the normative counterparts of, respectively, Propositions 14 and 15. These propositions characterized the equilibrium under, respectively, complete and incomplete information.

By assuming that information is complete, part (i) identifies in effect the first-best action. The possibility that the scalars $(\kappa_0^*, \kappa_1^*, \kappa_2^*)$ differ from their equilibrium counterparts reflects the familiar reasons why equilibrium and first-best allocations can differ in complete-information models. But as with the equilibrium, the key observation here is that only first-order beliefs matter when information is complete.

Part (ii) then shows how the constrained efficient action differs from the first-best one when information becomes incomplete. As with equilibrium, the planner wants each agent to align her action with her forecast of the average deviation in the population. By direct implication, higher-order beliefs matter. But unlike equilibrium, the extent of the desired alignment in actions and the associated role on higher-order beliefs is parameterized by the new scalar α^* , as opposed to the scalar α we encountered before. This new scalar encapsulates the social value of coordination.

Angeletos and Pavan (2007, 2009) analyze how this value can be understood in terms of two more familiar concepts: the welfare cost of inefficiency volatility in the aggregate outcome K ; and the welfare cost of inefficient dispersion in individual actions. In a business-cycle context, these objects map to the volatility of the “output gap” and the inefficient component of the cross-sectional dispersion in relative prices; see Angeletos, Iovino and La’O (2016b). To understand the general logic, note that if the planner induces the agents to coordinate more, then he also induces them to rely more on correlated sources of information, which in turn means more aggregate “mistakes” in actions (i.e., more volatile gaps between the equilibrium and the first-best level of aggregate activity) but also fewer idiosyncratic mistakes (i.e., less inefficient dispersion in the cross section). It follows that varying the level of coordination entails a trade-off between volatility and dispersion. The higher the social cost of the latter relative to the former, the higher the socially optimal degree of coordination (i.e., the higher α^*).

9.3 Policy Implications

What does the result above imply for policy? If α and α^* happen to coincide, then the equilibrium and the efficient allocation exhibit the same sensitivity to higher-order beliefs. Therefore, although the equilibrium inertia and the animal spirits that we repeatedly encountered in Sections 7 and 8 are necessarily a symptom of departure from the first best, neither of them is necessarily a call for policy stabilization, at least insofar as the policy maker is unable to eliminate the incompleteness of information.

Now contrast this observation with conventional wisdom. In standard, complete-information models, animal spirits are tied to multiple, and often Pareto-ranked, equilibria. This has led to the view that animal spirits are *prima-facia* rationale for policy intervention. For instance, the Keynesian tradition dictates that, if the business cycle is driven by animal spirits, then monetary and fiscal policy should be used to stabilize the economy. But if animal spirits are the product of incomplete information, and if α happens to coincide with α^* , then the aforementioned policy prescription is misguided.

Angeletos and La’O (2010) prove, in essence, that $\alpha = \alpha^*$ applies in the RBC model we studied in Subsection 8.1. In other words, the “aggregate demand externality” introduced by product differentiation gives rise to a private value for coordination that is perfectly aligned with its social counterpart

Angeletos and La’O (2012) and Angeletos, Iovino and La’O (2016b) consider a richer framework that allows, *inter alia*, the type of real rigidity studied in Subsection 8.1 to coexist with the type of nominal rigidity studied in Subsection 8.3 and featured in Woodford (2003), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009). Within that framework, it is shown that $\alpha = \alpha^*$ applies when monetary policy replicates flexible prices, as well as that such a policy is optimal when the business cycle is driven either by technology shocks or by shocks to the belief hierarchy about technology. The coincidence of the equilibrium and the efficient degrees of coordination is therefore an important benchmark for business-cycle analysis.

Angeletos and La’O (2012) further show that, similarly to the baseline New-Keynesian framework, the optimal monetary policy replicates flexible-price allocations as long as monetary policy does not have to substitute for missing tax instruments. But unlike that framework, replicating flexible-price allocations does not mean targeting price stability any more. Instead, because efficiency requires that both the quantity and the price of each firm move with its private information about the state of the economy, efficiency also requires that the aggregate price level moves with the average “sentiment” in the economy. This gives another concrete example of how accommodating realistic frictions in coordination can upset existing policy lessons.

Paciello and Wiederholt (2014) discuss additional implications in a variant setting in which firms are rationally inattentive. Monetary policy must substitute for missing tax instruments—there are markup shocks and no subsidies to correct them. For this kind of situation, the standard New-Keynesian model predicts that monetary policy should give up on price stability in order to mimic the missing counter-cyclical subsidy that would have offset the markup shocks. Paciello and Wiederholt (2014) instead show that the optimal policy may now target price stability in order to reduce the incentives of the firms to acquire information about the underlying markup shocks. Once again, existing policy lessons are upset.

A thorough analysis of the policy implications of higher-order uncertainty is beyond the scope of this chapter. Furthermore, the relevant literature is still relatively immature. For instance, all the aforementioned papers study settings in which firms are informationally constrained, but consumers are not; they also rule out the type of asset- and labor-market frictions that seem to be important in understanding business cycles. With the aforementioned examples we therefore only wish to indicate the significance of further investigations of the policy implications of incomplete information within the context of business cycles and monetary policy.

9.4 Welfare Effects of Public Information (and Central-Bank Transparency)

We conclude this section with a brief discussion of another topic, the welfare effects of the public information provided by policy makers, the media, or other sources.

As noted before, public signals can have a disproportionate effect on equilibrium outcomes in the presence of strategic uncertainty, because they become focal points and serve a role akin to coordination devices. In an influential article, Morris and Shin (2002) relied on this observation to show that the provision of public information can contribute to more volatility and thereby also to lower welfare. They then argued that this raises questions about the social value of the information disseminated by the financial media, as well as that it justifies “constructive ambiguity” in central bank communications.

Subsequent work has questioned the applicability of Morris and Shin’s welfare result to a macroeconomic context. Angeletos and Pavan (2004) document the opposite result in an investment game with production spillovers, illustrating how the welfare effects of information depend on the form of externalities, not just the form of complementarity. Angeletos and Pavan (2007) define and characterize the socially optimal degree of coordination in a flexible-class of linear-quadratic games and use this to show that Morris and Shin’s welfare result hinges on assuming that coordination motives are socially unwarranted—a property that need not hold in workhorse macroeconomic models. Finally, starting with Hellwig (2005), a number of more applied works have studied the welfare effects of different types of information in micro-founded monetary models in which nominal rigidity originates from incomplete information; see, e.g., Roca (2008), Walsh (2007), Baeriswyl and Cornand (2010), and Lorenzoni (2010). This has led to a variety of welfare lessons, some pro and some against central-bank transparency.

In a more recent paper, Angeletos, Iovino and La’O (2016b) develop a taxonomy of the welfare effects of information within a micro-founded business-cycle framework that encompasses the aforementioned applied works and disentangles the separate roles played by informational frictions and monetary policy. Importantly, the framework allows the incompleteness of information to be a source of *both* real and nominal rigidity, as in, respectively, Subsections 8.1 and 8.3. Accordingly, the welfare effects of either public or private information can be decomposed into two channels: one working through the real rigidity, and another working through the nominal rigidity.

The first channel is present regardless of the conduct of monetary policy. It also contains a particularly sharp answer to the question of interest: through this channel, more information unambiguously contributes to higher welfare when the business cycle is driven by efficient forces such as technology shocks, and to lower welfare when the business cycle is driven by distortionary forces such as markup shocks.

The second channel hinges on the conduct of monetary policy. As in the baseline New-Keynesian framework, there is a policy that neutralizes the nominal rigidity; this is the same as replicating flexible-price allocations, recast in the context of models where the nominal rigidity originates from an informational friction as opposed to a Calvo-like friction. At this benchmark, the welfare effects of information are shaped solely by the real-rigidity channel. Away from it, they hinge on whether the provision of more information dampens or amplifies the deviation from the flexible-price benchmark and on whether that deviation was desirable to begin with.

When the business cycle is driven by technology shocks, the policy that replicates flexible prices is optimal. When, instead, the business cycle is driven by markup shocks, a deviation from this policy benchmark is desirable. In the latter case, more information in the hands of private agents can decrease welfare, not only because it exacerbates the inefficiency of the underlying flexible-price fluctuations, but also because it curtails the monetary authority’s ability to combat these fluctuations.

These findings indicate that the welfare effects of information in baseline RBC and New-Keynesian models are closely connected to more familiar normative properties of RBC and New-Keynesian models, and have little, if anything, to do with the mechanism in Morris and Shin (2002). That said, Morris and Shin’s result may well apply to asset-markets applications insofar as speculative trading is a zero-sum game. If this is true, it would upset the conventional wisdom that the informational role of financial markets is welfare improving. It is also unclear—and therefore interesting to explore—whether the result applies to macroeconomic models that give a central role to financial markets, such as those discussed by Brunnermeier and Sannikov (2016), Gertler, Kiyotaki, and Prestipino (2016), and Uhlig and Guerrieri (2016) in this *Handbook*.

Amador and Weill (2010) identify a different reason for why the provision of public information can reduce welfare: the endogeneity of the information role of the price system. More specifically, they study the welfare effects of releasing public information about productivity and/or monetary shocks in a micro-founded model in which agents learn from the distribution of nominal prices. The release of such information induces private agents to rely less on their private sources of information. As this happens, the informational efficiency of the price system is reduced. The release of public information therefore has two competing effects: a beneficial direct one, and an adverse indirect one through the informativeness of the price system. Under certain conditions, the second effect dominates, yielding a negative overall welfare effect.

Vives (1993, 1997) and Amador and Weill (2012) develop related results within the context of a class of dynamic games with social learning. A key result is that, as long as some of the learning is private, the release of public information at some point may increase the overall level of information in the short run at the expense of lowering information in the long run. When agents are patient enough, this opens again the door to the possibility that public information is welfare deteriorating.

Whereas the above papers focus on the *aggregation* of information,¹⁰⁸ Burguet and Vives (2000) focus on the *acquisition* of information. They study a model in which agents can collect private information at a cost and show that the release of public information may reduce the incentives to do so. This finding is reminiscent of a key observation made by Grossman and Stiglitz (1980), namely that the information revealed by asset prices may reduce the incentives for individual traders to collect valuable private information.

Notwithstanding the potential relevance of this last set of results, it is worth noting two facts. First, these results require that private and public information are *substitutes* for one another, which is not necessarily the case; if instead they are complements, the provision of public information can raise both the acquisition and the aggregation of private information.¹⁰⁹ Second, these results operate even in settings in which the payoff of each agent is independent of the actions of other agents; it is an open question whether there is sufficiently interesting interaction between these results and the mechanisms our chapter is preoccupied with.

Last but not least, Baeriswyl and Cornand (2010) identify a potential trade off between the stabilization and the signaling roles of monetary policy. Suppose that the business cycle is driven by inefficient forces such as markup shocks. Suppose further that the monetary authority has some private information about these shocks and/or the level of economic activity.¹¹⁰ To be more concrete, think of the monetary authority knowing something about either the sources or the severity of an inefficient recession. In such a situation,

¹⁰⁸Closely related is also the literature on herds and information cascades (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992; Chamley, 2004). The results described above can indeed be seen as a smooth variant of a key lesson from that literature, namely the possibility that agents can *completely* cease to react to their private information after the revelation of sufficient information about the choices of other agents.

¹⁰⁹This relates to the issue of whether there is strategic substitutability or complementarity in the acquisition of *private* information with the asset-market context; see Barlevy and Veronesi (2000, 2007), Veldkamp (2006), and the discussion in Chapter 4 in Vives (2010).

¹¹⁰This does not require the monetary authority’s information to be more precise than that of any private agent; it only requires that the monetary authority’s information is not a priori publicly available to the private agents.

it is desirable for the monetary authority to withhold its information from the public, because disclosing it would only exacerbate the inefficient fluctuations. Clearly, it is also desirable to act on the basis of such information, by lowering interest rates or, perhaps, by engaging in unconventional policy measures. But note that acting means disclosing: the public will be able to extract a signal from the observed policy action about the state of the economy. It follows that the attempt to stabilize could backfire, which explains why there is a trade off between stabilization and signaling. Baeriswyl and Cornand (2010) proceed to characterize the optimal resolution of this trade off under the assumption that the policy maker can commit on a state-contingent policy rule prior to the arrival of any information. An interesting open question is how this trade off plays out in shaping the optimal policy in the absence of such commitment.¹¹¹

We conclude with an obvious disclaimer. Although many of the papers we have reviewed here touch upon the topic of central-bank communication and transparency, there is a vast literature on the topic that is simply beyond the scope of our chapter. Important early contributions include Cukierman and Meltzer (1986) and Stein (1989); see also the review in Geraats (2002).

10 Conclusion

Since the publication of the first volume of the *Handbook of Macroeconomics*, a growing literature has emerged on the macroeconomic effects of various forms of informational frictions. In this chapter, we have tried to present and synthesize the results of this literature as it relates to coordination, strategic uncertainty, and higher-order beliefs. In the light of this synthesis, we invite the reader to consider two complementary interpretations of the particular departure from standard macroeconomic models undertaken in this chapter:

1. Insofar as signals represent “hard information”, the theoretical advances we have reviewed in this chapter help understand the positive and normative implications of different kinds of information frictions within a variety of contexts, including financial crises and business cycles.
2. Insofar as signals represent “states of mind”, incomplete information is more broadly a vehicle for operationalizing the notion that coordination is imperfect and for developing a more flexible framework of how agents form expectations about endogenous economic outcomes.

The mechanisms we have reviewed apply regardless of the interpretation. The interpretation, however, matters for the mapping of the theory to the real world. For certain issues, such as those that have to do with the aggregation of information or the welfare effects of public information, the first interpretation seems most appropriate. For other issues, such as interpreting the observed variation in expectations and economic outcomes, we would favor the second interpretation.

Underlying the mechanisms we have reviewed in this chapter is the difference between the strategic and the decision-theoretic aspects of informational frictions, a theme that we repeatedly visited in this chapter. We have argued that strategic uncertainty helps formalize frictions in coordination and have explored its distinct positive implications.

Throughout this chapter, we maintained the standard solution concept of rational-expectations equilibrium. While this is a standard practice, it is not necessarily the best one. We are sympathetic to other

¹¹¹Note here a certain parallel to the policy-traps setting of Angeletos, Hellwig and Pavan (2006), which we reviewed in Subsection 5.8. In both papers, there is a trade off between manipulating the incentives of the private agents and disclosing information to them. Apart from the apparent difference in the setting, what distinguishes the two papers is the assumption about commitment: in Angeletos, Hellwig and Pavan (2006), policy is determined in the interim state, after the policy maker has received her information.

approaches that investigate plausible deviations from this practice. In fact, we are open to re-interpreting the departure considered in this chapter as a proxy for relaxing the equilibrium concept.

All in all, the following lesson emerges. Workhorse macroeconomic models, especially those used in the study of business cycles, abstract from the type of frictions in the coordination of expectations and behavior that we have studied in this chapter. As we have shown, this abstraction plays a central role in existing structural interpretations of the data and in the associated policy debate that is based on such structural interpretations. Once realistic frictions in coordination and expectations have been taken into account, our structural interpretation of the data, our view of the real world, and our prescriptions for policy can be significantly altered.

In this chapter we presented various applications of these ideas. We did not, however, offer a detailed empirical and quantitative assessment. This is because the relevant literature is still young and often confounds the decision-theoretic and the strategic aspects of informational frictions. By elucidating this distinction and the mechanisms that operate in a variety of applications, we hope to have offered, not only a sharper understanding and a certain synthesis of the literature, but also some guidance to future empirical and quantitative work.

Appendix: Proofs

Proof of Propositions 1, 2, and 3. See the main text.

Proof of Proposition 4. The continuity of G together with the compactness of D_k guarantees that there always exists a solution to $K = G(K, \mathbf{B})$.

If the economy features either strategic substitutability or weak complementarity, G_K is uniformly bounded from above by 1, which guarantees that the solution is unique for all \mathbf{B} . With K thus determined, the optimal action of any agent i is pinned down by $k_i = g(K, b_i)$. It follows that there exists a unique equilibrium strategy, that is, a unique mapping from the realizations of b_i to those of k_i . By the same token, there is also a unique mapping from the realizations of the cross-sectional distribution of fundamentals, \mathbf{B} , to those of the cross-sectional distribution of actions, \mathbf{K} .

If instead the economy features strong complementarity, let K^* denote the solution of fixed point problem in the definition of strong complementarity. In a small neighborhood to the left of K^* , we have that $G(K, \mathbf{B}) < K$ due to the fact that the derivative $G_K(K, \mathbf{B})$ is locally higher than one. Let k_{low} be the lower bound of D_k . At $K = k_{low}$, we necessarily have that $G(K, \mathbf{B}) \geq K$. By the continuity of G , there must then exist another fixed point to G . In either case, once K is determined as a solution to $K = G(K, \mathbf{B})$, the optimal action of any agent i is given by $k_i = g(K, b_i)$. It follows that there exists multiple equilibrium strategies, each one associated with a different mapping from \mathbf{B} to \mathbf{K} .

Proof of Proposition 5. See the main text.

Proof of Proposition 6. See Subsection 4.5.

Proof of Proposition 7. We start by studying the set of *monotone* or *threshold* equilibria, that is, equilibria in which the strategy of an agent satisfies the following property: for any realization of z , there is a threshold $x^*(z)$ such that an agent attacks if and only if $x \leq x^*(z)$.

When the agents follow such monotone strategies, the aggregate size of the attack is decreasing in θ , so that there is also a threshold $\theta^*(z)$ such that the status quo is abandoned if and only if $\theta \leq \theta^*(z)$. A monotone equilibrium is identified by thus identified by threshold functions x^* and θ^* .

The proof then proceeds in four steps. In step 1, we characterize the equilibrium θ^* for given x^* . In step 2, we characterize the equilibrium x^* for given θ^* . In step 3, we combine the two conditions to establish existence and to study the determinacy of monotone equilibria. In step 4, we conclude the proof by noting that, whenever the monotone equilibrium is unique, this gives also the unique rationalizable strategy.

Step 1. For given realizations of θ and z , the aggregate size of the attack is given by the mass of agents who receive signals $x \leq x^*(z)$. That is, letting $K(\theta, z)$ denote the size of attack when the fundamental is θ and the signal is z , we have

$$K(\theta, z) = Prob(x \leq x^*(z)|\theta) = \Phi(\sqrt{\alpha_\epsilon}(x^*(z) - \theta)),$$

where $\alpha_\epsilon = \sigma_\epsilon^{-2}$ and $\alpha_\zeta = \sigma_\zeta^{-2}$ denote, respectively, the precision of private and public information and Φ denotes the c.d.f. of the standardized Normal distribution. Note that $K(\theta, z)$ is decreasing in θ , so that regime change occurs if and only if $\theta \leq \theta^*(z)$, where $\theta^*(z)$ is the unique solution to

$$K(\theta^*(z), z) = \theta^*(z).$$

Rearranging the above gives following relation between the thresholds x^* and θ^* :

$$x^*(z) = \theta^*(z) + \frac{1}{\sqrt{\alpha_\epsilon}} \Phi^{-1}(\theta^*(z)). \quad (46)$$

Step 2. Given that regime change occurs if and only if $\theta \leq \theta^*(z)$, the payoff of an agent is

$$\mathbb{E} [U(k, K(\theta, z), \theta) \mid x, z] = k(b \Pr[\theta \leq \theta^*(z) \mid x, z] - c).$$

The posterior of the agent is a Normal distribution with mean $\delta x + (1 - \delta)z$ and precision α , namely

$$\theta \mid x, z \sim \mathcal{N}(\delta x + (1 - \delta)z, \alpha^{-1}),$$

where $\delta \equiv \alpha_\epsilon / (\alpha_\epsilon + \alpha_\zeta)$ captures the relative precision of private information and $\alpha \equiv \alpha_\epsilon + \alpha_\zeta$ captures the overall precision of information. Hence, the posterior probability of regime change is

$$\Pr[\theta \leq \theta^*(z) \mid x, z] = 1 - \Phi(\sqrt{\alpha}(\delta x + (1 - \delta)z - \theta^*(z))),$$

which is monotonic in x . It follows that the agent attacks if and only if $x \leq x^*(z)$, where $x^*(z)$ solves the indifference condition

$$b \Pr[\theta \leq \theta^*(z) \mid x^*(z), z] = c.$$

Substituting the expression for the posterior and the definition of δ and α , we obtain:

$$\Phi\left(\sqrt{\alpha_\epsilon + \alpha_\zeta} \left(\frac{\alpha_\epsilon}{\alpha_\epsilon + \alpha_\zeta} x^*(z) + \frac{\alpha_\zeta}{\alpha_\epsilon + \alpha_\zeta} z - \theta^*(z)\right)\right) = \frac{b - c}{b}. \quad (47)$$

Step 3. Combining (46) and (47), we conclude that $\theta^*(z)$ can be sustained in equilibrium if and only if it solves

$$G(\theta^*(z)) = g(z), \quad (48)$$

where

$$G(\theta) \equiv -\frac{\alpha_\zeta}{\sqrt{\alpha_\epsilon}} \theta + \Phi^{-1}(\theta) \quad \text{and} \quad g(z) \equiv \sqrt{1 + \frac{\alpha_\zeta}{\alpha_\epsilon}} \Phi^{-1}\left(1 - \frac{c}{b}\right) - \frac{\alpha_\zeta}{\sqrt{\alpha_\epsilon}} z.$$

With $\theta^*(z)$ given by (48), $x^*(z)$ is then given by (46).

We are now in a position to establish existence and determinacy of the equilibrium by considering the properties of the function G .

Note that $G(\theta)$ is continuous in θ , with $G(\underline{\theta}) = -\infty$ and $G(\bar{\theta}) = \infty$, which implies that there necessarily exists a solution and any solution satisfies $\theta^*(z) \in (\underline{\theta}, \bar{\theta})$. This establishes existence; we now turn to uniqueness.

Next, note that

$$\frac{\partial G(\theta)}{\partial \theta} = \frac{1}{\phi(\Phi^{-1}(\theta))} - \frac{\alpha_\zeta}{\sqrt{\alpha_\epsilon}}.$$

Since $\max_{w \in \mathbb{R}} \phi(w) = 1/\sqrt{2\pi}$, the following properties are true. If $\alpha_\zeta/\sqrt{\alpha_\epsilon} \leq \sqrt{2\pi}$, we have that G is strictly increasing in θ , which implies a unique solution to (48) for *all* values of z . If, instead, $\alpha_\zeta/\sqrt{\alpha_\epsilon} > \sqrt{2\pi}$, then G is non-monotonic in θ and there is an interval (\underline{z}, \bar{z}) such that (46) admits multiple solutions $\theta^*(z)$ whenever $z \in (\underline{z}, \bar{z})$ and a unique solution otherwise.

We conclude that the monotone equilibrium is unique if and only if $\alpha_\zeta/\sqrt{\alpha_\epsilon} \leq \sqrt{2\pi}$, or equivalently $\sigma_\epsilon \leq \sqrt{2\pi}\sigma_\zeta^2$.

Step 4. To complete the proof, we need to establish that, when the monotone equilibrium is unique, there is no other (non-monotone) equilibrium. This follows from a similar argument as in the proof of Proposition 6 in the main text.

To see this, let $h(x, z)$ denote the threshold that an agent finds it optimal to follow when all other agents follow a threshold x and the the public signal is z ; this is the same as the function h used in the proof of Proposition 6, modulo the presence of the public signal. Next, note that there is a one-to-one mapping between the thresholds θ^* that solve the equation $G(\theta) = g(z)$ and the fixed points of h : the monotone equilibria are defined by the fixed points of h . Finally, fix an arbitrary z and re-consider the type of sequences we constructed in the proof of Proposition 6. One can always show that the sequence “from below” converges to the lowest fixed point of h , whereas the sequence “from above” converges to the higher fixed point of h . When h admits a unique fixed point, corresponding to a unique monotone equilibrium, then the same kind of argument as that in the proof of Proposition 6 implies that the unique monotone equilibrium is also the unique rationalizable outcome. And when h admits multiple fixed points, the same argument implies that the set of rationalizable outcomes is contained within the two extreme monotone equilibria.

Proof of Propositions 8 and 9. See Morris and Shin (2003).

Proof of Proposition 11. Similar to the proof of Proposition 5.

Proof of Proposition 12. See the argument in the main text and Theorem 1 in Frankel and Pauzner (2000).

Proof of Proposition 13. Part (i) follows from Theorem 3 in Frankel and Pauzner (2000) along with the following fact: that the limit considered in Frankel-Pauzner and the limit considered in Morris-Shin coincide with the risk-dominance criterion of Harsanyi and Selten (1988).

Part (ii) follows from the fact that, in the limit as $\lambda \rightarrow 0$, the drift in K_t explodes to plus infinity whenever (θ_t, K_t) is on the left of κ^* and to minus infinity on the right of it.

Proof of Proposition 10. See Angeletos and Werning (2006).

Proof of Proposition 14. First, consider the case of perfect information, which means that the agent knows $(\theta_i, \bar{\theta})$ and, in equilibrium, also knows K . His best response is pinned down by the following first-order condition:

$$U_k(k_i, K, \theta_i) = 0.$$

Using the fact that U is quadratic, we have:

$$U_k(0, 0, 0) + U_{kk}k_i + U_{kK}K + U_{k\theta}\theta_i = 0. \tag{49}$$

Aggregating gives

$$U_k(0, 0, 0) + (U_{kk} + U_{kK})K + U_{k\theta}\bar{\theta} = 0. \tag{50}$$

which proves that the equilibrium value of K is a linear function of $\bar{\theta}$:

$$K = -\frac{U_k(0,0,0)}{U_{kk} + U_{kK}} - \frac{U_{k\theta}}{U_{kk} + U_{kK}}\bar{\theta}.$$

Substituting this back into condition (49) gives the equilibrium value of k_i as a linear function of θ_i and $\bar{\theta}$:

$$k_i = \kappa(\theta_i, \bar{\theta}) \equiv \kappa_0 + \kappa_1\theta_i + \kappa_2\bar{\theta},$$

and by implication $K = \kappa(\bar{\theta}, \bar{\theta})$, where

$$\kappa_0 \equiv -\frac{U_k(0,0,0)}{U_{kk} + U_{kK}}, \quad \kappa_1 \equiv -\frac{U_{k\theta}}{U_{kk}}, \quad \kappa_2 \equiv \frac{U_{k\theta}U_{kK}}{U_{kk}(U_{kk} + U_{kK})}. \quad (51)$$

Next, consider the case that information is complete but not necessarily perfect. Relative to the previous case, θ_i and $\bar{\theta}$ are not necessarily known any more, yet K remains known in equilibrium. The preceding proof therefore continues to work if we simply replace θ_i with $\mathbb{E}_i[\theta_i]$ and we accordingly replace $\bar{\theta}$ with the cross-sectional average of these first-order forecasts. That is, we can now express the equilibrium action as

$$k_i = \kappa(\mathbb{E}_i[\theta_i], \vartheta^1),$$

where $\vartheta^1 \equiv \int \mathbb{E}_j[\theta_j]dj$. Because complete information means that every agent knows $\mathbf{\Omega}$, and because $\mathbf{\Omega}$ is a sufficient statistic for $\bar{\theta}$ with respect to $(\omega_i, \mathbf{\Omega})$, the following is true:

$$\mathbb{E}_i[\bar{\theta}] = \mathbb{E}[\bar{\theta}|\mathbf{\Omega}] = \mathbb{E}[\theta_j|\mathbf{\Omega}] = \mathbb{E}[\mathbb{E}_j[\theta_j|\omega_j, \mathbf{\Omega}]|\mathbf{\Omega}] = \mathbb{E}[\mathbb{E}_j[\theta_j|\omega_j]|\mathbf{\Omega}] = \int \mathbb{E}_j[\theta_j]dj = \vartheta^1 \quad \forall i. \quad (52)$$

It follows that we can express the equilibrium action as:

$$k_i = \kappa(\mathbb{E}_i[\theta_i], \mathbb{E}_i[\bar{\theta}]) = \mathbb{E}_i[\kappa(\theta_i, \bar{\theta})],$$

which proves the proposition.

Proof of Proposition 15. Let us start with a preliminary observation. From the definition of $\kappa(\cdot)$, the following is trivially true for all θ_i and $\bar{\theta}$:

$$U_k(0,0,0) + U_{kk}\kappa(\theta_i, \bar{\theta}) + U_{kK}\kappa(\bar{\theta}, \bar{\theta}) + U_{k\theta}\theta_i = 0.$$

Taking the expectation of both sides, we have the following is also true, no matter the information of the agent:

$$U_k(0,0,0) + U_{kk}\mathbb{E}_i[\kappa(\theta_i, \bar{\theta})] + U_{kK}\mathbb{E}_i[\kappa(\bar{\theta}, \bar{\theta})] + U_{k\theta}\mathbb{E}_i[\theta_i] = 0. \quad (53)$$

Turning now the agent's optimal choice of k_i under incomplete information, note that this is characterized by the following first-order condition:

$$\mathbb{E}_i[U_k(k_i, K, \theta_i)] = 0.$$

Using the fact that U is quadratic, the above equation can be rewritten as follows:

$$U_k(0,0,0) + U_{kk}k_i + U_{kK}\mathbb{E}_i[K] + U_{k\theta}\mathbb{E}_i[\theta_i] = 0 \quad (54)$$

Subtracting equation (53) from (54), we get:

$$U_{kk}\mathbb{E}_i [k_i - \kappa(\theta_i, \bar{\theta})] + U_{kK}\mathbb{E}_i [K - \kappa(\bar{\theta}, \bar{\theta})] = 0.$$

We thus arrive at equation (15) in Proposition (15), namely,

$$k_i = \mathbb{E}_i [\kappa(\theta_i, \bar{\theta})] + \alpha \cdot \mathbb{E}_i [K - \kappa(\bar{\theta}, \bar{\theta})],$$

with

$$\alpha \equiv -\frac{U_{kK}}{U_{kk}} = \frac{\kappa_2}{\kappa_1 + \kappa_2}. \quad (55)$$

Proof of Proposition 16. Substituting equation (14) and (55) into equation (15), we can write an agent's optimal action k_i under incomplete information as

$$k_i = \kappa_0(1 - \alpha) + \kappa_1\theta_i + \alpha\mathbb{E}_i [K] + (\kappa_2 - \alpha(\kappa_1 + \kappa_2))\mathbb{E}_i [\bar{\theta}] = \kappa_0(1 - \alpha) + \kappa_1\theta_i + \alpha\mathbb{E}_i [K]. \quad (56)$$

Aggregating over i , we have:

$$K = \kappa_0(1 - \alpha) + \kappa_1\bar{\theta} + \alpha\bar{\mathbb{E}}[K].$$

Iterating, we have:

$$K = \kappa_0 + \kappa_1 \left\{ \sum_{h=0}^{\infty} \alpha^h \bar{\mathbb{E}}^h[\bar{\theta}] \right\}.$$

where $\bar{\mathbb{E}}^0[\bar{\theta}] = \bar{\theta}$. Substituting the above into (56) gives:

$$k_i = \kappa_0 + \kappa_1\theta_i + \alpha\mathbb{E}_i \left[\kappa_1 \left\{ \sum_{h=0}^{\infty} \alpha^h \bar{\mathbb{E}}^h[\bar{\theta}] \right\} \right]. \quad (57)$$

From the definition of α in condition (55), we have $\alpha\kappa_1 = (1 - \alpha)\kappa_2$. Substituting this into (57), we arrive at condition (16), namely:

$$k_i = \kappa_0 + \kappa_1\theta_i + \kappa_2\mathbb{E}_i \left\{ \sum_{h=0}^{\infty} (1 - \alpha)\alpha^h \bar{\mathbb{E}}^h[\bar{\theta}] \right\}.$$

Proof of Proposition 17. The posterior for θ conditional only on the public signal z is given by

$$\theta|z \sim N\left(\mu_{\theta|z}, \sigma_{\theta|z}^2\right), \quad (58)$$

where $\mu_{\theta|z} \equiv \eta_z z \equiv \frac{\sigma_{\zeta}^{-2}}{\sigma_{\theta}^{-2} + \sigma_{\zeta}^{-2}} z$ and $\sigma_{\theta|z}^{-2} \equiv \sigma_{\theta}^{-2} + \sigma_{\zeta}^{-2}$. Then we can derive the posterior for θ given both agent i 's private signal x_i and public signal z as follows:

$$\theta|x_i, z \sim N\left(\lambda_x x_i + \lambda_{\mu} \mu_{\theta|z}, \sigma_{\theta|x,z}^2\right), \quad (59)$$

where $\lambda_x \equiv \frac{(\sigma_u^2 + \sigma_{\epsilon}^2)^{-1}}{\sigma_{\theta|z}^{-2} + (\sigma_u^2 + \sigma_{\epsilon}^2)^{-1}} = \frac{(\sigma_u^2 + \sigma_{\epsilon}^2)^{-1}}{\sigma_{\theta|x,z}^{-2}} > 0$, $\lambda_{\mu} \equiv \frac{\sigma_{\theta|z}^{-2}}{\sigma_{\theta|z}^{-2} + (\sigma_u^2 + \sigma_{\epsilon}^2)^{-1}} = \frac{\sigma_{\theta}^{-2} + \sigma_{\zeta}^{-2}}{\sigma_{\theta|x,z}^{-2}} > 0$ and $\sigma_{\theta|x,z}^{-2} \equiv \left(\sigma_{\theta|z}^2\right)^{-1} + (\sigma_u^2 + \sigma_{\epsilon}^2)^{-1}$.

We can also derive the posterior for the correlated error in private signals u given both agent i 's private

signal x_i and public signal z , distributing as:

$$u|x_i, z \sim N\left(\Lambda(x_i - \mu_{\theta|z}), \sigma_{u|x,z}^2\right), \quad (60)$$

where $\Lambda \equiv \frac{(\sigma_{\theta|z}^2 + \sigma_\epsilon^2)^{-1}}{\sigma_u^{-2} + (\sigma_{\theta|z}^2 + \sigma_\epsilon^2)^{-1}} = \frac{((\sigma_\theta^{-2} + \sigma_\zeta^{-2})^{-1} + \sigma_\epsilon^2)^{-1}}{\sigma_u^{-2} + ((\sigma_\theta^{-2} + \sigma_\zeta^{-2})^{-1} + \sigma_\epsilon^2)^{-1}} = \frac{(\sigma_\theta^{-2} + \sigma_\zeta^{-2})(\sigma_u^2 + \sigma_\epsilon^2)^{-1}}{\sigma_u^{-2}((\sigma_u^2 + \sigma_\epsilon^2)^{-1} + \sigma_\theta^{-2} + \sigma_\zeta^{-2})} > 0$ and $\sigma_{u|x,z}^{-2} \equiv \sigma_u^{-2} + (\sigma_{\theta|z}^2 + \sigma_\epsilon^2)^{-1}$.

We now state two useful properties about λ_x, λ_μ , and Λ , which will be used later:

$$\lambda_x + \Lambda = \frac{(\sigma_u^2 + \sigma_\epsilon^2)^{-1}}{(\sigma_u^2 + \sigma_\epsilon^2)^{-1} + \sigma_\theta^{-2} + \sigma_\zeta^{-2}} \left(\frac{\sigma_u^{-2} + \sigma_\theta^{-2} + \sigma_\zeta^{-2}}{\sigma_u^{-2}} \right) < 1, \quad (61)$$

$$\lambda_\mu - \Lambda = \frac{\sigma_\theta^{-2} + \sigma_\zeta^{-2}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2} + (\sigma_u^2 + \sigma_\epsilon^2)^{-1}} \left(1 - \frac{(\sigma_u^2 + \sigma_\epsilon^2)^{-1}}{\sigma_u^{-2}} \right) > 0.. \quad (62)$$

To express in closed form the best response of the agent given her information, we calculate each term in equation (16). From equation (59), we know

$$\mathbb{E}_i[\theta] = \lambda_x x_i + \lambda_\mu \mu_{\theta|z}.$$

Aggregating over i gives

$$\bar{\mathbb{E}}[\theta] = \lambda_x (\theta + u) + \lambda_\mu \mu_{\theta|z}.$$

From equation (59) and equation (60), we have:

$$\mathbb{E}_i[\bar{\mathbb{E}}[\theta]] = \lambda_x (\lambda_x x_i + \lambda_\mu \mu_{\theta|z} + \Lambda(x_i - \mu_{\theta|z})) + \lambda_\mu \mu_{\theta|z}.$$

Aggregating over i , we have:

$$\bar{\mathbb{E}}[\bar{\mathbb{E}}[\theta]] = \lambda_x (\lambda_x + \Lambda) (\theta + u) + (\lambda_x (\lambda_\mu - \Lambda) + \lambda_\mu) \mu_{\theta|z}.$$

Similarly, for any $h \geq 0$, we have:

$$\begin{aligned} \mathbb{E}_i[\bar{\mathbb{E}}^h[\theta]] &= \lambda_x (\lambda_x + \Lambda)^h x_i + \left(\lambda_x (\lambda_\mu - \Lambda) \left(1 + (\lambda_x + \Lambda) + \dots + (\lambda_x + \Lambda)^{h-1} \right) + \lambda_\mu \right) \mu_{\theta|z} \quad (63) \\ &= \lambda_x (\lambda_x + \Lambda)^h x_i + \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^h}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \mu_{\theta|z}, \end{aligned}$$

and

$$\bar{\mathbb{E}}^{h+1}[\theta] = \bar{\mathbb{E}}[\bar{\mathbb{E}}^h[\theta]] = \lambda (\mathbb{E}_i[\theta] + \mathbb{E}_i[u]) = \lambda_x (\lambda_x + \Lambda)^h (\theta + u) + \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^h}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \mu_{\theta|z}. \quad (64)$$

As a result, for any $h \geq 1$, we have:

$$\begin{aligned}\omega_u^h &= \lambda_x (\lambda_x + \Lambda)^{h-1}, \\ \omega_\zeta^h &= \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^{h-1}}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \eta_z, \\ \omega_\theta^h &= \lambda_x (\lambda_x + \Lambda)^{h-1} + \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^{h-1}}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \eta_z = \omega_u^h + \omega_\zeta^h.\end{aligned}$$

It is then straightforward to check that

$$0 < \omega_u^h < 1, \quad \omega_\zeta^h > 0, \quad \text{and} \quad \omega_\theta^h = \omega_u^h + \omega_\zeta^h > 0.$$

Now we use the fact that $\lambda_x + \lambda_\mu = 1$ to prove a useful property about the coefficients on u and ζ :

$$\omega_u^h + \frac{\omega_\zeta^h}{\eta_z} = 1 \tag{65}$$

This comes from

$$\begin{aligned}\omega_u^h + \frac{\omega_\zeta^h}{\eta_z} &= \lambda_x (\lambda_x + \Lambda)^{h-1} + \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^{h-1}}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \\ &= \frac{\lambda_x \left((\lambda_x + \Lambda)^{h-1} - (\lambda_x + \Lambda)^h + (1 - (\lambda_x + \Lambda)) \left(1 - (\lambda_x + \Lambda)^{h-1} \right) \right)}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \\ &= \frac{\lambda_x (1 - (\lambda_x + \Lambda))}{1 - (\lambda_x + \Lambda)} + \lambda_\mu = 1.\end{aligned}$$

Together with the fact that $0 < \omega_u^h < 1$ and $0 < \eta_z < 1$, we have:

$$\omega_\theta^h = \eta_z \omega_u^h + \omega_\zeta^h + (1 - \eta_z) \omega_u^h = \eta_z + (1 - \eta_z) \omega_u^h < 1. \tag{66}$$

To prove

$$\omega_\theta^h > \omega_\theta^{h+1},$$

note that for any $h \geq 1$, we have:

$$\omega_u^h = \lambda_x (\lambda_x + \Lambda)^{h-1} > \lambda_x (\lambda_x + \Lambda)^h = \omega_u^{h+1}.$$

Together with equation (66) and the fact that $0 < \eta_z < 1$, we then have that $\omega_\theta^h > \omega_\theta^{h+1}$, $\forall h \geq 1$.

To prove

$$0 < \frac{\text{Var}(\bar{\mathbb{E}}^h[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^h[\theta])} < \frac{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta])} < 1,$$

note that

$$\frac{\text{Var}(\bar{\mathbb{E}}^h[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^h[\theta])} = \frac{(\omega_\zeta^h)^2 \sigma_\zeta^2 + (\omega_u^h)^2 \sigma_u^2}{(\omega_\zeta^h)^2 \sigma_\zeta^2 + (\omega_u^h)^2 \sigma_u^2 + (\omega_\theta^h)^2 \sigma_\theta^2},$$

we obviously have

$$0 < \frac{\text{Var}(\bar{\mathbb{E}}^h[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^h[\theta])} < 1.$$

We thus only need to prove

$$\frac{(\omega_\zeta^h)^2 \sigma_\zeta^2 + (\omega_u^h)^2 \sigma_u^2}{(\omega_\theta^h)^2 \sigma_\theta^2} < \frac{(\omega_\zeta^{h+1})^2 \sigma_\zeta^2 + (\omega_u^{h+1})^2 \sigma_u^2}{(\omega_\theta^{h+1})^2 \sigma_\theta^2} \quad \forall h \geq 1.$$

Define $t(\omega) = \frac{\omega}{1-\omega} > 0$, $\forall \omega \in (0, 1)$, and use equation (65) and (66), we have:

$$\frac{(\omega_\zeta^h)^2 \sigma_\zeta^2 + (\omega_u^h)^2 \sigma_u^2}{(\omega_\theta^h)^2 \sigma_\theta^2} = \frac{t(\omega_u^h)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\omega_u^h) + \eta_z)^2 \sigma_\theta^2}.$$

Next, note that $1 > \omega_u^h > \omega_u^{h+1} > 0$ and $t(x)$ increases with $x \in (0, 1)$, we have:

$$t(\omega_u^{h+1}) < t(\omega_u^h) \leq t(\omega_u^1) = \frac{\lambda_x}{1-\lambda_x} = \frac{(\sigma_u^2 + \sigma_\epsilon^2)^{-1}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} < \frac{\sigma_u^{-2}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}}.$$

As a result, for $t \in (0, t(\omega_u^1)]$, we have:

$$t\sigma_u^2 - \eta_z\sigma_\zeta^2 = t\sigma_u^2 - \frac{1}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} < 0; \quad (67)$$

$$\frac{\partial \left[\frac{t^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t + \eta_z)^2 \sigma_\theta^2} \right]}{\partial t} = \left(\frac{2t\sigma_u^2}{(t + \eta_z)^2 \sigma_\theta^2} - \frac{2(t^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2)}{(t + \eta_z)^3 \sigma_\theta^2} \right) = 2\eta_z \left(\frac{t\sigma_u^2 - \eta_z\sigma_\zeta^2}{(t + \eta_z)^3 \sigma_\theta^2} \right) < 0. \quad (68)$$

Together with the fact that $t(\omega_u^{h+1}) < t(\omega_u^h)$, we have

$$\frac{t(\omega_u^h)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\omega_u^h) + \eta_z)^2 \sigma_\theta^2} < \frac{t(\omega_u^{h+1})^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\omega_u^{h+1}) + \eta_z)^2 \sigma_\theta^2}, \quad \forall h \geq 1,$$

and therefore

$$\frac{\text{Var}(\bar{\mathbb{E}}^h[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^h[\theta])} < \frac{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta]|\theta)}{\text{Var}(\bar{\mathbb{E}}^{h+1}[\theta])}, \quad \forall h \geq 1.$$

Proof of Proposition 18. From equation (17), we know that the individual's optimal action is given by:

$$k_i = \mathbb{E}_i \left\{ \sum_{h=0}^{\infty} (1-\alpha)\alpha^h \bar{\mathbb{E}}^h[\bar{\theta}] \right\}.$$

Substituting $\mathbb{E}_i[\bar{\mathbb{E}}^h[\theta]]$ from equation (63), we obtain

$$\begin{aligned}
k_i &= (1 - \alpha) \sum_{h=0}^{+\infty} \alpha^h \left(\lambda_x (\lambda_x + \Lambda)^h x_i + \left(\lambda_x (\lambda_\mu - \Lambda) \frac{1 - (\lambda_x + \Lambda)^h}{1 - (\lambda_x + \Lambda)} + \lambda_\mu \right) \mu_{\theta|z} \right) \\
&= \frac{\lambda_x (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)} x_i + (1 - \alpha) \left(\frac{\lambda_\mu}{1 - \alpha} + \frac{\lambda_x (\lambda_\mu - \Lambda)}{1 - (\lambda_x + \Lambda)} \left(\frac{1}{1 - \alpha} - \frac{1}{1 - \alpha (\lambda_x + \Lambda)} \right) \right) \mu_{\theta|z} \\
&= \frac{\lambda_x (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)} x_i + \left(\lambda_\mu + \frac{\alpha \lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \mu_{\theta|z}.
\end{aligned}$$

Aggregating over i , we get

$$K = \frac{\lambda_x (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)} (\theta + u) + \left(\lambda_\mu + \frac{\alpha \lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \mu_{\theta|z}.$$

As a result,

$$\phi_u = \frac{\lambda_x (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)}, \quad (69)$$

$$\phi_\zeta = \left(\lambda_\mu + \frac{\alpha \lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \eta_z, \quad (70)$$

$$\phi_\theta = \phi_u + \phi_\zeta. \quad (71)$$

Similarly to the proof of condition (65), using the fact that $\lambda_x + \lambda_\mu = 1$, we can prove:

$$\phi_u + \frac{\phi_\zeta}{\eta_z} = 1. \quad (72)$$

Together with the fact that $\lambda_x + \Lambda < 1, \lambda_\mu - \Lambda > 0$, and $0 < \eta_z < 1$, we have:

$$0 < \phi_u, \phi_\zeta < 1$$

$$0 < \phi_\theta = \eta_z + (1 - \eta_z) \phi_u < 1 \quad (73)$$

We can also express the agents' average expectation of K as follows:

$$\begin{aligned}
\bar{\mathbb{E}}[K] &= \frac{\lambda_x (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)} (\bar{\mathbb{E}}[\theta] + \bar{\mathbb{E}}[u]) + \left(\lambda_\mu + \frac{\alpha \lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \mu_{\theta|z} \\
&= \frac{\lambda_x (1 - \alpha) (\lambda_x + \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} (\theta + u) + \left(\lambda_\mu + \frac{\lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \mu_{\theta|z}.
\end{aligned}$$

As a result,

$$\psi_u = \frac{\lambda_x (1 - \alpha) (\lambda_x + \Lambda)}{1 - \alpha (\lambda_x + \Lambda)}, \quad (74)$$

$$\psi_\zeta = \left(\lambda_\mu + \frac{\lambda_x (\lambda_\mu - \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right) \eta_z, \quad (75)$$

$$\psi_\theta = \psi_u + \psi_\zeta. \quad (76)$$

Similarly to the proof of equation (65) and (72), using the fact that $\lambda_x + \lambda_\mu = 1$, we can prove the

following:

$$\psi_u + \frac{\psi_\zeta}{\eta_z} = 1. \quad (77)$$

Together with the fact that $\lambda_x + \Lambda < 1, \lambda_\mu - \Lambda > 0$, and $\eta_z > 0$, we have:

$$0 < \psi_u < \phi_u < 1,$$

$$0 < \phi_\zeta < \psi_\zeta < 1,$$

$$0 < \psi_\theta = \eta_z + (1 - \eta_z) \psi_u < \phi_\theta < 1.$$

To prove

$$0 < \frac{\text{Var}(K|\theta)}{\text{Var}(K)} < \frac{\text{Var}(\bar{\mathbb{E}}[K]|\theta)}{\text{Var}(\bar{\mathbb{E}}[K])} < 1,$$

we proceed in the same way as the proof of Proposition 17. First, note that

$$\frac{\text{Var}(K|\theta)}{\text{Var}(K)} = \frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2 + \phi_\theta^2 \sigma_\theta^2} \quad \text{and} \quad \frac{\text{Var}(\bar{\mathbb{E}}[K]|\theta)}{\text{Var}(\bar{\mathbb{E}}[K])} = \frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2 + \psi_\theta^2 \sigma_\theta^2}.$$

We thus have

$$0 < \frac{\text{Var}(K|\theta)}{\text{Var}(K)}, \frac{\text{Var}(\bar{\mathbb{E}}[K]|\theta)}{\text{Var}(\bar{\mathbb{E}}[K])} < 1.$$

We only need to prove

$$\frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_\theta^2 \sigma_\theta^2} < \frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_\theta^2 \sigma_\theta^2}.$$

Defining again $t(x) = \frac{x}{1-x} > 0, \forall x \in (0, 1)$, we have

$$\frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_\theta^2 \sigma_\theta^2} = \frac{t(\phi_u)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\phi_u) + \eta_z)^2 \sigma_\theta^2},$$

$$\frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_\theta^2 \sigma_\theta^2} = \frac{t(\psi_u)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\psi_u) + \eta_z)^2 \sigma_\theta^2}.$$

$$\frac{\sigma_u^{-2}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} > \frac{(\sigma_u^2 + \sigma_\zeta^2)^{-1}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} = t(\lambda_x) > t(\phi_u) > t(\psi_u) > 0,$$

where we use the fact that $\phi_u = \frac{\lambda_x(1-\alpha)}{1-\alpha(\lambda_x+\Lambda)} < \lambda_x$, which comes from $\lambda_x + \Lambda < 1$, to prove $t(\lambda_x) > t(\phi_u)$. Similar as the proof of Proposition 17, for $t \in (0, t(\lambda_x)]$, we have equations (67) and (68) hold. As a result,

$$\frac{t(\phi_u)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\phi_u) + \eta_z)^2 \sigma_\theta^2} < \frac{t(\psi_u)^2 \sigma_u^2 + \eta_z^2 \sigma_\zeta^2}{(t(\psi_u) + \eta_z)^2 \sigma_\theta^2},$$

and therefore

$$\frac{\text{Var}(K|\theta)}{\text{Var}(K)} < \frac{\text{Var}(\bar{\mathbb{E}}[K]|\theta)}{\text{Var}(\bar{\mathbb{E}}[K])}.$$

Proof of Proposition 19. Let $w \equiv \lambda_x u + (1 - \lambda_x) \eta_z \zeta$ be the the residual of projecting $\bar{\mathbb{E}}[\theta]$ on θ . Next, let $R \equiv -(1 - \lambda_x) \eta_z \sigma_\zeta^2 u + \lambda_x \sigma_u^2 \zeta$. Clearly, we have both $w \perp \theta$ and $R \perp \theta$. Furthermore, we have $R \perp w$.

To see this, note that $cov(R, w) = -(1 - \lambda_x) \lambda_x \eta_z \sigma_\zeta^2 \sigma_u^2 + (1 - \lambda_x) \lambda_x \eta_z \sigma_\zeta^2 \sigma_u^2 = 0$, and both R and w are distributed Normally. It follows that θ , w , and R are an orthogonal basis of the space spanned by random variable θ , ζ and u .

Using the above, we can express u and ζ as linear transformations of w and R . Replacing them in to $K = \phi_\theta \theta + \phi_u u + \phi_\zeta \zeta$, we can then express K as a linear combination of θ , w , and R :

$$K = \phi_\theta \theta + \frac{\phi_u \lambda_x \sigma_u^2 + (1 - \phi_u) (1 - \lambda_x) \eta_z^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} w + \frac{(\lambda_x - \phi_u) \eta_z}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} R. \quad (78)$$

Similarly, we can express $\bar{\mathbb{E}}[K] = \psi_\theta \theta + \psi_u u + \psi_\zeta \zeta$ as a linear combination of θ , w , and R :

$$\bar{\mathbb{E}}[K] = \phi_\theta \theta + \frac{\psi_u \lambda_x \sigma_u^2 + (1 - \psi_u) (1 - \lambda_x) \eta_z^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} w + \frac{(\lambda_x - \psi_u) \eta_z}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} R. \quad (79)$$

From the above, we then have

$$\begin{aligned} Var(K | \theta, \bar{\mathbb{E}}\theta) &= Var(K | \theta, w) = \left[\frac{(\lambda_x - \phi_u) \eta_z}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} \right]^2 Var(R) \\ &= \frac{(\lambda_x - \phi_u)^2 \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} \\ &= \left[\frac{\alpha (1 - (\lambda_x + \Lambda))}{1 - \alpha (\lambda_x + \Lambda)} \right]^2 \frac{\lambda_x^2 \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} > 0, \end{aligned} \quad (80)$$

$$\begin{aligned} Var(\bar{\mathbb{E}}K | \theta, \bar{\mathbb{E}}\theta) &= Var(\bar{\mathbb{E}}K | \theta, w) = \left[\frac{(\lambda_x - \psi_u) \eta_z}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} \right]^2 Var(R) \\ &= \frac{(\lambda_x - \psi_u)^2 \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} \\ &= \left[\frac{1 - (\lambda_x + \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right]^2 \frac{\lambda_x^2 \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} > 0, \end{aligned} \quad (81)$$

and finally

$$\begin{aligned} Cov(K, \bar{\mathbb{E}}K | \theta, \bar{\mathbb{E}}\theta) &= Cov(K, \bar{\mathbb{E}}K | \theta, w) = \frac{(\lambda_x - \phi_u) (\lambda_x - \psi_u) \eta_z^2}{\left[\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2 \right]^2} Var(R) \\ &= \frac{(\lambda_x - \phi_u) (\lambda_x - \psi_u) \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} \\ &= \alpha \left[\frac{1 - (\lambda_x + \Lambda)}{1 - \alpha (\lambda_x + \Lambda)} \right]^2 \frac{\lambda_x^2 \eta_z^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1 - \lambda_x)^2 \eta_z^2 \sigma_\zeta^2} > 0, \end{aligned} \quad (82)$$

where we have used the facts that

$$\lambda_x - \phi_u = \lambda_x \left[1 - \frac{1 - \alpha}{1 - \alpha (\lambda_x + \Lambda)} \right] = \lambda_x \left[\frac{1 - \alpha (\lambda_x + \Lambda) - (1 - \alpha)}{1 - \alpha (\lambda_x + \Lambda)} \right] = \lambda_x \left[\frac{\alpha (1 - (\lambda_x + \Lambda))}{1 - \alpha (\lambda_x + \Lambda)} \right]$$

and

$$\lambda_x - \psi_u = \lambda_x \left[1 - \frac{(1 - \alpha)(\lambda_x + \Lambda)}{1 - \alpha(\lambda_x + \Lambda)} \right] = \frac{\lambda_x(1 - (\lambda_x + \Lambda))}{1 - \alpha(\lambda_x + \Lambda)}.$$

Proof of Proposition 20. Using the the fact that $\lambda_x + \Lambda < 1$, shown in condition (61), we have:

$$\begin{aligned} \frac{\partial \phi_u}{\partial \alpha} &= \frac{-\lambda_x}{1 - \alpha(\lambda_x + \Lambda)} + \frac{\lambda_x(1 - \alpha)(\lambda_x + \Lambda)}{(1 - \alpha(\lambda_x + \Lambda))^2} \\ &= \frac{\lambda_x}{1 - \alpha(\lambda_x + \Lambda)} \left(\frac{\lambda_x + \Lambda - 1}{1 - \alpha(\lambda_x + \Lambda)} \right) < 0. \end{aligned}$$

From condition (73), we have

$$\frac{\partial \phi_\theta}{\partial \alpha} = (1 - \eta_z) \frac{\partial \phi_u}{\partial \alpha} < 0.$$

Similarly,

$$\begin{aligned} \frac{\partial \psi_u}{\partial \alpha} &= (\lambda_x + \Lambda) \left(\frac{-\lambda_x}{1 - \alpha(\lambda_x + \Lambda)} + \frac{\lambda_x(1 - \alpha)(\lambda_x + \Lambda)}{(1 - \alpha(\lambda_x + \Lambda))^2} \right) \\ &= \frac{\lambda_x(\lambda_x + \Lambda)}{1 - \alpha(\lambda_x + \Lambda)} \left(\frac{\lambda_x + \Lambda - 1}{1 - \alpha(\lambda_x + \Lambda)} \right) < 0 \end{aligned}$$

and therefore also

$$\frac{\partial \psi_\theta}{\partial \alpha} = (1 - \eta_z) \frac{\partial \psi_u}{\partial \alpha} < 0.$$

Note that

$$Cov(K, \theta) = \phi_\theta \sigma_\theta^2 \quad \text{and} \quad Cov(\bar{\mathbb{E}}K, \theta) = \psi_\theta \sigma_\theta^2.$$

It follows that

$$\frac{\partial Cov(K, \theta)}{\partial \alpha} < 0 \quad \text{and} \quad \frac{\partial Cov(\bar{\mathbb{E}}K, \theta)}{\partial \alpha} < 0.$$

To prove

$$\frac{\partial}{\partial \alpha} \left(\frac{Var(K|\theta)}{Var(K)} \right) > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} \left(\frac{Var(\bar{\mathbb{E}}K|\theta)}{Var(\bar{\mathbb{E}}K)} \right) > 0,$$

note that

$$0 < \frac{Var(K|\theta)}{Var(K)} = \frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2 + \phi_\theta^2 \sigma_\theta^2} < 1 \quad \text{and} \quad 0 < \frac{Var(\bar{\mathbb{E}}K|\theta)}{Var(\bar{\mathbb{E}}K)} = \frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2 + \psi_\theta^2 \sigma_\theta^2} < 1,$$

To get the desired result, we therefore only need to prove

$$\frac{\partial}{\partial \alpha} \left(\frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_\theta^2 \sigma_\theta^2} \right) > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} \left(\frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_\theta^2 \sigma_\theta^2} \right) > 0.$$

Defining again $t(x) = \frac{x}{1-x} > 0, \forall x \in (0, 1)$ we have

$$\frac{\phi_u^2 \sigma_u^2 + \phi_\zeta^2 \sigma_\zeta^2}{\phi_\theta^2 \sigma_\theta^2} = \frac{t(\phi_u)^2 \sigma_u^2 + \sigma_\zeta^2}{(t(\phi_u) + \eta_z)^2 \sigma_\theta^2},$$

$$\frac{\psi_u^2 \sigma_u^2 + \psi_\zeta^2 \sigma_\zeta^2}{\psi_\theta^2 \sigma_\theta^2} = \frac{t(\psi_u)^2 \sigma_u^2 + \sigma_\zeta^2}{(t(\psi_u) + \eta_z)^2 \sigma_\theta^2},$$

$$\frac{\sigma_u^{-2}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} > \frac{(\sigma_u^2 + \sigma_\zeta^2)^{-1}}{\sigma_\theta^{-2} + \sigma_\zeta^{-2}} = t(\lambda_x) > t(\phi_u) > t(\psi_u) > 0,$$

where we used the fact that $\phi_u = \frac{\lambda_x(1-\alpha)}{1-\alpha(\lambda_x+\Lambda)} < \lambda_x$, which comes from $\lambda_x + \Lambda < 1$, to prove $t(\lambda_x) > t(\phi_u)$. Because $\frac{\partial \phi_u}{\partial \alpha} < 0$ and $\frac{\partial \psi_u}{\partial \alpha} < 0$, we also have

$$\frac{\partial}{\partial \alpha} (t(\phi_u)) < 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} (t(\psi_u)) < 0.$$

Similar as the proof of Proposition 17 and 18, for $t \in (0, t(\lambda_x)]$, we have equations (67) and (68) hold. As a result,

$$\frac{\partial}{\partial \alpha} \left(\frac{\text{Var}(K|\theta)}{\text{Var}(K)} \right) > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} \left(\frac{\text{Var}(\bar{\mathbb{E}}K|\theta)}{\text{Var}(\bar{\mathbb{E}}K)} \right) > 0,$$

We finally prove

$$\frac{\partial \text{Var}(K|\theta, \bar{\mathbb{E}}\theta)}{\partial \alpha} > 0, \quad \frac{\partial \text{Var}(K|\theta, \bar{\mathbb{E}}\theta)}{\partial \alpha} > 0 \quad \text{and} \quad \frac{\partial (\text{Cov}(K, \bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta))}{\partial \alpha} > 0.$$

From conditions (80), (81), and (82), we have

$$\begin{aligned} \text{Var}(K|\theta, \bar{\mathbb{E}}\theta) &= \left[\frac{\alpha(1-(\lambda_x+\Lambda))}{1-\alpha(\lambda_x+\Lambda)} \right]^2 \frac{\lambda_x^2 \eta_y^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1-\lambda_x)^2 \eta_y^2 \sigma_\zeta^2}, \\ \text{Var}(\bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta) &= \left[\frac{1-(\lambda_x+\Lambda)}{1-\alpha(\lambda_x+\Lambda)} \right]^2 \frac{\lambda_x^2 \eta_y^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1-\lambda_x)^2 \eta_y^2 \sigma_\zeta^2}, \\ \text{Cov}(K, \bar{\mathbb{E}}K|\theta, \bar{\mathbb{E}}\theta) &= \alpha \left[\frac{1-(\lambda_x+\Lambda)}{1-\alpha(\lambda_x+\Lambda)} \right]^2 \frac{\lambda_x^2 \eta_y^2 \sigma_u^2 \sigma_\zeta^2}{\lambda_x^2 \sigma_u^2 + (1-\lambda_x)^2 \eta_y^2 \sigma_\zeta^2}. \end{aligned}$$

Note that

$$\frac{\partial}{\partial \alpha} \left\{ \frac{1-(\lambda_x+\Lambda)}{1-\alpha(\lambda_x+\Lambda)} \right\} > 0 \quad \text{and} \quad \frac{\partial}{\partial \alpha} \left\{ \frac{\alpha[1-(\lambda_x+\Lambda)]}{1-\alpha(\lambda_x+\Lambda)} \right\} > 0.$$

The result then follows.

Proof of Proposition 21. Note that for any random variable X , and any information set I , according to law of total variance, we have:

$$\text{Var}(\mathbb{E}[X|I]) \leq \text{Var}(X).$$

As a result, we have:

$$\text{Var}(\bar{\mathbb{E}}[\theta]) = \text{Var}(\mathbb{E}[\mathbb{E}[\theta|\omega]|\mathbf{\Omega}]) \leq \text{Var}(\mathbb{E}[\theta|\omega]) \leq \text{Var}(\theta).$$

Similarly, for any $h \geq 2$, we have:

$$\text{Var}(\bar{\mathbb{E}}^h[\theta]) \leq \text{Var}(\mathbb{E}[\mathbb{E}[\bar{\mathbb{E}}^{h-1}[\theta]|\omega]|\mathbf{\Omega}]) \leq \text{Var}(\mathbb{E}[\bar{\mathbb{E}}^{h-1}[\theta]|\omega]) \leq \text{Var}(\bar{\mathbb{E}}^{h-1}[\theta]).$$

It follows that $Var(\bar{\mathbb{E}}^h[\theta]) \leq Var(\theta) \forall h \geq 1$, and therefore

$$Var(K) = Var\left(\sum_{h=1}^{\infty} (1-\alpha)\alpha^{h-1}\bar{\mathbb{E}}^h[\theta]\right) \leq \left(\sum_{h=1}^{\infty} (1-\alpha)\alpha^{h-1}\sqrt{Var(\theta)}\right)^2 = Var(\theta).$$

For the last result, we used the fact that, for any random variable X, Y and scalars $a, b \geq 0$, the following inequality is true:

$$\begin{aligned} Var(aX + bY) &= a^2Var(X) + 2abCov(X, Y) + b^2Var(Y) \\ &\leq a^2Var(X) + 2ab\sqrt{Var(X)Var(Y)} + b^2Var(Y) \\ &= \left(a\sqrt{Var(X)} + b\sqrt{Var(Y)}\right)^2. \end{aligned}$$

Proof of Proposition 22. See the main text.

Proof of Proposition 23. The result follows essentially from Woodford (2003); see Angeletos and La'O (2010) for the particular version stated in Proposition 23. The exact characterization of the scalars $(\gamma_K, \gamma_\theta, \gamma_v)$ cannot be found in these earlier works, but can be obtained with the method developed in Huo and Takayama (2015a,b).

Proof of Proposition 24. See Proposition 2 in Myatt and Wallace (2012), and Proposition 1 and Corollary 1 in Pavan (2015).

Proof of Proposition 25. See Proposition 4 in Myatt and Wallace (2012).

Proof of Proposition 26. Under complete information, $\mathbb{E}_{it}Y_t = Y_t$ for all i . As a result, optimal output level in island i in equation (29) can be written as:

$$y_{it} = (1-\alpha)\chi a_{i,t} + \alpha Y_t.$$

Aggregating the above equation, we have $Y_t = \chi A_t$, and thus $N_t = (\chi - 1)A_t$, which completes the proof.

Proof of Proposition 27. Aggregating equation (29), we have:

$$Y_t = (1-\alpha)\chi A_t + \alpha\bar{\mathbb{E}}_t Y_t.$$

Iterating, we have:

$$Y_t = \chi \sum_{h=0}^{\infty} (1-\alpha)\alpha^h \bar{\mathbb{E}}_t^h[A_t].$$

Proof of Proposition 28. Given the information structure introduced in Assumption 9, we have:

$$\bar{\mathbb{E}}_t^h[A_t] = \left(\frac{\sigma_\epsilon^{-2}}{\sigma_u^{-2} + \sigma_\epsilon^{-2}}\right)^h u_t + A_{t-1} \quad \forall h \in \mathbb{N}.$$

Substituting it into equation (30) gives

$$Y_t = \chi A_{t-1} + \chi \frac{(1-\alpha)(\sigma_u^{-2} + \sigma_\epsilon^{-2})}{\sigma_u^{-2} + (1-\alpha)\sigma_\epsilon^{-2}} u_t \equiv \chi A_{t-1} + \phi u_t,$$

where

$$\phi \equiv \chi \frac{(1-\alpha)(\sigma_u^{-2} + \sigma_\epsilon^{-2})}{\sigma_u^{-2} + (1-\alpha)\sigma_\epsilon^{-2}}$$

measures the response of aggregate output to a innovation in the current fundamental. Note that $\lim_{\alpha \rightarrow 1^-} \phi = 0$. This means that, no matter how precise the information is, we can always make the aforementioned response to be arbitrarily small by boosting enough the degree of strategic complementarity.

Proof of Proposition 29. See Section 4 in Woodford (2003).

Proof of Proposition 30. For the first property, see argument in subsection 8.3. For the second property, see Figure 7.

Proof of Proposition 31. Because m_{t-1} is common knowledge in the beginning of period t , it has no real effect in that period, which means that p_{it} must move one-to-one with m_{t-1} . Because local demand varies only with the sum $v_t + \xi_{it}$ and young agents have no information about the individual components of this sum, the variation in p_{it} conditional on m_{t-1} must be spanned by the variation in that sum. Along with the assumption that p_{it} is Normal, this means that $p_{it} - m_{t-1}$ is linear in the sum $v_t + \xi_{it}$. In what follows, we thus guess and verify the existence of an equilibrium in which

$$p_{it} = m_{t-1} + \phi(v_t + \xi_{it}), \tag{83}$$

for some scalar ϕ (which remains to be determined).

Taking the first-order condition of the problem of a young agent, we obtain the optimal labor supply, up to a constant, as follows:

$$n_{it} = \frac{1}{\kappa} (\mathbb{E}_{it} [p_{it} - m_t] - \mathbb{E}_{it} [p_{j,t+1} - m_{t+1}]). \tag{84}$$

Because $\mathbb{E}_{it} [m_{t+1}] = \mathbb{E}_{it} [m_t]$, we can rewrite the above as

$$n_{it} = \frac{1}{\kappa} (p_{it} - \mathbb{E}_{it} [p_{j,t+1}]).$$

Intuitively, labor supply depends on the perceived *relative* price. Using (83), we get that $\mathbb{E}_{it} [p_{j,t+1}] = \mathbb{E}_{it} [m_t + \phi(v_{t+1} + \xi_{it+1})] = \mathbb{E}_{it} [m_t]$, simply because no agent in period t has any information about the next-period shocks v_{t+1} and $\xi_{i,t+1}$. We conclude that the optimal labor supply can be express as follows:

$$n_{it} = \frac{1}{\kappa} (p_{it} - \mathbb{E}_{it} [m_t]).$$

Note next that the local demand is given by

$$d_{it} = m_t + \xi_{it} = m_{t-1} + v_t + \xi_{it},$$

and that market clearing imposes

$$p_{it} + n_{it} = d_{it}.$$

Combining the above we get

$$y_{it} = n_{it} = \frac{1}{\kappa + 1} (\xi_{i,t} + v_t - \mathbb{E}_{it}[v_t]),$$

and

$$p_{it} = v_t + \xi_{it} - \frac{1}{\kappa + 1} (\xi_{i,t} + v_t - \mathbb{E}_{it}[v_t]) = m_{t-1} + \frac{\kappa}{\kappa + 1} (\xi_{i,t} + v_t) + \frac{1}{\kappa + 1} \mathbb{E}_{it}[v_t].$$

Aggregating the last equation, and letting $\bar{\mathbb{E}}_t$ denotes the average expectation of young agents in the cross section of islands, we reach the following result for the aggregate price level:

$$p_t = m_{t-1} + \frac{\kappa}{\kappa + 1} v_t + \frac{1}{\kappa + 1} \bar{\mathbb{E}}[v_t] = \frac{\kappa}{\kappa + 1} m_t + \frac{1}{\kappa + 1} \bar{\mathbb{E}}[m_t].$$

Using the fact that $m_t = m_{t-1} + v_t$ and that m_{t-1} is commonly known, we can rewrite the above as

$$p_t = (1 - \beta)m_t + \beta \bar{\mathbb{E}}[m_t],$$

where $\beta \equiv \frac{1}{\kappa + 1}$. This gives part (i) in the proposition.

Real output can then be expressed as

$$y_t = m_t - p_t = \beta \{m_t - \bar{\mathbb{E}}[m_t]\},$$

which gives part (ii) of the proposition. The above condition simply means that real output is pinned down by the average forecast error of the underlying money supply (also known as the “unanticipated” component of money growth).

Part (iii) follows from solving the signal-extraction problem of the young agents. In particular, condition (83) implies that the observation of p_{it} contains the same information as the observation of the signal

$$x_{it} \equiv \frac{1}{\phi} (p_{it} - m_{t-1}) = \frac{1}{\phi} (v_t + \xi_{it}),$$

which in turn means that $\mathbb{E}_{it}[m_t] = m_{t-1} + \delta x_{it}$ and therefore $\bar{\mathbb{E}}_t[m_t] = m_{t-1} + \delta v_t$, where $\delta \equiv \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\xi^2} \in (0, 1)$. By implication, we can write the price level as $p_t = m_{t-1} + (1 - \beta + \beta\delta)v_t$, which in turn verifies our initial guess in (83) with $\phi = 1 - \beta + \beta\delta \in (0, 1)$. This completes the characterization of the equilibrium; it also proves condition (37) with $\lambda \equiv 1 - \delta = \frac{\sigma_\xi^2}{\sigma_v^2 + \sigma_\xi^2}$.

Proof of Proposition 32. The average first-order belief of θ_t is a simple weighted average of beliefs of the firms that updated their information at different lags:

$$\bar{\mathbb{E}}_t[\theta_t] = \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j \mathbb{E}_{t-j}[\theta_t],$$

where $\mathbb{E}_{t-j}[\theta_t]$ is the expectation conditional on all the information that is available at $t - j$. This proves that the following condition holds for $h = 1$:

$$\bar{\mathbb{E}}_t^h[\theta_t] = \sum_{j=0}^{+\infty} \left[\left(1 - (1 - \lambda)^{j+1}\right)^h - \left(1 - (1 - \lambda)^j\right)^h \right] \mathbb{E}_{t-j}[\theta_t]. \quad (85)$$

We then prove that the condition holds for all $h \geq 1$ by induction.

Thus suppose that (85) holds for some arbitrary h . It follows that

$$\mathbb{E}_{t-j} [\bar{\mathbb{E}}_t^h [\theta_t]] = \left(1 - (1 - \lambda)^{j+1}\right)^h \mathbb{E}_{t-j} [\theta_t] + \sum_{l=j+1}^{\infty} \left[\left(1 - (1 - \lambda)^{l+1}\right)^h - \left(1 - (1 - \lambda)^l\right)^h \right] \mathbb{E}_{t-l} [\theta_t].$$

Then it also follows that

$$\begin{aligned} \bar{\mathbb{E}}_t^{h+1} [\theta_t] &= \sum_{j=0}^{+\infty} \left[\lambda (1 - \lambda)^j \left(1 - (1 - \lambda)^{j+1}\right)^h + \sum_{l=0}^{j-1} \lambda (1 - \lambda)^l \left[\left(1 - (1 - \lambda)^{j+1}\right)^h - \left(1 - (1 - \lambda)^j\right)^h \right] \right] \mathbb{E}_{t-j} [\theta_t] \\ &= \sum_{j=0}^{+\infty} \left[\lambda (1 - \lambda)^j \left(1 - (1 - \lambda)^{j+1}\right)^h + \left[1 - (1 - \lambda)^j\right] \left[\left(1 - (1 - \lambda)^{j+1}\right)^h - \left(1 - (1 - \lambda)^j\right)^h \right] \right] \mathbb{E}_{t-j} [\theta_t] \\ &= \sum_{j=0}^{+\infty} \left[\left(1 - (1 - \lambda)^{j+1}\right)^{h+1} - \left(1 - (1 - \lambda)^j\right)^{h+1} \right] \mathbb{E}_{t-j} [\theta_t], \end{aligned}$$

which means that condition (85) holds for $h + 1$ and proves the claim.

Now let us focus in the case in which $\rho = 0$, that is, $\theta_t = \theta_{t-1} + v_t$. As a result,

$$\mathbb{E}_{t-j} [\theta_t] - \mathbb{E}_{t-j-1} [\theta_t] = v_{t-j}, \quad \forall j \geq 0.$$

Then condition (85) can be written as

$$\bar{\mathbb{E}}_t^h [\theta_t] = \sum_{j=0}^{+\infty} \left\{ \left(1 - (1 - \lambda)^{j+1}\right)^h \left(\mathbb{E}_{t-j} [\theta_t] - \mathbb{E}_{t-j-1} [\theta_t] \right) \right\} = \sum_{j=0}^{+\infty} \left\{ \left(1 - (1 - \lambda)^{j+1}\right)^h v_{t-j} \right\}.$$

Proof of Propositions 33 and 34. See, respectively, Propositions 2 and 1 in Allen, Morris and Shin (2006).

Proof of Proposition 35. This follows from the same argument as in Step 1 in the proof of Propositions 1 and 2 in Angeletos and Pavan (2009). The only non-essential differences are (i) that Angeletos and Pavan (2009) uses a different notation and (ii) that their framework allows for externalities to obtain, not only from the mean action, but also from the dispersion of actions.

Proof of Proposition 36. See Propositions 1 and 2 in Angeletos and Pavan (2009).

References

- ABEL, A. B., J. C. EBERLY, AND S. PANAGEAS (2007): “Optimal inattention to the stock market,” *The American Economic Review*, 97, 244–249.
- (2013): “Optimal inattention to the stock market with information costs and transactions costs,” *Econometrica*, 81, 1455–1481.
- ABREU, D. AND M. K. BRUNNERMEIER (2003): “Bubbles and crashes,” *Econometrica*, 71, 173–204.
- ACHARYA, S. (2013): “Dispersed beliefs and aggregate demand management,” *University of Maryland mimeo*.
- ADAM, K. (2007): “Optimal monetary policy with imperfect common knowledge,” *Journal of Monetary Economics*, 54, 267–301.
- AKERLOF, G. A. AND R. J. SHILLER (2010): *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*, Princeton University Press.
- ALBAGLI, E., C. HELLOWIG, AND A. TSYVINSKI (2014): “Risk-taking, rent-seeking, and investment when financial markets are noisy,” *Yale mimeo*.
- (2015): “A Theory of Asset Prices based on Heterogeneous Information,” *Yale mimeo*.
- ALLEN, F., S. MORRIS, AND A. POSTLEWAITE (1993): “Finite bubbles with short sale constraints and asymmetric information,” *Journal of Economic Theory*, 61, 206–229.
- ALLEN, F., S. MORRIS, AND H. S. SHIN (2006): “Beauty contests and iterated expectations in asset markets,” *Review of financial Studies*, 19, 719–752.
- ALVAREZ, F. AND F. LIPPI (2014): “Price setting with menu cost for multiproduct firms,” *Econometrica*, 82, 89–135.
- ALVAREZ, F., F. LIPPI, AND L. PACIELLO (2011): “Optimal Price Setting With Observation and Menu Costs,” *The Quarterly Journal of Economics*, 126, 1909–1960.
- (2015): “Phillips curves with observation and menu costs,” *Uchicago mimeo*.
- AMADOR, M. AND P.-O. WEILL (2010): “Learning from Prices: Public Communication and Welfare,” *Journal of Political Economy*, 118, 866–907.
- (2012): “Learning from private and public observations of others’ actions,” *Journal of Economic Theory*, 147, 910–940.
- AMATO, J. D. AND H. S. SHIN (2006): “Imperfect common knowledge and the information value of prices,” *Economic theory*, 27, 213–241.
- ANDRADE, P. AND H. LE BIHAN (2013): “Inattentive professional forecasters,” *Journal of Monetary Economics*, 60, 967–982.
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2015): “Quantifying Confidence,” *NBER Working Paper Series*.

- ANGELETOS, G.-M., C. HELLWIG, AND A. PAVAN (2006): “Signaling in a global game: Coordination and policy traps,” *Journal of Political Economy*, 114, 452–484.
- (2007): “Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks,” *Econometrica*, 75, 711–756.
- ANGELETOS, G.-M., L. IOVINO, AND J. LA’O (2016a): “Efficiency and policy with endogenous learning,” *MIT mimeo*.
- (2016b): “Real Rigidity, Nominal Rigidity, and the Social Value of Information,” *The American Economic Review*, 106, 200–227.
- ANGELETOS, G.-M. AND J. LA’O (2009): “Incomplete information, higher-order beliefs and price inertia,” *Journal of Monetary Economics*, 56, 19–37.
- (2010): “Noisy business cycles,” in *NBER Macroeconomics Annual 2009, Volume 24*, University of Chicago Press, 319–378.
- (2012): “Optimal Monetary Policy with Informational Frictions,” *MIT mimeo*.
- (2013): “Sentiments,” *Econometrica*, 81, 739–779.
- ANGELETOS, G.-M. AND C. LIAN (2016a): “Dampening General Equilibrium: From Micro Elasticities to Macro Effects,” *MIT mimeo*.
- (2016b): “A (Real) Theory of Aggregate Demand,” *MIT mimeo*.
- ANGELETOS, G.-M., G. LORENZONI, AND A. PAVAN (2010): “Beauty contests and irrational exuberance: a neoclassical approach,” *MIT mimeo*.
- ANGELETOS, G.-M. AND A. PAVAN (2004): “Transparency of Information and Coordination in Economies with Investment Complementarities,” *The American Economic Review*, 94, 91–98.
- (2007): “Efficient use of information and social value of information,” *Econometrica*, 75, 1103–1142.
- (2009): “Policy with dispersed information,” *Journal of the European Economic Association*, 7, 11–60.
- (2013): “Selection-free predictions in global games with endogenous information and multiple equilibria,” *Theoretical Economics*, 8, 883–938.
- ANGELETOS, G.-M. AND I. WERNING (2006): “Crises and Prices: Information Aggregation, Multiplicity, and Volatility,” *The American Economic Review*, 96, 1720–1736.
- ARAÚJO, L. AND B. GUIMARAES (2015): “Intertemporal coordination with delay options,” *Journal of Economic Theory*, 157, 793–810.
- ARELLANO, C., Y. BAI, AND P. J. KEHOE (2012): “Financial Frictions and Fluctuations in Volatility,” *University of Minnesota mimeo*.
- ATKESON, A. (2000): “Discussion of Morris and Shin’s ‘Rethinking multiple equilibria in macroeconomic modelling’,” *NBER Macroeconomics Annual*, 15, 162–171.

- AUMANN, R. J. (1974): "Subjectivity and correlation in randomized strategies," *Journal of mathematical Economics*, 1, 67–96.
- (1987): "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica*, 1–18.
- AUMANN, R. J. AND A. HEIFETZ (2002): "Incomplete information," *Handbook of game theory with economic applications*, 3, 1665–1686.
- AUMANN, R. J., J. PECK, AND K. SHELL (1988): "Asymmetric information and sunspot equilibria: a family of simple examples," *CAE mimeo*.
- AZARIADIS, C. (1981): "Self-fulfilling prophecies," *Journal of Economic Theory*, 25, 380–396.
- BACCHETTA, P. AND E. VAN WINCOOP (2006): "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *The American Economic Review*, 96, 552–576.
- BAERISWYL, R. AND C. CORNAND (2010a): "Optimal Monetary Policy in Response to Cost-Push Shocks: The Impact of Central Bank Communication," *International Journal of Central Banking*, 6, 31–52.
- (2010b): "The signaling role of policy actions," *Journal of Monetary Economics*, 57, 682–695.
- BANERJEE, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 797–817.
- BANNIER, C. E. (2006): "The role of information disclosure and uncertainty in the 1994/95 Mexican peso crisis: Empirical evidence," *Review of International Economics*, 14, 883–909.
- BARLEVY, G. AND P. VERONESI (2000): "Information acquisition in financial markets," *The Review of Economic Studies*, 67, 79–90.
- (2007): "Information acquisition in financial markets: a correction," *University of Chicago mimeo*.
- BARRO, R. J. (1976): "Rational expectations and the role of monetary policy," *Journal of Monetary Economics*, 2, 1–32.
- (1978): "Unanticipated money, output, and the price level in the United States," *The Journal of Political Economy*, 549–580.
- BARSKY, R. B. AND E. R. SIMS (2011): "News shocks and business cycles," *Journal of Monetary Economics*, 58, 273–289.
- (2012): "Information, animal spirits, and the meaning of innovations in consumer confidence," *The American Economic Review*, 1343–1377.
- BEAUDRY, P. AND F. PORTIER (2006): "Stock Prices, News, and Economic Fluctuations," *The American Economic Review*, 96, 1293–1307.
- BEBCHUK, L. A. AND I. GOLDSTEIN (2011): "Self-fulfilling credit market freezes," *Review of Financial Studies*, 24, 3519–3555.
- BENHABIB, J. AND R. E. FARMER (1994): "Indeterminacy and increasing returns," *Journal of Economic Theory*, 63, 19–41.
- (1999): "Indeterminacy and sunspots in macroeconomics," *Handbook of macroeconomics*, 1, 387–448.

- BENHABIB, J., X. LIU, AND P. WANG (2016): “Sentiments, financial markets, and macroeconomic fluctuations,” *Journal of Financial Economics*, forthcoming.
- BENHABIB, J. AND P. WANG (2015): “Endogenous information acquisition and countercyclical uncertainty,” *NYU mimeo*.
- BENHABIB, J., P. WANG, AND Y. WEN (2015): “Sentiments and aggregate demand fluctuations,” *Econometrica*, 83, 549–585.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015): “Information and volatility,” *Journal of Economic Theory*, 158, 427–465.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust predictions in games with incomplete information,” *Econometrica*, 81, 1251–1308.
- BIAIS, B. AND P. BOSSAERTS (1998): “Asset prices and trading volume in a beauty contest,” *The Review of Economic Studies*, 65, 307–340.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, 992–1026.
- BLANCHARD, O. J., J.-P. L’HUILIER, AND G. LORENZONI (2013): “News, Noise, and Fluctuations: An Empirical Exploration,” *American Economic Review*, 103, 3045–3070.
- BLOOM, N. (2009): “The impact of uncertainty shocks,” *Econometrica*, 77, 623–685.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA EKSTEN, AND S. TERRY (2014): “Really uncertain business cycles,” *Stanford mimeo*.
- BRANCH, W. A. (2007): “Sticky information and model uncertainty in survey data on inflation expectations,” *Journal of Economic Dynamics and Control*, 31, 245–276.
- BRUNNERMEIER, M. K. AND Y. SANNIKOV (2016): “Macro, Money and Finance: A Continuous Time Approach,” *Handbook of Macroeconomics*, 2, forthcoming.
- BRUNNERMEIER, M. K., A. SIMSEK, AND W. XIONG (2014): “A Welfare Criterion For Models With Distorted Beliefs,” *The Quarterly Journal of Economics*, 129, 1753–1797.
- BRYANT, J. (1983): “A simple rational expectations Keynes-type model,” *The Quarterly Journal of Economics*, 98, 525–528.
- BURDZY, K. AND D. M. FRANKEL (2005): “Shocks and business cycles,” *Advances in Theoretical Economics*, 5.
- BURDZY, K., D. M. FRANKEL, AND A. PAUZNER (2001): “Fast equilibrium selection by rational players living in a changing world,” *Econometrica*, 69, 163–189.
- BURGUET, R. AND X. VIVES (2000): “Social learning and costly information acquisition,” *Economic theory*, 15, 185–205.
- CABRALES, A., R. NAGEL, AND R. ARMENTER (2007): “Equilibrium selection through incomplete information in coordination games: an experimental study,” *Experimental Economics*, 10, 221–234.

- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12, 383–398.
- (1988): “Servicing the public debt: The role of expectations,” *The American Economic Review*, 647–661.
- CARLSSON, H. AND E. VAN DAMME (1993a): “Equilibrium selection in stag hunt games,” *Frontiers of Game Theory*, 237–253.
- (1993b): “Global games and equilibrium selection,” *Econometrica*, 989–1018.
- CARVALHO, C. AND F. NECHIO (2014): “Do people understand monetary policy?” *Journal of Monetary Economics*, 66, 108–123.
- CASS, D. AND K. SHELL (1983): “Do sunspots matter?” *The Journal of Political Economy*, 193–227.
- CAVALLO, A., G. CRUCES, AND R. PEREZ-TRUGLIA (2015): “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *NBER Working Paper Series*.
- CESPA, G. AND X. VIVES (2012): “Dynamic trading and asset prices: Keynes vs. Hayek,” *The Review of Economic Studies*, 79, 539–580.
- (2015): “The Beauty Contest and Short-Term Trading,” *The Journal of Finance*, 70, 2099–2154.
- CHAHROUR, R. (2014): “Public Communication and Information Acquisition,” *American Economic Journal: Macroeconomics*, 6, 73–101.
- CHAHROUR, R. AND G. GABALLO (2015): “On the Nature and Stability of Sentiments,” *Boston College mimeo*.
- CHAMLEY, C. (1999): “Coordinating regime switches,” *Quarterly Journal of Economics*, 869–905.
- (2004): *Rational herds: Economic models of social learning*, Cambridge University Press.
- CHASSANG, S. (2010): “Fear of miscoordination and the robustness of cooperation in dynamic global games with exit,” *Econometrica*, 78, 973–1006.
- CHEN, Q., I. GOLDSTEIN, AND W. JIANG (2010): “Payoff complementarities and financial fragility: Evidence from mutual fund outflows,” *Journal of Financial Economics*, 97, 239–262.
- CHRISTIANO, L., C. ILUT, R. MOTTO, AND M. ROSTAGNO (2008): “Monetary Policy and Stock Market Boom-Bust Cycles,” *Northwestern mimeo*.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary policy shocks: What have we learned and to what end?” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, 65–148.
- (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy,” *Journal of Political Economy*, 113, 1–45.
- CHRISTIANO, L. J., M. EICHENBAUM, AND R. VIGFUSSON (2003): “What happens after a technology shock?” *NBER Working Paper Series*.

- COIBION, O. AND Y. GORODNICHENKO (2012): “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy*, 120, 116–159.
- (2015): “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 105, 2644–78.
- COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2015): “How Do Firms Form Their Expectations? New Survey Evidence,” *NBER Working Paper Series*.
- COLOMBO, L., G. FEMMINIS, AND A. PAVAN (2014): “Information Acquisition and Welfare,” *Review of Economic Studies*, 81, 1438–1483.
- COOPER, R. (1999): *Coordination Games: Complementarities and Macroeconomics*, Cambridge University Press, Cambridge.
- COOPER, R. AND A. JOHN (1988): “Coordinating coordination failures in Keynesian models,” *The Quarterly Journal of Economics*, 441–463.
- CORNAND, C. (2006): “Speculative attacks and informational structure: an experimental study,” *Review of International Economics*, 14, 797–817.
- CORNAND, C. AND F. HEINEMANN (2009): “Speculative Attacks with Multiple Sources of Public Information,” *Scandinavian Journal of Economics*, 111, 73–102.
- CORSETTI, G., A. DASGUPTA, S. MORRIS, AND H. S. SHIN (2004): “Does one Soros make a difference? A theory of currency crises with large and small traders,” *The Review of Economic Studies*, 71, 87–113.
- CORSETTI, G., B. GUIMARAES, AND N. ROUBINI (2006): “International lending of last resort and moral hazard: a model of IMF’s catalytic finance,” *Journal of Monetary Economics*, 53, 441–471.
- COSTAIN, J. S. ET AL. (2007): “A Herding Perspective on Global Games and Multiplicity,” *The BE Journal of Theoretical Economics*, 7, 1–55.
- CUKIERMAN, A. AND A. H. MELTZER (1986): “A theory of ambiguity, credibility, and inflation under discretion and asymmetric information,” *Econometrica*, 1099–1128.
- DANIELSSON, J. AND F. PEÑARANDA (2011): “On the impact of fundamentals, liquidity, and coordination on market stability,” *International Economic Review*, 52, 621–638.
- DASGUPTA, A. (2007): “Coordination and delay in global games,” *Journal of Economic Theory*, 134, 195–225.
- DASGUPTA, A., J. STEINER, AND C. STEWART (2012): “Dynamic coordination with individual learning,” *Games and Economic Behavior*, 74, 83–101.
- DAVID, J. M., H. A. HOPENHAYN, AND V. VENKATESWARAN (2014): “Information, misallocation and aggregate productivity,” *NYU mimeo*.
- DAVILA, E. (2012): “Does Size Matter? Bailouts with Large and Small Banks,” *Harvard mimeo*.
- DEKEL, E., M. SINISCALCHI, ET AL. (2015): “Epistemic Game Theory,” *Handbook of Game Theory with Economic Applications*, 4, 619–702.

- DENTI, T. (2015): “Games with Unrestricted Information Acquisition,” *MIT mimeo*.
- DIAMOND, D. W. AND P. H. DYBVIK (1983): “Bank runs, deposit insurance, and liquidity,” *The Journal of Political Economy*, 401–419.
- DIAMOND, P. A. (1982): “Aggregate demand management in search equilibrium,” *The Journal of Political Economy*, 881–894.
- DUFFY, J. AND J. OCHS (2012): “Equilibrium selection in static and dynamic entry games,” *Games and Economic Behavior*, 76, 97–116.
- EDMOND, C. (2013): “Information Manipulation, Coordination, and Regime Change,” *The Review of economic studies*, 80, 1422–1458.
- EISENBACH, T. (2016): “Rollover risk as market discipline: a two-sided inefficiency,” *Federal Reserve Bank of New York mimeo*.
- ENNIS, H. M. AND T. KEISTER (2009): “Bank runs and institutions: The perils of intervention,” *The American Economic Review*, 99, 1588–1607.
- EVANS, G. W. AND S. HONKAPOHJA (1999): “Learning dynamics,” *Handbook of macroeconomics*, 1, 449–542.
- (2001): *Learning and expectations in macroeconomics*, Princeton University Press.
- (2009): “Learning and Macroeconomics,” *Annual Review of Economics*, 1, 421–451.
- FARHI, E. AND J. TIROLE (2012): “Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts,” *American Economic Review*, 102, 60–93.
- FARMER, R. E. (1996): “A theory of business cycles,” *Finnish Economic Papers*, 9, 91–109.
- FARMER, R. E. AND M. WOODFORD (1997): “Self-fulfilling prophecies and the business cycle,” *Macroeconomic Dynamics*, 1, 740–769.
- FRANKEL, D. M. (2016): “Optimal Insurance for Users of Network Goods,” *Iowa State University mimeo*.
- FRANKEL, D. M., S. MORRIS, AND A. PAUZNER (2003): “Equilibrium selection in global games with strategic complementarities,” *Journal of Economic Theory*, 108, 1–44.
- FRANKEL, D. M. AND A. PAUZNER (2000): “Resolving indeterminacy in dynamic settings: the role of shocks,” *Quarterly Journal of Economics*, 285–304.
- FUSTER, A., D. LAIBSON, AND B. MENDEL (2010): “Natural Expectations and Macroeconomic Fluctuations,” *Journal of Economic Perspectives*, 24, 67–84.
- FUTIA, C. A. (1981): “Rational Expectations in Stationary Linear Models,” *Econometrica*, 49, 171–192.
- GABAIX, X. (2014): “A Sparsity-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 129, 1661–1710.
- (2015): “Behavioral Macroeconomics Via Sparse Dynamic Programming,” *NYU mimeo*.

- GABALLO, G. (2015): “Price Dispersion, Private Uncertainty and Endogenous Nominal Rigidities,” *Banque De France mimeo*.
- GALI, J. (1999): “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, 89, 249–271.
- GALI, J. AND P. RABANAL (2005): “Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar US Data?” in *NBER Macroeconomics Annual 2004, Volume 19*, MIT Press, 225–318.
- GEANAKOPOLOS, J. (2010): “The leverage cycle,” in *NBER Macroeconomics Annual 2009, Volume 24*, University of Chicago Press, 1–65.
- GENNAIOLI, N., Y. MA, AND A. SHLEIFER (2015): “Expectations and investment,” in *NBER Macroeconomics Annual 2015, Volume 30*, University of Chicago Press.
- GERAATS, P. M. (2002): “Central Bank Transparency,” *The Economic Journal*, 112, 532–565.
- GERTLER, M., N. KIYOTAKI, AND A. PRESTIPINO (2016): “Wholesale Banking and Bank Runs in Macroeconomic Modelling of Financial Crises,” *Handbook of Macroeconomics*, 2, forthcoming.
- GIANNITSAROU, C. AND F. TOXVAERD (2006): “Recursive global games,” *University of Cambridge mimeo*.
- GOLDSTEIN, I. (2005): “Strategic Complementarities and the Twin Crises,” *The Economic Journal*, 115, 368–390.
- GOLDSTEIN, I., E. OZDENOREN, AND K. YUAN (2011): “Learning and complementarities in speculative attacks,” *The Review of Economic Studies*, 78, 263–292.
- (2013): “Trading frenzies and their impact on real investment,” *Journal of Financial Economics*, 109, 566–582.
- GOLDSTEIN, I. AND A. PAUZNER (2004): “Contagion of self-fulfilling financial crises due to diversification of investment portfolios,” *Journal of Economic Theory*, 119, 151–183.
- (2005): “Demand–deposit contracts and the probability of bank runs,” *the Journal of Finance*, 60, 1293–1327.
- GREENWOOD, R. AND A. SHLEIFER (2014): “Expectations of returns and expected returns,” *Review of Financial Studies*, 27, 714–746.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the impossibility of informationally efficient markets,” *The American Economic Review*, 70, 393–408.
- GUESNERIE, R. (1992): “An exploration of the eductive justifications of the rational-expectations hypothesis,” *The American Economic Review*, 1254–1278.
- (2008): “Macroeconomic and monetary policies from the eductive viewpoint,” *Central Banking, Analysis, and Economic Policies Book Series*, 13, 171–202.
- GUESNERIE, R. AND M. WOODFORD (1993): “Endogenous fluctuations,” in *Advances in Economic Theory*, ed. by J.-J. Laffont, Cambridge University Press, vol. 2, 289–412.

- GUIMARAES, B. (2006): “Dynamics of currency crises with asset market frictions,” *Journal of International Economics*, 68, 141–158.
- GUIMARAES, B., C. MACHADO, AND M. RIBEIRO (2014): “A model of the confidence channel of fiscal policy,” *Journal of Money, Credit and Banking*, forthcoming.
- GUIMARAES, B. AND S. MORRIS (2007): “Risk and wealth in a model of self-fulfilling currency attacks,” *Journal of Monetary Economics*, 54, 2205–2230.
- HAMILTON, J. D. (1996): “This is what happened to the oil price-macroeconomy relationship,” *Journal of Monetary Economics*, 38, 215–220.
- HANSEN, L. P. AND T. J. SARGENT (2008): *Robustness*, Princeton university press.
- HARRISON, J. M. AND D. M. KREPS (1978): “Speculative investor behavior in a stock market with heterogeneous expectations,” *The Quarterly Journal of Economics*, 323–336.
- HARSANYI, J. C. (1967-1968): “Games with Incomplete Information Played by Bayesian Players, Parts I, II, and III,” *Management Science*, 14, 159–182, 320–334, 486–502.
- HARSANYI, J. C. AND R. SELTEN (1988): “A general theory of equilibrium selection in games,” *MIT Press Books*, 1.
- HASSAN, T. A. AND T. M. MERTENS (2011): “Market sentiment: A tragedy of the commons,” *The American Economic Review*, 101, 402–405.
- (2014a): “Information Aggregation in a Dynamic Stochastic General Equilibrium Model,” *NBER Macroeconomics Annual*, 29, 159–207.
- (2014b): “The Social Cost of Near-Rational Investment,” *UChicago mimeo*.
- HAYEK, F. A. (1945): “The use of knowledge in society,” *The American economic review*, 519–530.
- HE, Z., A. KRISHNAMURTHY, AND K. MILBRADT (2016): “A model of safe asset determination,” *UChicago mimeo*.
- HE, Z. AND A. MANELA (2014): “Information Acquisition in Rumor-Based Bank Runs,” *The Journal of Finance*.
- HE, Z. AND W. XIONG (2012): “Dynamic debt runs,” *Review of Financial Studies*, 25, 1799–1843.
- HEIDHUES, P. AND N. MELISSAS (2006): “Equilibria in a dynamic global game: the role of cohort effects,” *Economic Theory*, 28, 531–557.
- HEINEMANN, F. (2000): “Unique equilibrium in a model of self-fulfilling currency attacks: Comment,” *The American Economic Review*, 90, 316–318.
- HEINEMANN, F., R. NAGEL, AND P. OCKENFELS (2004): “The theory of global games on test: Experimental analysis of coordination games with public and private information,” *Econometrica*, 72, 1583–1599.
- (2009): “Measuring strategic uncertainty in coordination games,” *The Review of Economic Studies*, 76, 181–221.

- HELLWIG, C. (2005): “Heterogeneous information and the welfare effects of public information disclosures,” *UCLA mimeo*.
- HELLWIG, C., A. MUKHERJI, AND A. TSYVINSKI (2006): “Self-Fulfilling Currency Crises: The Role of Interest Rates,” *American Economic Review*, 96, 1769–1787.
- HELLWIG, C. AND L. VELDKAMP (2009): “Knowing what others know: Coordination motives in information acquisition,” *The Review of Economic Studies*, 76, 223–251.
- HELLWIG, C. AND V. VENKATESWARAN (2009): “Setting the right prices for the wrong reasons,” *Journal of Monetary Economics*, 56, 57–77.
- HELLWIG, M. F. (1980): “On the aggregation of information in competitive markets,” *Journal of economic theory*, 22, 477–498.
- HOLDEN, S., G. JAMES, AND N. A. VIGIER (2014): “An Equilibrium Theory of Credit Rating,” *University of Oslo mimeo*.
- HOWITT, P. AND R. P. MCAFEE (1992): “Animal spirits,” *The American Economic Review*, 493–507.
- HUO, Z. AND N. TAKAYAMA (2015a): “Higher Order Beliefs, Confidence, and Business Cycles,” *Yale mimeo*.
- (2015b): “Rational expectations models with higher order beliefs,” *Yale mimeo*.
- IACHAN, F. S. AND P. T. NENOV (2015): “Information quality and crises in regime-change games,” *Journal of Economic Theory*, 158, 739–768.
- ILUT, C. L. AND M. SCHNEIDER (2014): “Ambiguous business cycles,” *American Economic Review*, 104, 2368–2399.
- IZMALKOV, S. AND M. YILDIZ (2010): “Investor sentiments,” *American Economic Journal: Microeconomics*, 2, 21–38.
- JACKSON, M. AND J. PECK (1991): “Speculation and price fluctuations with private, extrinsic signals,” *Journal of Economic Theory*, 55, 274–295.
- JAIMOVICH, N. AND S. REBELO (2009): “Can News about the Future Drive the Business Cycle?” *American Economic Review*, 99, 1097–1118.
- KACPERCZYK, M. T., S. VAN NIEUWERBURGH, AND L. VELDKAMP (2015): “A rational theory of mutual funds’ attention allocation,” *NYU mimeo*.
- KAJII, A. AND S. MORRIS (1997a): “Commonp-Belief: The General Case,” *Games and Economic Behavior*, 18, 73–82.
- (1997b): “The robustness of equilibria to incomplete information,” *Econometrica*, 1283–1309.
- KASA, K. (2000): “Forecasting the forecasts of others in the frequency domain,” *Review of Economic Dynamics*, 3, 726–756.
- KASA, K., T. B. WALKER, AND C. H. WHITEMAN (2007): “Asset prices in a time series model with perpetually disparately informed, competitive traders,” *University of Iowa mimeo*.

- KILEY, M. T. (2007): “A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting,” *Journal of Money, Credit and Banking*, 39, 101–125.
- KIYOTAKI, N. AND J. MOORE (1997): “Credit Cycles,” *The Journal of Political Economy*, 105, 211–248.
- KLENOW, P. J. AND J. L. WILLIS (2007): “Sticky information and sticky prices,” *Journal of Monetary Economics*, 54, 79–99.
- KONDOR, P. (2012): “The more we know about the fundamental, the less we agree on the price,” *The Review of Economic Studies*, 79, 1175–1207.
- KOVÁČ, E. AND J. STEINER (2013): “Reversibility in dynamic coordination problems,” *Games and Economic Behavior*, 77, 298–320.
- KOZENIAUSKAS, N., A. ORLIK, AND L. VELDKAMP (2014): “Black Swans and the many shades of uncertainty,” *NYU mimeo*.
- KUMAR, S., H. AFROUZI, O. COIBION, AND Y. GORODNICHENKO (2015): “Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand,” *NBER Working Paper Series*.
- KURLAT, P. (2015): “Optimal stopping in a model of speculative attacks,” *Review of Economic Dynamics*, 18, 212–226.
- KURLAT, P. AND L. VELDKAMP (2015): “Should we regulate financial information?” *Journal of Economic Theory*, 158, 697–720.
- KURZ, M. (1994): “On the structure and diversity of rational beliefs,” *Economic Theory*, 4, 877–900.
- (2012): “A new keynesian model with diverse beliefs,” *Stanford mimeo*.
- LAFFONT, J.-J. (1985): “On the Welfare Analysis of Rational Expectations Equilibria with Asymmetric Information,” *Econometrica*, 53, 1–29.
- LEVIN, J. (2001): “A note on global games with overlapping generations,” *Stanford mimeo*.
- LLOSA, L.-G. AND V. VENKATESWARAN (2015): “Efficiency with Endogenous Information Choice,” *NYU mimeo*.
- LORENZONI, G. (2009): “A Theory of Demand Shocks,” *American Economic Review*, 99, 2050–84.
- (2010): “Optimal monetary policy with uncertain fundamentals and dispersed information,” *The Review of Economic Studies*, 77, 305–338.
- LUCAS, R. E. (1972): “Expectations and the Neutrality of Money,” *Journal of economic theory*, 4, 103–124.
- (1973): “Some international evidence on output-inflation tradeoffs,” *The American Economic Review*, 63, 326–334.
- LUO, Y. (2008): “Consumption dynamics under information processing constraints,” *Review of Economic dynamics*, 11, 366–385.
- LUO, Y., J. NIE, AND E. R. YOUNG (2015): “Slow Information Diffusion and the Inertial Behavior of Durable Consumption,” *Journal of the European Economic Association*, 13, 805–840.

- MAĆKOWIAK, B. AND M. WIEDERHOLT (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769–803.
- (2015): "Business Cycle Dynamics under Rational Inattention," *The Review of Economic Studies*, 82, 1502–1532.
- MAKAROV, I. AND O. RYTCHKOV (2012): "Forecasting the forecasts of others: Implications for asset pricing," *Journal of Economic Theory*, 147, 941–966.
- MANKIW, N. G. AND R. REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 1295–1328.
- (2007): "Sticky information in general equilibrium," *Journal of the European Economic Association*, 5, 603–613.
- (2011): "Imperfect Information and Aggregate Supply," *Handbook of Monetary Economics*.
- MANKIW, N. G., R. REIS, AND J. WOLFERS (2004): "Disagreement about inflation expectations," in *NBER Macroeconomics Annual 2003, Volume 18*, The MIT Press, 209–270.
- MATEJKA, F. (2015a): "Rationally inattentive seller: Sales and discrete pricing," *The Review of Economic Studies*, forthcoming.
- (2015b): "Rigid pricing and rationally inattentive consumer," *Journal of Economic Theory*, 158, 656–678.
- MATEJKA, F. AND A. MCKAY (2015): "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105, 272–98.
- MATEJKA, F. AND C. A. SIMS (2011): "Discrete Actions in Information-Constrained Tracking Problems," *CERGE-EI mimeo*.
- MATHEVET, L. (2010): "A contraction principle for finite global games," *Economic Theory*, 42, 539–563.
- (2014): "Beliefs and rationalizability in games with complementarities," *Games and Economic Behavior*, 85, 252–271.
- MATHEVET, L. AND J. STEINER (2013): "Tractable dynamic global games and applications," *Journal of Economic Theory*, 148, 2583–2619.
- MATSUYAMA, K. (1991): "Increasing returns, industrialization, and indeterminacy of equilibrium," *The Quarterly Journal of Economics*, 617–650.
- (1995): "Complementarities and cumulative processes in models of monopolistic competition," *Journal of Economic Literature*, 33, 701–729.
- MCGRATTAN, E. R. (2004): "Comment on Gali and Rabanal's "Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar US Data?,"" *NBER Macroeconomics annual*, 19, 289–308.
- MEESE, R. A. AND K. ROGOFF (1983): "Empirical exchange rate models of the seventies: Do they fit out of sample?" *Journal of international economics*, 14, 3–24.

- MELOSI, L. (2014): “Estimating models with dispersed information,” *American Economic Journal: Macroeconomics*, 6, 1–31.
- MERTENS, J.-F. AND S. ZAMIR (1985): “Formulation of Bayesian analysis for games with incomplete information,” *International Journal of Game Theory*, 14, 1–29.
- MESSNER, S. AND X. VIVES (2005): “Informational and Economic Efficiency in REE with Asymmetric Information,” *IESE mimeo*.
- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 128, 1687–1726.
- MIAN, A. AND A. SUFI (2014): “What explains the 2007–2009 drop in employment?” *Econometrica*, 82, 2197–2223.
- MILANI, F. (2007): “Expectations, learning and macroeconomic persistence,” *Journal of Monetary Economics*, 54, 2065–2082.
- (2011): “Expectation shocks and learning as drivers of the business cycle,” *The Economic Journal*, 121, 379–401.
- MONDERER, D. AND D. SAMET (1989): “Approximating common knowledge with common beliefs,” *Games and Economic Behavior*, 1, 170–190.
- (1996): “Proximity of information in games with incomplete information,” *Mathematics of Operations Research*, 21, 707–725.
- MONDRIA, J. (2010): “Portfolio choice, attention allocation, and price comovement,” *Journal of Economic Theory*, 145, 1837–1864.
- MORRIS, S., A. POSTLEWAITE, AND H. S. SHIN (1995): “Depth of knowledge and the effect of higher order uncertainty,” *Economic Theory*, 6, 453–467.
- MORRIS, S. AND H. S. SHIN (1997): “Approximate common knowledge and co-ordination: Recent lessons from game theory,” *Journal of Logic, Language and Information*, 6, 171–190.
- (1998): “Unique equilibrium in a model of self-fulfilling currency attacks,” *American Economic Review*, 587–597.
- (2001): “Rethinking multiple equilibria in macroeconomic modeling,” in *NBER Macroeconomics Annual 2000, Volume 15*, MIT Press, 139–182.
- (2002): “Measuring Strategic Uncertainty,” *Princeton mimeo*.
- (2003): “Global games: theory and applications,” in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, Cambridge University Press.
- (2004a): “Coordination risk and the price of debt,” *European Economic Review*, 48, 133–153.
- (2004b): “Liquidity black holes,” *Review of Finance*, 8, 1–18.
- (2006): “Catalytic finance: When does it work?” *Journal of International Economics*, 70, 161–177.

- (2012): “Contagious adverse selection,” *American Economic Journal: Macroeconomics*, 1–21.
- MORRIS, S., H. S. SHIN, AND M. YILDIZ (2016): “Common belief foundations of global games,” *Journal of Economic Theory*, forthcoming.
- MORRIS, S. AND M. YANG (2016): “Coordination and the Relative Cost of Distinguishing Nearby States,” *Princeton mimeo*.
- MURPHY, K. M., A. SHLEIFER, AND R. W. VISHNY (1989): “Industrialization and the Big Push,” *The Journal of Political Economy*, 1003–1026.
- MYATT, D. P. AND C. WALLACE (2012): “Endogenous information acquisition in coordination games,” *The Review of Economic Studies*, 79, 340–374.
- NAGAR, V. AND G. YU (2014): “Accounting for Crises,” *American Economic Journal: Macroeconomics*, 6, 184–213.
- NAKAMURA, E. AND J. STEINSSON (2014): “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 104, 753–792.
- NIMARK, K. (2008): “Dynamic pricing and imperfect common knowledge,” *Journal of Monetary Economics*, 55, 365–382.
- (2011): “Dynamic higher order expectations,” *Universitat Pompeu Fabra mimeo*.
- NIMARK, K. P. AND S. PITSCHNER (2015): “Beliefs, coordination and media focus,” *Cornell mimeo*.
- OBSTFELD, M. (1986): “Rational and Self-fulfilling Balance-of-Payments Crises,” *American Economic Review*, 76, 72–81.
- (1996): “Models of currency crises with self-fulfilling features,” *European economic review*, 40, 1037–1047.
- OZDENOREN, E. AND K. YUAN (2008): “Feedback effects and asset prices,” *The journal of finance*, 63, 1939–1975.
- PACIELLO, L. AND M. WIEDERHOLT (2014): “Exogenous Information, Endogenous Information, and Optimal Monetary Policy,” *Review of Economic Studies*, 81, 356–388.
- PAVAN, A. (2015): “Attention, Coordination, and Bounded Recall,” *Northwestern mimeo*.
- PEARLMAN, J. G. AND T. J. SARGENT (2005): “Knowing the forecasts of others,” *Review of Economic Dynamics*, 8, 480–497.
- PECK, J. AND K. SHELL (1991): “Market uncertainty: correlated and sunspot equilibria in imperfectly competitive economies,” *The Review of Economic Studies*, 58, 1011–1029.
- PRATI, A. AND M. SBRACIA (2002): “Currency Crises and Uncertainty about Fundamentals,” *IMF mimeo*.
- RADNER, R. (1962): “Team decision problems,” *The Annals of Mathematical Statistics*, 33, 857–881.
- RAMEY, V. A. (2016): “Macroeconomic Shocks and Their Propagation,” *Handbook of Macroeconomics*, 2, forthcoming.

- REIS, R. (2006): "Inattentive producers," *The Review of Economic Studies*, 73, 793–821.
- (2009): "Optimal monetary policy rules in an estimated sticky-information model," *American Economic Journal: Macroeconomics*, 1–28.
- ROCA, M. (2005): "Transparency and Monetary Policy with Imperfect Common Knowledge," *Columbia mimeo*.
- ROCHET, J.-C. AND X. VIVES (2004): "Coordination failures and the lender of last resort: was Bagehot right after all?" *Journal of the European Economic Association*, 2, 1116–1147.
- RONDINA, G. AND T. WALKER (2014): "Dispersed Information and Confounding Dynamics," *Indiana University mimeo*.
- RUBINSTEIN, A. (1989): "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *The American Economic Review*, 385–391.
- SAKOVICS, J. AND J. STEINER (2012): "Who matters in coordination problems?" *The American Economic Review*, 102, 3439–3461.
- SARGENT, T. J. (1991): "Equilibrium with signal extraction from endogenous variables," *Journal of Economic Dynamics and Control*, 15, 245–273.
- (2008): "Evolution and intelligent design," *The American Economic Review*, 98, 3–37.
- SARTE, P.-D. (2014): "When Is Sticky Information More Information?" *Journal of Money, Credit and Banking*, 46, 1345–1379.
- SCHAAL, E. AND M. TASCHEREAU-DUMOUCHEL (2015): "Coordinating Business Cycles," *NYU mimeo*.
- SCHEINKMAN, J. A. AND W. XIONG (2003): "Overconfidence and speculative bubbles," *Journal of Political Economy*, 111, 1183–1220.
- SHELL, K. (1977): "Monnaie et allocation intertemporelle," *CNRS mimeo*.
- (1989): "Sunspot equilibrium," in *The New Palgrave: General Equilibrium*, Springer, 274–280.
- SHURCHKOV, O. (2013): "Coordination and learning in dynamic global games: experimental evidence," *Experimental Economics*, 16, 313–334.
- SIMS, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–690.
- (2006): "Rational inattention: Beyond the linear-quadratic case," *The American Economic Review*, 96, 158–163.
- (2010): "Rational inattention and monetary economics," *Handbook of Monetary Economics*, 3, 155–181.
- SIMSEK, A. (2013): "Belief disagreements and collateral constraints," *Econometrica*, 81, 1–53.
- SINGLETON, K. J. (1987): "Asset prices in a time-series model with disparately informed, competitive traders," in *New approaches to monetary economics*, ed. by W. A. Barnett and K. J. Singleton, Cambridge University Press, 249–272.

- SOCKIN, M. AND W. XIONG (2015): “Informational frictions and commodity markets,” *The Journal of Finance*, 70, 2063–2098.
- SOLOMON, R. H. (2003): “Anatomy of a twin crisis,” *Bank of Canada mimeo*.
- STEIN, J. C. (1989): “Cheap talk and the Fed: A theory of imprecise policy announcements,” *The American Economic Review*, 32–42.
- STEINER, J. (2008): “Coordination cycles,” *Games and Economic Behavior*, 63, 308–327.
- STEINER, J., C. STEWART, AND F. MATEJKA (2015): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” *CERGE-EI mimeo*.
- STEVENS, L. (2015): “Coarse Pricing Policies,” *Federal Reserve Bank of Minneapolis mimeo*.
- SZKUP, M. AND I. TREVINO (2015): “Information acquisition in global games of regime change,” *Journal of Economic Theory*, 160, 387–428.
- TARASHEV, N. A. (2007): “Speculative attacks and the information role of the interest rate,” *Journal of the European Economic Association*, 5, 1–36.
- TAYLOR, J. B. (1979): “Staggered wage setting in a macro model,” *The American Economic Review*, 69, 108–113.
- (1980): “Aggregate dynamics and staggered contracts,” *The Journal of Political Economy*, 1–23.
- TILLMANN, P. (2004): “Disparate information and the probability of currency crises: empirical evidence,” *Economics Letters*, 84, 61–68.
- TOWNSEND, R. M. (1983): “Forecasting the forecasts of others,” *The Journal of Political Economy*, 546–588.
- TOXVAERD, F. (2008): “Strategic merger waves: A theory of musical chairs,” *Journal of Economic Theory*, 140, 1–26.
- TUTINO, A. (2013): “Rationally inattentive consumption choices,” *Review of Economic Dynamics*, 16, 421–439.
- UHLIG, H. AND V. GUERRIERI (2016): “Macroeconomics and Financial Markets,” *Handbook of Macroeconomics*, 2, forthcoming.
- VAN NIEUWERBURGH, S. AND L. VELDKAMP (2009): “Information immobility and the home bias puzzle,” *The Journal of Finance*, 64, 1187–1215.
- (2010): “Information acquisition and under-diversification,” *The Review of Economic Studies*, 77, 779–805.
- VAN ZANDT, T. AND X. VIVES (2007): “Monotone equilibria in Bayesian games of strategic complementarities,” *Journal of Economic Theory*, 134, 339–360.
- VELDKAMP, L. L. (2006): “Media frenzies in markets for financial information,” *The American Economic Review*, 577–601.
- (2011): *Information choice in macroeconomics and finance*, Princeton University Press.

- VENKATESWARAN, V. (2014): “Heterogeneous Information and Labor Market Fluctuations,” *NYU mimeo*.
- VERRECCHIA, R. E. (1982): “Information acquisition in a noisy rational expectations economy,” *Econometrica*, 1415–1430.
- VIVES, X. (1988): “Aggregation of Information in Large Cournot Markets,” *Econometrica*, 56, 851–76.
- (1993): “How fast do rational agents learn?” *The Review of Economic Studies*, 60, 329–347.
- (1997): “Learning from others: a welfare analysis,” *Games and Economic Behavior*, 20, 177–200.
- (2005): “Complementarities and games: New developments,” *Journal of Economic Literature*, 43, 437–479.
- (2010): *Information and learning in markets: the impact of market microstructure*, Princeton University Press.
- (2014): “Endogenous public information and welfare in market games,” *IESE mimeo*.
- WALSH, C. E. (2007): “Optimal Economic Transparency,” *International Journal of Central Banking*, 3, 5–36.
- WEINSTEIN, J. AND M. YILDIZ (2007a): “Impact of higher-order uncertainty,” *Games and Economic Behavior*, 60, 200–212.
- (2007b): “A structure theorem for rationalizability with application to robust predictions of refinements,” *Econometrica*, 75, 365–400.
- WOODFORD, M. (1986): “Stationary sunspot equilibria in a finance constrained economy,” *Journal of Economic Theory*, 40, 128–137.
- (1991): “Self-Fulfilling Expectations and Fluctuations in Aggregate Demand,” in *The New Keynesian Macroeconomics*, M.I.T. Press.
- (2003): “Imperfect Common Knowledge and the Effects of Monetary Policy,” *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*.
- (2009): “Information-constrained state-dependent pricing,” *Journal of Monetary Economics*, 56, 100–124.
- (2013): “Macroeconomic Analysis Without the Rational Expectations Hypothesis,” *Annual Review of Economics*, 5, 303–346.
- YANG, M. (2015): “Coordination with flexible information acquisition,” *Journal of Economic Theory*, 158, 721–738.