

NBER WORKING PAPER SERIES

NEAR-RATIONAL EQUILIBRIA IN HETEROGENEOUS-AGENT MODELS:
A VERIFICATION METHOD

Leonid Kogan
Indrajit Mitra

Working Paper 30111
<http://www.nber.org/papers/w30111>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2022

We thank Santiago Bazdresch (discussant), Jack Favilukis (discussant), Ken Judd, Aubhik Khan, Jonathan Parker, Tony Smith (discussant), Julia Thomas, and seminar participants at City University of Hong Kong, Hong Kong University, INSEAD, MIT, Nanyang Technological University, Ohio State University, UC Berkeley, UCLA, University of Michigan, and at the Duke AP conference, Minnesota Macro-Asset Pricing conference, MIT Capital Markets Workshop, Society for Economic Dynamics Annual Meeting, SITE Stanford, and the Western Finance Association Meeting for comments. We thank Anton Petukhov for excellent research assistance. The views expressed here are ours and not necessarily those of the Federal Reserve Bank of Atlanta, the Federal Reserve System, or 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Leonid Kogan and Indrajit Mitra. 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.

Near-Rational Equilibria in Heterogeneous-Agent Models: A Verification Method
Leonid Kogan and Indrajit Mitra
NBER Working Paper No. 30111
June 2022
JEL No. C02,C18,C63,C68,E00,E37,G1

ABSTRACT

We propose a general simulation-based procedure for estimating quality of approximate policies in heterogeneous-agent equilibrium models, which allows to verify that such approximate solutions describe a near-rational equilibrium. Our procedure endows agents with superior knowledge of the future path of the economy, while imposing a suitable penalty for such foresight. The relaxed problem is more tractable than the original, and results in an upper bound on agents' welfare. Our method is general, straightforward to implement, and can be used in conjunction with various solution algorithms. We illustrate our approach in two applications: the incomplete-markets model of Krusell and Smith (1998) and the heterogeneous firm model of Khan and Thomas (2008).

Leonid Kogan
MIT Sloan School of Management
100 Main Street, E62-636
Cambridge, MA 02142
and NBER
lkogan@mit.edu

Indrajit Mitra
Federal Reserve Bank of Atlanta
1000 Peachtree Street N.E.
Atlanta, GA 30309
indrajit.mitra@atl.frb.org

1 Introduction

Cross-sectional heterogeneity among households and firms is at the heart of many important economic phenomena. In general, it is impossible to aggregate a cross-section of agent characteristics in dynamic heterogeneous-agent economies, especially in the presence of constraints (e.g., financial constraints) or un-hedgeable sources of risk (e.g., idiosyncratic labor income shocks). In such models, endogenous quantities, such as risk premia, depend on the cross-sectional distribution of agent characteristics, such as household wealth or firm capital. Since the cross-sectional distribution is an infinite dimensional state-variable, it is typically impossible to solve exactly for equilibrium in this class of models. Under the common solution approach, introduced in the highly influential paper by Krusell and Smith (1998), agents follow relatively simple approximate policies that avoid the burden of solving a dynamic optimization problem with a high-dimensional state space. Specifically, agents summarize the state of the economy by a low-dimensional state vector, typically keeping track of only a few cross-sectional moments.¹ The approximate solution of the original model can be viewed as an exact equilibrium in a near-rational economy, in which agents pursue suboptimal policies (see Krusell and Smith (1998), page 874). If agents suffer small welfare losses from failing to fully optimize, expanding further resources on improving the policies is unproductive, and approximate policies are plausible as a description of near-rational behavior. This argument is in the spirit of modeling economic agents as satisficing rather than optimizing, as in Simon (1978).

Do approximate solutions thus constructed describe near-rational equilibria? A significant limitation of the commonly used approaches, including the one by Krusell and Smith, is that currently there are no reliable general methods for verifying the degree of welfare loss under candidate near-rational equilibria in heterogeneous-agent models. We discuss the limitations of one common approach, based on Euler equation errors, in Section 2. In this paper we propose a general method for bounding welfare losses due to suboptimality of policies under the approximate solution. Our technique allows one to compute provable bounds on the

¹This approximation method is widely used, see e.g., Heathcote, Storesletten, and Violante (2009) and Guvenen (2011) for a survey of the solution methodology and applications.

degree of welfare loss under approximate policies. It is straightforward to implement and has general applicability, being usable in conjunction with various approximation algorithms.

The key to tractability of our approach is that it establishes an upper bound on the agents' welfare loss without computing the optimal policies. This is essential when dealing with infinite-dimensional models, where optimal individual policies are infeasible to compute, even in a candidate near-rational equilibrium. The main idea of our approach is the following. We alter the original problem of an agent by enlarging her information set to allow for perfect knowledge of the future path of prices (more generally, the aggregate state process of the economy), while simultaneously penalizing the agent's objective for such foresight. The modified problem is much more tractable than the original problem, because the aggregate state of the economy in the modified problem follows a deterministic process. Moreover, if the penalty for perfect foresight is chosen properly (we discuss the precise requirement in the main body of the paper), the value function of the modified problem is always higher than the value function of the original problem. We thus obtain an upper bound on the agent's welfare. The lower bound results from following the sub-optimal policy prescribed by the approximate solution. The gap between the two bounds limits the agent's welfare loss from above. A narrow gap indicates that the degree of sub-optimality is economically small, and the approximate equilibrium is indeed near-rational. A large gap does not necessarily imply that the sub-optimal policy is grossly inefficient, as it may result from the value function of the modified problem being significantly higher than the value function of the original problem.

To illustrate the potential of our method, we apply it to two well-known models, which feature an approximate equilibrium with aggregate uncertainty. First, we consider the incomplete markets model of Krusell and Smith (1998). This is a stochastic growth model in which individual agents face uninsurable labor income risk as well as aggregate shocks to the productivity of capital. Krusell and Smith compute an approximate equilibrium by summarizing the cross-sectional distribution of wealth among the agents using only the average per capita level of wealth. Our second application of the information relaxation approach is to the model of Khan and Thomas (2008). Their model features a heterogeneous

cross-section of firms in general equilibrium. In the approximate equilibrium, Khan and Thomas summarize the cross-sectional distribution by the mean capital stock of all firms in the economy. We provide accompanying code which shows our approach applied to these two models and to the simpler illustrative examples in Section 2 and Section 3 of this paper.²

We quantify the degree of sub-optimality of agents' policies under both models. We establish that in both settings the original solutions imply relatively low individual welfare losses for most initial configurations of the economy. Thus, for the calibrated models under consideration, our method confirms that their approximate solutions describe a near-rational equilibrium. This is especially important for the model of Khan and Thomas (2008), because we are able to show that the key finding of that paper that non-linearities in individual firm policies do not have a quantitatively large effect on aggregate dynamics is not a result of firms adopting grossly sub-optimal policies.

Next, we stress-test the approximation algorithms in the above applications by introducing transitional dynamics in an economy perturbed away from its steady state. Since the standard solution methods, such as that of Krusell-Smith, are not intended to approximate equilibrium dynamics accurately when the economy is away from its steady state, it is not clear a-priori how well such methods may perform. For both models, we consider the following two transitional dynamics experiments. Starting from the steady-state of the model, we consider: (i) an unanticipated permanent increase in the volatility of aggregate productivity shocks (we consider two-fold and five-fold increases); or, (ii) an unanticipated 50% reduction in capital stock of all agents in the Krusell-Smith economy and a similar reduction in capital stock of every firm in the Khan-Thomas economy. In the first case, the economy transitions to a new steady state following a permanent regime shift. In the second case, the economy reverts to the original steady state following a large transient shock. Our methodology shows the welfare bound in each experiment to be larger than in the steady-state case, in some instances rising by more than an order of magnitude, thus indicating potentially large welfare losses.

²See <https://www.dropbox.com/sh/rqe859kstso6vk0/AACBZNNuxCIqUY7BZRJfqt4a?dl=0>.

Related literature

The basic idea of using information relaxations and martingale multipliers to formulate a dual stochastic optimization problem can be traced back to Bismut (1973) (in a continuous-time setting) and Rockafellar and Wets (1976) and Pliska (1982) (in the discrete-time finite horizon setting). Back and Pliska (1987) apply this technique to single-agent problems in financial economics. Most of the existing applications of information relaxations deal with the optimal stopping problems, typically in the context of pricing American or Bermudan options, e.g., Davis and Karatzas (1994), Rogers (2002), Haugh and Kogan (2004), and Andersen and Broadie (2004). Rogers (2007) and Brown, Smith, and Sung (2010) extend the information-relaxation idea to general dynamic optimization problems. We use the formulation in Brown et al. (2010), which incorporates both perfect and partial information relaxations and derives penalty processes from the value function of the original problem. Our paper is the first to apply the information relaxation approach to approximate solutions of heterogeneous-agent equilibrium models.

The existing literature on approximate solutions of equilibrium models uses several approaches to evaluating approximate solutions. One common approach is to compare forecasts of aggregate states of the economy with actual realizations from the simulation, and to judge the approximation quality by the accuracy of the forecasts, e.g., their R^2 . A well-known limitation of this approach is that a high forecast R^2 does not guarantee that the approximation quality is high (den Haan, 2010). Krusell and Smith (1998) evaluate multi-period forecasts as a more stringent test, and den Haan (2010) develops a yet more stringent procedure for comparing the law of motion used to formulate agents' policy functions to the true law of motion implied by the approximate solution of the model. These approaches have two main limitations relative to the method we propose in this paper: they do not provide a guarantee of approximation quality and do not describe the welfare cost of approximation errors.

Another popular approach, due to den Haan and Marcet (1994), evaluates Euler equation errors of the approximate solution along the simulated path of the economy. Under the null hypothesis that the agent's policies are optimal, the \mathcal{L}^2 norm of the Euler equation errors is

distributed as a χ^2 random variable, and a standard hypothesis test can be carried out. The limitation of this method is that small Euler equation errors do not imply low welfare loss. As we show in Section 2, Euler equation errors can be small while sub-optimal policies are infinitely costly in welfare terms. The inadequacy of Euler equation errors as a measure of the approximation quality of equilibrium solutions is also highlighted in Kubler and Schmedders (2003).

Santos (2000) shows that small Euler equation errors do imply small policy function errors for a restricted class of models – importantly, this result is limited to the models in which equilibria correspond to the solution of the central planner’s problem. In more general models, theoretical guarantees on the accuracy of policy functions are not available. In such situations, our approach allows one to compute a generally applicable bound on the approximation accuracy of agents’ policies. Kubler and Schmedders (2005) propose a method of error analysis where the quality of the approximation to the equilibrium is judged by its proximity to an exact equilibrium in a close-by economy. In contrast to this approach, our method establishes an upper bound on the welfare loss in the original economy, which is due to agents following suboptimal policies.

In the language of numerical analysis, estimates of the errors in equilibrium policies, as in Santos (2000), represent forward error analysis – where the quality of the approximation is judged by how close the approximate solution (including agents’ policies and endogenous processes, such as prices) is to the exact solution. In comparison, Kubler and Schmedders (2005) use the logic of backward error analysis, where one evaluates how much the inputs of the model need to be modified to make the approximate solution satisfy all equilibrium conditions exactly. One can view our method as a form of forward error analysis with provable guarantees of approximation quality, where the distance between the approximate and exact solution is measured in economic terms – in terms of individual welfare loss under approximate policies. Judd, Maliar, and Maliar (2017) provide a complementary view of the solution quality. They establish a lower bound on the (forward) error of an approximate solution to an equilibrium model. While our approach provides a sufficient condition for the accuracy of an approximate solution, the lower bound provides a necessary condition since

the true errors are larger than the lower bound.

The rest of the paper is organized as follows. In Section 2 we show that the method of Euler equation errors can fail to detect sub-optimal policies with large utility losses. In Section 3 we formulate the relaxed problem and outline the construction of penalty functions. To illustrate our approach, we apply it to a model for which the optimal policy is known in closed form. In Sections 4 and 5, we apply our method to the Krusell-Smith model and the model of Khan-Thomas, respectively. Section 6 concludes.

2 Shortcomings of the Euler equation errors approach

We use two examples to show that the method of Euler equation errors can fail to detect sub-optimal policies with large utility losses. In the first example, the agent incurs infinite loss in expected utility from adopting a sub-optimal policy; however, the Euler equation errors remain finite. In the second example, a finite sample test based on Euler equation errors fails to reject the null hypothesis of policy optimality, even though the welfare loss associated with the policy is substantial. We describe the setting for both of our examples next.

Consider an investor with time-0 wealth w_0 and log utility over consumption. The agent has access to a risk-free bond with a constant rate of return R^B and a stock whose distribution of time $t + 1$ return R_{t+1}^S depends on the time t value of a state variable X . Assume that X follows an n state Markov process and takes n possible values X_1, X_2, \dots, X_n with a time independent transition probability $P_{ij} \equiv \text{Prob}(X_{t+1} = X_j | X_t = X_i)$. The investor solves

$$\sup_{c_t > 0, \phi_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log c_t, \quad (1)$$

subject to the budget constraint:

$$w_{t+1} = (w_t - c_t) (\phi_t R_{t+1}^S + (1 - \phi_t) R^B). \quad (2)$$

In (1), β is the time-preference parameter, c_t is time- t consumption, and ϕ_t is the share of the investor's wealth in the stock at time- t .

The optimal consumption policy c_t^* is to consume a constant fraction of wealth:

$$c_t^* = (1 - \beta)w_t. \quad (3)$$

The optimal portfolio policy is to choose ϕ_t to maximize the certainty equivalent of the one-period return on wealth, that is,

$$\phi_t^* = \arg \sup B(\phi_t, X_t), \quad (4)$$

where $B(\phi_t, X_t)$ is the certainty equivalent of the one-period return on wealth:

$$B(\phi_t, X_t) \equiv \mathbb{E} [\log (\phi_t R_{t+1}^S + (1 - \phi_t)R^B) | X_t] . \quad (5)$$

The expectation in equation (5) is over the distribution of stock returns R_{t+1}^S conditional on the time- t realization of the state variable X . We derive the policies (3) and (4) in Appendix A.1.

Next, consider the suboptimal policy resulting from the investor having incorrect beliefs about the distribution of stock returns. Specifically, instead of the true transition probabilities P_{ij} , the investor supposes that the transition probability is $\hat{P}_{ij} = \text{Prob}(X_{t+1} = X_j | X_t = X_i)$. The investor solves

$$\sup_{(\phi_u, c_u)_{u \geq 0}} \hat{\mathbb{E}}_0 \sum_{t=0}^{\infty} \beta^t \log c_t, \quad (6)$$

subject to the budget constraint (2), where the expectation $\hat{\mathbb{E}}_0$ in (6) is taken under the investor's beliefs. Since the agent's consumption policy is independent of portfolio returns (the agent has log-utility), the investor still consumes the same fraction of wealth as in equation (3), that is,

$$\hat{c}_t = (1 - \beta)c_t. \quad (7)$$

The investor's portfolio policy $\hat{\phi}_t$ is

$$\hat{\phi}_t = \arg \sup_{\phi} \hat{\mathbb{E}} [\log (\phi R_{t+1}^S + (1 - \phi) R^B) | X_t] . \quad (8)$$

Euler equation errors. To quantify deviations from optimality, den Haan and Marcet (1994), henceforth DM, proposed a hypothesis test based on deviations from the first-order optimality conditions. In our example, the investor's first-order optimality equations are $1 = \mathbb{E} [R^B (c_{t+1}/c_t)^{-1} | X_t]$ and $1 = \mathbb{E} [R_{t+1}^S (c_{t+1}/c_t)^{-1} | X_t]$. Deviations from these first-order conditions, that is, the Euler equation errors, are defined as:

$$\epsilon_t^B \equiv 1 - E [R^B / (\hat{\phi} R_{t+1}^S + (1 - \hat{\phi}) R^B) | X_t], \quad \epsilon_t^S \equiv 1 - E [R_{t+1}^S / (\hat{\phi}_t R_{t+1}^S + (1 - \hat{\phi}_t) R^B) | X_t], \quad (9)$$

where we used equation (2) to express consumption growth in terms of the portfolio return in the Euler equation. DM show that under the null hypothesis that the policy is optimal, the test statistics constructed from the \mathcal{L}^2 norm of the Euler equation errors

$$s^B \equiv \sum_{t=1}^T (\epsilon_t^B)^2 / T, \quad s^S \equiv \sum_{t=1}^T (\epsilon_t^S)^2 / T, \quad (10)$$

approach a χ^2 distribution with one degree of freedom as T approaches infinity.

We follow DM in implementing their test: we simulate N paths of stock returns (each of length T) and compute the values of the statistics s^B and s^S along each sample path. Using many paths minimizes the likelihood of not rejecting the null hypothesis due to luck. Finally, we compute the fraction of times these statistics fall in the upper and lower critical 5% region of a χ^2 distribution with one degree of freedom. If these realized fractions are substantially different from 5%, we have evidence that the policy being examined is not optimal.

To test the reliability of the DM approach, we explicitly compute the investor's loss in expected utility from adopting the policies (7) and (8) relative to adopting the optimal policies (3) and (4). We report this welfare loss as a fractional certainty equivalent loss, which we define as follows. Let $U((\phi_t, c_t)_{t \geq 0}; W_0, X_0)$ be the investor's expected utility from adopting policies ϕ_t and c_t for $t \geq 0$, for initial wealth W_0 and initial state

X_0 . For each W_0 and X_0 , we first compute the expected utility under the optimal policy $U((\phi_t^*, c_t^*)_{t \geq 0}; W_0, X_0)$. Next, we compute the initial wealth \widehat{W}_0 that is needed to achieve this utility from adopting the potentially suboptimal policy. That is, we solve for \widehat{W}_0 from the equation $U((\phi_t^*, c_t^*)_{t \geq 0}; W_0, X_0) = U((\widehat{\phi}_t, \widehat{c}_t)_{t \geq 0}; \widehat{W}_0, X_0)$. We define the fractional certainty equivalent loss as:

$$\eta \equiv \frac{\widehat{W}_0 - W_0}{W_0}. \quad (11)$$

2.1 Finite Euler equation residuals, unbounded welfare loss

Consider the special case in which stock returns are independent and identically distributed. In this case, the certainty equivalent of the single period ahead portfolio return is a constant $B(\phi) = \mathbb{E} \log(\phi R_{t+1}^S + (1 - \phi)R^B)$. Therefore, the optimal portfolio is a constant ϕ^* , where $\phi^* = \arg \sup B(\phi)$. The suboptimal portfolio policy (8) also implies a constant $\widehat{\phi}$ where

$$\widehat{\phi} = \arg \sup_{\phi} \widehat{\mathbb{E}} [\log(\phi R_{t+1}^S + (1 - \phi)R^B)]. \quad (12)$$

We prove in Appendix A.2, that the fractional certainty equivalent loss from adopting the suboptimal policy is

$$\eta = \exp\left(\frac{\beta}{1 - \beta} (B(\phi^*) - B(\widehat{\phi}))\right) - 1. \quad (13)$$

Equation (13) shows that if $\widehat{\phi} \neq \phi^*$, then the welfare loss becomes arbitrarily large as the time preference parameter β approaches unity. However, we see from equation (9) that the magnitude of Euler equation errors does not increase as β approaches one; indeed, these errors are independent of β . The Euler equation errors do not blow up because these errors are based on deviations from the one-period ahead Euler equations, which fail to aggregate the effect of such deviations over multiple periods. In terms of utility loss, however, a sufficiently patient investor who puts non-negligible weight to utility loss far into the future suffers a very large utility loss, even though the single-period deviation from the suboptimal policy appears small.

2.2 Low-probability persistent disasters

In this section we use the setting from Section 2 to show that the finite-sample test based on Euler equation errors may fail to reject suboptimal policies. The key feature of our example is the presence of a rare but persistent disaster state. The investor underestimates the persistence of this state. Since the state is rare, the investor's mistake is infrequently realized, and therefore a finite-sample test may fail to detect a policy associated with significant welfare loss.

Specifically, we consider the case in which the state variable X takes three values X_1 , X_2 , and X_3 ; the realized stock returns in these states are 1.3, 0.91, and 0.7, respectively. The true transition probability matrix is shown in Panel A of Table 1. The investor believes (incorrectly) that the transition probability matrix is the one shown in Panel B of Table 1. Comparing the transition matrices, we see that the investor underestimates the persistence of the disaster state X_3 . Note that even though the investor makes large errors in the conditional dynamics of the state variable X , these errors are much milder unconditionally. For example, while the true relative frequencies of occurrence of X_1 , X_2 , and X_3 are 49.45%, 49.45%, and 1.1%, respectively, the investor believes these frequencies to be 49.5%, 49.5%, and 1%, respectively. Similarly, while the true average stock return and volatility are 10.05% and 19.84%, respectively, the investor believes these quantities to be 10.09% and 19.82%, respectively. Finally, we choose the investor's time preference parameter $\beta = 0.99$ and the risk-free rate $R^B = 1.04$.

As a result of underestimating the persistence of X_3 , the investor overinvests in the stock in this state relative to the optimal portfolio by a significantly large amount. While the optimal policy is to invest $\phi^*(X_3) = 0.52$, the investor chooses $\hat{\phi}(X_3) = 1.73$. The first row of Table 2 shows the investor's welfare loss in each state as a result of adopting the suboptimal policy (12). We see that the fractional certainty equivalent loss is substantial: 4.09% in states X_1 and X_2 and 8.39% in the state X_3 .

Next, we use the DM test using Euler equation errors to assess evidence against the null hypothesis (that the investor's policy is optimal). The stringency of the DM test depends on the length of the time series T used. As DM point out, a low value of T increases the

likelihood of a Type II error (i.e., the test fails to detect a suboptimal policy), whereas if T is sufficiently large, every approximate solution will be rejected. DM suggest choosing T to be “substantially bigger than the length of the empirical series...”. In their example, they choose the series to be 20 times the length of the empirical series. We assume that the model of Section 2 is being used to analyze consumption-portfolio choices using the post-War data. Similar to DM, we choose $T = 6000$, which is close to twenty times the length of the quarterly consumption series in post-War data. We simulate $N = 10,000$ paths of stock returns and we compute the values of s^B and s^S along each sample path. Finally, we compute the fraction of times s^B and s^S fall in the upper and lower critical 5% region of a χ^2 distribution with one degree of freedom.

Rows 2 through 4 in Panel A of Table 2 show the results. We see that no entry is substantially different from 5%; hence according to this test, we would fail to reject the null in spite of the fact that the investor incurs substantial welfare loss from adopting these policies. In order to detect evidence against the null, we have to choose a much longer time series. This can be seen from Panel B of Table 2, where we use $T = 100,000$. In this case, both s^B and s^S appear in the 5% critical region 43.12% of the time when the initial state is X_3 .

Our example highlights a weakness of the approach of using Euler equation errors (see also the related discussion in den Haan and Marcet (1994)). Both examples in Section 2.1 and Section 2.2 underscore the need for a reliable test to detect suboptimal policies. We propose and discuss such a test next.

3 Information relaxation: The main idea

In our analysis of approximate equilibria, we apply the information relaxation method proposed in Brown et al. (2010). We introduce the main ideas of this method in this section, and refer the readers to Brown et al. (2010) for full technical details.

Consider a standard finite-horizon consumption-savings problem. Time is discrete, $t = (0, \dots, T)$. Each period the agent receives a random labor income which takes two possible values $\{y_H, y_L\}$ with $y_H > y_L$. The probability of receiving y_H is p each period. The agent

chooses consumption c_t and stores the rest in a risk-free asset with constant total return R . We denote the agent's feasible consumption policy by $C = (c_0, c_1, \dots, c_T)$.

At each date t , the agent observes the history of income shocks realized up to and including this date (but not the future shocks). We denote this history by $y^t = (y_0, y_1, \dots, y_t)$. All feasible consumption policies must be adapted to the information structure of the agent, i.e., consumption choices are functions of the observed past histories of income shocks. Thus, making the agent's information structure explicit, $C = (c_0(y^0), c_1(y^1), \dots, c_T(y^T))$.

The agent has a time-separable constant relative risk aversion utility function with curvature γ . Let w_t denote the agent's wealth at the beginning of period t , including the realized income in the current period. The agent solves the dynamic optimization problem:

$$\sup_{\{C: c_t \leq w_t\}} \mathbb{E}_0 \left[\sum_{t=0}^T \beta^t \frac{c_t^{1-\gamma}}{1-\gamma} \right], \quad (14)$$

where the agent's wealth and consumption satisfy the dynamic budget constraint:

$$w_t = (w_{t-1} - c_{t-1})R + y_t. \quad (15)$$

We denote the value function of the above problem by $V_t(w_t)$:

$$V_t(w_t) = \mathbb{E}_t \left[\sum_{s=t}^T \beta^{s-t} \frac{(c_s^*)^{1-\gamma}}{1-\gamma} \right], \quad (16)$$

where C^* is the optimal consumption policy.

We formulate a relaxed problem by allowing the agent to have access to information about the future realizations of income shocks. The name ‘‘information relaxation’’ reflects the notion that this formulation relaxes information constraints placed on the agent. Specifically, consider a *complete information relaxation*, whereby we allow the agent to condition her consumption choices on the knowledge of the entire future sequence of income shocks. To distinguish the feasible policies of the relaxed problem from those of the original problem, we denote the former by $C^{\mathcal{R}} = (c_0^{\mathcal{R}}(y^T), c_1^{\mathcal{R}}(y^T), \dots, c_T^{\mathcal{R}}(y^T))$.

While providing the agent with such informational advantage compared to that in the

original problem, we impose a penalty on the objective function, designed to offset the effect of information relaxation. The penalty is a stochastic process λ_t , which depends on the consumption policy and the entire path of income shocks, y^T : $\lambda_t(C^{\mathcal{R}}, y^T)$. The only requirement we impose on the penalty process is that if the consumption process is chosen to depend only on the information available to the agent, the resulting penalty is non-positive in expectation, i.e.,

$$\mathbb{E}_0 [\lambda_t (C^{\mathcal{R}}, y^T)] \leq 0. \quad (17)$$

It is easy to see that the value function of the relaxed problem:

$$V_0^{\mathcal{R}}(w_0) = \sup_{\{C^{\mathcal{R}}: c_t^{\mathcal{R}} \leq w_t^{\mathcal{R}}\}} \mathbb{E}_0 \left[\sum_{t=0}^T \beta^t \left(\frac{(c_t^{\mathcal{R}})^{1-\gamma}}{1-\gamma} - \lambda_t(C^{\mathcal{R}}, y^T) \right) \right], \quad (18)$$

subject to the dynamic budget constraint:

$$w_t^{\mathcal{R}} = (w_{t-1}^{\mathcal{R}} - c_{t-1}^{\mathcal{R}})R + y_t, \quad (19)$$

is at least as high as the value function of the original problem. The reason is that the consumption policy C^* , optimal under the agent's original problem (14-15), is also a feasible policy for the relaxed problem (18-19), and the expected penalty under such policy adds a non-negative term to the agent's expected utility, according to (17). Thus, we establish that:

$$V_0^{\mathcal{R}}(w_0) \geq V_0(w_0). \quad (20)$$

Next, consider a feasible but suboptimal consumption policy \widehat{C} . Under this suboptimal policy, the expected utility of the agent is given by:

$$\widehat{V}_0(w_0) = \mathbb{E}_0 \left[\sum_{t=0}^T \beta^t \frac{(\widehat{c}_t)^{1-\gamma}}{1-\gamma} \right], \quad (21)$$

which results in a welfare loss of $V_0 - \widehat{V}_0$. To estimate this welfare loss resulting from a suboptimal strategy, we use the inequality (20) to conclude that the agent's welfare loss is bounded above by the difference between the value function of the relaxed problem (18-19)

and the expected utility under the suboptimal policy $\widehat{c}_t(y^t)$:

$$V_0(w_0) - \widehat{V}_0(w_0) \leq V_0^{\mathcal{R}}(w_0) - \widehat{V}_0(w_0). \quad (22)$$

We thus have a framework for computing bounds on welfare loss resulting from sub-optimal strategies: define a valid penalty process for the relaxed problem, and then compare the expected utility of the agent under information relaxation with her expected utility under the suboptimal policy of interest. This approach is especially useful where it is infeasible to solve for the optimal policies of the original problem.³

While this formulation is rather general, it is only useful as long as the resulting bound is relatively tight, i.e., as long as the value function of the relaxed problem $V_0^{\mathcal{R}}$ is not much higher than the expected utility of the agent under the optimal consumption policy, V_0 . Brown et al. show that it is possible to make the difference $V_0^{\mathcal{R}} - V_0$, i.e. the *duality gap* arbitrarily small. In particular, they show (using backwards induction) that under an *ideal* penalty, $V_0^{\mathcal{R}} - V_0 = 0$.

The definition of an ideal penalty by Brown et al. adopted to our example is as follows. The penalty is defined for each possible sequence of income shocks, and each possible sequence of consumption choices, without requiring the consumption policy to be non-anticipating. Specifically, consider an arbitrary path of income shocks y^T , and a budget-feasible positive sequence of consumption choices $c^T = (c_0, c_1, \dots, c_T)$. Note that c^T is not a consumption policy, it denotes a sequence of positive real numbers representing a particular path of consumption. The corresponding values of agent's wealth (w_0, w_1, \dots, w_T) satisfy the dynamic budget constraint equation (19). To develop intuition of how an ideal penalty may look, we first consider a somewhat restricted version of information relaxation: suppose that at time t , the decision maker is able to anticipate perfectly the state next period, and must use the original information structure starting from the following period, i.e., perfect foresight is there for a single period only. In this case, the Bellman equation of the relaxed problem at

³For complete information relaxation, the relaxed problem is deterministic and hence easy to solve.

time t takes the form:

$$V_t^{\mathcal{R}}(w_t) = \sup_{\{c_t^{\mathcal{R}}: c_t^{\mathcal{R}} \leq w_t\}} \frac{(c_t^{\mathcal{R}})^{1-\gamma}}{1-\gamma} - \lambda_t^*(c_t^{\mathcal{R}}, y^{t+1}) + \beta V_{t+1}((w_t - c_t^{\mathcal{R}})R + y_{t+1}). \quad (23)$$

In the above expression, V_{t+1} is the value function of the original problem, which coincides with the value function of the relaxed problem starting at $t+1$, given our assumption of a single-period relaxation at time t . There is no expectation operator in the Bellman equation, because the agent anticipates income at time $t+1$, y_{t+1} , perfectly. $\lambda_t^*(c_t^{\mathcal{R}}, y^{t+1})$ is the penalty term, to be determined, which depends on the realization of income next period, and the current choice of consumption under the single-period perfect foresight, $c_t^{\mathcal{R}}$, which is also allowed to dependent on next-period income. Compare the above expression to the Bellman equation of the original consumption choice problem:

$$V_t(w_t) = \sup_{\{c_t: c_t \leq w_t\}} \frac{(c_t)^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}[V_{t+1}((w_t - c_t)R + \tilde{y}_{t+1})|y^t]. \quad (24)$$

The second term on the right in (24) is the standard expectation of V_{t+1} over the possible values of \tilde{y}_{t+1} , taking the realizations of income shocks y_0, \dots, y_t as given. In order for the optimal consumption policy of the relaxed problem to be the same as for the original problem, we set the penalty to equalize expressions in equations (23) and (24):

$$\lambda_t^*(c_t^{\mathcal{R}}, y^{t+1}) = \beta V_{t+1}((w_t - c_t^{\mathcal{R}})R + y_{t+1}) - \beta \mathbb{E}[V_{t+1}((w_t - c_t^{\mathcal{R}})R + \tilde{y}_{t+1})|y^t]. \quad (25)$$

Brown et al. (2010) show, using mathematical induction, that the above expression for the ideal penalty is valid in general, and not just for single-period information relaxations (for completeness, we reproduce their proof in the context of our problem in Appendix A.3). Given the current portfolio value of the relaxed problem, $w_t^{\mathcal{R}}$, the general form of the ideal penalty is

$$\lambda_t^*(C^{\mathcal{R}}, y^T) = \beta V_{t+1}((w_t^{\mathcal{R}} - c_t^{\mathcal{R}})R + y_{t+1}) - \beta \mathbb{E}[V_{t+1}((w_t^{\mathcal{R}} - c_t^{\mathcal{R}})R + \tilde{y}_{t+1})|y^t]. \quad (26)$$

y_{t+1} in the above expression is not a random variable. It is the value of the income shock at

time $t + 1$ from the particular path of shocks y^T for which we are defining the ideal penalty. In contrast, the labor income shock at time $t + 1$, i.e., \tilde{y}_{t+1} , is a random variable. The second term on the right in (26) is a conditional expectation of V_{t+1} over the possible values of \tilde{y}_{t+1} , taking the realizations of income shocks y_0, \dots, y_t as given. Thus, the second term depends only on $(c^{\mathcal{R}})^t$ and y^t , and so the penalty λ_t^* depends on $(c^{\mathcal{R}})^t$ and y^{t+1} . In particular, $\lambda_t^*(C^{\mathcal{R}}, y^T)$ depends on the contemporaneous consumption choice $c_t^{\mathcal{R}}$ and the future income shock y_{t+1} explicitly, and on the earlier consumption choices $(c^{\mathcal{R}})^{t-1}$ and income shocks y^t implicitly, through $w_t^{\mathcal{R}}$ and the dynamic budget constraint. Going forward, we use more concise notation for the penalty:

$$\lambda_t^*(C^{\mathcal{R}}, y^T) = \beta (V_{t+1}(w_{t+1}^{\mathcal{R}}) - \mathbb{E}_t[V_{t+1}(w_{t+1}^{\mathcal{R}})]) . \quad (27)$$

To visualize how the penalty affects the solution of the relaxed problem, consider the dependence of the ideal penalty on the contemporaneous consumption choice c_t . For this numerical example, we choose the parameters shown in Table 3. The time-1 penalty λ_1^* , is shown in Figure 1. It depends on c_1 , the income in the following period y_2 , and current wealth w_1 , where w_1 captures the dependence of the penalty on prior consumption choices and income shocks. We plot the penalty for two different levels of current wealth, $w_1 = 4$ and $w_1 = 5$, in Panels A and B of Figure 1, respectively. Each line in these figures shows the penalty as a function of c_1 and the two possible values of y_2 : the solid line corresponds to $y_2 = y_H$, while the dash-dot line corresponds to $y_2 = y_L$.

To see how the penalty discourages the agent from using information about future income, consider a relaxed problem with the agent observing the time-2 income shock in advance and using this information in his time-1 consumption decision. Without the penalty, the agent can take advantage of his knowledge of the future. In particular, if the agent knows that the time-2 income shock is high, i.e., $y_2 = y_H$, it is optimal to choose higher time-1 consumption than if $y_2 = y_L$. Figure 1 shows that the ideal penalty discourages such behavior. If the agent chooses higher consumption in the $y_2 = y_H$ state relative to the $y_2 = y_L$ state, the expected penalty is positive (shown by the solid line). An ideal penalty has the property that the

benefit of perfect foresight is exactly offset by the negative effect of the penalty, and the agent finds it optimal to choose a non-anticipative consumption policy while knowing future realizations of income shocks. As long as the consumption choice is non-anticipating, i.e., c_1 does not depend on y_2 , the expected value of the penalty is zero (shown by the dash line), and the welfare of the agent is not impaired by the penalty.

Next we show how the ideal penalty discourages the agent from conditioning the time-0 consumption choice on knowledge of y_2 to achieve a smoother consumption path. To see this compare panels A and B of Figure 1. From these figures we see that for both realizations of y_2 , the gradient of the penalty λ_1^* is larger in absolute value for $w_1 = 4$ than for $w_1 = 5$. Selecting higher c_0 in the $y_2 = y_H$ state relative to the $y_2 = y_L$ state raises the expected penalty term λ_1^* , making it positive even if the consumption choice at time 1 is non-anticipative. This illustrates the inter-temporal connections between various penalty terms and consumption choices.

An ideal penalty is as difficult to compute as the solution of the original problem. We therefore define the penalty based on an approximation to the value function:

$$\lambda_t(C^{\mathcal{R}}, y^T) = \beta \left(\widehat{V}_{t+1}(w_{t+1}^{\mathcal{R}}) - \mathbb{E}_t \left[\widehat{V}_{t+1}(w_{t+1}^{\mathcal{R}}) \right] \right), \quad (28)$$

where we define \widehat{V}_t to be the expected utility resulting from the agent's consumption policy. This is a feasible penalty because it satisfies equation (17). However, this penalty choice results in an upward biased estimate of the actual welfare loss of the agent. We see this from Panel A of Figure 2. In this figure, the value function of the relaxed problem $V^{\mathcal{R}}(w)$ (shown by the dot-dash line), is greater than the value function $V(w)$ (shown by the dashed line). This duality gap arises from a sub-optimal choice of penalty; had we used the ideal penalty, λ^* , the two value functions would have coincided and the duality gap would have been zero. The solid line in this figure is the expected utility of the agent $\widehat{V}(w)$, from adopting the sub-optimal policy. We estimate it by simulating many paths of shocks and computing the sample mean of realized utilities from adopting the agent's policy. The information relaxation approach provides us with an estimate of the difference $V^{\mathcal{R}}(w) - \widehat{V}(w)$; this difference is an upper bound on the actual welfare loss $V(w) - \widehat{V}(w)$ (see equation (22)).

Panel B of Figure 2 shows an estimate of the upper bound on welfare loss of an agent using sub-optimal policies and compares this with the actual welfare loss. The agent uses a consumption policy based on the optimal solution of the model with the probability of the high state equal to $\hat{p} = 0.89$, whereas the true probability is $p = 0.9$. We report welfare loss of an agent as a fractional certainty equivalent loss, η , defined similar to equation (11):

$$\eta(w_0) = \frac{w'_0 - w_0}{w'_0}, \quad (29)$$

where w'_0 is computed by solving:

$$\widehat{V}_0(w'_0) = V_0^{\mathcal{R}}(w_0). \quad (30)$$

The numerator of η in equation (29) is therefore the additional amount of time-0 wealth, w'_0 , needed by an agent following a sub-optimal policy to obtain the level of expected utility equal to the value function of the relaxed problem with time-0 wealth equal to w_0 , keeping all other state variables the same. The solid line in Figure 2 shows an estimate of the upper bound on welfare loss of the agent $\bar{\eta}$, computed using the information relaxation approach. We obtain this estimate in three steps: (i) by simulating 500 paths of income shocks, (ii) solving the relaxed problem (i.e., equations (18) and (19)) along each path thereby obtaining an estimate of $V_0^{\mathcal{R}}(w, y^T)$ along each path, and (iii) taking the sample mean, $V_0^{\mathcal{R}}(w)$ over all 500 paths.⁴ For this simple, partial equilibrium example, since we are able to solve for the actual welfare loss, we plot it in the same figure (the dot-dash line) alongside the upper-bound in panel B of Figure 2. We define the actual welfare loss, η^* , using the value function $V_0(w)$ instead of $V_0^{\mathcal{R}}(w)$, i.e. w'_0 in equation (29) is the root of the equation $\widehat{V}_0(w'_0) = V_0(w_0)$ instead of equation (30). In this example, the agent's welfare loss relative to the optimal policy is less than 0.27% in certainty equivalent terms. Information relaxation bounds the maximum welfare loss at less than 0.3%.

In order for information relaxation to be useful, it is necessary that the duality gap

⁴We do not show the 5% and 95% confidence intervals for the estimated upper bound on welfare loss in this example because this estimate is so precise that the confidence band is not visible separately from the estimated mean.

between $V_0^{\mathcal{R}}(w)$ and $V_0(w)$ be small. Put differently, applying a more sub-optimal penalty results in a larger duality gap and a less informative upper bound on welfare loss. To see this, consider a penalty function that is identically zero, for example. In other words, the agent is not penalized for foresight. Note that a zero penalty is a feasible penalty since it satisfies equation (17). We see from Panel C of Figure 2 that when we re-estimate the maximum welfare loss with a zero penalty, the upper bound on welfare loss is quite inflated. With this sub-optimal penalty choice, information relaxation bounds the maximum welfare loss at less than 65%, whereas the actual welfare loss is less than 0.27%.

Next, we vary the degree of sub-optimality of policies adopted by the agent and we estimate the welfare loss for each of these policies. In particular, we compare the welfare loss to the agent from adopting policies corresponding to $\hat{p} = 0.87$, $\hat{p} = 0.88$, and $\hat{p} = 0.89$ and report both the upper bound (solid line) and the actual welfare loss (dot-dash line) in panels A, B, and C, respectively of Figure 3. From these figures we see that the upper bound on welfare loss declines as the agent's policy approaches the optimal policy. For example, the maximum value of the upper bound on welfare loss is less than 2.3% when $\hat{p} = 0.87$, while it is less than 0.3% when $\hat{p} = 0.89$.

Comparing across Panels A through C of Figure 3, we see that the duality gap between $V^{\mathcal{R}}(w)$ and $V(w)$ decreases as the agent's policy approaches the optimal policy. This is because when the agent adopts a policy that is closer to the optimal policy, the expected utility $\hat{V}(w)$ improves and approaches the value function $V(w)$. This, in turn, implies that the penalty λ approaches the ideal penalty λ^* , and therefore, the duality gap decreases. In other words, the estimated upper bound on welfare loss declines as the agent's policy approaches the optimal policy both because of a decline in the actual welfare loss and also because of a decline in the duality gap.

The general information relaxation approach follows the same logic as the basic example above, with a multivariate state vector potentially replacing the wealth of the agent as an argument in the value function, and multiple choice variables potentially replacing the single choice variable, c . In addition, the general approach allows for partial information relaxations, where the agent receives some but not complete information about the future. Formally,

we describe the structure of the agent’s information as a filtration $\mathbb{F} = \{\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_T\}$, and the information set of the relaxed problem as a finer filtration $\mathbb{G} = \{\mathcal{G}_0, \mathcal{G}_1, \dots, \mathcal{G}_T\}$, where $\mathcal{F}_t \subseteq \mathcal{G}_t \subseteq \mathcal{F}_T$. Then, we define the relaxed problem under the information structure \mathbb{G} , and we define the penalty process as:

$$\lambda_t = \beta \left(\mathbb{E}_t[V(x_{t+1})|\mathcal{G}_t] - \mathbb{E}_t[V(x_{t+1})|\mathcal{F}_t] \right), \quad (31)$$

where the two expectation operators above are conditional on the corresponding information sets, and x_{t+1} denotes the time- $(t + 1)$ state vector (to avoid introducing more notation, we suppress the dependence of the penalty process on the choice variables and the exogenous shocks).

Before we apply the information relaxation approach to specific models, we note that, in general, there are two potential sources of numerical error in the solution of an equilibrium model. First, errors are introduced by approximating a high dimensional state space of a problem with a lower dimensional proxy. A second source of error arises when a continuous state space is approximated by a finite number of points on a discrete grid. Our focus in this paper is on the former, i.e., we provide an approach to obtain an upper bound on errors introduced by the curse of dimensionality.⁵

4 Application 1: imperfect insurance with aggregate uncertainty

We demonstrate the potential of the information relaxation methodology by computing bounds on the welfare loss of individual agents in the incomplete market model of Krusell and Smith (1998) (henceforth, KS). This model is a canonical example of a model with an

⁵Discretization of a continuous state variable introduces a subtlety. Strictly speaking, our method produces an upper bound only if the dual problem can be solved exactly. Otherwise, the welfare loss resulting from discretization no longer guarantees that the value function of the discrete approximation of the relaxed problem is an upward biased estimate of the value function of the original problem. Therefore, the dual problem has to be solved exactly to obtain a provable upper bound of the welfare loss from adopting heuristic policies. This can be achieved by providing the agent with perfect foresight and solving the resulting deterministic problem exactly.

infinite dimensional state space. We review the model, the equilibrium concept, and their solution approach briefly. We refer the reader to the original paper for details.

4.1 The model

The KS model is the Aiyagari model (Aiyagari, 1994) with aggregate uncertainty. Time is discrete, $t = (0, 1, \dots, \infty)$. There is a continuum of agents of unit measure with identical constant relative risk aversion preferences:

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma} \right]. \quad (32)$$

There is a single consumption good produced using a Cobb-Douglas production function:

$$y_t = z_t k_t^\alpha l_t^{1-\alpha}, \quad (33)$$

where the capital share parameter, $0 < \alpha < 1$, k and l are capital and labor inputs, respectively, and z is aggregate productivity. All agents are exposed to the aggregate shock z , which takes one of two values $z = \{z_h, z_l\}$ (with $z_h > z_l$) and follows a Markov chain. Each period's output is partly used for consumption and partly added to the next-period capital stock, resulting in the capital accumulation constraint:

$$k_t = (1 - \delta)k_{t-1} + y_{t-1} - c_{t-1}. \quad (34)$$

where δ is the capital depreciation rate.

Households collect capital rent and labor income each period. Individual labor income is exposed to idiosyncratic employment shocks, ε_t . We assume that each agent supplies \bar{l} units of labor if employed ($\varepsilon_t = 1$), and zero units if unemployed ($\varepsilon_t = 0$). Employment shocks are cross-sectionally independent, conditionally on the aggregate productivity shock. Thus, based on the law of large numbers, the unemployed fraction of the population depends only on the aggregate state. We denote the equilibrium unemployment rate conditional on z_h and z_l by u_h and u_l , respectively. Then the aggregate labor supply in the two states u_h and u_l are given by $L_h = (1 - u_h)\bar{l}$ and $L_l = (1 - u_l)\bar{l}$, respectively.

We look for a competitive recursive equilibrium. Let $\psi_t(k, \epsilon)$ denote the cross-sectional distribution function at time t , defined over individual capital stock and employment status. Aggregate output depends on the aggregate capital stock, $K_t = \int \psi(k, \epsilon) dk d\epsilon$, and the aggregate supply of labor. Input prices in competitive equilibrium are determined by their marginal product, hence the capital rent r and the wage rate w are given by:

$$r(K_t, L_t, z_t) = \alpha z_t \left(\frac{K_t}{L_t} \right)^{\alpha-1}, \quad w(K_t, L_t, z_t) = (1 - \alpha) z_t \left(\frac{K_t}{L_t} \right)^{\alpha}. \quad (35)$$

Individuals optimize their consumption-investment policies under rational expectations about market prices, i.e., we assume that they correctly forecast the law of motion of the equilibrium cross-sectional distribution of agents, denoted by:

$$\psi_t = \mathcal{H}(\psi_{t-1}, z_{t-1}). \quad (36)$$

Thus, optimal individual policies depend on the cross-sectional distribution of capital.

The value function of the agents satisfies the Bellman equation:

$$\begin{aligned} V_t(k_t, \varepsilon_t, z_t, \psi_t) &= \sup_{c_t \geq 0, k_{t+1} \geq 0} \left[\frac{c_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t [V_{t+1}(k_{t+1}, \varepsilon_{t+1}, z_{t+1}, \psi_{t+1})] \right] \\ \text{where } k_{t+1} &= (1 - \delta + r_t)k_t + w_t \bar{l} \varepsilon_t - c_t, \\ \psi_{t+1} &= \mathcal{H}(\psi_t, z_t), \end{aligned} \quad (37)$$

The main difficulty in computing the competitive equilibrium arises because of the dependence of equilibrium prices on the cross-sectional distribution of agents. Thus, to solve for the equilibrium, we must determine the law of motion, $\psi_{t+1} = \mathcal{H}(\psi_t, z_t)$.

Solution approach. To make the problem tractable, KS use a low-dimensional approximation to the infinite-dimensional cross-sectional distribution ψ_t . This approach, introduced in Krusell and Smith (1998), approximates all relevant information about the cross-sectional distribution by its first K moments, $\{m_1, m_2, \dots, m_K\}$. This results in a low-dimensional fixed-point problem. In particular, in their analysis, KS restrict their attention to the cross-sectional mean m_1 . For notational simplicity, from now on, we omit the sub-script in m_1 and

denote the distribution's mean simply by m . To speed up computation further, KS posit an approximate law of motion for m that is log-linear:

$$\widehat{\mathcal{H}} : \log m_{t+1} = a^z + b^z \log m_t, \quad z = \{z_g, z_b\}. \quad (38)$$

In an approximate equilibrium, the equilibrium dynamics of the cross-sectional distribution of capital in the economy must conform closely to the assumed law of motion, equation (38).

To solve the fixed-point problem, we start with an initial guess for the four constants $\{a^z, b^z\}$. The individual optimization problem, i.e., equation (37), is solved for optimal policies $k_{t+1}(k_t, \epsilon_t, z_t, m_t)$. With these policies, we simulate a long time series of the cross-sectional distribution using a large sample cross-section of agents and compute the time-series of the cross-sectional mean, m .⁶ Next, we compute the ordinary least squares regression estimates of $\{a^z, b^z\}$ based on equation (38). We re-solve the individual problem, i.e., equation (37), with these updated estimates of $\{a^z, b^z\}$, and the new optimal policies are used to simulate a new time-series of m . This is used to update $\{a^z, b^z\}$, and this process is continued till the system converges.⁷ In solving for the approximate equilibrium, we use parts of the code in Maliar, Maliar, and Valli (2010).

4.2 The relaxed problem

We apply the approach that we described in Section 3 to compute an upper bound on the welfare loss of individual agents by relaxing the information set of the agents. In particular, we start with an initial cross-sectional distribution of capital across agents and draw a sequence of aggregate shocks z_0, \dots, z_{T-1} and a panel of idiosyncratic employment shocks. Given the aggregate and individual shocks in each period, we use the approximate equilibrium policies of the agents and compute their choice of capital stock for the following period, which, in turn, gives us the cross-sectional mean of the distribution of capital stock in the following

⁶See also, Young (2010), who uses a histogram over the capital grid to obtain the time-series of m , instead of the simulating the path of m using a cross-section. While this approach speeds up the computation, the equilibrium is approximate since agents approximate the cross-sectional distribution by its mean.

⁷We stop when the maximum change in the policy function k_{t+1} between successive iterations is less than 10^{-8} , and the change in the norm of the vector $\{a^z, b^z\}$ between successive iterations is less than 10^{-8} .

period, m_{t+1} . Using equation (35), we compute the prices, i.e., capital returns r and wages w , corresponding to the realized sequence of aggregate shocks. To minimize the gap between the value function of the relaxed problem and the value function of the original problem of the agent, we apply a partial relaxation, revealing only future aggregate shocks but not the agent's idiosyncratic employment shocks.

We define the penalty according to equation (31):

$$\lambda_t = \beta \left(\mathbb{E}_t [\widehat{V}_{t+1}(k_{t+1}, \tilde{\varepsilon}_{t+1}, z_{t+1}, m_{t+1}) | \mathcal{G}_t] - \mathbb{E}_t [\widehat{V}_{t+1}(k_{t+1}, \tilde{\varepsilon}_{t+1}, \tilde{z}_{t+1}, m_{t+1}) | \mathcal{F}_t] \right), \quad (39)$$

where \mathcal{G}_t denotes the information set of the agent. In equation (39), \mathcal{G}_t contains the following period's realization of the aggregate shock z_{t+1} and the cross-sectional mean of capital m_{t+1} , but not $\tilde{\varepsilon}_{t+1}$. Therefore, we average over possible employed and un-employed future states. Knowledge of the transition probabilities for z and ε allows us to compute both the terms in (39) above explicitly as a function of the decision variables of the relaxed problem, $(k_{t+1}^{\mathcal{R}}, c_t^{\mathcal{R}})$. Finally, we use the budget constraint to eliminate $k_{t+1}^{\mathcal{R}}$. Along a particular path, the penalty λ_t is then a function of consumption $c_t^{\mathcal{R}}$ only.

4.3 Results

We carry out our baseline analysis using the same parameters as in Krusell and Smith (1998). The preference parameters are $\beta = 0.99$, and $\gamma = 1$. On the production side, the parameters are: the capital share $\alpha = 0.36$, the depreciation rate $\delta = 0.025$, aggregate productivity shocks take values $z_h = 1.01$, $z_l = 0.99$, and the corresponding aggregate unemployment rates are $u_h = 0.04$, $u_l = 0.10$. The transition probability matrix for (z, ε) is in Table 4.

All of our simulation results use a sample cross-section of $N = 10^5$ agents. Sample paths are $T = 10^3$ long, and we average over 500 paths to compute unbiased estimates of the value function of the relaxed problem, $V^{\mathcal{R}}$. In choosing the number of paths, we face a trade-off—using more paths provides more precise estimates of $V^{\mathcal{R}}$, but increases the computational time since the relaxed problem is solved path-by-path. We use a simulation-based approach to estimate the expected utility, \widehat{V} , under the sub-optimal policy of the agent in the approximate equilibrium. In particular, we simulate 10^5 future paths of aggregate and idiosyncratic shocks.

For each path, we compute realized prices r and w , as well as the realized utility. The sample mean of the realized utilities is our estimate of \widehat{V} .

As in Section 3, we report the welfare loss as a fractional certainty equivalent. The definition is analogous to equation (11):

$$\eta(k_0) = \frac{k'_0 - k_0}{k'_0}, \quad (40)$$

where k'_0 is the root of the equation $\widehat{V}_0(k'_0) = V_0^{\mathcal{R}}(k_0)$. The numerator of η is the extra capital needed by an agent following a sub-optimal policy to obtain the level of expected utility equal to the value function of the relaxed problem with initial capital k_0 , keeping all other state variables the same.

Baseline results

The welfare loss of an agent depends on the current state: the agent's capital stock, employment status, and the state of the aggregate economy captured by the realization of the aggregate shock and the cross-sectional distribution of capital across the agents. In our baseline results, we report welfare loss for an agent in a typical state of the economy, i.e., where the cross-sectional distribution of capital corresponds to the stochastic steady state.

Figure 4 summarizes the results. Panels A and B correspond to the state in which the agent is unemployed (i.e., $\epsilon = 0$) and employed (i.e., $\epsilon = 1$), respectively. The aggregate shock z is low. We see from both figures that individual welfare losses are small, especially for high levels of initial capital. For example, an agent who is unemployed (Panel A) and has capital stock equal to the bottom 5% of the distribution of capital stock, suffers a welfare loss that is at most 0.13% in fractional certainty equivalent terms. This number drops to 0.04% for an agent whose capital stock is equal to the distribution's median. These numbers are similar for an agent who is currently employed (Panel B). Thus, our results verify that the approximation of Krusell and Smith based on moment truncation, produces an approximate equilibrium in which agents come very close to fully optimizing their objectives when the economy is in a typical initial state.

Transitional dynamics

Next, we consider how accurately the approximate solutions describe the transitional dynamics of the economy following an unanticipated aggregate shock. The transitional dynamics of equilibrium models is often of great interest. Yet the standard solution methods, such as that of KS, are not intended to approximate equilibrium dynamics accurately when the economy is away from its steady state. It is therefore unclear a priori how well such approximation approaches perform under such experiments. This becomes an important issue when drawing conclusions about transitional dynamics of the economy based on approximate numerical solutions. The information relaxation approach is useful in this context because it provides a provable upper bound which quantifies the approximation accuracy of numerical solutions in such experiments.

We illustrate this approach using two examples in which we compute the transitional dynamics of the KS economy following two kinds of unanticipated shocks. In our first experiment, the economy experiences a large transitory shock: a sudden loss of 50% of capital stock of every agent. Such large shocks are considered in the disaster risk literature (see e.g., Gourio 2012). The second shock is a regime change: the economy gradually transitions from its baseline equilibrium to the new stochastic steady state following an unanticipated permanent increase in the volatility of the aggregate shock z . We consider two values for this increase: a two-fold and a five-fold increase in volatility.

Figure 5 shows the result for the scenario in which all agents suddenly lost half of their capital stock. We assume that every agent knows that the economy has experienced the shock which has depleted the aggregate capital stock of the economy to half its value. In other words, they observe the mean m decline. Accordingly, they adopt the policy corresponding to the new value of m , in the period in which the shock is realized. Since all structural parameters of the economy have stayed unchanged, agents rebuild their capital stock over the next several periods. Panels A and B correspond to the welfare loss of an agent who is unemployed and employed, respectively, in the period in which the unanticipated shock arrives. From these two figures we see that the upper bound to the welfare loss is larger relative to the baseline scenario. For example, an agent who is unemployed (Panel A) and

has capital stock equal to the bottom 5% of the distribution, suffers a welfare loss that is at most 0.30% in fractional certainty equivalent terms. The corresponding value was 0.13% in the stochastic steady-state. Comparing Panels A and B, we see that the welfare losses are similar in magnitude. Thus, our information relaxation method establishes that, following the shock to capital stock, the moment truncation approximation generates relatively low individual welfare loss.

Next we use the information-relaxation algorithm to quantify how well the Krusell-Smith algorithm performs following a permanent increase in the volatility of the aggregate shock, z . We assume that agents re-optimize and immediately switch to new policies following the regime change. Figure 6 shows the result. Panels A and B correspond to the welfare loss of an agent following a two-fold and five-fold increase in aggregate volatility, respectively. In both cases the agent is initially employed and the aggregate state of the economy is low. Panel A shows that for the two-fold increase in volatility, the upper bound to the welfare loss is not much larger than the baseline scenario. However, panel B shows that the welfare loss is potentially larger by an order of magnitude for the five-fold increase in volatility compared to the welfare loss in the stochastic state. For example, an agent who has capital stock equal to the bottom 5% of the distribution, suffers a welfare loss that is at most 3.0% in fractional certainty equivalent terms. The corresponding value was 0.13% in the stochastic steady-state. Thus, our information relaxation method establishes that the quality of the KS approximation approach continues to be good following an increase in aggregate volatility, unless the shock is extremely large.

5 Application 2: heterogeneous firms with aggregate uncertainty

In this section we apply our method to the general equilibrium model of Khan and Thomas (2008). This model features a heterogeneous cross-section of firms facing non-convex adjustment costs and exposed to persistent aggregate and firm-specific productivity shocks. We use information relaxation to estimate the upper bound to the loss in firm value from following

suboptimal investment policies. We briefly outline the model using the notation of Khan and Thomas (2008). We refer the reader to the original paper for more details.

5.1 The model

A continuum of firms of unit mass use a decreasing returns to scale technology with production function:

$$y_t = z_t \varepsilon_t k_t^\alpha n_t^\nu, \quad (41)$$

where $0 < \alpha < 1$ and $\alpha + \nu < 1$. Productivity has an aggregate component z and a firm-specific component ε . Both follow a Markov process, where $z_t \in \{z^1, \dots, z^{N_z}\}$ and $P(z_{t+1} = z^j | z_t = z^i) = \pi_{ij}$, $\varepsilon_t \in \{\varepsilon^1, \dots, \varepsilon^{N_\varepsilon}\}$ and $P(\varepsilon_{t+1} = \varepsilon^j | \varepsilon_t = \varepsilon^i) = \pi_{ij}^\varepsilon$.⁸ The firm hires labor n_t in period t after observing that period's productivity. The capital accumulation constraint is:

$$\gamma k_{t+1} = (1 - \delta)k_t + i_t, \quad (42)$$

where i_t is investment, δ is the depreciation rate of capital, and γ is a constant which captures the growth rate of labor-augmenting technological progress. Firms face non-convex adjustment cost of capital: there are no adjustment costs if investment is within a small range $i_t \in [ak_t, bk_t]$, where the parameters $a \leq 0 \leq b$. However, if investment is outside this range, then the adjustment cost is equal to $\xi_t \omega_t$, where ω_t is the real wage rate and ξ_t is a uniformly distributed random variable $\xi \sim \mathcal{U}[0, \bar{\xi}]$ that is independent across firms and over time.

⁸These transition probabilities are computed using the discretization method of Rouwenhorst (1995) applied to the AR(1) processes: $\log z' = \rho_z \log z + \eta'_z$ and $\log \varepsilon' = \rho_\varepsilon \log \varepsilon + \eta'_\varepsilon$, where $\eta'_z \sim N(0, \sigma_{\eta_z}^2)$ and $\eta'_\varepsilon \sim N(0, \sigma_{\eta_\varepsilon}^2)$.

Each period a firm chooses labor and investment to maximize its net present value:

$$\begin{aligned}
V^1(\varepsilon_t, k_t, \xi_t; z_t, \mu_t) &= \sup_{i_t \in \mathbb{R}, n_t \geq 0} \left[(z_t \varepsilon_t k_t^\alpha n_t^\nu - \omega_t n_t - i_t - \omega_t \xi_t I\{i_t \notin [ak_t, bk_t]\}) p_t \right. \\
&\quad \left. + \beta \sum_{i=1}^{N_z} \pi_{ij} \sum_{l=1}^{N_\varepsilon} \pi_{lm}^\varepsilon V(\varepsilon^m, k_{t+1}; z^j, \mu_{t+1}) \right] \\
\mu_{t+1} &= \mathcal{H}(\mu_t, z_t) \\
V(\varepsilon_t, k_t; z_t, \mu_t) &= \int_0^{\bar{\xi}} \bar{\xi}^{-1} V^1(\varepsilon_t, k_t, \xi; z_t, \mu_t) d\xi, \tag{43}
\end{aligned}$$

where $I\{\}$ is the indicator function, $V^1(\varepsilon_t, k_t, \xi_t; z_t, \mu_t)$ is the present value of a firm with idiosyncratic productivity ε_t and realized adjustment cost ξ_t , $V(\varepsilon_t, k_t; z_t, \mu_t)$ is the present value of the firm prior to the realization of the adjustment cost ξ_t , p_t is the price at which current output is valued, and μ is the cross-sectional distribution over individual firms' capital stock and idiosyncratic shocks. The presence of non-convex adjustment costs and persistent differences in firm-productivity lead to non-linear firm policies. This prevents aggregation of the cross-section of firms into a representative firm. Equilibrium prices, therefore, depend on the cross-sectional distribution of capital stock and idiosyncratic shocks μ_t of all firms in the economy, in addition to current aggregate productivity z_t . Firms optimize under rational expectations about market prices, i.e., we assume that firms correctly forecast the equilibrium law of motion of the distribution, $\mu_{t+1} = \mathcal{H}(\mu_t, z_t)$.

The model is closed by assuming that a representative household owns all firms in the economy. The household maximizes expected lifetime utility over consumption C_t and labor L_t :

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(C_t, L_t) \right] = \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (\log C_t - \varphi L_t) \right]. \tag{44}$$

An exact solution of equilibrium policies is infeasible in this model because of the dependence of prices on the cross-sectional distribution of capital μ . Khan and Thomas (2008) adopt the solution approach of Krusell and Smith (1998) and approximate the cross-sectional distribution by its mean. To solve the model, they make two more approximations. First,

they assume that the mean m follows a log-linear law of motion:

$$\widehat{H} : \log m_{t+1} = a^z + b^z \log m_t, \quad z = \{z^1, \dots, z^{N_z}\}. \quad (45)$$

Second, they assume that the dependence of the price p_t on the mean takes the form:

$$\log \widehat{p}_t = a_p^z + b_p^z \log m_t, \quad z = \{z^1, \dots, z^{N_z}\}. \quad (46)$$

The solution algorithm iteratively solves for optimal firm policies and the constants $\{a^z, b^z, a_p^z, b_p^z\}_{z=1}^{N_z}$ such that the decision rules of individual firms are consistent with the aggregate savings and leisure decisions of the representative household. For exact details see Khan and Thomas (2008).

5.2 The relaxed problem

To compute the value function $V_0^{\mathcal{R}}(\varepsilon, k; z_0, \mu_0)$ of the relaxed problem for a particular initial state of the economy (z_0, μ_0) , we simulate paths of the economy starting from this state. The relaxed information set contains the realization of all future aggregate productivity shocks; however, it does not include future realizations of idiosyncratic productivity and adjustment costs shocks. We define the penalty in the relaxed problem according to equation (31):

$$\begin{aligned} \lambda_t &= \beta \left(\mathbb{E}_t [\widehat{V}_{t+1}(\varepsilon_{t+1}, k_{t+1}; z_{t+1}, \widehat{\mu}_{t+1}) | \mathcal{G}_t] - \mathbb{E}_t [\widehat{V}_{t+1}(\varepsilon_{t+1}, k_{t+1}; z_t^l, \widehat{\mu}_{t+1}) | \mathcal{F}_t] \right) \\ &= \beta \left(\sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon \widehat{V}(\varepsilon^l, k_{t+1}; z_{t+1}, m_{t+1}) - \sum_{j=1}^{N_z} \pi_{ij} \sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon \widehat{V}(\varepsilon^l, k_{t+1}; z^i, m_{t+1}) \right). \quad (47) \end{aligned}$$

Along a particular path z_0, z_1, \dots, z_T of the aggregate shock z , the time- t penalty λ_t is therefore a function of the control $k_{t+1}^{\mathcal{R}}$. We solve for the optimal control of this relaxed problem and obtain a single realization of the value function of the relaxed problem. We repeat this over many simulated paths of aggregate shocks; the sample mean over these paths is our estimate of $V_0^{\mathcal{R}}(\varepsilon, k; z_0, \mu_0)$.

5.3 Results

We carry out our baseline analysis using the same parameters as in Khan and Thomas (2008) (see Table 5 of this paper). We choose the number of states for the aggregate and idiosyncratic shocks equal $N_z = 11$ and $N_\varepsilon = 15$, respectively, as in their original paper. All of our simulation results use a sample cross-section of $N = 10^5$ firms. Sample paths are $T = 10^3$ long, and we average over 15000 paths to compute unbiased estimates of the value function of the relaxed problem, $V^{\mathcal{R}}(\varepsilon, k; z_0, \psi_0)$. Similar to our approach for the KS model in Section 4, we use a simulation-based approach to estimate the expected utility, $\widehat{V}(\varepsilon, k; z_0, \psi_0)$, under the sub-optimal policy of the firm in the approximate equilibrium. In particular, we simulate 10^5 future paths of aggregate and idiosyncratic shocks, and compute the realized utility along each path. The sample mean of the realized utilities is our estimate of \widehat{V} . Since there are no optimizations involved, this step is fast.

As in Section 3, we report the welfare loss as a fractional certainty equivalent. The definition is analogous to equation (29):

$$\eta = \frac{k'_0 - k_0}{k'_0}, \quad (48)$$

where k'_0 is the root of the equation $\widehat{V}_0(k'_0) = V_0^{\mathcal{R}}(k_0)$. The numerator of η in equation (48), is the extra capital needed by a firm following a sub-optimal policy to obtain the level of expected utility equal to the value function of the relaxed problem with initial capital k_0 , keeping all other state variables the same.

Baseline Results

The loss in firm value depends on the aggregate state of the economy, i.e. the realization of the aggregate shock and the cross-sectional distribution μ , as well as the firm's current state, i.e. the capital stock k and the idiosyncratic shock ε . In our baseline results, we report the loss in firm value in a typical state of the economy, i.e., where the cross-sectional distribution corresponds to the stochastic steady state. In Figure 7 we present upper bounds on the loss in firm value in three different idiosyncratic productivity states: the lowest state $\varepsilon = \varepsilon^1$ in

Panel A, the middle state $\varepsilon = \varepsilon^8$ in Panel B, and the highest state $\varepsilon = \varepsilon^{15}$ in Panel C, all as functions of the firm’s capital stock, k .

We see from Figure 7 that information relaxation bounds the welfare losses from following sub-optimal policies in the near-rational equilibrium to be economically negligible. For example, comparing across Panels A through C, we see that the upper bound is largest in Panel A, corresponding to the state $\varepsilon = \varepsilon^1$. Even in this state, a firm with capital stock equal to the bottom 5% of the distribution suffers a welfare loss that is at most 0.02% in fractional certainty equivalent terms. In the same figure we see that a firm with capital stock equal to the median of the distribution suffers a welfare loss that is less than 0.01%. For higher firm-specific shocks, the upper bound is even smaller. In Panel C which corresponds to $\varepsilon = \varepsilon^{15}$, for example, the loss is essentially zero. In sum, this shows that firms do not incur significant loss by following suboptimal policies under circumstances they typically encounter in the near-rational equilibrium.

Transitional Dynamics

Next, we test the efficiency of the suboptimal policies by considering transitional dynamics of the economy away from the stochastic steady-state. As in Section 4, we consider two unanticipated shocks. In our first experiment, the economy experiences a large transitory shock: a sudden destruction of 50% of capital stock of every firm in the economy. The second shock is a regime change: the economy gradually transitions from its baseline equilibrium to the new stochastic steady state following an unanticipated permanent increase in the volatility of the aggregate shock z . We consider two shock sizes—a two-fold increase and also an extremely large shock corresponding to a five-fold increase in aggregate volatility.

Figure 8 shows the upper bound on welfare loss for the unanticipated destruction of 50% of capital stock of firms. We see from these figures that the upper bound is now larger than in the stochastic steady-state. For example, in Panel A which corresponds to the state $\varepsilon = \varepsilon^1$, a firm with capital stock equal to the bottom 5% of the distribution suffers a welfare loss that is bounded at 1% in fractional certainty equivalent terms. The corresponding value in the stochastic steady-state was 0.02%. The potential losses to firm value are even larger in

Panels B and C of Figure 8.

Figure 9 shows the upper bound on welfare loss for the unanticipated increase in the volatility of aggregate productivity. Panel A corresponds to a doubling of aggregate volatility; we see that the upper bound is similar to the stochastic steady-state in this case. Panel B corresponds to a five-fold increase in aggregate volatility. We see from these figures that the upper bound is much larger than in the stochastic steady-state. For example, a firm with capital stock equal to the bottom 5% of the distribution suffers a welfare loss which is bounded from above by our information relaxation approach at 0.5% in fractional certainty equivalent terms. The corresponding value in the stochastic steady-state is about an order of magnitude smaller at 0.07%. The results of these two transitional dynamics experiments therefore indicate potentially large losses in firm value in the Khan and Thomas economy when the economy experiences very large, unanticipated shocks.

6 Conclusion

Our analysis shows that information relaxation techniques can be effectively used to establish the accuracy of approximate solutions for equilibria in heterogeneous-agent models. This methodology is general, easy to implement, and has a wide range of potential applications beyond the scope of this paper. For instance, information relaxation could be used to evaluate the accuracy of solutions obtained using perturbation techniques. The latter approach is widely used for DSGE models (for a recent application, see Mertens and Judd (2012) and Mertens (2011)) because of its ability to handle models with high-dimensional state vectors, and is supported by the computational software Dynare. More recently, Boppart, Krusell, and Mitman (2018) and Auclert, Bardóczy, Rognlie, and Straub (2020) introduce an efficient method to solve for the approximate equilibrium of a large class of heterogeneous agent models using a first-order perturbation around the stochastic steady-state. They assume that the equilibrium is well approximated as a linear system in the space of perfect-foresight shock sequences of finite length. Our method can be used to guarantee the quality of the linearization assumption when their method is applied to solve for the equilibrium of a

particular model. Yet another natural application of our approach is to evaluating welfare loss resulting from heuristic policies motivated by behavioral biases of the agents.

Finally, our objective has been to establish near-rationality of individual policies under approximate solutions of equilibrium models, as measured by the associated welfare loss. A small welfare loss implies that individual agents have little to gain by refining their strategies. However, there is no guarantee that price dynamics in such a near-rational economy is similar to that in an exact equilibrium of the original model. Small mistakes by individual agents may potentially lead to large differences in equilibrium outcomes (e.g., Akerlof and Yellen (1985), Jones and Stock (1987), Naish (1993), Krusell and Smith (1996), Hassan and Mertens (2011)). An important and challenging task for future research is to develop general quantitative tools for evaluating the effect of small deviations from individual rationality on equilibrium price dynamics.

References

- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109(3), 659 – 684.
- Akerlof, G. A. and J. L. Yellen (1985). Can small deviations from rationality make significant differences to economic equilibria? *The American Economic Review* 75, 708 – 720.
- Andersen, L. and M. Broadie (2004). Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science* 50(9), 1222 – 1234.
- Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2020). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. Working Paper 26123, National Bureau of Economic Research.
- Back, K. and S. R. Pliska (1987). The shadow price of information in continuous time decision problems. *Stochastics* 22(2), 151 – 186.
- Bismut, J. M. (1973). Conjugate convex functions in optimal stochastic control. *Journal of Math. Anal. Appl.* 44, 384 – 404.
- Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control* 89, 68–92.
- Brown, D., J. E. Smith, and P. Sung (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58(4), 785 – 801.
- Davis, M. H. A. and I. Karatzas (1994). A deterministic approach to optimal stopping. In F. P. Kelly (Ed.), *Probability, Statistics and Optimisation*, pp. 455–466. John Wiley & Sons Ltd.
- den Haan, W. J. (2010). Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. *Journal of Economic Dynamics and Control* 34(1), 79 – 99.
- den Haan, W. J. and A. Marcet (1994). Accuracy in simulations. *Review of Economic Studies* 61(1), 3 – 17.
- Gourio, F. (2012, May). Disaster risk and business cycles. *American Economic Review* 102(6), 2734–66.
- Guvenen, F. (2011). Macroeconomics with heterogeneity: A practical guide. Working Paper 17622, National Bureau of Economic Research.
- Hassan, T. A. and T. M. Mertens (2011, May). The social cost of near-rational investment. NBER Working Papers 17027, National Bureau of Economic Research, Inc.
- Haugh, M. B. and L. Kogan (2004). Pricing American options: A duality approach. *Operations Research* 52(2), 258 – 270.

- Heathcote, J., K. Storesletten, and G. L. Violante (2009). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics* 1(1), 319–354.
- Jones, S. R. G. and J. H. Stock (1987, March). Demand disturbances and aggregate fluctuations: The implications of near rationality. *Economic Journal* 97(385), 49–64.
- Judd, K. L., L. Maliar, and S. Maliar (2017). Lower bounds on approximation errors to numerical solutions of dynamic economic models. *Econometrica* 85(3), 991 – 1012.
- Khan, A. and J. K. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica* 76(2), 395–436.
- Krusell, P. and A. A. Smith (1996). Rules of thumb in macroeconomic equilibrium: A quantitative analysis. *Journal of Economic Dynamics and Control* 20, 527–558.
- Krusell, P. and A. A. Smith (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106(5), 867 – 896.
- Kubler, F. and K. Schmedders (2003). Stationary equilibria in asset-pricing models with incomplete markets and collateral. *Econometrica* 71(6), 1767–1793.
- Kubler, F. and K. Schmedders (2005). Approximate versus exact equilibria in dynamic economies. *Econometrica* 73(4), 1205–1235.
- Maliar, L., S. Maliar, and F. Valli (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell-smith algorithm. *Journal of Economic Dynamics and Control* 34(1), 42–49.
- Mertens, T. M. (2011). Excessively volatile stock markets: Equilibrium computation and policy analysis. *Working Paper*.
- Mertens, T. M. and K. L. Judd (2012). Equilibrium existence and approximation for incomplete market models with substantial heterogeneity. *Working Paper*.
- Naish, H. F. (1993, January). The near optimality of adaptive expectations. *Journal of Economic Behavior & Organization* 20(1), 3–22.
- Pliska, S. R. (1982). Duality theory for some stochastic control models. *Lecture Notes in Control and Information Sciences* 43(2), 329 – 337.
- Rockafellar, R. T. and R. J. B. Wets (1976). Nonanticipativity and l^1 martingales in stochastic optimization problems. *Mathematical Programming Studies* 6, 170 – 187.
- Rogers, L. C. G. (2002). Monte Carlo valuation of American options. *Mathematical Finance* 12(3), 271 – 286.
- Rogers, L. C. G. (2007). Pathwise stochastic optimal control. *SIAM J. Control Optim.* 46(3), 1116–1132.

- Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. Cooley (Ed.), *Frontiers of Business Cycle Research*. Princeton, NJ: Princeton University Press.
- Santos, M. S. (2000). Accuracy of numerical solutions using the euler equation residuals. *Econometrica* 68(6), 1377–1402.
- Simon, H. A. (1978). Rationality as process and as product of thought. *The American Economic Review* 68, 1 – 16.
- Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control* 34(1), 36 – 41.

Appendix

A Proofs related to Section 2

A.1 Optimal policies

The investor's problem (1) implies the Bellman equation:

$$V(w_t, X_t) = \sup_{c_t > 0, \phi_t} \log c_t + \beta \mathbb{E}_t V(w_{t+1}, X_{t+1}), \quad (\text{A.1})$$

where $V(w_t, X_t)$ is the investor's value function, w_t is the time- t wealth of the investor, and the investor's wealth evolves according to equation (2). The value function satisfies

$$V(w_t, X_t) = \frac{1}{1 - \beta} \log w_t + v(X_t). \quad (\text{A.2})$$

Equations (A.1) and (A.2), together with the budget constraint (2), imply

$$V(w_t, X_t) = \sup_{c_t > 0, \phi_t} \log c_t + \beta \mathbb{E}_t \left(\frac{1}{1 - \beta} \log (w_t - c_t) + \log (\phi_t R_{t+1}^S + (1 - \phi_t) R^B) + v(X_{t+1}) \right). \quad (\text{A.3})$$

The first order conditions for optimal consumption (3) and optimal portfolio choice (4) follow from (A.3).

A.2 Derivation of expected utility loss equation (13)

The optimal and suboptimal consumption policies in Section 2.1 both have the form

$$c_t = (1 - \beta)w_t. \quad (\text{A.4})$$

That is $c_t = c_t^* = (1 - \beta)w_t$ for the optimal policy and $c_t = \hat{c}_t = (1 - \beta)w_t$ for the suboptimal policy. This implies that the time- t consumption c_t is given by

$$c_t = (1 - \beta)w_0 \prod_{s=0}^{t-1} \frac{w_{s+1}}{w_s}. \quad (\text{A.5})$$

In the setting of Section 2.1, since stock returns are independent and identically distributed, both the optimal and the suboptimal portfolio are time-invariant. Let us denote the share of the investor's wealth in the stock for both the optimal and the suboptimal portfolio choices by ϕ (i.e., $\phi = \phi^*$ for the optimal portfolio and $\phi = \hat{\phi}$ for the suboptimal portfolio). Combining the budget constraint (2) with the consumption policy (A.4), the return on wealth under both the optimal and the suboptimal policies is given by

$$\frac{w_{s+1}}{w_s} = \beta (\phi R_{s+1}^S + (1 - \phi) R^B). \quad (\text{A.6})$$

Combining equations (A.5) and (A.6), taking logs we get

$$\log c_t = \log(1 - \beta) + \log w_0 + t \log \beta + \sum_{s=0}^{t-1} \log (\phi R_{s+1}^S + (1 - \phi) R^B). \quad (\text{A.7})$$

Equation (A.7) implies that

$$\sum_{t=0}^{\infty} \beta^t \log c_t = \frac{\log(1-\beta)}{1-\beta} + \frac{\log w_0}{1-\beta} + \frac{\beta \log \beta}{(1-\beta)^2} + \sum_{t=0}^{\infty} \beta^t \sum_{s=0}^{t-1} \log(\phi R_{s+1}^S + (1-\phi)R^B) \quad (\text{A.8})$$

where we used the identities $\sum_{t=0}^{\infty} \beta^t = 1/(1-\beta)$ and $\sum_{t=0}^{\infty} t\beta^t = \beta/(1-\beta)^2$. Finally, taking the time-0 expectation of both sides of equation (A.8), and using the definition of the certainty equivalent of the one-period return on wealth equation (5), we get

$$\begin{aligned} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \log c_t \right] &= \frac{\log(1-\beta)}{1-\beta} + \frac{\log w_0}{1-\beta} + \frac{\beta \log \beta}{(1-\beta)^2} + \sum_{t=0}^{\infty} \beta^t t B(\phi), \\ &= \frac{\log(1-\beta)}{1-\beta} + \frac{\log w_0}{1-\beta} + \frac{\beta \log \beta}{(1-\beta)^2} + \frac{\beta}{(1-\beta)^2} B(\phi), \end{aligned} \quad (\text{A.9})$$

where we used the identity $\sum_{t=0}^{\infty} \beta^t t = \beta/(1-\beta)^2$ in going from the first to the second equality.

Equation (A.9) implies that the expected utility from adopting the optimal consumption and portfolio policy is

$$U(\phi^*, (c_t^*)_{t \geq 0}; w_0, X_0) = \frac{\log w_0}{1-\beta} + \frac{\log(1-\beta)}{1-\beta} + \frac{\beta \log \beta}{(1-\beta)^2} + B(\phi^*) \frac{\beta}{(1-\beta)^2}, \quad (\text{A.10})$$

where ϕ^* is given by equation (8). Similarly, the expected utility under the suboptimal policy is

$$U(\hat{\phi}, (\hat{c}_t)_{t \geq 0}; \hat{W}_0, X_0) = \frac{\log w_0}{1-\beta} + \frac{\log(1-\beta)}{1-\beta} + \frac{\beta \log \beta}{(1-\beta)^2} + B(\hat{\phi}) \frac{\beta}{(1-\beta)^2}, \quad (\text{A.11})$$

where $\hat{\phi}$ is given by equation (12). Note that the true distribution of returns is used in computing $B(\hat{\phi})$ in equation (A.11). Using the definition of the fractional certainty equivalent welfare loss equation (11), we get that the fractional certainty equivalent loss from adopting the suboptimal policy is

$$\eta = \exp \left(\frac{\beta}{1-\beta} \left(B(\phi^*) - B(\hat{\phi}) \right) \right) - 1.$$

A.3 Duality gap for the ideal penalty

Proposition 1 (Ideal penalty) *Let $V_t(w_t)$ and $V_0^{\mathcal{R}}(w_t^{\mathcal{R}})$ be the time- t value functions of the original problem (14) and the relaxed problem (18), respectively. Then,*

$$V_0(w_0) = V_0^{\mathcal{R}}(w_0), \quad (\text{A.12})$$

if the penalty process λ_t in equation (18) is chosen to be equal to λ_t^* , where

$$\lambda_t^*(C^{\mathcal{R}}, y^T) = \beta (V_{t+1}((w_t^{\mathcal{R}} - c_t^{\mathcal{R}})R + y_{t+1}) - \mathbb{E}[V_{t+1}((w_t^{\mathcal{R}} - c_t^{\mathcal{R}})R + \tilde{y}_{t+1})|y^t]). \quad (\text{A.13})$$

Proof. We prove Proposition 1 by adapting the proof in Brown et al. (2010) in the context of the problem in Section 3. As in that section, we assume complete information relaxation.

The Bellman equation for the original problem is

$$V_t(w_t) = \sup_{c_t: 0 < c_t \leq w_t} \frac{c_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E} [V_{t+1}(w_{t+1})|y^t], \quad (\text{A.14})$$

where w_{t+1} is given by equation (15). The Bellman equation for the relaxed problem is

$$V_t^{\mathcal{R}}(w_t^{\mathcal{R}}) = \sup_{c_t^{\mathcal{R}}: 0 < c_t^{\mathcal{R}} \leq w_t^{\mathcal{R}}} \frac{(c_t^{\mathcal{R}})^{1-\gamma}}{1-\gamma} - \lambda_t^*(C^{\mathcal{R}}, y^T) + \beta V_{t+1}^{\mathcal{R}}(w_{t+1}^{\mathcal{R}}) \quad (\text{A.15})$$

where $w_{t+1}^{\mathcal{R}}$ and $\lambda_t^*(C^{\mathcal{R}}, y^T)$ are given by equations (19) and (A.13), respectively.

The penalty (A.13) is feasible since it trivially satisfies the condition (17). The proof of equation (A.12) follows from using mathematical induction. To see this, note that the terminal value functions of the original and relaxed problems are both identical and equal to zero, that is, $V_{T+1}(w_{T+1}) = 0$ and $V_{T+1}^{\mathcal{R}}(w_{T+1}^{\mathcal{R}}) = 0$. For the inductive step, we show that if the value functions are the same function of wealth at time $t + 1$, that is, if

$$V_{t+1}(w_{t+1}) = V_{t+1}^{\mathcal{R}}(w_{t+1}), \quad (\text{A.16})$$

then the time- t value functions for the original and relaxed problems must also be the same function of wealth: $V_t(w_t) = V_t^{\mathcal{R}}(w_t)$. Indeed, this follows from noting that the right-hand side of equation (A.15) reduces to the right-hand side of (A.14), after using equation (A.16) and the expression for the ideal penalty (A.13) in (A.15). ■

Table 1: Transition matrix for the portfolio choice problem in Section 2.2.

A. Objective probabilities				B. Subjective probabilities			
	X_1	X_2	X_3		X_1	X_2	X_3
X_1	0.495	0.495	0.01	X_1	0.495	0.495	0.01
X_2	0.495	0.495	0.01	X_2	0.495	0.495	0.01
X_3	0.45	0.45	0.1	X_3	0.495	0.495	0.01

Table 2: Welfare loss and Euler equation errors for the portfolio choice problem in Section 2.2 for two different lengths of time series T .

A. $T = 6000$				B. $T = 100,000$			
	X_1	X_2	X_3		X_1	X_2	X_3
Fractional C.E. loss η	4.09	4.09	8.39	Fractional C.E. loss η	4.09	4.09	8.39
Lower 5%, Bond	4.76	5.31	4.59	Lower 5%, Bond	4.12	4.54	1.02
Upper 5%, Bond	5.14	4.8	6.49	Upper 5%, Bond	5.31	3.71	43.12
Lower 5%, Stock	4.76	5.31	4.59	Lower 5%, Stock	4.12	4.54	1.02
Upper 5%, Stock	5.14	4.8	6.49	Upper 5%, Stock	5.31	3.71	43.12

Table 3: Parameter values used in the consumption-saving problem of Section 3.

Parameter	Symbol	Value
Probability of high state	p	0.9
Labor income, high state	y_H	4
Labor income, low state	y_L	1
Risk-free interest rate	R	1.02
Agent's time-preference parameter	β	0.9
Agent's risk aversion	γ	5
Horizon	T	100

Table 4: Transition matrix for the model of Krusell and Smith in Section 4.

$z, \epsilon/z', \epsilon'$	$(z_b, 0)$	$(z_b, 1)$	$(z_g, 0)$	$(z_g, 1)$
$(z_b, 0)$	0.525000	0.350000	0.031250	0.093750
$(z_b, 1)$	0.038889	0.836111	0.002083	0.122917
$(z_g, 0)$	0.093750	0.031250	0.291667	0.583333
$(z_g, 1)$	0.009115	0.115885	0.024306	0.850694

Table 5: Baseline parameters of the Khan and Thomas (2008) model.

γ	β	δ	α	ν	ϕ	ρ_z	σ_{η_z}	ρ_ε	$\sigma_{\eta_\varepsilon}$	b	ξ
1.016	0.977	0.069	0.256	0.640	2.400	0.859	0.014	0.859	0.022	0.011	0.0083




Figure 1: Ideal penalty. The ideal penalty λ_1^* for the consumption-saving problem of Section 3, plotted as a function of time-1 consumption choice c_1 . We use the parameters shown in Table 3. The agent believes the probability of y_H is 0.89. Panel A corresponds to the time-1 wealth $w_1 = 4$, while Panel B corresponds to $w_1 = 5$. The solid line is the penalty function in the state $y_2 = y_H$, while the dash-dot line is the penalty in the state $y_2 = y_L$. The dash line shows the expected value of the penalty over the two possible realizations of y_2 . This expectation is identically equal to zero.




Figure 2: Welfare loss, upper bound. Panel A shows the relative ordering of the upper bound $V^{\mathcal{R}}$, the agent’s value function under the optimal policy V , and an unbiased estimate of the agent’s expected utility from adopting sub-optimal policies \hat{V} . While the true probability p of a high income shock, y_H is 0.9, the sub-optimal policies correspond to the optimal policy for realization of y_H that equals $\hat{p} = 0.89$. The solid line in Panel B shows the upper bound of certainty equivalent loss from adopting the sub-optimal policy, while the dot-dash line shows the actual certainty equivalent loss. The upper bound is computed using \hat{V} . Panel C shows the upper bound and the actual certainty equivalent loss with no penalty for foresight. We use the parameters shown in Table 3.




Figure 3: Certainty equivalent loss for different sub-optimal policies. Panels A, B, and C correspond to the certainty equivalent loss from adopting policies with varying degrees of sub-optimality. While the true probability p of a high income shock, y_H is 0.9, the sub-optimal policies in panels A, B, and C correspond to the optimal policy where the agent’s belief for realization of y_H equals $\hat{p} = 0.87$, $\hat{p} = 0.88$, and $\hat{p} = 0.89$, respectively. We use the parameters shown in Table 3.

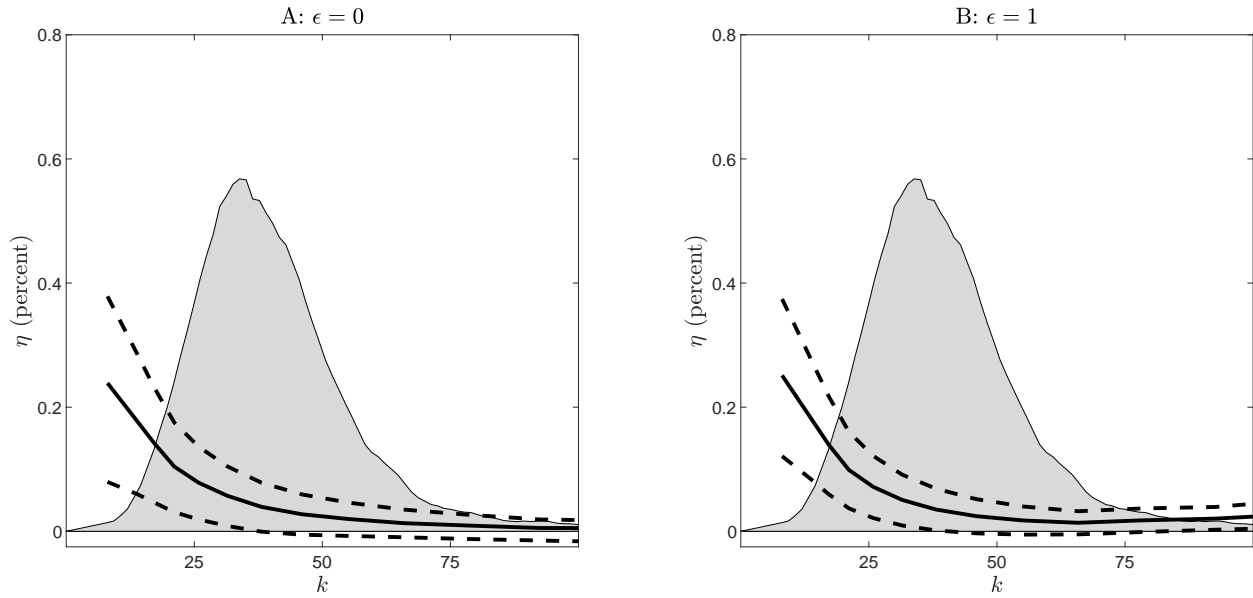


Figure 4: Upper bound on welfare loss: Model of Krusell and Smith, stochastic steady-state. Panels A and B show the upper bound of an agent’s welfare loss as a function of his capital stock, k , when he is unemployed and employed, respectively. The welfare loss is measured as a fractional certainty equivalent loss η , which is defined in equation (40). The aggregate state of the economy is low. The value function of the relaxed problem is estimated by averaging over 500 paths of aggregate shocks. The shaded area shows the cross-sectional distribution of capital and corresponds to the stochastic steady-state distribution. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Krusell and Smith (1998).

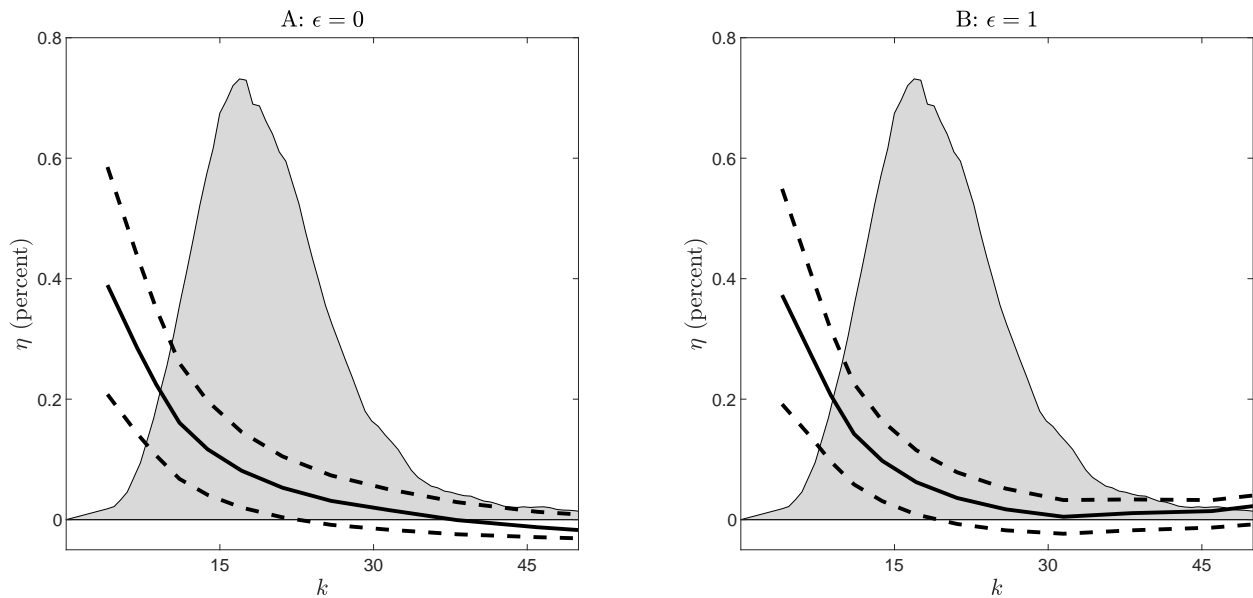


Figure 5: Upper bound on welfare loss: Model of Krusell and Smith, transitional dynamics following 50% destruction of capital stock of all agents. Panels A and B show the upper bound of an agent's welfare loss as a function of his capital stock, k , when he is unemployed and employed, respectively. The welfare loss is measured as a fractional certainty equivalent loss η , which is defined in equation (40). The aggregate state of the economy is low. The area under the shaded curve shows the cross-sectional distribution of capital after the permanent shock is realized. The value function of the relaxed problem is estimated by averaging over 500 paths of aggregate shocks. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Krusell and Smith (1998).

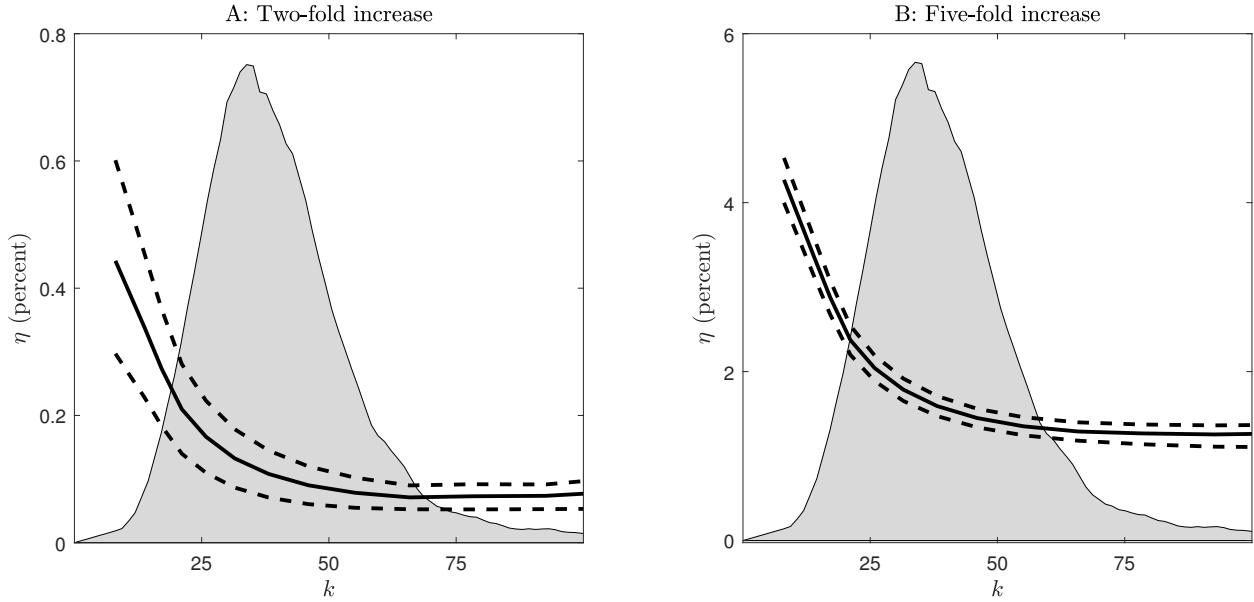


Figure 6: Upper bound on welfare loss: Model of Krusell and Smith, transitional dynamics following a permanent increase in aggregate volatility. Panels A and B show the upper bound of an agent's welfare loss as a function of his capital stock, k , following a two-fold and five-fold increase in aggregate volatility, respectively. In both cases the agent is initially employed and the aggregate state of the economy is low. The welfare loss is measured as a fractional certainty equivalent loss η , which is defined in equation (40). The area under the shaded curve shows the cross-sectional distribution of capital in the stochastic steady-state. The value function of the relaxed problem is estimated by averaging over 500 paths of aggregate shocks. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Krusell and Smith (1998).




Figure 7: Upper bound on welfare loss: Model of Khan and Thomas, stochastic steady-state. Panels A, B, and C show the upper bound of loss in firm value as a function of the firm’s capital stock, k , for three different idiosyncratic shocks, $\epsilon = 1$, $\epsilon = 8$, and $\epsilon = 15$, respectively. The loss is measured as a fractional certainty equivalent loss η , which is defined in equation (48). The aggregate state of the economy is low. The area under the shaded curve shows the cross-sectional distribution of capital in the stochastic steady-state, conditional on ϵ . The value function of the relaxed problem is estimated by averaging over 15000 paths of aggregate shocks. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Khan and Thomas (2008) and are also reported in Table 5.




Figure 8: Upper bound on welfare loss: Model of Khan and Thomas, transitional dynamics following 50% destruction of capital stock of all firms. Panels A, B, and C show the upper bound of loss in firm value as a function of the firm’s capital stock, k , for three different idiosyncratic shocks, $\epsilon = 1$, $\epsilon = 8$, and $\epsilon = 15$, respectively. The loss is measured as a fractional certainty equivalent loss η , which is defined in equation (48). The aggregate state of the economy is low. The area under the shaded curve shows the cross-sectional distribution of capital immediately after the economy-wide capital loss is realized. The value function of the relaxed problem is estimated by averaging over 15000 paths of aggregate shocks. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Khan and Thomas (2008) and are also reported in Table 5.




Figure 9: Upper bound on welfare loss: Model of Khan and Thomas, transitional dynamics following a permanent increase in aggregate volatility. Panels A and B show the upper bound of loss in firm value as a function of the firm's capital stock, k , for a two-fold and five-fold unanticipated increase in aggregate volatility, respectively. In both cases, the idiosyncratic shock $\epsilon = 8$. The loss is measured as a fractional certainty equivalent loss η , which is defined in equation (48). The aggregate state of the economy is low. The area under the shaded curve shows the cross-sectional distribution of capital in the stochastic steady-state, conditional on ϵ . The value function of the relaxed problem is estimated by averaging over 15000 paths of aggregate shocks. Dashed lines are 95% Monte Carlo confidence bounds. All parameters values are identical to those in Khan and Thomas (2008) and are also reported in Table 5.