

NBER WORKING PAPER SERIES

THE VALUE OF ARBITRAGE

Eduardo Dávila
Daniel D. Graves
Cecilia Parlatore

Working Paper 29744
<http://www.nber.org/papers/w29744>

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

We would like to thank Manuel Amador, Markus Brunnermeier, Ricardo Caballero, John Campbell, John Cochrane, Wenxin Du, Winston Dou, Darrel Duffie, Emmanuel Farhi, Xavier Gabaix, Joel Hasbrouck, Ben Hébert, Urban Jermann, Ralph Koijen, Anton Korinek, Karen Lewis, Guido Lorenzoni, Martin Oehmke, Stavros Panageas, Adriano Rampini, Alexi Savov, Jesse Schreger, Jeremy Stein, Harald Uhlig, Dimitri Vayanos, and Adrien Verdelhan for helpful comments and discussions. Jack Kelly, Jalen Li, and Flint O'Neil provided excellent research assistance. This material is based upon work supported by the NSF Graduate Research Fellowship Program under Grant No. DGE-1752134. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Eduardo Dávila, Daniel D. Graves, and Cecilia Parlatore. 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.

The Value of Arbitrage
Eduardo Dávila, Daniel D. Graves, and Cecilia Parlato
NBER Working Paper No. 29744
February 2022
JEL No. D61,G12,G18

ABSTRACT

This paper studies the social value of closing price differentials in financial markets. We show that arbitrage gaps (price differentials between markets) exactly correspond to the marginal social value of executing an arbitrage trade. We further show that arbitrage gaps and measures of price impact are sufficient to compute the total social value from closing an arbitrage gap. Theoretically, we show that, for a given arbitrage gap, the total social value of arbitrage is higher in more liquid markets. We apply our framework to compute the welfare gains from closing arbitrage gaps in the context of covered interest parity violations and several dual-listed companies. The estimates of the value of closing arbitrage gaps vary substantially across applications.

Eduardo Dávila
Yale University
Department of Economics
28 Hillhouse Avenue, Room 304
New Haven, CT 06520-8268
and NBER
eduardo.davila@yale.edu

Cecilia Parlato
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012
and NBER
cps272@stern.nyu.edu

Daniel D. Graves
Yale University
Department of Economics
New Haven, CT-C 06511
d.graves@yale.edu

1 Introduction

Arbitrage, the practice of making a sure profit off a price difference between two or more assets, is one of the bedrocks of modern finance. In particular, the absence of arbitrage opportunities provides a unifying principle to derive tight restrictions among asset prices. Conventional arbitrage logic relies on the ability to trade frictionlessly. However, different frictions generate arbitrage gaps in practice, with a growing body of evidence showing that deviations from the law of one price are widespread. An often-heard critique of the existing empirical literature that identifies these deviations is that it is hard to know whether a given arbitrage gap is associated with large or small welfare costs. In this paper, we tackle this issue by providing a framework to measure the welfare gains associated with closing an arbitrage gap; that is, we study the value of arbitrage.

We initially derive our results within a stylized model of trading in segmented financial markets. We consider an environment without uncertainty in which two sets of investors trade identical risk-free assets in two segmented markets: type A investors trade a risk-free asset in market A , while type B investors trade an identical risk-free asset in market B . If both markets were fully integrated, the prices of the risk-free asset in both markets would be equal in equilibrium, and no arbitrage opportunities would exist. Instead, we introduce an arbitrageur sector that can conduct an arbitrage trade — buying in the underpriced market and selling in the overpriced one — and explore how individual and social welfare vary as a function of the scale of such a trade.¹ This approach allows us to avoid taking a stance on the exact frictions that prevent the arbitrageur sector from fully equalizing asset prices across markets. Conceptually, our approach can be interpreted as a smooth way of moving from an autarky equilibrium to a fully integrated equilibrium.

Our first main result shows that the marginal social welfare gain of an arbitrage trade, i.e., the marginal value of arbitrage, is given by the arbitrage gap, defined as the price differential between markets. That is, we show that arbitrage gaps are sufficient statistics for the marginal social value of arbitrage. While one would expect the marginal value of arbitrage to increase with the scale of the arbitrage gap, our analysis shows that the arbitrage gap *exactly* corresponds to the marginal social value of arbitrage. Intuitively, this price differential fully captures the difference in willingness to pay between investors in different markets.

While the arbitrage gap measures the marginal social value of arbitrage, an arbitrage trade creates distributive pecuniary externalities that cancel out in the aggregate. Accounting for these externalities, we can also characterize the marginal individual value of arbitrage and show that i) the welfare of both type A and type B investors increases with the scale of the arbitrage trade and ii) the welfare of arbitrageurs — given by the arbitrage profits — initially increases with the scale of the arbitrage trade, reaches a maximum, and then decreases, becoming zero when the arbitrage gap is closed.

Our second main result shows that the total social value of arbitrage can be recovered by measuring the arbitrage gap for different levels of the arbitrage trade m . While prices are sufficient

¹In the Appendix, we explicitly model multiple frictions that can endogenously generate a positive arbitrage gap in equilibrium and show that increasing the scale of the arbitrage trade has equivalent implications to relaxing a particular friction.

to compute the marginal gain, it is necessary to use quantity measures — in particular, measures of price impact — to understand the total value of arbitrage. From a practical perspective, price differentials/arbitrage gaps along with measures of price impact become sufficient statistics for the value of arbitrage. The main upshot of our approach is that, conditional on measuring arbitrage gaps and price impact, our results are valid in a large class of environments.

Our third main result shows that the value of arbitrage depends on the degree of market liquidity. For a given arbitrage gap, the total social value of arbitrage is higher in more liquid markets, in which price impact is lower. Intuitively, observing an arbitrage gap that is very easy to close because the market is illiquid implies that the welfare gains from closing such a gap are likely to be small. Alternatively, finding an arbitrage gap in a liquid market, in which prices are not very sensitive to the quantities traded, implies that the welfare gains from closing such a gap are potentially large.

Before presenting our empirical results, we extend the baseline model to a general dynamic stochastic environment that features many agents and assets, and that allows for rich patterns of asset segmentation. This extension allows us to show how the insights from the baseline model extend to complex real-world scenarios, which is particularly relevant to correctly interpret the welfare estimates that we present in the empirical applications.

Our first empirical application examines violations of covered interest parity (CIP). Using high-frequency data from the Chicago Mercantile Exchange (CME), we compute price impact estimates for the FX futures market. Combining these estimates with the cross-currency bases from Bloomberg, following [Du, Tepper and Verdelhan \(2018\)](#), we provide estimates of the welfare gains associated with eliminating violations of covered interest parity, as well as the size of the required gap-closing trades.² We find that the welfare gains from closing CIP arbitrage gaps, even during the March 2020 start of the COVID-19 crisis in the US, never exceed \$1.2B for the five studied currency pairs, and never exceed \$300M outside of the Yen-Dollar basis. As we explain, in relation to the size of volume in FX markets, these welfare gains can be considered small. However, when normalized by combined daily GDP, the welfare gains become non-negligible, of up to 0.3% of daily GDP in the EUR/USD case. We find that the largest estimated gap-closing arbitrage trades do not exceed \$1.22T and never exceed \$455B outside of the Yen-Dollar basis. We document that periods of market distress, in which CIP violations are at their peak (e.g., in March 2020) are typically associated with illiquidity (large price impact), which underlies our estimates of welfare gains and gap-closing trades. The fact that the CIP deviations, while significant, are never larger than 64bps (quarterly) throughout our sample also contributes to our quantitative findings.

Our second empirical application provides estimates of the welfare gains associated with closing arbitrage gaps in the case of dual-listed companies (also referred to as “Siamese twin stocks”).

²As we explain in Remark 8, the measures of welfare gains that we report in our empirical applications should be interpreted as the gains from closing the arbitrage gap at a given point in time. If closing the arbitrage gap at that point in time closes all future gaps, all present and future direct welfare gains are captured by the number that we report. If closing the arbitrage gap at a given point in time does not narrow future arbitrage gaps, and it would be necessary to implement an arbitrage trade every day, the estimates that we report correspond to the daily flow of direct welfare gains.

We compute the welfare gains from closing arbitrage gaps in three particular scenarios: i) Royal Dutch/Shell, the canonical dual-listed company that featured arbitrage opportunities for nearly a century; ii) Smithkline Beecham; and iii) Rio Tinto. In this application, we combine measures of the arbitrage gaps with the price impact estimates for global equities from [Frazzini, Israel and Moskowitz \(2018\)](#), who use a global database of a hedge fund’s \$1.7T in stock transactions spanning two decades. We find that the welfare gains from closing arbitrage gaps in the Royal Dutch/Shell case are substantial, peaking at approximately \$2B in 1996 when converted to USD. This finding is due to the extreme magnitude of the arbitrage gap for Royal Dutch/Shell, with a price deviation at the time of maximum welfare gains of over 20%, despite the fact that the company had a market cap of over \$100B at the time. Because of the significant size of the arbitrage gap, a trade size equal to around \$36B is required to close the arbitrage gap. The persistence of extreme deviations from parity over many years results in an extended period over which the welfare gains from closing arbitrage gaps in Royal Dutch/Shell are over \$1B.

Finally, we show that the value of closing arbitrage gaps is not large for all dual-listed arbitrages. In the case of Smithkline Beckham (US-based) and Beecham (UK-based), we find minuscule welfare gains due to the limited liquidity of the UK-traded shares. Similarly, for the case of Rio Tinto, which experiences the smallest mean absolute divergence of all the studied twin shares, we find minimal welfare gains from closing arbitrage gaps at most times. The combination of lower liquidity than Royal Dutch/Shell, as measured by dollar trade volume, and limited divergence from parity yields a smaller magnitude of welfare gains, as implied by our theory.

Related Literature The absence of arbitrage opportunities is considered by many to be the Fundamental Theorem of Asset Pricing, and every modern finance textbook, e.g., [Duffie \(2001\)](#), [Cochrane \(2005\)](#), and [Campbell \(2017\)](#) develops the implications of no-arbitrage pricing. However, following [Shleifer and Vishny \(1997\)](#), there is a growing literature that studies the frictions that impose limits to arbitrage. For example, limits to arbitrage are the result of collateral constraints in [Gromb and Vayanos \(2002\)](#) and [Liu and Longstaff \(2004\)](#); of funding constraints in [Xiong \(2001\)](#) and [Kondor \(2009\)](#); and of a non-competitive arbitrageur sector in [Oehmke \(2009\)](#), [Duffie and Strulovici \(2012\)](#), and [Fardeau \(2016\)](#). Most of this literature has a positive focus and studies how the behavior of arbitrageurs impacts price dynamics.

The work by [Gromb and Vayanos \(2002\)](#), which has a normative emphasis, is perhaps the most closely related to ours. They show that the competitive equilibrium in an environment in which financially constrained investors arbitrage between segmented markets is constrained inefficient. More recently, [Hébert \(2020\)](#) argues that arbitrage gaps can be part of the optimal regulation that tackles pecuniary or aggregate demand externalities — see [Dávila and Korinek \(2018\)](#) and [Farhi and Werning \(2016\)](#). In contrast to these papers, we do not seek to determine whether a given arbitrage gap implied by an equilibrium with financial frictions is efficient. Instead, in the spirit of [Lucas \(1987\)](#) and [Alvarez and Jermann \(2004\)](#), we seek to understand the welfare impact of a hypothetical experiment that eliminates the underlying frictions that prevent arbitrage gaps from closing. Methodologically, our paper is closest to the work of [Alvarez and Jermann \(2004\)](#),

who tackle the question of what the potential gains from reducing business cycles are through a hypothetical experiment in which business cycle fluctuations could be eliminated. To our knowledge, we provide the first framework to quantify the potential welfare gains from closing arbitrage gaps.³

Our first empirical application is motivated by the recent evidence documenting and rationalizing the systematic breakdown of covered interest parity (CIP). [Du, Tepper and Verdelhan \(2018\)](#) find persistent deviations from CIP for many currency pairs starting in the global financial crisis of 2007–2008 and show that such deviations cannot be attributed solely to transaction costs or credit risk differences between the arbitrage legs. The consistent presence of these arbitrage opportunities is connected with the implications of the Supplementary Leverage Ratio (SLR) Rule explored by [Duffie and Krishnamurthy \(2016\)](#): with post-crisis regulatory constraints on banks’ balance sheets, arbitrage opportunities arise as intermediation of large arbitrage trades becomes infeasible. This is directly explored in [Boyarchenko et al. \(2020\)](#), who show that CIP arbitrage offers limited returns in the recent regulatory regime and use hedge fund return data to explore the impact of regulation, particularly the SLR, on leverage-dependent hedge funds’ returns. Relatedly, [Amador et al. \(2020\)](#) develop a model in which CIP violations arise as a direct result of the zero lower bound constraints on monetary policy in conjunction with exchange-rate targeting policies, as illustrated by the Swiss National Bank throughout much of the past decade.

The application of our framework to dual-listed companies is motivated by the persistent violations of arbitrage relations in dual-listed companies documented by [Froot and Dabora \(1999\)](#) and [De Jong, Rosenthal and Van Dijk \(2009\)](#), among others. [Froot and Dabora \(1999\)](#) find that when a company is listed on two different exchanges with a fixed cash flow claim ratio across the two instruments, the violations of parity with that ratio can be extreme, sometimes exceeding 30%. This finding holds even for extraordinarily liquid and large equities like Royal Dutch/Shell, one of the largest oil majors. [De Jong, Rosenthal and Van Dijk \(2009\)](#) show that despite this fact, trading strategies that would take advantage of twin share divergences possess low Sharpe ratios and significant left tails.

Our measurement of price impact is based on the extensive theoretical and empirical literature that stresses the square root power law of price impact, which states that prices react to large signed orders with a change approximately proportional to the square root of the order size. On the theory side, [Gabaix et al. \(2006\)](#) use a model of large trader dynamics with first-order risk-averse liquidity providers to show how a Zipf’s law for institutional trader size results in a square root price impact function. Several empirical studies have supported this conclusion using different datasets and varied empirical approaches. In particular, [Gabaix et al. \(2003, 2006\)](#) find support for a square root power law using large-cap French stock data from the 1990s as well as US TAQ data, while [Frazzini, Israel and Moskowitz \(2018\)](#) use two decades of actual stock execution data from a large hedge fund and find the same power law behavior. Finally, our results are also related to the recent work by [Gabaix and Koijen \(2021\)](#), who study the role played by trading flows in shaping

³Conceptually, our approach is also related to the work of [Gourinchas and Jeanne \(2006\)](#), who measure the gains from international financial integration within a calibrated neoclassical model. They find that such gains are small for plausible calibrations.

asset prices. Section G of the Appendix connects our price impact estimates to theirs.

Outline Section 2 introduces the baseline model, while Section 3 presents the welfare implications of closing arbitrage gaps. Section 4 extends the results from the baseline model to a general dynamic stochastic environment. Sections 5 and 6 develop the empirical applications and Section 7 concludes. All proofs and derivations are in the Appendix. The Appendix also includes extensions of our baseline model and robustness results for our applications.

2 Baseline Model

In Sections 2 and 3, we characterize the value of arbitrage in a stylized model of trading in segmented financial markets. The simplicity of the model allows us to transparently illustrate how to measure the welfare gains associated with closing an arbitrage gap. In Section 4, we generalize our results to environments with uncertainty, multiple trading dates, multiple assets, and general patterns of trading segmentation.

2.1 Investors' and Arbitrageurs' Problems

Environment There are two dates $t \in \{0, 1\}$ and no uncertainty. There is a single consumption good (dollar) at each date, which serves as numeraire. There are two types of investors, each in unit measure, indexed by $i = \{A, B\}$. There are two identical risk-free assets traded in two segmented markets, with both assets and markets also indexed by $i = \{A, B\}$. Type A investors exclusively trade risk-free asset A in market A , while type B investors exclusively trade risk-free asset B in market B . Each risk-free asset has an identical payoff $d_1 > 0$ at date 1. We denote the price of risk-free asset i , traded in market i , by p^i .

First, we describe the problem of investors. Subsequently, we describe the problem of arbitrageurs.

Investors Investors in each market have time-separable utility, with a flow utility of consumption $u_i(\cdot)$, which satisfies standard regularity conditions, and a discount factor β_i . Type i investors have dollar endowments n_0^i and n_1^i , and initial asset holdings q_{-1}^i . Investors choose asset holdings q_0^i at date 0, taking prices as given. Note that we allow different investor types to have different preferences, endowments, and initial asset holdings.

Hence, the demand of type i investors for asset i is given by the solution to

$$\max_{q_0^i} u_i(c_0^i) + \beta_i u_i(c_1^i),$$

subject to the following budget constraints:

$$\begin{aligned} c_0^i &= n_0^i - p^i \Delta q_0^i \\ c_1^i &= n_1^i + d_1 q_0^i, \end{aligned}$$

where $\Delta q_0^i = q_0^i - q_{-1}^i$ and where c_0^i and c_1^i denote the consumption of type i investors at dates 0 and 1, respectively. Therefore, type i investors optimally choose asset holdings according to

$$p^i = \frac{\beta_i u'_i(c_1^i)}{u'_i(c_0^i)} d_1.$$

Arbitrageurs Arbitrageurs (indexed by α) are the only agents who can simultaneously trade in both market A and market B . For simplicity, we assume that arbitrageurs have no initial endowments of dollars or assets and that their flow utility is linear.⁴ The budget constraints of arbitrageurs at dates 0 and 1, which parallel those of investors, are given by

$$c_0^\alpha = -p^A q_0^{\alpha A} - p^B q_0^{\alpha B} \quad (1)$$

$$c_1^\alpha = d_1 q_0^{\alpha A} + d_1 q_0^{\alpha B}, \quad (2)$$

where c_0^α and c_1^α denote the consumption of arbitrageurs at dates 0 and 1, respectively, and $q_0^{\alpha A}$ and $q_0^{\alpha B}$ denote the asset holdings of arbitrageurs in markets A and B , respectively.

In equilibrium, arbitrageurs implement an arbitrage trading strategy with zero cash flow at date 1 and consume the date 0 revenue generated by such a strategy.⁵ Therefore, as will become clear shortly, it is helpful to reformulate arbitrageurs' budget constraints in terms of the *size/scale* of the arbitrage trade, denoted by m , and the *direction* of such a trade, denoted by x_0^α . Using these notions, we can rewrite the budget constraints of arbitrageurs — introduced in Equations (1) and (2) — as follows:

$$c_0^\alpha = -\left(p^A + p^B x_0^\alpha\right) m \quad (3)$$

$$c_1^\alpha = d_1 (1 + x_0^\alpha) m, \quad (4)$$

where $m = q_0^{\alpha A}$ and $x_0^\alpha = \frac{q_0^{\alpha B}}{q_0^{\alpha A}}$. Equation (4), when combined with the zero cash flow restriction at

⁴Given our assumptions and the notion of equilibrium introduced below, it is irrelevant whether arbitrageurs in our model represent a single investor or many. Moreover, every formal result remains valid when arbitrageurs have concave utility and positive dollar endowments, n_0^α and n_1^α . In that case, the utility of arbitrageurs — introduced in Equation (6) below — would take the form

$$V^\alpha(m, p^A, p^B) = u_\alpha(n_0^\alpha + (p^B - p^A)m) + \beta_\alpha u_\alpha(n_1^\alpha),$$

and their date 0 marginal utility of consumption, λ_0^α , which is necessary to make welfare assessments in money-metric terms, would be $\lambda_0^\alpha = u'_\alpha(n_0^\alpha + (p^B - p^A)m)$.

⁵Some authors refer to trading strategies with a positive initial cash flow and nonnegative payoff under all future scenarios as type A arbitrages. In contrast, trading strategies with zero initial cost, nonnegative payoffs under all future scenarios, but a strictly positive expected payoff are often referred to as type B arbitrages. For simplicity, we focus on type A arbitrages throughout the paper. Our results remain valid if we considered type B arbitrages.

date 1, implies that

$$c_1^\alpha = 0 \quad \Rightarrow \quad x_0^\alpha = -1 \quad \Rightarrow \quad m = q_0^{\alpha A} = -q_0^{\alpha B}. \quad (5)$$

Equation (5) shows that an arbitrage strategy in this environment involves a one-for-one long/short strategy combining assets A and B . Consequently, the restriction that $c_1^\alpha = 0$ pins down the direction — but not the scale — of the arbitrage trade. Equation (5) also implies that whenever $m > 0$, arbitrageurs are buyers in market A and sellers in market B , and vice versa when $m < 0$.

Therefore, we are able to express the utility of arbitrageurs — which we denote by $V^\alpha(m, p^A, p^B)$ and is simply given by their date 0 consumption c_0^α — as a function of the scale of the arbitrage trade m as follows:

$$V^\alpha(m, p^A, p^B) = (p^B - p^A) m. \quad (6)$$

As we describe next, instead of making assumptions on the behavior of arbitrageurs (e.g., whether they are competitive or strategic, or whether they face different types of financing constraints or trading costs, etc.), we make the scale of the arbitrage trade m a primitive of the model and focus on the impact of how varying m affects social welfare. Note that we refer to $V^\alpha(m, p^A, p^B)$ as the utility of arbitrageurs, instead of their indirect utility or value function since m is not optimally chosen by arbitrageurs.

2.2 Equilibrium

Arbitrage equilibrium We are now ready to define an arbitrage equilibrium, which is a particular notion of competitive equilibrium. In an arbitrage equilibrium, equilibrium prices and allocations in both markets A and B are a function of the scale of the arbitrage trade m , which we take as a primitive.

Definition. (*Arbitrage equilibrium*) An arbitrage equilibrium, parameterized by the size/scale of the arbitrage trade $m \equiv q_0^{\alpha A}$, is defined as a set of consumption allocations, asset holdings, and asset prices $p^A(m)$ and $p^B(m)$ such that *i*) investors maximize utility subject to their budget constraints, *ii*) arbitrageurs follow an arbitrage trading strategy with zero cash flow at date 1 and a positive cash flow at date 0, and *iii*) asset markets A and B clear, that is,

$$\Delta q_0^A + q_0^{\alpha A} = 0 \quad \text{and} \quad \Delta q_0^B + q_0^{\alpha B} = 0.$$

Assumptions Going forward, we proceed under the following two assumptions. First, without loss of generality, we assume that the equilibrium price in market B is higher than in market A when there is no arbitrage activity, i.e., $m = 0$. Formally,

$$p^B(0) - p^A(0) > 0. \quad (7)$$

Second, we restrict our attention to scenarios in which the equilibrium price in market A (B) is an increasing (decreasing) function of m . Formally, we suppose that the equilibrium price functions

$p^A(m)$ and $p^B(m)$ satisfy

$$\frac{dp^A(m)}{dm} > 0 \quad \text{and} \quad \frac{dp^B(m)}{dm} < 0. \quad (8)$$

The second assumption makes the model well-behaved by avoiding scenarios in which prices in an asset market fall (increase) when there is higher (lower) demand for the asset. Our empirical results in Sections 5 and 6 are consistent with this assumption. It is straightforward to provide conditions on primitives that guarantee that this assumption is satisfied.

Arbitrage gap and gap-closing trade An important object for our purposes is the price differential between the two identical assets when written as a function of the scale of the arbitrage trade m . We refer to this price differential, which we denote by $\mathcal{G}_{BA}(m)$, as the *arbitrage gap*, or arbitrage basis. Formally, $\mathcal{G}_{BA}(m)$ is given by

$$\mathcal{G}_{BA}(m) \equiv p^B(m) - p^A(m). \quad (\text{arbitrage gap/basis}) \quad (9)$$

In terms of the newly defined arbitrage gap, $\mathcal{G}_{BA}(m)$, our second assumption implies that the arbitrage gap narrows as the size of the arbitrage trade m increases, whenever $m < m^*$:

$$\mathcal{G}'_{BA}(m) = \frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm} < 0, \text{ for all } m < m^*,$$

where we denote by m^* the scale of the arbitrage trade that closes the gap, i.e., $p^A(m^*) = p^B(m^*)$. We refer to this level of trade as the *gap-closing* arbitrage trade. Going forward, to simplify the exposition, we always suppose that m lies in $[0, m^*]$.

Figure 1 illustrates the behavior of the allocations and prices of an arbitrage equilibrium as a function of m . The left panel shows the equilibrium asset holdings of type A and type B investors as a function of the scale of the arbitrage trade m . The right panel shows the equilibrium prices in markets A and B as a function of the scale of the arbitrage trade m , illustrating how the arbitrage gap converges to 0 as m approaches the gap-closing trade m^* .

2.3 Remarks on Equilibrium

Before analyzing the welfare implications of our model, we conclude the characterization of the equilibrium with two remarks.

Remark 1. (Unspecified limits to arbitrage) A fully unconstrained perfectly competitive sector would compete away any arbitrage profits until $p^B(m^*) = p^A(m^*)$.⁶ Therefore, when indexing the equilibrium by the scale of the arbitrage trade m , we implicitly assume that there is some

⁶As shown in Section D of the Appendix, if unconstrained competitive arbitrageurs freely chose the scale of the arbitrage trade m , the solution to the arbitrageurs' problem would take the form:

$$m = \begin{cases} -\infty, & \text{if } p^A > p^B \\ \infty, & \text{if } p^A < p^B, \end{cases}$$

with $p^A = p^B$ and $m = m^*$ being the only outcome compatible with the existence of a competitive equilibrium.

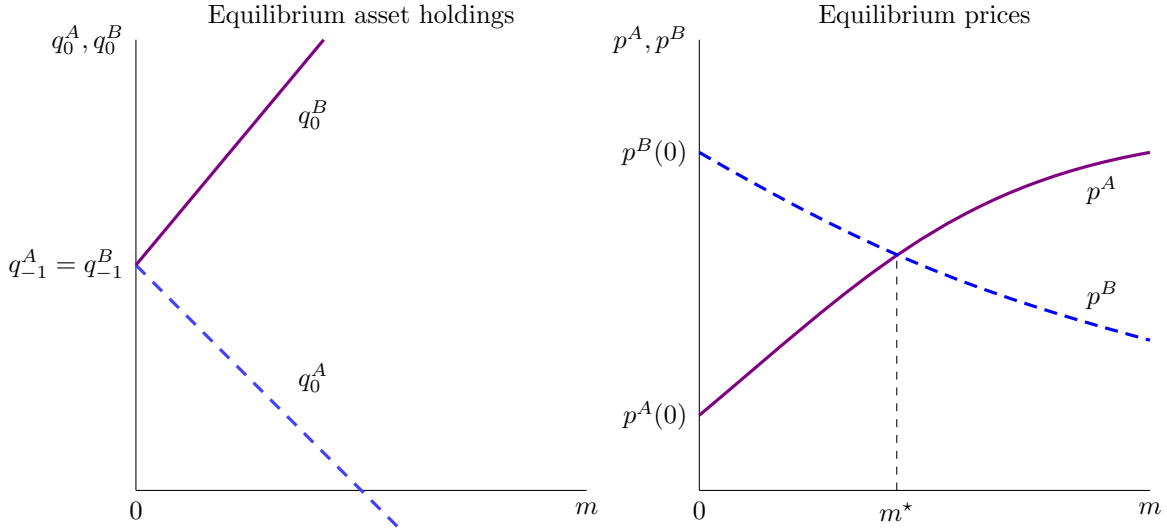


Figure 1: Equilibrium asset holdings and prices

Note: The left panel of Figure 1 shows the equilibrium asset holdings of type A and type B investors as a function of the size/scale of the arbitrage trade m . For simplicity, this illustration assumes that $q_{-1}^A = q_{-1}^B$. The right panel of Figure 1 shows the equilibrium price in market A and market B as a function of the size/scale of the arbitrage trade m . The right panel also illustrates the gap-closing arbitrage trade, m^* .

friction in the background that makes arbitrageurs unwilling or unable to reach the unconstrained competitive benchmark. Since our goal is to reach conclusions that are valid regardless of the specific limits to arbitrage that are relevant in a particular scenario, we purposefully avoid modeling frictions explicitly. However, there is a direct mapping between varying the scale of arbitrage m and loosening/tightening the friction that prevents complete price equalization. In Section C of the Online Appendix, we formally explain — in the context of multiple microfounded models — how trading frictions or departures from competitive behavior can endogenously generate positive arbitrage gaps of the form considered here.

Remark 2. (Autarky equilibrium and integrated equilibrium as special cases of arbitrage equilibrium)

By varying the scale of the arbitrage trade m we can capture two related type of equilibria: an autarky equilibrium, in which $m = 0$, and an integrated equilibrium, in which both assets in markets A and B trade under a single integrated market clearing constraint, $p^A = p^B$ by assumption, and arbitrageurs are redundant. While in general it is hard to compare the outcomes of an autarky equilibrium and an integrated equilibrium, the approach that we develop in this paper can be interpreted as a smooth way to connect both equilibria.

3 The Value of Arbitrage

Since our main objective is to understand the welfare properties of an arbitrage equilibrium, we need to characterize the welfare of investors and arbitrageurs. We denote the indirect utilities of type A and type B investors by $V^A(p^A(m))$ and $V^B(p^B(m))$, respectively. We denote the utility

of arbitrageurs in equilibrium by $V^\alpha (m, p^A(m), p^B(m))$, introduced in Equation (6).

3.1 Marginal Value of Arbitrage

First, we characterize the marginal value of arbitrage, that is, we show how individual and social welfare vary with the scale of the arbitrage trade m . We express all individual welfare assessments in money-metric terms, measured in date 0 dollars. That is, we normalize individual utility changes by date 0 marginal utility of consumption, denoted by λ_0^A , λ_0^B , and λ_0^α , respectively, which allows us to aggregate money-metric welfare changes across all agents.⁷

Formally, the (money-metric) change in social welfare induced by increasing the scale of the arbitrage trade m , which we denote by $\frac{dW(m)}{dm}$, takes the form

$$\frac{dW(m)}{dm} = \frac{dV^A(m)}{\lambda_0^A} + \frac{dV^B(m)}{\lambda_0^B} + \frac{dV^\alpha(m)}{\lambda_0^\alpha},$$

where $\frac{dW}{dm}$ and all its constituents are measured in date 0 dollars. Proposition 1 introduces the first main result of the paper.

Proposition 1. (Marginal value of arbitrage)

a) *(Marginal individual value of arbitrage)* The marginal individual value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade $m \in [0, m^*]$, measured in date 0 dollars, for type A investors, type B investors, and arbitrageurs, is respectively given by

$$\frac{dV^A(m)}{\lambda_0^A} = \frac{dp^A(m)}{dm} m > 0 \tag{10}$$

$$\frac{dV^B(m)}{\lambda_0^B} = \frac{dp^B(m)}{dm} (-m) > 0 \tag{11}$$

$$\frac{dV^\alpha(m)}{\lambda_0^\alpha} = \left(\frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm} \right) m + p^B(m) - p^A(m) \gtrless 0. \tag{12}$$

b) *(Marginal social value of arbitrage)* The marginal social value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade $m \in [0, m^*]$, aggregated and measured in date 0 dollars, is given by

$$\frac{dW(m)}{dm} = \underbrace{p^B(m) - p^A(m)}_{\mathcal{G}_{BA}(m)} > 0. \tag{13}$$

Proposition 1a) shows that increasing the scale of the arbitrage trade has two types of first-order welfare effects: *direct effects* and *pecuniary effects*, which affect the types of agents in this economy differently. The direct effects, which correspond to the arbitrage gap $p^B(m) - p^A(m)$, only affect the welfare of arbitrageurs directly and are zero for investors. Intuitively, increasing the scale of

⁷This approach can be interpreted as selecting, for all agents, a set of uniform generalized social marginal welfare weights of the form described in Saez and Stantcheva (2016). This approach is analogous to measuring welfare in terms of consumer surplus — see Dávila and Schaab (2021) for details.

the arbitrage trade by a unit increases the net profit/welfare gain of arbitrageurs at the margin by exactly $p^B(m) - p^A(m)$ dollars.

The pecuniary effects, which include the terms in which $\frac{dp^A(m)}{dm}$ and $\frac{dp^B(m)}{dm}$ appear, correspond to the distributive pecuniary externalities of increasing the scale of the arbitrage trade, using the terminology in [Dávila and Korinek \(2018\)](#).⁸ Intuitively, increasing the scale of the arbitrage trade by a unit changes prices in both markets, with price increases making buyers of an asset better off and sellers worse off, and vice versa. In particular, type A investors, who are net sellers of the type A asset, are better off when $\frac{dp^A(m)}{dm} > 0$. Similarly, type B investors, who are net buyers of the type B asset, are worse off when $\frac{dp^B(m)}{dm} > 0$. Arbitrageurs, on the other side of both trades, are worse off when $\frac{dp^A(m)}{dm} > 0$ and better off when $\frac{dp^B(m)}{dm} > 0$, as we further discuss below. Note that the strength of the pecuniary effects increases with the size of net trades, here given by m , and is zero when $m \rightarrow 0$.

Proposition 1b) leverages the fact that the pecuniary effects cancel out in the aggregate for any m , due to market clearing, to show that the direct effects are the single source of marginal social welfare gains. This result shows, in a general equilibrium environment, that the arbitrage gap corresponds to the marginal social value of arbitrage. Importantly, while one would expect the value of closing an arbitrage gap to be increasing in the size of the arbitrage gap, our analysis shows that the arbitrage gap *exactly* corresponds to the marginal social value of arbitrage.

While our paper is focused on characterizing and measuring social welfare, after aggregating welfare gains/losses across agents, Proposition 1a) shows that arbitrage trades have distributional welfare consequences. We summarize those in the following corollary.

Corollary 1. (Distributional consequences of arbitrage) *An increase in the scale of the arbitrage trade m always makes type A and type B investors better off. Starting from $m = 0$, arbitrageurs are initially better off as the scale of the arbitrage trade m increases, but there is a level of m such that further increases in m make arbitrageurs worse off. Hence, increasing the scale of arbitrage is Pareto improving for low values of m .*

The left panel of Figure 2 illustrates how the marginal individual value of arbitrage and the marginal social value of arbitrage vary with the size of the arbitrage trade. While the welfare of both types of investors is increasing in m , the welfare of arbitrageurs is hump-shaped. At all times, the pecuniary effects of arbitrage improve the welfare of investors in markets A and B . This occurs because increasing the scale of the arbitrage trade increases the price in market A and type A investors are net sellers of that asset (to arbitrageurs), so they profit from selling at higher prices. Symmetrically, increasing the scale of the arbitrage trade reduces the price in market B and type B investors are net buyers of that asset (from arbitrageurs), so they profit from buying at lower prices. While an increase in m always nets arbitrageurs the arbitrage gap, which is positive whenever $m < m^*$, the pecuniary effects are always negative for arbitrageurs, for the opposite

⁸Distributive pecuniary externalities are well understood. [Gromb and Vayanos \(2002\)](#) are the first to identify this type of pecuniary externality in models of arbitrage. See also [Geanakoplos and Polemarchakis \(1986\)](#), [Lorenzoni \(2008\)](#), or [Dávila and Korinek \(2018\)](#), among others, for examples of this type of pecuniary externality in different contexts.

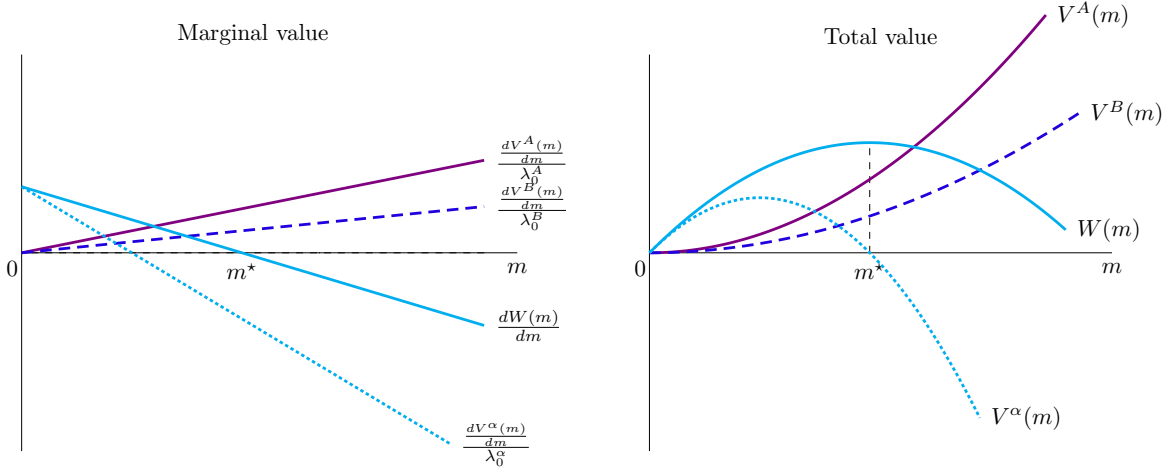


Figure 2: Marginal/total individual and social value of arbitrage

Note: The left panel of Figure 2 shows the money-metric marginal individual value of arbitrage for type A and B investors, $\frac{dV^A(m)}{\lambda_0^A \frac{dm}{dm}}$ and $\frac{dV^B(m)}{\lambda_0^B \frac{dm}{dm}}$, and for arbitrageurs, $\frac{dV^\alpha(m)}{\lambda_0^\alpha \frac{dm}{dm}}$, as a function of the scale of the arbitrage trade m , as characterized in Proposition 1. It also shows the marginal social value of arbitrage, $\frac{dW(m)}{dm}$, whose zero defines the gap-closing trade, m^* . The right panel of Figure 2 shows the indirect utility of type A and B investors and the utility of arbitrageurs — all expressed in money-metric terms after integrating the respective marginal individual values of arbitrage — as a function of the scale of the arbitrage trade m . It also shows the total social value of arbitrage, as characterized in Proposition 2. For simplicity, we assume that $V^A(0) = V^B(0) = W(0) = V^\alpha(0)$.

reasons that make investors better off: increasing m increases the price in market A , in which arbitrageurs buy, and lowers the price in market B , in which arbitrageurs sell. In other words, arbitrageurs move prices against them with each additional unit arbitrated.⁹

While the distributional consequences of closing arbitrage gaps are interesting and worth studying further, our paper focuses on social welfare, after aggregating welfare gains/losses across agents, for two reasons. First, in more complex environments in which a given investor trades in many asset markets — like the one considered in Section 4 — there are no clear predictions for the welfare of a given individual investor, since the pecuniary effects can potentially take different signs in different markets for any particular investor. Second, measuring individual gains/losses from arbitrage would require information that is not readily available, in particular detailed information on the net trades and asset holdings of every market participant.

3.2 Total Value of Arbitrage

Next, we characterize the total social value of arbitrage and study its properties. Proposition 2 introduces the second main result of the paper.

Proposition 2. (Total value of arbitrage)

⁹As shown in Section D of the Appendix, a monopolist arbitrageur would set m as to solve $\frac{dV^\alpha}{dm} = 0$. Interestingly, introducing a monopolist arbitrageur generates a Pareto improvement in the model considered here. As discussed above, competitive arbitrageurs would trade until $p^B = p^A$, in turn eliminating all arbitrage profits.

a) (Total social value of arbitrage) The total social value of arbitrage, that is, the total value associated with a change in the scale of the arbitrage trade from m_0 (any initial arbitrage trade) to m^* (the gap-closing arbitrage trade), aggregated across all agents and measured in date 0 dollars, is given by

$$W(m^*) - W(m_0) = \int_{m_0}^{m^*} W'(\tilde{m}) d\tilde{m} = \int_{m_0}^{m^*} \mathcal{G}_{BA}(\tilde{m}) d\tilde{m}. \quad (14)$$

b) (Sufficient statistics) It is sufficient to know i) the initial arbitrage gap, $\mathcal{G}_{BA}(m_0)$, and ii) measures of price impact in both markets A and B, that is, $\frac{dp^A(m)}{dm}$ and $\frac{dp^B(m)}{dm}$, to exactly compute the social value of arbitrage, since

$$\mathcal{G}_{BA}(\tilde{m}) = p^B(m_0) - p^A(m_0) + \int_{m_0}^{\tilde{m}} \left(\frac{dp^B(\hat{m})}{d\hat{m}} - \frac{dp^A(\hat{m})}{d\hat{m}} \right) d\hat{m}.$$

Proposition 2a) combines the Fundamental Theorem of Calculus with our characterization in Proposition 1b). Intuitively, if the marginal social value of arbitrage is given by $\mathcal{G}_{BA}(m)$, by summing over arbitrage gaps for different values of m , we can recover the total social value of arbitrage. Proposition 2b) shows that knowing the existing arbitrage gap and how this gap evolves with m is sufficient to compute the total social value of arbitrage. Since the existing arbitrage gap is easily observable/measurable, Proposition 2b) implies that measuring price impact is the main hurdle to compute the total social value of arbitrage in particular scenarios.

It is worth highlighting that the exact characterization of the total social value of arbitrage is expressed as a function of equilibrium objects (prices and arbitrage trades). In this sense, our characterization is valid regardless of the specific assumptions made on the primitives of the model. In other words, as long as it is possible to come up with measures of price gaps and price impact, it is possible to compute the social value of arbitrage without the need to fully specify the primitives of the economy. Thus, Proposition 2 provides the foundation of our empirical applications in Sections 5 and 6.¹⁰

Finally, note that it is possible to find a simple upper bound for the social value of arbitrage that depends exclusively on the initial arbitrage and the gap-closing arbitrage trade:

$$W(m^*) - W(m_0) \leq \mathcal{G}_{BA}(m_0)(m^* - m_0).$$

However, since finding the gap-closing arbitrage trade m^* implicitly involves making assumptions about price impact in both markets, there is little to gain by using this approximation instead of

¹⁰If price impact follows a power law $p^i(m) = p^i(m_0) + \alpha^i \text{sgn}(m - m_0) |m - m_0|^\beta$ for $i = \{A, B\}$ — see Section 5 for more details — the gap-closing trade and the total social value of arbitrage can be computed in closed form in terms of $p^B(m_0) - p^A(m_0)$, $\alpha^A + \alpha^B$ and β , as follows:

$$m^* - m_0 = \left(\frac{p^B(m_0) - p^A(m_0)}{\alpha^A + \alpha^B} \right)^{1/\beta}$$

$$W(m^*) - W(m_0) = \frac{\beta}{1 + \beta} (p^B(m_0) - p^A(m_0))^{\frac{1+\beta}{\beta}} \left(\frac{1}{\alpha^A + \alpha^B} \right)^{1/\beta}.$$




Figure 3: The value of arbitrage and market liquidity

Note: Figure 3 illustrates how the social value of closing an arbitrage gap varies with the degree of market liquidity by comparing two different economies that have an identical arbitrage gap when $m = m_0$ but for which price impact is different. The left panel of Figure 3 corresponds to an economy in which markets are illiquid, i.e., price impact is high, so a given trade moves prices substantially. In this case, the area with horizontal lines defines the total social value of arbitrage for that economy. The right panel of Figure 3 corresponds to an economy in which markets are liquid, i.e., price impact is low, so a given trade barely impacts prices. In this case, the area with vertical lines defines the total social value of arbitrage for that economy. It is evident that the area with vertical lines is larger than the area with horizontal lines, as characterized in Proposition 3. Note that, matching our empirical estimates, Figure 3 implies that price impact is concave.

the exact characterization in Equation (14).

3.3 The Value of Arbitrage and Market Liquidity

We now explore how the social value of arbitrage varies as a function of market liquidity. In particular, Proposition 3 shows that the value of arbitrage depends on the degree of market liquidity for a given arbitrage gap. Here we use a conventional notion of liquidity: we say that a market is liquid (illiquid) when price impact in that market is low (high).

Proposition 3. (Market liquidity and the value of arbitrage) *For a given arbitrage gap, $\mathcal{G}_{BA}(m) = p^B(m) - p^A(m)$, the social value of arbitrage is higher (lower) in more liquid (illiquid) markets, in which price impact is lower (higher).*

Proposition 3 concludes that finding arbitrage gaps in very liquid markets has the potential to be associated with large welfare losses. On the contrary, finding arbitrage gaps in illiquid markets is likely to yield small welfare losses. Intuitively, observing a large arbitrage gap that is very easy to close because price impact is large implies that welfare losses are small. Alternatively, finding an arbitrage gap in a market in which prices are not particularly sensitive to the quantities traded implies that the value of arbitrage is potentially large.

Figure 3 illustrates Proposition 3 by comparing two economies that have the same arbitrage gap when $m = m_0$, which implies that the marginal social value of arbitrage is initially the same, but for which the price impact of trades is different. In this context, following Proposition 2, the

total social value of arbitrage, $W(m^*) - W(m_0)$, corresponds to the triangle-like area between the equilibrium prices. Figure 3 shows how economies with higher price impact, i.e., associated with more illiquid markets and steeper pricing functions, have lower welfare costs from closing arbitrage gaps than economies with lower price impact, i.e., associated with more liquid markets and flatter pricing functions. In practical terms, Proposition 3 points out that arbitrage gaps in highly liquid markets have the potential to be associated with large welfare losses.

It is worth highlighting that Proposition 3 fixes a given arbitrage gap and then compares scenarios in which markets are more or less liquid. Even though this is a natural way to theoretically describe the impact of market liquidity on the value of arbitrage, one may conjecture that large arbitrage gaps are more likely to emerge in illiquid markets. In fact, in the context of the CIP application that we present in Section 5, we empirically show that the magnitude of arbitrage gaps is larger when price impact measures are also larger — see Figure 5. The welfare measures that we report in that section account for this correlation.

3.4 Remarks on Welfare Assessments

The following remarks further explain and qualify our approach to making welfare assessments.

Remark 3. (Constrained efficient benchmark/Absence of additional externalities) The economy studied in this section is constrained efficient whenever the arbitrage gap is zero. It is straightforward to show that if there were non-pecuniary externalities associated with the arbitrage trade, the marginal social value of arbitrage would have to account for how changing m affects those. For instance, if closing an arbitrage gap exacerbated a pollution externality somewhere in the economy, this would have to be accounted for as part of the marginal welfare effect of closing a gap. As shown in Section 4, pecuniary externalities in dynamic environments would also have to be accounted for when computing marginal welfare effects. Importantly, what calls for augmenting our characterization of welfare gains is not the existence of additional economic decisions, but whether these decisions are associated with externalities, i.e., are constrained inefficient or not. In any case, the direct marginal welfare effect of closing the arbitrage gap, in Equation (13), never vanishes. For this reason, our welfare calculations should be understood as capturing the direct welfare gains from closing arbitrage gaps.

Remark 4. (Interpretation of welfare gains as potential gains) It is worth highlighting that our characterization of the total social value of arbitrage does not include potential costs associated with implementing a gap-closing trade. For instance, our characterization does not include physical costs associated with implementing the gap-closing trade. Relatedly, in specific scenarios, one may argue that arbitrage gaps emerge as part of an optimal regulation. In that case, closing an arbitrage gap may be costly in terms of eliminating a desirable regulation. Therefore, our welfare calculations should be understood as capturing the potential gains from closing arbitrage gaps. In order to determine whether policies or technologies that eliminate arbitrage gaps are socially desirable, one would need to trade off the gains that we identify against potential costs.¹¹

¹¹In this sense, our approach parallels the literature that measures the welfare cost of business cycles (Lucas (1987));

Remark 5. (Invariance of welfare gains to underlying frictions) It is important to highlight how making the scale of the arbitrage trade m a primitive of the model impacts our welfare calculations. First, from a theoretical perspective, our characterization captures the welfare change associated with changing the allocation of m , regardless of why this change took place. However, relaxing the friction that limits arbitrage in a fully specified model may not only change the allocation of m , but also have a direct welfare impact either through its interaction with constrained inefficient choices or directly, as discussed in the previous two remarks. Therefore, our welfare calculations should be understood as capturing the welfare gains that accrue directly by varying the actual amount arbitrated. Second, from a measurement/empirical perspective, in our applications we estimate price impact measures that directly link prices to quantities traded. While different underlying frictions may have different implications for price impact, by measuring price impact directly, we avoid taking a stance on the exact underlying frictions. If one wanted to consider the welfare gains from relaxing a particular friction, it would be necessary to translate how relaxing that friction i) changes price impact directly and ii) affects the actual amount of arbitrage, to ultimately determine how equilibrium prices change.

Remark 6. (Relation to welfare gains from trade) At an abstract level, our approach is connected to the literature in international trade that studies the welfare gains from trade (Dixit and Norman (1980); Arkolakis, Costinot and Rodríguez-Clare (2012)). In fact, by relabeling consumption in different dates as consumption of different goods, our baseline model can be reinterpreted as an international trade model. However, that relabeling exercise is not valid more generally since the economic structure of arbitrage problems is very different from the structure of international trade environments, as shown in the next section and several of the extensions discussed in the Appendix. Interestingly, while in international trade information on quantities traded is readily available but prices are hard to gather, in financial markets prices are easy to obtain but not information on transactions.

Remark 7. (Relation to Harberger triangle) Figure 3 suggests that there exists a connection between our baseline model and a Marshallian supply/demand environment. In particular, our characterization of the total social value of arbitrage resembles a Harberger triangle (Harberger, 1964), while price impact measures can be mapped to supply/demand elasticities. Both our characterization and Harberger's are certainly based on aggregating marginal valuations. However, there are clear differences. For instance, supply is fixed in our baseline model, which features two excess demands (the pricing functions). Once again, more generally, the analogies cease when considering more complex environments, as shown in the next section and several of the extensions discussed in the Appendix.

Alvarez and Jermann (2004)). That body of work characterizes the potential gains from eliminating business cycles, which must be traded off against the costs of doing so.

4 General Model

In this section, we extend the results from the baseline model to a general dynamic stochastic environment that features many agents and assets, and that allows for rich patterns of asset segmentation. The main purpose of this section is to show how the insights from the baseline model extend to more realistic environments. This is particularly relevant to correctly interpret the welfare estimates that we present in the empirical applications in Sections 5 and 6. A reader more interested in the empirical implementation of our results can skip this section and jump to Sections 5 and 6, with the caveat that the interpretation of the empirical estimates of welfare gains partly relies on the insights from this section.¹²

4.1 Environment

We consider an environment with $T + 1$ dates indexed by $t = \{0, 1, \dots, T\}$ and a single consumption good (dollar), which serves as numeraire. Starting at date $t = 1$, a state $s_t \in S$ is realized, where $S = \{1, \dots, |S|\}$. The initial state at date 0 is s_0 .¹³ We denote the history of states up to date t (inclusive) by $s^t \equiv \{s_0, s_1, \dots, s_t\}$ and the probability of each history s^t by $\pi(s^t) \equiv \pi(s^t | s_0)$. The economy is populated by investors and arbitrageurs, whom we describe next.

Investors There is a finite number of investor types (investors, for short), indexed by $i \in I$, where $I = \{1, \dots, |I|\}$. There is a finite number of assets that investors can trade, indexed by $j \in J$, where $J = \{1, \dots, |J|\}$. Asset j has payoffs $d_t^j(s^t) \geq 0$, where $d_0^j(s^0) = 0, \forall j$, and trades at a price $p_t^j(s^t)$, where $p_T^j(s^T) = 0, \forall j$. At each date t , investor i chooses holdings of asset j , denoted by $q_t^{ij}(s^t)$. The initial holdings of asset j by investor i are given by $q_{-1}^{ij}(s^{-1})$.

Investor i has preferences of the form

$$u^i(c_0^i(s^0)) + \sum_{t=1}^T \sum_{s^t} \pi(s^t) (\beta_i)^t u^i(c_t^i(s^t)), \quad (15)$$

where $c_t^i(s^t)$ denotes the consumption of investor i at date t given a history s^t , β_i denotes investor i 's discount factor, and $u^i(\cdot)$ is an increasing and concave flow utility function.¹⁴ We denote the dollar endowment of investor i at time t given a history s^t by $n_t^i(s^t)$.

Consequently, investor i faces the following budget constraints:

$$c_t^i(s^t) = n_t^i(s^t) + \sum_j (d_t^j(s^t) + p_t^j(s^t)) q_{t-1}^{ij}(s^{t-1}) - p_t^j(s^t) q_t^{ij}(s^t), \quad \forall t, \forall s^t. \quad (16)$$

Even if the available assets do not span all possible contingencies (markets are incomplete) — a

¹²In Section D of the Online Appendix, we include several special cases of the general model studied here that may be helpful to build economic intuition.

¹³At times, to simplify the exposition, we drop the explicit dependence of date 0 variables on s_0 ; for instance, we write m_0 instead of $m_0(s_0)$.

¹⁴Our results easily extend to environments in which utility is not time-separable and in which $u^i(\cdot)$ is time- and/or state-dependent.

possibility that our environment allows for — if all investors can frictionlessly trade all assets at all times, it is well-known that no arbitrage opportunities exist — see e.g., [Duffie \(2001\)](#) or [Cochrane \(2005\)](#). For that reason, we assume that investors choose asset holdings subject to a potentially vector valued constraint of the form:

$$\Phi^i \left(\left\{ q_t^{ij}(s^t) \right\}_{t \in \{0, \dots, T-1\}} \right) \leq 0. \quad (17)$$

This constraint captures rich forms of market segmentation. For instance, if investor $i = 1$ cannot trade asset $j = 1$, then such an investor faces the following constraint: $q_t^{1,1}(s^t) = 0, \forall s^t, \forall t$. Consequently, the extreme segmentation assumed in the baseline model can be seen as a special case of the environment described here.¹⁵

Arbitrageurs As in the baseline model, we assume, for simplicity, that arbitrageurs have no initial endowments of dollars or assets, and that their flow utility is linear. Therefore, arbitrageurs face the following budget constraints, which parallel those of investors:

$$c_t^\alpha(s^t) = \sum_j \left(d_t^j(s^t) + p_t^j(s^t) \right) q_{t-1}^{\alpha j}(s^{t-1}) - \sum_j p_t^j(s^t) q_t^{\alpha j}(s^t), \quad \forall t, \forall s^t, \quad (18)$$

where $c_t^\alpha(s^t)$ denotes the consumption of arbitrageurs at date t given a history s^t , while $q_t^{\alpha j}(s^t)$ denotes the arbitrageurs' asset holdings of asset j at date t given a history s^t , where $q_{-1}^{\alpha j}(s_1) = 0, \forall j$.

To simplify the exposition, we assume i) that arbitrageurs can trade all assets, and ii) that the number of assets equals one more than the number of states, that is, $|J| = |S| + 1$.¹⁶ These assumptions simply reduce the set of cases to consider and imply that generically there is a unique solution to the problem of characterizing the directions of arbitrage trades at each date/state, as we show next.

4.2 Equilibrium

First, we define the notion of arbitrage equilibrium in this more general environment. Next, to simplify the exposition, we only present a minimal set of results that allows us to properly describe [Proposition 4](#), which is the main result of this section.

Definition. (*Arbitrage equilibrium*) An arbitrage equilibrium, parameterized by the size/scales of the arbitrage trades, $m_0 \equiv q_0^{\alpha 1}$ at date 0 and $m_t(s^t) \equiv q_t^{\alpha 1}(s^t), \forall t = \{1, \dots, T-1\}, \forall s^t$, is defined as a set of consumption allocations, asset holdings, and asset prices such that i) investors maximize utility in [Equation \(15\)](#), subject to their budget constraints in [Equation \(16\)](#), and trading constraints

¹⁵Equation (17) can also capture borrowing or short-sale restrictions. Similar insights would emerge if the constraint directly depended on other equilibrium objects, such as prices.

¹⁶This condition is verified in the baseline model, where we assumed that there was a single state, so $|S| = 1$, but two different assets, $|J| = 2$. Our results can be extended to the $|J| > |S| + 1$ case, in which there may be multiple arbitrage opportunities, or the $|J| < |S| + 1$, in which arbitrage opportunities only emerge in specific scenarios.

in Equation (17); ii) arbitrageurs follow an arbitrage trading strategy with zero cash flow at all dates and histories and a positive cash flow at date 0; and iii) asset markets clear for all assets, dates, and histories, that is,

$$\sum_i \Delta q_t^{ij} (s^t) di + \Delta q_t^{\alpha j} (s^t) = 0, \quad \forall j, \forall t \in \{0, \dots, T-1\}, \forall s^t. \quad (19)$$

It is straightforward to characterize the optimality conditions of investors — see Equation A.4 in the Appendix. Understanding how arbitrageurs choose the directions of their trades, as we describe next, is important to understand Proposition 4.

First, as in the baseline model, note that we can reformulate the arbitrageurs' budget constraints in terms of a size/scale of the arbitrage trade and a direction for such a trade. In this dynamic stochastic setup, we must define a size/scale and direction for each date and each history. Formally, we can rewrite the budget constraints of arbitrageurs as follows:

$$c_t^\alpha (s^t) = m_{t-1} (s^{t-1}) \sum_j \left(d_t^j (s^t) + p_t^j (s^t) \right) x_{t-1}^{\alpha j} (s^{t-1}) - m_t (s^t) \sum_j p_t^j (s^t) x_t^{\alpha j} (s^t), \quad \forall t, \forall s^t, \quad (20)$$

where we define sizes/scales $m_t (s^t) = q_t^{\alpha 1} (s^t)$, $\forall t, \forall s^t$, and trading directions $x_t^{\alpha j} (s^t) = \frac{q_t^{\alpha j} (s^t)}{q_t^{\alpha 1} (s^t)}$, $\forall t, \forall s^t$.

Second, when combining Equation (20) with the fact that arbitrageurs' consumption must be zero at all times after date 0, we can decompose the directions of arbitrageurs trades at history s^{t-1} into a static component and a dynamic component, as follows:

$$\underbrace{\tilde{x}_{t-1}^\alpha (s^{t-1})}_{\text{direction}} = - \underbrace{\left(\tilde{D}_{t|s^{t-1}} + \tilde{P}_{t|s^{t-1}} \right)^{-1} \left(D_{t|s^{t-1}}^1 + P_{t|s^{t-1}}^1 \right)}_{\text{static component}} + \underbrace{\left(\tilde{D}_{t|s^{t-1}} + \tilde{P}_{t|s^{t-1}} \right)^{-1} \frac{M_{t|s^{t-1}}}{m_{t-1} (s^{t-1})} \left(P_{t|s^{t-1}}^1 + \text{diag} \left(\tilde{P}_{t|s^{t-1}} \left(\tilde{X}_{t|s^{t-1}}^\alpha \right)' \right) \right)}_{\text{dynamic component}}, \quad (21)$$

where $\tilde{x}_{t-1}^\alpha (s^{t-1})$ corresponds to a vector of trading directions of dimension $J-1$, since we have normalized $x_{t-1}^{\alpha 1} = 1$, $D_{t|s^{t-1}}$ and $\tilde{D}_{t|s^{t-1}}$ are payoff matrices, $P_{t|s^{t-1}}$ and $\tilde{P}_{t|s^{t-1}}$ are price matrices, and $M_{t|s^{t-1}}$ and $\tilde{X}_{t|s^{t-1}}^\alpha$, are matrices of trade sizes/scales and trade directions, respectively. We provide explicit formal definitions of these matrices in Section B of the Appendix.

The static component in the first line of Equation (21) captures the one-period-forward direction of the trade purely based on asset payoffs and prices. For instance, in our baseline model, the static component simply takes the value -1 — see Equation (5) — since the two assets in that case have identical payoffs.

The dynamic component in the second line of Equation (21) captures how arbitrageurs adjust the direction of their arbitrage trade at a given date/state because future arbitrage trades are non-zero (formally, when $M_{t|s^{t-1}} \neq 0$). As we show next, only the static component, which can

be easily computed in specific applications, is necessary to determine the direct welfare gains from closing arbitrage gaps.

Finally, note that restrictions on the size of future arbitrage trades, for instance those highlighted in [Shleifer and Vishny \(1997\)](#), can be mapped to constraints on $M_{t|s^{t-1}}$, which in turn affect the dynamic component. Therefore, our results are based on the assumption that is actually possible to build an arbitrage strategy at s^{t-1} , for given values of $M_{t|s^{t-1}}$, as explained in detail in Section [D](#) of the Online Appendix.

4.3 The Value of Arbitrage

We now characterize the marginal social value of arbitrage in this general model in [Proposition 4](#). This proposition is the counterpart of part b) of [Proposition 2](#). In [Section B](#) of the Appendix, we also characterize the marginal individual value of arbitrage for investors and arbitrageurs, and include several intermediate results.

Proposition 4. (Marginal social value of arbitrage: general model) *The marginal social value of arbitrage at date 0, that is, the marginal value of increasing the scale of the arbitrage trade m_0 , aggregated across all agents and measured in date 0 dollars, is given by*

$$\begin{aligned} \frac{dW(m_0)}{dm_0} &= \sum_i \frac{dV^i(m_0)}{\lambda_0^i} + \frac{dV^\alpha(m_0)}{\lambda_0^\alpha} \\ &= \underbrace{-p_0^1 + \tilde{p}_0 \cdot (\tilde{D}_1 + \tilde{P}_1)^{-1} (D_1^1 + P_1^1)}_{\text{direct welfare gains}} - \underbrace{\sum_j \sum_{t=0}^{T-1} \sum_{s^{t+1}} \Xi_{t+1}^j(s^{t+1}) \frac{dp_{t+1}^j(s^{t+1})}{dm_0}}_{\text{distributive pecuniary effects}}, \end{aligned}$$

where $\Xi_{t+1}^j(s^{t+1})$, formally defined in [Equation \(A.11\)](#) of the Appendix, is a term that depends on *i*) marginal rates of substitutions of investors and arbitrageurs between history s^{t+1} and date 0, and *ii*) net trades of asset j at history s^{t+1} .

[Proposition 4](#) shows that the insights of the baseline model extend to much more general environments. In particular, [Proposition 4](#) shows that the marginal social value of arbitrage can be decomposed into two components: the direct welfare gains, at date 0, and the distributive pecuniary effects of the arbitrage trade, starting from date 1 onward.¹⁷ Our baseline model exclusively features the direct welfare gains, which correspond to the arbitrage gap, after appropriately computing the correct static direction of the arbitrageurs trade, given by $(\tilde{D}_1 + \tilde{P}_1)^{-1} (D_1^1 + P_1^1)$ and defined in [Equation \(21\)](#). In particular, in the baseline model, which featured two assets with identical payoffs, the direct welfare gains term collapses to the price difference between both assets. Interestingly, even though arbitrage gaps may remain open at future dates, the direct welfare gains of the arbitrage trade at date 0 only include the date 0 arbitrage gap.

Since financial markets may be incomplete in this general model, it is well-known that this economy is in principle constrained inefficient. Therefore, consistent with our discussion of [Remark](#)

¹⁷As in the baseline model, the date 0 distributive pecuniary effects of the arbitrage trade cancel out.

3, one would expect the marginal social value of arbitrage to include additional terms. In particular, changing the size of the amount arbitrated at date 0 impacts the prices of all assets at all future dates, which is captured by the term $\frac{dp_{t+1}^j(s^{t+1})}{dm_0}$. These price changes translate into welfare via i) the differences in marginal rates of substitutions among agents, and ii) the size of net trades at all future dates and states, which define the term in brackets that multiplies the price changes. This result is fully consistent with the characterization of distributive pecuniary externalities in [Dávila and Korinek \(2018\)](#), who show that these externalities are zero-sum at each specific history. In general, from a date 0 perspective, the distributive pecuniary effects can be welfare improving or decreasing, and it is possible to find scenarios in which they are zero, for instance, when markets become complete, or when all agents follow buy-and-hold strategies. We conclude the analysis of the general model with a remark of practical relevance.

Remark 8. (Implications for empirical applications) Proposition 4 implies that the measures of welfare gains that we report in Sections 5 and 6 should be interpreted as the gains from closing the arbitrage gap at a given point in time, that is, they capture the direct welfare gains of closing an arbitrage gap at a given date. If closing the arbitrage gap at that point in time closes all future gaps, all present and future direct welfare gains are captured by the number that we report. If closing the arbitrage gap at a given point in time does not narrow future arbitrage gaps, and it would be necessary to implement an arbitrage trade every day, the estimates that we report correspond to the daily flow of direct welfare gains.

5 Application 1: Covered Interest Parity

In our first empirical application, we use our theoretical results to provide estimates of the welfare gains associated with eliminating violations of covered interest parity (CIP) for multiple currency pairs. Following the work of [Du, Tepper and Verdelhan \(2018\)](#), a substantial body of research has studied violations of CIP — see [Du and Schreger \(2021\)](#) for a recent survey.

CIP is a no-arbitrage condition that relates spot foreign exchange rates, forward exchange rates, and interest rates. In particular, an investor can exchange present dollars for dollars three months later in two different ways. First, the investor could invest in a three-month US T-Bill. Alternatively, the investor could exchange dollars for euros in the spot market, use the proceeds to purchase a German three-month zero coupon bond, and additionally sell a three-month forward contract in the exact amount of the face value of the German zero coupon bond to convert the payoff of the bond into dollars. Assuming no differential sovereign risk, these two strategies have exactly the same payoffs, so the CIP condition holds when the price of each strategy is the same. Should the CIP condition not hold, the investor could purchase the cheaper leg and sell short the more expensive leg to generate risk-free profits.

[Du, Tepper and Verdelhan \(2018\)](#) show that during and after the global financial crisis of 2007–2008, CIP went from consistently holding within a narrow band to being systematically violated for many currency pairs, at times with significant magnitudes. They present evidence that explains negative cross-currency bases with respect to the Dollar in the post-crisis regime by

banking regulations and demand imbalances for currencies. We take the existence of CIP deviations as a starting point and, building on the theoretical results derived in the previous two sections, proceed to estimate the potential welfare gains associated with closing such deviations.

5.1 Measurement Approach

To simplify the exposition, so far we have presented our theoretical results in a single-currency environment. In Proposition 5, we characterize the marginal social value of arbitrage in the context of a multi-currency model — fully developed in the Appendix — that nests our baseline model and that is suitable to study CIP deviations.

Proposition 5. (CIP: Marginal social value of arbitrage) *The marginal social value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade, m , in the CIP context, aggregated and measured in date 0 units of domestic currency (dollars), is given by*

$$\frac{dW(m)}{dm} = \frac{S(m)}{F(m)} \frac{1}{1+r^f(m)} - \frac{1}{1+r^d(m)}, \quad (22)$$

where $S(m)$ denotes the spot foreign exchange rate, $F(m)$ denotes the forward exchange rate, $r^f(m)$ denotes the foreign interest rate, and $r^d(m)$ denotes the domestic interest rate. The arbitrage trade m is defined as purchasing the domestic leg and selling the foreign leg.

If the CIP basis/deviation, defined in Equation (22), is positive, the return on the foreign leg is lower than the return on the domestic leg. In that case, there is scope to sell the foreign leg and buy the domestic leg to close the arbitrage gap.¹⁸ If the CIP condition holds, then

$$\frac{S(m)}{F(m)} \frac{1}{1+r^f(m)} = \frac{1}{1+r^d(m)} \quad \text{and} \quad \frac{dW(m)}{dm} = 0.$$

Once we have characterized the marginal social value of arbitrage in Proposition 5, we can use Proposition 2 to find the total social value of arbitrage. Therefore, conceptually, to find the empirical counterpart of the total social value of arbitrage, we must i) measure the cross-currency basis, which is easily observable, and ii) estimate price impact for each of the markets involved in the arbitrage trade, which is substantially harder. To achieve this second step, we use a high-frequency dataset of the FX futures market, whose centralized order book and trade and quote transparency allow us to provide credible measures of price impact for quarterly FX futures. However, due to limitations associated with measurement, we must base our welfare calculations on some assumptions, which we describe next. At a high level, these assumptions reflect the impossibility of credibly estimating multidimensional, rich, and complex price impact functions.

First, we assume that (foreign and domestic) interest rates are not impacted by the arbitrage trade. That is, we assume that the price impact of arbitrage trades on the bond markets is zero. This assumption can be seen as an appropriate approximation for the following reason. The futures

¹⁸Note that in a single currency economy, $S(m) = F(m) = 1$, and Equation (22) collapses to Equation (13).

contract that corresponds to the three-month USD LIBOR rate — the CME’s Eurodollar futures — regularly features front-month resting liquidity at the bid and offer of over \$100B. Given i) that the largest gap-closing trades that we find are of the same order of magnitude as the Eurodollar resting liquidity, ii) that the front-month bid-ask spread is only a quarter of a basis point, and iii) that resting futures liquidity is a lower bound for the true available liquidity, the impact of executing the gap-closing trades on interest rates is extremely small, justifying our assumption. Note that this assumption biases our estimate of welfare gains upward. In other words, if we allowed trading in bond markets to be subject to price impact, we would find smaller welfare gains from closing arbitrage gaps and smaller gap-closing arbitrage trades.

Second, we assume that the estimated price impact function of both the spot and three-month forward markets can be approximated by the price impact function of the CME’s front-month currency futures contracts.¹⁹ We argue that this is an appropriate approximation because front-month futures contracts are regularly used interchangeably with spot FX by traders and because the three-month forward and spot foreign exchange markets are very closely related to the futures market (in fact, they both are at some point in time equivalent to the futures).

Third, we must make an assumption on cross-price impact. Because the spot and forward prices are highly correlated and fundamentally connected, significant purchases in one market spill over into the other market, which we refer to as cross-price impact. As cross-price impact is notoriously difficult to estimate, we use our estimates of directional price impact from the futures market and we assume that cross-price impact of simultaneous transactions reduces the price impact by 90%. As explained in detail in Section F.1 of the Appendix, under the baseline square root specification of price impact, this is the same as assuming that \$100B of transactions have the same effect that \$1B of transactions would have under the alternate assumption of zero cross-price impact. We consider this a conservative assumption, which again biases our results towards finding large welfare gains.

Fourth, we abstract from cross-currency spillovers, for which it is very hard to find credible estimates. That is, we provide welfare estimates for the thought exercise of closing CIP deviations for each currency pair one at a time. Hence, since we do not account for the impact that closing the CIP deviation for a given currency pair has on other CIP deviations, simply adding up the welfare gains across the five currency pairs that we consider is likely to overestimate the welfare gains of closing all arbitrage gaps at once.

Overall, we purposefully adopt these assumptions to bias our results towards finding large welfare gains. Therefore, our results should be interpreted as upper bounds for the potential direct welfare gains from closing CIP deviations.

5.2 Estimating Price Impact

Here, we briefly describe the data used to estimate price impact in the FX futures market. Section E of the Appendix includes summary statistics, a more detailed description of the data sources, and

¹⁹The front-month contract is defined as the quarterly contract that is closest to delivery. Hence, the front-month contract is equivalent to a forward contract whose maturity varies between three months and several days, depending on the date.

some useful institutional background on FX markets. Next, we describe our estimation procedure and discuss the price impact estimates.

Data Description In order to estimate price impact in the FX futures market, we use high-frequency transaction data and top-of-the-book data from the Chicago Mercantile Exchange (CME). Our dataset contains the universe of transactions, bid changes, and offer changes, recorded with a millisecond timestamp corresponding to the time that the CME’s matching engine processed the order book change or trade.

Our dataset starts on December 15, 2019 and ends on February 26, 2021 and corresponds to five contract months for each of five different futures contracts. Specifically, our dataset covers the March 2020, June 2020, September 2020, December 2020, and March 2021 futures contracts for the following currency pairs: Australian Dollar/USD, British Pound/USD, Canadian Dollar/USD, Euro/USD, and Yen/USD.

Estimation Procedure As shown in the previous sections, our objective is to estimate the impact of trading a given quantity of contracts/shares on the price of the asset of interest. To do so, we build on the literature that has studied the relationship between order size and asset prices over the past several decades. In particular, [Gabaix et al. \(2003\)](#), [Almgren et al. \(2005\)](#), [Bouchaud \(2010\)](#), [Frazzini, Israel and Moskowitz \(2018\)](#), and [Graves \(2021\)](#), among others, have empirically found that price impact satisfies an approximate square root power law relation for a wide variety of countries, time periods, and financial instruments from the 1980s to the present.²⁰

Therefore, we estimate the following functional form for price impact in the FX futures markets:

$$F_{\tau+1} - F_{\tau} = \theta + \alpha \operatorname{sgn}(Q_{\tau}) |Q_{\tau}|^{\beta} + \varepsilon_{\tau}, \quad (23)$$

where F_{τ} , whose imputation is described in Equation (24) below, denotes the price of a futures contract right before a given transaction takes place; Q_{τ} denotes the actual size of the transaction; ε_{τ} is an error term; and θ , α , and β are parameters. Three features of our nonlinear least squares estimation procedure are worth highlighting.

First, note that the unit of observation in our nonlinear regression is a transaction. Hence, every time a transaction takes place, we record the difference in prices right before the transaction to right before the next transaction, as well as the size of the transaction, to generate an observation that informs the estimation of price impact. In our dataset, transactions may occur within microseconds of each other.

Second, since we want to find the impact of a trade on the price, we do not use the prices at which transactions take place. Instead, we impute the price F_{τ} (analogously, $F_{\tau+1}$) using the

²⁰[Gabaix et al. \(2006\)](#) and [Donier et al. \(2015\)](#), among others, have developed theoretical models that justify the power law functional form.

following weighted average of the bid and ask prices:

$$F_\tau = \omega_\tau F_\tau^B + (1 - \omega_\tau) F_\tau^A, \quad \text{and} \quad \omega_\tau = \frac{M_\tau^A}{M_\tau^B + M_\tau^A}, \quad (24)$$

where F_τ^B and F_τ^A respectively denote the bid and ask prices and M_τ^B and M_τ^A respectively denote the size of the bid and the ask right before a transaction takes place.²¹

Third, to be consistent with the empirical literature described above and to improve the efficiency of our estimation, in our baseline estimation we assume that the power law coefficient is predetermined at $\beta = \frac{1}{2}$. Section F.5 of the Appendix shows that our conclusions are almost identical when β can be freely estimated.

Price Impact Estimates Given the high-frequency nature of our data, we are able to compute precise estimates of price impact on a daily basis between December 15, 2019 and February 26, 2021, which allows us to capture the time-varying nature of price impact. Figure 4 illustrates our price impact estimates in the FX futures market for the five currency pairs in our dataset.

Figure 4a shows the price impact functions of the form described in Equation (23) for the five currency pairs that we study, estimated using all transactions in our sample. Since our measure of price impact is nonlinear, it is more informative to present the whole estimated function, rather than simply reporting the estimated α coefficients. Table OA-5 of the Appendix includes summary statistics of the estimated daily price impact coefficients. In terms of magnitudes, our estimates imply that a trade of \$10B moves the FX futures rate by roughly 0.15% in the case of the Euro and by roughly 0.35% in the case of the Australian Dollar, with the other currency pairs in between. The nonlinear (concave) nature of our price impact estimates implies that the price impact of larger trades does not scale proportionally. For instance, a trade of \$100B moves the FX futures rate by roughly 0.45% in the case of the Euro and by roughly 1.1% in the case of the Australian Dollar, with the other currency pairs in between. As we discuss below in the context of our welfare computations, these price impact estimates are consistent with ancillary evidence we provide on the behavior of FX markets.

Figure 4b shows the evolution over time of the estimates of the price impact α coefficients for each day of our sample. We draw three conclusions from this figure. First, price impact spikes in periods of market distress, particularly the COVID-19 crisis in March 2020, which are associated with high uncertainty, limited liquidity, and significant CIP violations — as we show next. This is an important observation for our welfare calculations, as the period of elevated price impacts corresponds precisely to the greatest violations of covered interest parity since the global financial crisis of 2007–2008, as we further highlight below. Second, daily price impact coefficients are remarkably stable outside of the pandemic-related surge in March 2020. Finally, it seems that the average daily estimates of the price impact α coefficients have not recovered to the pre-pandemic levels.

²¹The choice of weights in Equation (24) follows Burghardt, Hanweck and Lei (2006). See Graves (2021) for a detailed discussion of these and other imputation methods, e.g., midpoint.

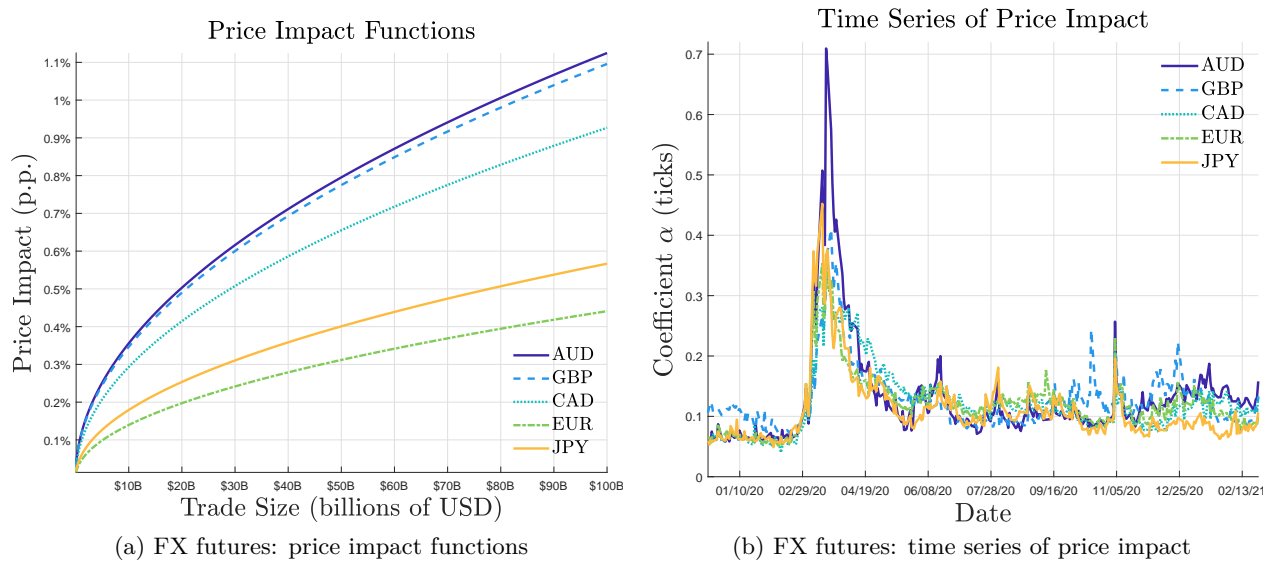


Figure 4: FX futures: price impact estimation

Note: Figure 4a shows the price impact functions of the form introduced in Equation (23) for the five currency pairs that we study, estimated using all transactions in our sample over the period between December 15, 2019 and February 26, 2021. To report price impact in percentage points, we use the mean price over the dataset as reference. Figure 4b shows daily estimates of the α coefficients from the regression introduced in Equation (23). The α coefficients reported in Figure 4b here are expressed in ticks, the minimum amount that a given futures contract is allowed to move, which is closely tied to the median bid-ask spread — see Table OA-1 of the Appendix for tick values. By using ticks, we can display all estimates in the same figure, since the α coefficients have the same order of magnitude. As explained in the text, price impact increases dramatically during the COVID-19 pandemic crisis in March 2020, then rapidly subsides to a new, more elevated equilibrium. Outside of the crisis episode, estimates are stable.

Before using these price estimates to make welfare calculations, we would like to make three observations regarding our estimation of price impact. First, note that while our estimates of price impact rely on trade sizes ranging from hundreds of thousands of dollars to several hundred million dollars, the required gap-closing trades often run into the tens or hundreds of billions, making extrapolation unavoidable.²² Our specification of price impact, which builds on a large empirical literature, transparently illustrates how the needed extrapolation takes place. In particular, Section F.6 of the Appendix explores the robustness of our conclusions to imposing alternative values of the power law coefficient β .

Second, to address the potential concern that microstructure issues are driving our price impact estimates, in Section F of the Appendix we explore how our estimates change as we lengthen the number of transactions on the left hand side of Equation (23) from a single transaction to 5000 transactions. We actually find that price impact is increasing in the number of transactions, which suggests that our estimates are not driven by microstructure noise.

Finally, note that our approach mitigates but does not fully circumvent endogeneity concerns. By i) studying transaction time instead of clock time, ii) only using the interval between trades, and iii) using an imputed price instead of the midpoint or most recent transaction price, we reduce the risk that signed quantities are influenced by future prices (generally milliseconds forward in time). These concerns would be much more acute, for instance, when measuring price changes over longer horizons.

5.3 Measuring CIP Deviations

The next inputs necessary for our welfare calculations are the cross-currency bases. Here, we compute the relevant cross-currency bases using spot rates, three-month forward rates, foreign three-month interest rates, and US three-month interest rates. In the body of the paper, we use secured three-month lending rates whenever such instruments are available, while in Section F.4 of the Appendix, we show that our results are robust to using LIBOR rates.

Data Description We use data from Bloomberg to compute cross-currency bases between February 2, 2008 and February 26, 2021. Table OA-3 of the Appendix provides the summary statistics for the cross-currency bases, while Figure 5a plots the time-series evolution of the cross-currency bases for the five relevant currency pairs. We provide an exact list of the Bloomberg data series used in the Appendix.

CIP Deviations Figure 5a shows the three-month cross-currency bases for the USD versus the five currency pairs that we consider, which are some of the world’s most liquid currencies. As discussed extensively in Du, Tepper and Verdelhan (2018), the post-global financial crisis period contains substantially and persistently negative cross-currency bases. Two major events are immediately visible in the figure: the global financial crisis of 2007–2008 and the COVID-19

²²At the sizes of the largest trades in the dataset, our estimated price impact is of the order of several basis points. This is a meaningfully large fluctuation that is of the same order of magnitude as most CIP deviations in the dataset.




Figure 5: CIP deviations/Correlation with price impact estimates

Note: Figure 5a shows the three-month CIP deviations for the five currency pairs that we consider. The basis is calculated as $\frac{S}{F} \frac{1}{(1+r^f)^{1/4}} - \frac{1}{(1+r^d)^{1/4}}$, where r^f and r^d are annualized three-month foreign and domestic (USD) interest rates. The two largest CIP deviations take place in September 2008, during the global financial crisis of 2007–2008, and in March 2020, at the onset of the COVID-19 pandemic. Our results are comparable to those in Du, Tepper and Verdelhan (2018), with two minor expositional differences: i) they present ten-day moving averages of the gaps, while we show daily values, and ii) our values are not annualized, so they must be multiplied by 4 to be compared directly to the annualized value of the bases reported by Du, Tepper and Verdelhan (2018). In Section F.4 of the Appendix we include results using LIBOR. Figure 5b shows that there exists a negative relation between daily price impact estimates (the α coefficients) and the observed CIP deviations. This is equivalent to stating that the relation between price impact estimates and the magnitude of the observed CIP deviations is positive. Figure 5b shows the relation for the case of the Euro during the period between December 15, 2019 and February 26, 2021. The regression line and its R^2 shown in the figure are computed using only the observations with large impact (when $\alpha \geq 0.1$). A similar negative correlation emerges for the other four currency pairs: the analogous R^2 's are 0.289 (AUD), 0.683 (GBP), 0.792 (CAD), and 0.637 (JPY).

pandemic in March 2020, both of which appear as downward spikes. These events can be thought of as dollar-liquidity driven events wherein positive shocks to the demand for US dollars in foreign countries drove covered interest away from parity.

Figure 5b illustrates the relation between daily measures of cross-currency bases (CIP deviations) with the daily price impact estimates introduced in Figure 4. This figure shows that there exists a clear positive relation between the daily price impact estimates (the α coefficients) and the magnitude of observed CIP deviations for the case of the Euro (similar patterns emerge for all currency pairs). That is, days in which CIP deviations are large are days in which the FX market is illiquid, as measured by high price impact (a high α coefficient). The comovement between price impact measures and CIP deviations is relevant for our welfare conclusions, as described next.

5.4 Gap-closing Trades and Welfare Estimates

Armed with both price impact estimates and cross-currency bases, we are ready to leverage the theoretical framework introduced in the previous sections to compute gap-closing arbitrage trades and total social welfare gains from closing CIP deviations. First, given our specification of price impact for a given currency pair market, we compute the gap-closing arbitrage trade m^* as the size of the arbitrage that closes the arbitrage gap:

$$\frac{S(m^*)}{F(m^*)} \frac{1}{1+r^f(m^*)} = \frac{1}{1+r^d(m^*)}. \quad (25)$$

Second, given that we have estimates of $\frac{dW(m)}{dm}$, as defined in Equation (22), for different trade sizes, we compute the social welfare gain of closing the CIP gap for a given currency pair by integrating the value of the CIP deviation until the gap is closed, that is,

$$\int_0^{m^*} \frac{dW(\tilde{m})}{d\tilde{m}} d\tilde{m}. \quad (26)$$

Since the cross-currency bases are generally negative, Equation 6a is typically negative, but its absolute value reflects the welfare gains associated with closing the gap. As explained in the previous section, the measures of welfare gains that we report should be interpreted as the gains from closing the arbitrage gap at a given point in time. If closing the arbitrage gap at a given point in time closes all future gaps, the estimates that we report capture all present and future welfare gains. If closing the arbitrage gap at a given point in time does not impact future prices, and it would be necessary to implement an arbitrage trade every day, the estimates that we report correspond to the daily flow of direct welfare gains.

Figure 6 presents daily estimates of i) gap-closing arbitrage trades and ii) welfare gains from closing CIP deviations for the five currency pairs that we study. Figure 6a shows that the estimated gap-closing arbitrage trade for each of the five currency pairs never exceeds \$1.22T and never exceeds \$455B outside of the Yen-Dollar basis. One can readily see why gap-closing trades do not spike in periods with high CIP deviations: consistent with Figure 5b, at precisely the moment in which CIP

violations reach their peak in March 2020, price impact also surges, resulting in only a moderate increase in gap-closing trades. Since events of distress are commonly associated with both illiquidity and CIP deviations, the magnitude of the gap-closing arbitrage size does not fluctuate as much as price impact measures and CIP deviations independently.

Figure 6b shows that the estimated welfare gains from closing CIP deviations never exceed \$1.2B for the five studied pairs, and never exceed \$300M outside of the Yen-Dollar basis. It is useful to normalize these estimates to gain some perspective. Two normalizations seem reasonable. First, it is natural to normalize welfare gains by a measure of total daily trading volume. For instance, in the spot FX markets for the EUR/USD and the JPY/USD pairs, daily trading volumes are respectively \$416B and \$260B (BIS, 2019). Therefore, the estimated normalized welfare gains from closing an arbitrage gap for the EUR/USD is no larger than $\frac{\$300M}{\$416B} \approx 0.072\%$ of spot FX daily volume. Based on this number, one could argue that the gains from closing gaps are relatively small, despite taking place in one of the most highly liquid markets that exist. Alternatively, one could normalize by a measure of combined daily GDP — in the EUR/USD case, roughly $\frac{\$35T}{365} \approx \$100B$. In that case, the estimated normalized welfare gains from closing an arbitrage gap are no larger than $\frac{\$300M}{\$100B} \approx 0.3\%$ of combined daily GDP, which is a non-negligible amount.

There are two features of our analysis that explain these quantitative findings. First, the CIP deviations are significant, but not that large in magnitude. This is in contrast to Section 6, in which our approach finds large welfare gains in the case of some dual-listed companies for which arbitrage gaps are very large. Second, as described above, the fact that CIP deviations are larger in magnitude when price impact estimates are larger makes the size of the gap-closing trade small, lowering the scope for welfare gains.

Section F of the Appendix reports additional findings. First, by backward extrapolating our price impact estimates from 2019–2021, we can produce a sensible approximation of the gap-closing trades and the welfare gains from closing CIP deviations for the 2010–2019 period. A similar pattern to the one we observe in 2019–2021 is immediately apparent in Figure OA-2 of the Appendix — the greatest welfare gains from closing CIP deviations are in the Yen-Dollar basis and briefly spike to \$1.6B, but for all other currency pairs never even exceed \$250M. Second, we characterize the combinations of price impact coefficients and CIP violations that must exist to generate large welfare gains from closing CIP deviations. In Figure OA-3 of the Appendix, we show isoquants of CIP deviations and price impact estimates (the α coefficients) that yield the same level of welfare gains, corresponding to \$100M, \$1B, and \$10B in the EUR/USD case. Figure OA-3 clearly illustrates that large welfare gains from closing CIP deviations can only emerge when CIP deviations are extremely large and price impact estimates are low. To put our estimates in perspective, Figure OA-3 includes average and extreme measures of CIP deviations and price impact estimates in our sample. Taking \$10B as the desired target number, and assuming that price impact takes average values, Figure OA-3 of the Appendix shows that one would need to observe a EUR/USD CIP deviation of 250 basis points, a number which is an order of magnitude larger than any CIP deviation measured.

We conclude our welfare analysis of CIP deviations with two remarks. First, one might conjecture that closing five-year CIP deviations may be associated with higher welfare gains. Du,

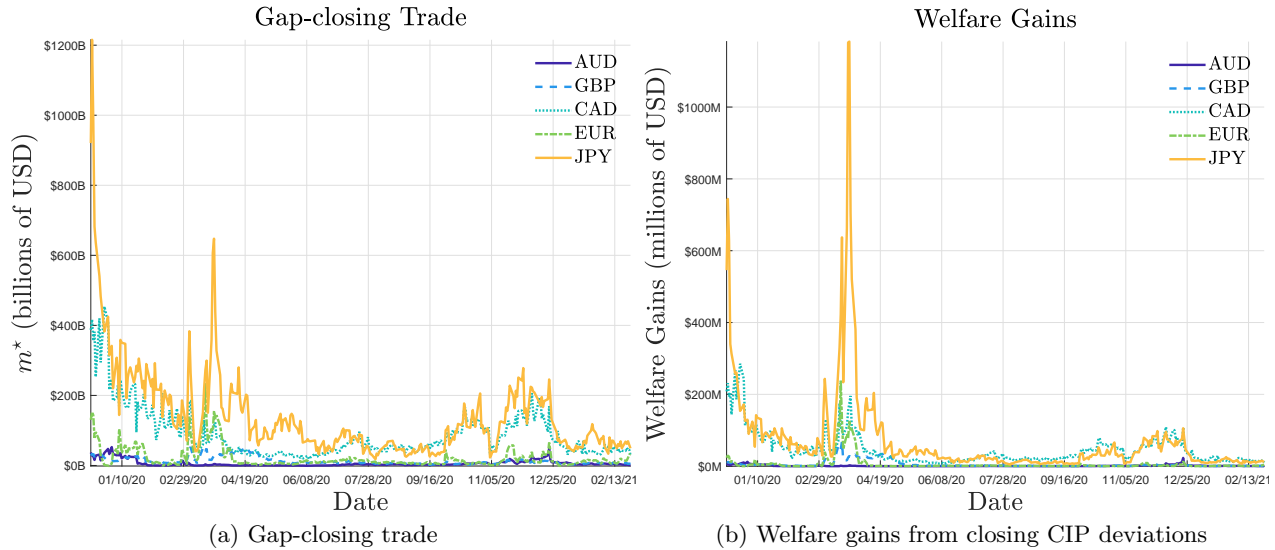


Figure 6: CIP: Welfare Gains/Gap-closing Trade

Note: Figure 6a shows the gap-closing arbitrage trade m^* that is needed for the covered interest parity condition to hold. Gap-closing trades rarely exceed \$500B. Figure 6b shows the daily estimates of the welfare gains from closing CIP deviations between December 15, 2019 and February 26, 2021. These are computed as described in Equation (26), reporting the absolute value of the gains. Estimated welfare gains from closing CIP deviations never exceed \$1B for the five studied pairs, and never exceed \$250M outside of the Yen-Dollar basis.

Tepper and Verdelhan (2018) and others find that such CIP violations are greater than their three-month counterparts, but have the same order of magnitude. However, the five-year T-Note and corresponding foreign interest rate markets are not nearly as liquid as the three-month T-Bill and Eurodollar markets. More importantly, the five-year forward market with its distant delivery date is far less liquid than the three-month forward market. If comprehensive five-year forward quote and trade data were available, the estimated price impact coefficients would be significantly larger than those estimated from front-month futures, which are among the most liquid markets in the world. Although the welfare gains from longer time horizons might seem larger in magnitude due to the large CIP deviations documented in the literature, the lack of liquidity would make the gap-closing trade small.

Finally, note that our conclusions regarding the magnitude of gap-closing trades as well as our assumptions on cross-price impact seem to be supported by the Federal Reserve’s quarterly reports and press announcements. In particular, on September 18, 2008, right as the CIP violations peaked, the Federal Reserve increased swap lines with other central banks by \$180B; on March 19, 2020, swap lines were extended to multiple additional central banks around the world while that same week saw a reduction in the pricing of US Dollar liquidity via swap lines between the Federal Reserve and each of the central banks of Canada, the EU, Japan, and Switzerland. The latter effort was with the intention of reducing dollar shortages and coincided with the precise week that the CIP violations peaked in 2020. Looking at quarterly Federal Reserve reports, we observe that central bank liquidity swaps went from functionally \$0 at the beginning of March 2020 to just

over \$350B by the end of March. The expansion of these swap lines coincided precisely with the elimination of most of the CIP arbitrage gaps. In addition, the two most utilized swap lines were with the BOJ (\$223B at its peak) and the ECB (\$145B at its peak) — see [Federal Reserve \(2020\)](#) — which aligns quite well with the magnitude of our estimates in [Figure 6a](#). While anecdotal, the precise coincidence of swap line usage and the elimination of most of the CIP arbitrage gaps at a time of scarce dollar liquidity lends support to the gap-closing sizes that we find.

6 Application 2: Dual-Listed Companies

In our second empirical application, we use our theoretical results to provide estimates of the welfare gains associated with closing arbitrage gaps for dual-listed companies, also referred to as “Siamese twin” stocks. We compute the welfare gains from closing arbitrage gaps in three particular scenarios: i) Royal Dutch/Shell, the canonical dual-listed company that featured arbitrage opportunities for nearly a century, see, e.g., [Shleifer \(2000\)](#); ii) Smithkline Beecham; and iii) Rio Tinto.²³

Under a dual-listing arrangement, two corporations function as a single operating business but retain separate stock exchange listings. An equalization agreement determines the distribution of cash flows and voting rights among the shareholders of the twin parents. For instance, in the case of Royal Dutch and Shell, Royal Dutch was a Dutch company traded out of Amsterdam and Shell was a British company traded out of London. From 1907 until their unification in July 2005, the stocks of Royal Dutch and Shell traded on different exchanges but with a 60%/40% fixed division of the joint cash flow and ownership rights. Given that the dividend streams were fixed at a 1.5:1 ratio, absence of arbitrage implies that the market capitalizations of both companies should have always respected such a ratio. As extensively documented in the finance literature — see e.g., [Froot and Dabora \(1999\)](#) and [De Jong, Rosenthal and Van Dijk \(2009\)](#) — the prices of Royal Dutch and Shell demonstrated consistent and extreme deviations from the ratio, sometimes amounting to as much as 20 to 30%. Like CIP, this represents a textbook arbitrage opportunity.²⁴

6.1 Measurement Approach

The main challenge associated with computing the welfare gains from closing arbitrage gaps in the case of dual-listed companies is to determine the price impact function of each involved stock, since only daily stock prices and trading volumes are readily available. Since we do not have access to detailed transaction data — as we do in the case of the futures FX market — we rely on the price impact estimates for global equities of [Frazzini, Israel and Moskowitz \(2018\)](#).

Using a global database of a hedge fund’s \$1.7T in stock transactions spanning two decades, [Frazzini, Israel and Moskowitz \(2018\)](#) estimate the price impact of trades as a function of trade

²³According to [Lowenstein \(2001\)](#), the Royal Dutch/Shell trade accounted for more than half of the losses incurred by Long Term Capital Management in equity pairs trading (of \$286 million).

²⁴As originally noted by [Shleifer and Vishny \(1997\)](#), the need to unwind arbitrage trades before prices converge can make arbitrage trades risky in practice. In [Section D](#) of the Appendix we describe how to incorporate this possibility in our framework in a version of the general model studied in [Section 4](#).

size and stock characteristics. Consistent with earlier literature, they find that price impact is well described by a square root functional form. The central component of their main price impact specification for a given stock j takes the form

$$\frac{\Delta p^j(m)}{p^j} = \alpha^j \operatorname{sgn}(m) |m|^{\frac{1}{2}},$$

where $\frac{\Delta p^j(m)}{p^j}$ denotes the proportional price change of stock j after a trade of size m — expressed as the fraction of daily volume, which is the trade size as a percentage of one-year average daily volume — takes place.²⁵ The value of α^j is the key parameter that modulates price impact.

In order to compute gap-closing trades and welfare estimates, we use the price impact coefficient that [Frazzini, Israel and Moskowitz \(2018\)](#) estimate using their full sample: $\alpha = 0.000889$. This is a moderate estimate — in particular in the Royal Dutch/Shell scenario — since they find lower price impact estimates for large-cap stocks and Royal Dutch/Shell was one of the largest companies in the world, with an average market capitalization of around \$100B in the 1990s (around \$160B today). Smithkline Beecham and Rio Tinto are similarly large companies.

In the three scenarios that we consider, we combine the price impact estimates of [Frazzini, Israel and Moskowitz \(2018\)](#) with the data on dual-listed stocks from [De Jong, Rosenthal and Van Dijk \(2009\)](#) to compute gap-closing trades and welfare estimates. We do so by building on the theoretical results derived in Sections 3 and 4, as we did in the CIP case. In particular, in the Royal Dutch/Shell scenario, we report gap-closing trades and welfare estimates for each trading day between December 8, 1987 and October 3, 2002. Specifically, given a share price of p^{RD} for Royal Dutch and p^S for Shell, one can build an arbitrage trade by purchasing m shares of Royal Dutch and selling short $\frac{3}{2}m$ shares of Shell, which generates a marginal social value welfare gain of

$$\frac{dW(m)}{dm} = p^{RD}(m) - \frac{3}{2}p^S(m). \quad (27)$$

As explained in Section 3, we can integrate Equation (27) over m to find the total social value of closing the arbitrage gap, and find its zero to compute the gap-closing trade. We follow an analogous procedure in the case of Smithkline Beecham and Rio Tinto.

Finally, it is worth pointing out that while the price impact estimates of [Frazzini, Israel and Moskowitz \(2018\)](#) are based on transactions that account for at most 13.1% of the daily volume of a stock, the required gap-closing trades are at times larger than that. Hence, as in the CIP case, this makes extrapolation unavoidable. The way in which we specify price impact transparently illustrates how extrapolation takes place and how our results are sensitive to using different estimates.

²⁵Because the price impact specification of [Frazzini, Israel and Moskowitz \(2018\)](#) relies on a normalization of trades that employs one-year averages of volume, we may find different welfare estimates in periods with identical arbitrage gaps.

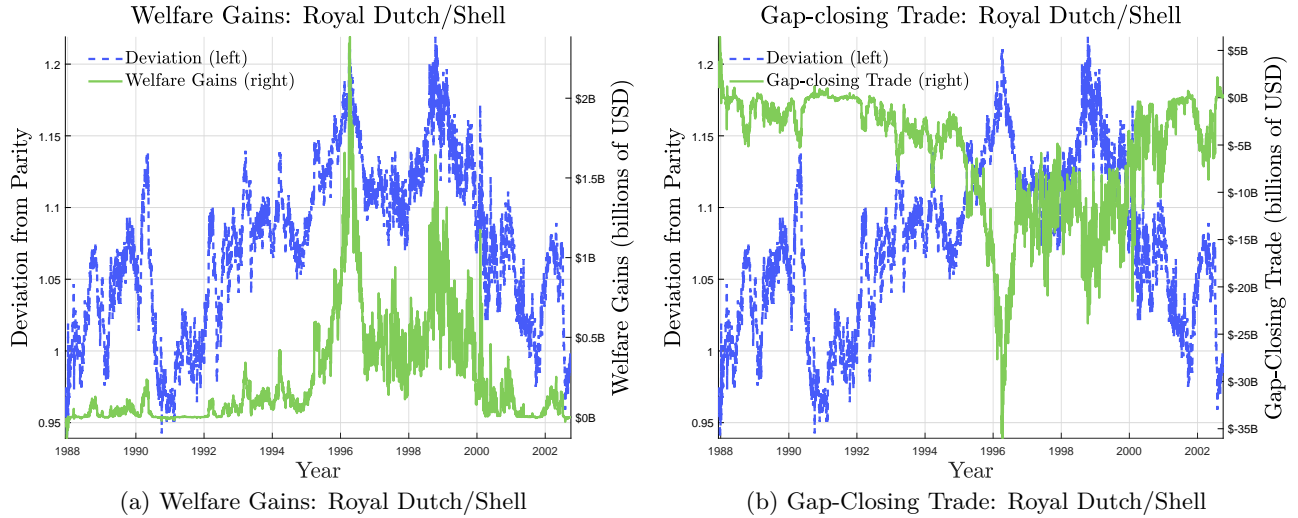


Figure 7: Royal Dutch/Shell

Note: Both Figures 7a and 7b show the price deviation relative to parity (normalized to 1). Figure 7a shows the welfare gains from closing the arbitrage gap for Royal Dutch/Shell. Peak deviations are associated with significant welfare gains, with much of the 1996–2000 period seeing welfare gains from closing the gap of over \$1B. Figure 7b shows the gap-closing arbitrage trade for Royal Dutch/Shell. The maximum gap-closing trade is around \$35B, which takes place in a date in which the deviation from parity is 20%. For reference, the overall market capitalization on that date (4/2/96) is around \$118B.

6.2 Gap-closing Trades and Welfare Estimates

6.2.1 Royal Dutch/Shell

Figures 7a and 7b show welfare gains and gap-closing trades, along with deviations from parity, for each trading day between December 8, 1987 and October 3, 2002. Figure 7a shows the welfare gains from closing the arbitrage gap in nominal USD, which peak at approximately \$2B in 1996. In contrast to our CIP results, this twin share divergence is associated with significant welfare gains. This finding can be explained by the fact that the deviations from parity in the Royal Dutch/Shell scenario are extreme. As detailed in Figure 7a, the divergence at the time of maximum welfare gains from closing arbitrage gaps is over 20% despite the fact that the company had a market cap of over \$100B at the time. Because of this significant deviation from parity, a trade size equal to around \$36B is required to close the arbitrage gap — as shown by Figure 7b. The persistence of extreme deviations from parity over many years results in an extended period over which the welfare gains from closing the arbitrage gap in Royal Dutch/Shell are over \$1B.

6.2.2 Smithkline Beecham

GlaxoSmithKline was created in 2000 as a combination of Glaxo Wellcome and Smithkline Beecham. Smithkline Beecham, which was one of the world’s major pharmaceutical companies until the merger, had dual-listed shares as a result of the “stapled stock structure” ensuing after the merger of its own precursors, Smithkline Beckham (US-based) and Beecham (UK-based).

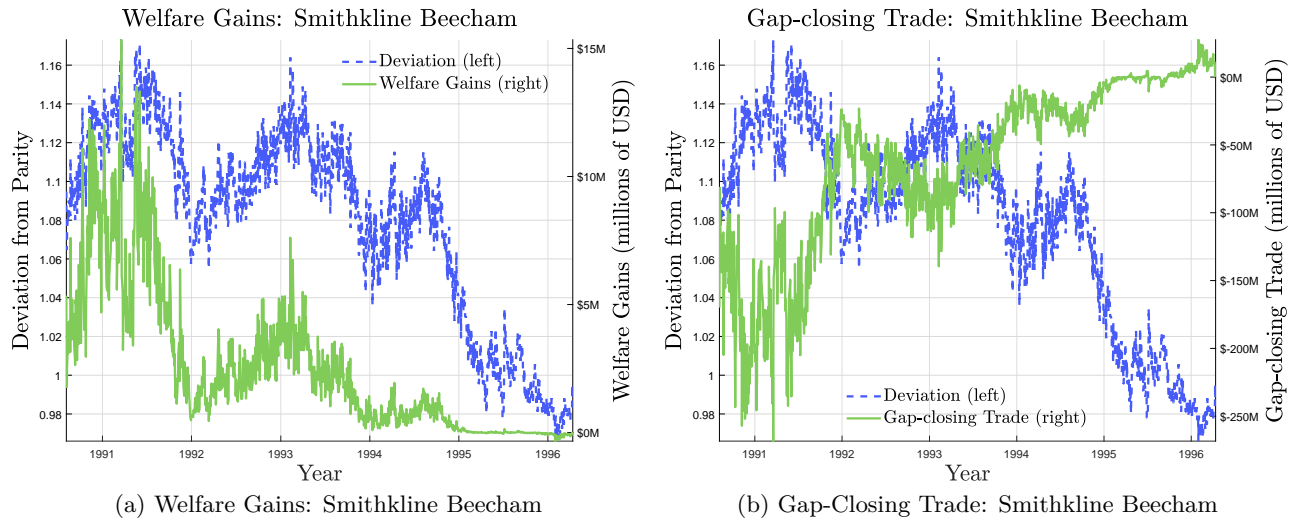


Figure 8: Smithkline Beecham

Note: Both Figures 8a and 8b show the price deviation relative to parity (normalized to 1). Figure 8a shows the welfare gains from closing the arbitrage gap for Smithkline Beecham. The welfare gains never exceed \$16M due to the limited liquidity of one of the two listings. Figure 8b shows the gap-closing arbitrage trade for Smithkline Beecham. The maximum gap-closing trade is around \$250M.

Figures 8a and 8b show welfare gains and gap-closing trades, along with deviations from parity, for each trading day between August 6, 1990 and April 12, 1996. Unlike Royal Dutch/Shell, the Smithkline Beecham twin share divergence ratio never reaches more than 20% and is vanishingly small in the later years of the sample. Figure 8a, which shows the welfare gains from closing the arbitrage gap in nominal USD, illustrates how the welfare gains are persistently small in general, and several orders of magnitude smaller than the maximal gains in Royal Dutch/Shell case. This can be seen as a direct result of the limited liquidity in the UK-traded H-share class.

6.2.3 Rio Tinto

Finally, we consider the case of Rio Tinto, a dual-listed company that represents one of the largest mining corporations in the world, with dual listings in Australia and the United Kingdom.

Figures 9a and 9b show welfare gains and gap-closing trades, along with deviations from parity, for each trading day between December 18, 1997 and October 3, 2002. Figure 9a clearly shows that the deviations from parity are less dramatic than in the Royal Dutch/Shell scenario. As noted in De Jong, Rosenthal and Van Dijk (2009), Rio Tinto experiences the smallest mean absolute divergence among the dual-listed companies studied in that paper. Largely as a result of this fact, we find minimal welfare gains from closing arbitrage gaps throughout most of the price history, except when the deviation from parity briefly diverges to its most extreme points with peak welfare gains of roughly \$150M and maximal gap-closing trades of approximately \$3B — see Figure 9b. When compared to our findings in the Royal Dutch/Shell scenario, our finding of low welfare gains in the Rio Tinto scenario can be explained by the combination of i) lower liquidity than Royal

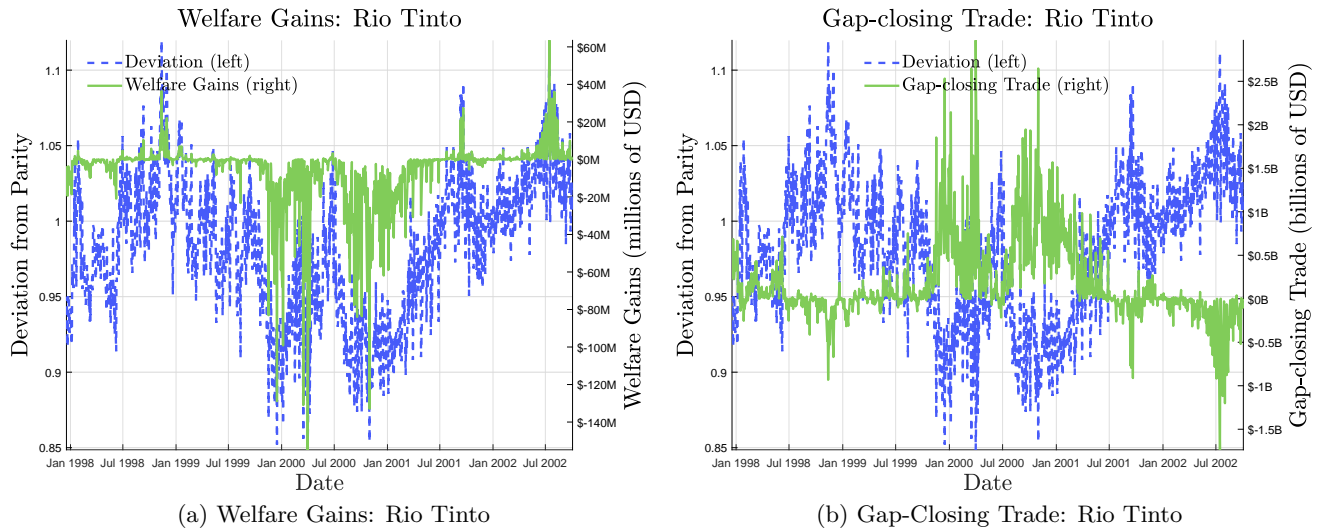


Figure 9: Rio Tinto

Note: Both Figures 9a and 9b show the price deviation relative to parity (normalized to 1). Figure 9a shows the welfare gains from closing the arbitrage gap for Rio Tinto. The welfare gains never exceed \$150M. Figure 9b shows the gap-closing arbitrage trade for Rio Tinto. The maximum gap-closing trade is around \$3B.

Dutch/Shell, as measured by dollar trade volume, and ii) smaller deviations from parity.

6.3 Final Remarks

As illustrated by our three applications, the welfare gains of closing arbitrage gaps for different dual-listed companies can be substantially different. Our findings can be traced directly to the theoretical connection between liquidity/price impact and welfare: arbitrage gaps that are resistant to closure because of minimal price impact from trades are associated with potentially large welfare gains from closing such gaps, whereas environments in which price gaps can be easily closed are associated with minimal social gains from arbitrage. Hence, the welfare gains from closing arbitrage gaps in dual-listed companies are directly tied to the liquidity of each of the dual-listed shares, as well, of course, to the magnitude of the deviations from parity.

Finally, it is worth highlighting that an intuitive prerequisite for finding large gains from closing arbitrage gaps in the case of dual-listed companies is a balanced distribution of trading volume across the two venues. For instance, consider a hypothetical stock that represents one of the largest listed companies on an exchange: if its underpriced dual-listed counterpart is thinly traded by comparison, any arbitrage gap can be closed with only a small transaction because the leg of the trade transacting in the illiquid stock would trigger a swift price change. Similarly, if one venue represents the primary trading venue, while the other venue is a foreign listing with limited trading volumes, then even extreme deviations from parity need not generate large welfare implications. On the other hand, if a company is dual-listed in two major trading venues that feature significant liquidity (like Royal Dutch/Shell), then the welfare implications can be surprisingly large, in particular if the company possesses an outsized market capitalization.

7 Conclusion

This paper shows theoretically and in practice that arbitrage gaps and measures of price impact are sufficient to compute the direct welfare gains associated with closing arbitrage gaps. The approach introduced in this paper can be applied to any environment in which violations of no-arbitrage conditions exist. Going forward, our results show that understanding the relation between quantities traded and asset prices (price impact) is as important as identifying arbitrage gaps for the purpose of understanding the value of arbitrage. While our welfare assessments are exact in models in which there are strict arbitrage opportunities, one may conjecture that our results remain approximately valid for quasi-arbitrages. Extending our approach to those situations is a fascinating topic for further research.

References

- Almgren, Robert, Chee Thum, Emmanuel Hauptmann, and Hong Li. 2005. “Direct Estimation of Equity Market Impact.” *Risk*, 18(7): 58–62.
- Alvarez, Fernando, and Urban J Jermann. 2004. “Using Asset Prices to Measure the Cost of Business Cycles.” *Journal of Political Economy*, 112(6): 1223–1256.
- Amador, Manuel, Javier Bianchi, Luigi Bocola, and Fabrizio Perri. 2020. “Exchange Rate Policies at the Zero Lower Bound.” *The Review of Economic Studies*, 87(4): 1605–1645.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. “New Trade Models, Same Old Gains?” *American Economic Review*, 102(1): 94–130.
- BIS. 2019. “Global Foreign Exchange Market turnover in 2019.” *Triennial Central Bank Survey*.
- Bouchaud, Jean-Philippe. 2010. “Price Impact.” *Encyclopedia of Quantitative Finance*.
- Boyarchenko, Nina, Thomas M Eisenbach, Pooja Gupta, Or Shachar, and Peter Van Tassel. 2020. “Bank-Intermediated Arbitrage.” *Working Paper*.
- Burghardt, Galen, Jerry Hanweck, and Lauren Lei. 2006. “Measuring Market Impact and Liquidity.” *The Journal of Trading*, 1(4): 70–84.
- Campbell, John Y. 2017. *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press.
- Cochrane, John. 2005. *Asset Pricing: Revised Edition*. Princeton University Press.
- Dávila, Eduardo, and Andreas Schaab. 2021. “Welfare Assessments with Heterogeneous Individuals.” *Working Paper*.
- Dávila, Eduardo, and Anton Korinek. 2018. “Pecuniary Externalities in Economies with Financial Frictions.” *The Review of Economic Studies*, 85(1): 352–395.
- De Jong, Abe, Leonard Rosenthal, and Mathijs A Van Dijk. 2009. “The Risk and Return of Arbitrage in Dual-Listed Companies.” *Review of Finance*, 13(3): 495–520.
- Dixit, A., and V. Norman. 1980. *Theory of International Trade: A Dual, General Equilibrium Approach*. Cambridge Economic Handbooks, Cambridge University Press.
- Donier, Jonathan, Julius Bonart, Iacopo Mastromatteo, and Jean-Philippe Bouchaud. 2015. “A Fully Consistent, Minimal Model for Non-linear Market Impact.” *Quantitative finance*, 15(7): 1109–1121.
- Duffie, Darrell. 2001. *Dynamic Asset Pricing Theory, Third Edition*. Princeton University Press.
- Duffie, Darrell, and Arvind Krishnamurthy. 2016. “Passthrough Efficiency in the Fed’s New Monetary Policy Setting.” 1815–1847.
- Duffie, Darrell, and Bruno Strulovici. 2012. “Capital Mobility and Asset Pricing.” *Econometrica*, 80(6): 2469–2509.
- Du, Wenxin, Alexander Tepper, and Adrien Verdelhan. 2018. “Deviations from Covered Interest Rate Parity.” *The Journal of Finance*, 73(3): 915–957.
- Du, Wenxin, and Jesse Schreger. 2021. “CIP Deviations, the Dollar, and Frictions in International Capital Markets.” *NBER Working Paper*.
- Fardeau, Vincent. 2016. “Dynamic Strategic Arbitrage.” *Working Paper*.

- Farhi, Emmanuel, and Iván Werning.** 2016. “A theory of macroprudential policies in the presence of nominal rigidities.” *Econometrica*, 84(5): 1645–1704.
- Federal Reserve.** 2020. “Report on the Federal Reserve’s Balance Sheet.”
- Frazzini, Andrea, Ronen Israel, and Tobias J Moskowitz.** 2018. “Trading Costs.” *Working Paper*.
- Froot, K.A., and E. Dabora.** 1999. “How Are Stock Prices Affected by the Location of Trade?” *Journal of Financial Economics*, 53(2): 189–216.
- Gabaix, Xavier, and Ralph Koijen.** 2021. “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis.” *NBER Working Paper*.
- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley.** 2006. “Institutional Investors and Stock Market Volatility.” *The Quarterly Journal of Economics*, 121(2): 461–504.
- Gabaix, X., P Gopikrishnan, V Plerou, and HE Stanley.** 2003. “A Theory of Power Law Distributions in Financial Market Fluctuations.” *Nature*, 423: 267–270.
- Geanakoplos, John, and Herakles Polemarchakis.** 1986. “Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market is Incomplete.” *Essays in Honor of Kenneth J. Arrow*, 3: 65–95.
- Gourinchas, P.O., and O. Jeanne.** 2006. “The Elusive Gains from International Financial Integration.” *Review of Economic Studies*, 73(3): 715–741.
- Graves, Daniel.** 2021. “Microstructure Sources of Volatility: The Price Impact Function and Power Laws of Return and Volume in High Frequency Financial Data.” *Working Paper*.
- Gromb, Denis, and Dimitri Vayanos.** 2002. “Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs.” *Journal of Financial Economics*, 66(2-3): 361–407.
- Harberger, Arnold C.** 1964. “The measurement of waste.” *The American Economic Review*, 54(3): 58–76.
- Hébert, Benjamin.** 2020. “Externalities as Arbitrage.” *NBER Working Paper*.
- Kondor, P.** 2009. “Risk in Dynamic Arbitrage: The Price Effects of Convergence Trading.” *The Journal of Finance*, 64(2): 631–655.
- Liu, J., and F.A. Longstaff.** 2004. “Losing Money on Arbitrage: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities.” *Review of Financial Studies*, 17(3): 611–641.
- Lorenzoni, Guido.** 2008. “Inefficient Credit Booms.” *Review of Economic Studies*, 75(3): 809–833.
- Lowenstein, R.** 2001. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House Publishing Group.
- Lucas, Robert E.** 1987. *Models of Business Cycles*. Basil Blackwell New York.
- Oehmke, Martin.** 2009. “Gradual Arbitrage.” *Working Paper*.
- Saez, Emmanuel, and Stefanie Stantcheva.** 2016. “Generalized Social Marginal Welfare Weights for Optimal Tax Theory.” *American Economic Review*, 106(1): 24–45.
- Shleifer, A., and R. Vishny.** 1997. “The Limits of Arbitrage.” *The Journal of Finance*, 52(1): 35–55.

- Shleifer, Andrei.** 2000. *Inefficient Markets: An Introduction to Behavioral Finance (Clarendon Lectures in Economics)*. Oxford University Press, USA.
- Xiong, W.** 2001. “Convergence Trading with Wealth Effects: An Amplification Mechanism in Financial Markets.” *Journal of Financial Economics*, 62(2): 247–292.

APPENDIX

A Proofs and Derivations: Section 3

Proof of Proposition 1. (Marginal value of arbitrage)

Proof. a) The indirect utility of an investor $i = \{A, B\}$ is given by

$$V^i(p^i) = u_i(c_0^{i*}) + \beta_i u_i(c_1^{i*}), \quad (\text{A.1})$$

where c_0^{i*} and c_1^{i*} are the consumption allocations implied by the optimal portfolio choice of the investor which satisfies the first order condition

$$p^i = \frac{\beta_i u'_i(c_1^{i*})}{u'_i(c_0^{i*})} d_1.$$

From the market clearing condition, we have that p^i is a function of m . Therefore, differentiating the indirect utility of investors in Equation (A.1) and applying the envelope theorem imply that

$$\begin{aligned} \frac{dV^i}{dm} &= -u'_i(c_0^{i*}) \frac{dp^i}{dm} \Delta q_0^i + \underbrace{(-u'_i(c_0^{i*}) p^i + \beta_i u'_i(c_1^{i*}) d_1)}_{=0} \frac{dq_0^i}{dm} \\ &= -u'_i(c_0^{i*}) \frac{dp^i}{dm} \Delta q_0^i. \end{aligned}$$

Defining $\lambda_0^i = u'_i(c_0^{i*})$ and using the market clearing condition in market i yield the first result in a).

The utility of arbitrageurs is given by

$$V^\alpha(m, p^A, p^B) = (p^B - p^A) m.$$

Differentiating this expression with respect to m and taking into account that the equilibrium prices depend on m through the market clearing conditions, we have

$$\frac{dV^\alpha}{dm} = p^B - p^A + \left(\frac{dp^B}{dm} - \frac{dp^A}{dm} \right) m.$$

Note that $\lambda_0^\alpha = 1$ proves the second result in a).

b) The marginal social value of arbitrage is given by

$$\frac{dW}{dm} = \frac{dV^A}{dm} \frac{1}{\lambda_0^A} + \frac{dV^B}{dm} \frac{1}{\lambda_0^B} + \frac{dV^\alpha}{dm} \frac{1}{\lambda_0^\alpha}.$$

Then, using the results from part a) of this proposition, we have

$$\begin{aligned} \frac{dW}{dm} &= p^B - p^A + \left(\frac{dp^B}{dm} - \frac{dp^A}{dm} \right) m - \frac{dp^A}{dm} \Delta q_0^A - \frac{dp^B}{dm} \Delta q_0^B \\ &= p^B - p^A - \frac{dp^A}{dm} \underbrace{(\Delta q_0^A + m)}_{=0} - \frac{dp^B}{dm} \underbrace{(\Delta q_0^B - m)}_{=0} \\ &= p^B - p^A, \end{aligned}$$

where the last line exploits the market clearing conditions $\Delta q_0^A + m = 0$ and $\Delta q_0^B - m = 0$. This proves the result. \square

Proof of Proposition 2. (Total value of arbitrage)

Proof. a) From the investors' first order conditions and market clearing, we have the equilibrium prices p^A and p^B are continuous in m . Then, $\frac{\partial W}{\partial m}$ is a continuous real function of m and using the Fundamental Theorem of Calculus we have that

$$W(m^*) - W(m_0) = \int_{m_0}^{m^*} W'(m) dm.$$

Using the result in Proposition 1b) proves the result.

b) From part a), we have that

$$W(m^*) - W(m_0) = \int_{m_0}^{m^*} \mathcal{G}_{BA}(\tilde{m}) d\tilde{m}.$$

Using the Fundamental Theorem of Calculus we can express the arbitrage gap at m as

$$\mathcal{G}_{BA}(\tilde{m}) \equiv p^B(\tilde{m}) - p^A(\tilde{m}) = p^B(0) - p^A(0) + \int_0^{\tilde{m}} \left(\frac{dp^B(\hat{m})}{d\hat{m}} - \frac{dp^A(\hat{m})}{d\hat{m}} \right) d\hat{m}.$$

Therefore, measures of the current arbitrage gap $p^B(m_0) - p^A(m_0)$ and price impacts $\frac{dp^A(\tilde{m})}{d\tilde{m}}$ and $\frac{dp^B(\tilde{m})}{d\tilde{m}}$ are sufficient to exactly compute the social value of arbitrage. \square

Proof of Proposition 3. (Market liquidity and the value of arbitrage)

Proof. From Proposition 2 we have that the social value of arbitrage is

$$W(m^*) - W(m_0) = \int_{m_0}^{m^*} \mathcal{G}_{BA}(\tilde{m}) d\tilde{m}, \tag{A.2}$$

where

$$\mathcal{G}_{BA}(\tilde{m}) \equiv p^B(\tilde{m}) - p^A(\tilde{m}) = p^B(0) - p^A(0) + \int_0^{\tilde{m}} \left(\frac{dp^B(\hat{m})}{d\hat{m}} - \frac{dp^A(\hat{m})}{d\hat{m}} \right) d\hat{m}.$$

Price impact is higher when the arbitrageur moves the price against him more, i.e., when $\left| \frac{dp^i(m)}{dm} \right|$ is higher. Therefore, since

$$\frac{dp^A(m)}{dm} > 0 \quad \text{and} \quad \frac{dp^B(m)}{dm} < 0,$$

more liquid markets have lower $\frac{dp^A(m)}{dm}$ and higher $\frac{dp^B(m)}{dm}$, which imply a higher $\mathcal{G}_{BA}(m)$. Therefore, our results follows from the expression for the social value of arbitrage in Equation (A.2). \square

B Proofs and Derivations: Section 4

Matrix definitions

Here we provide explicit definitions of each of the elements of Equation (20). The $|J| \times 1$ vector of arbitrageurs' trading directions at history s^t , $x_t^\alpha(s^t)$, is given by

$$x_t^\alpha(s^t) = \begin{pmatrix} x_t^{\alpha 1}(s^t) \\ \vdots \\ x_t^{\alpha j}(s^t) \\ \vdots \\ x_t^{\alpha |J|}(s^t) \end{pmatrix}.$$

The definition of $x_t^\alpha(s^t)$ allows us to express the matrix $X_{t|s^{t-1}}^\alpha$ of dimension $|S| \times |J|$, which collects the vectors $x_t^\alpha(s^t)$ that share a common predecessor s^{t-1} , as follows:

$$X_{t|s^{t-1}}^\alpha = \begin{pmatrix} x_t^\alpha(s^{t-1}, \underline{s}_t)' \\ \vdots \\ x_t^\alpha(s^{t-1}, s_t)' \\ \vdots \\ x_t^\alpha(s^{t-1}, \bar{s}_t)' \end{pmatrix} = \begin{pmatrix} x_t^{\alpha 1}(s^{t-1}, \underline{s}_t) & \cdots & x_t^{\alpha j}(s^{t-1}, \underline{s}_t) & \cdots & x_t^{\alpha |J|}(s^{t-1}, \underline{s}_t) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_t^{\alpha 1}(s^{t-1}, s_t) & & x_t^{\alpha j}(s^{t-1}, s_t) & & x_t^{\alpha |J|}(s^{t-1}, s_t) \\ \vdots & & \vdots & \ddots & \vdots \\ x_t^{\alpha 1}(s^{t-1}, \bar{s}_t) & \cdots & x_t^{\alpha j}(s^{t-1}, \bar{s}_t) & \cdots & x_t^{\alpha |J|}(s^{t-1}, \bar{s}_t) \end{pmatrix},$$

where we denote the first and last realizations of the state s_t by \underline{s}_t and \bar{s}_t , respectively. Analogously to $X_{t|s^{t-1}}^\alpha$, we define the matrices $D_{t|s^{t-1}}$ and $P_{t|s^{t-1}}$ of dimension $|S| \times |J|$, which collect asset payoffs and asset prices for all assets as follows:

$$D_{t|s^{t-1}} = \begin{pmatrix} d_t^1(s^{t-1}, \underline{s}_t) & \cdots & d_t^j(s^{t-1}, \underline{s}_t) & \cdots & d_t^{|J|}(s^{t-1}, \underline{s}_t) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ d_t^1(s^{t-1}, s_t) & & d_t^j(s^{t-1}, s_t) & & d_t^{|J|}(s^{t-1}, s_t) \\ \vdots & & \vdots & \ddots & \vdots \\ d_t^1(s^{t-1}, \bar{s}_t) & \cdots & d_t^j(s^{t-1}, \bar{s}_t) & \cdots & d_t^{|J|}(s^{t-1}, \bar{s}_t) \end{pmatrix}$$

and

$$P_{t|s^{t-1}} = \begin{pmatrix} p_t^1(s^{t-1}, \underline{s}_t) & \cdots & p_t^j(s^{t-1}, \underline{s}_t) & \cdots & p_t^{|J|}(s^{t-1}, \underline{s}_t) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ p_t^1(s^{t-1}, s_t) & & p_t^j(s^{t-1}, s_t) & & p_t^{|J|}(s^{t-1}, s_t) \\ \vdots & & \vdots & \ddots & \vdots \\ p_t^1(s^{t-1}, \bar{s}_t) & \cdots & p_t^j(s^{t-1}, \bar{s}_t) & \cdots & p_t^{|J|}(s^{t-1}, \bar{s}_t) \end{pmatrix}.$$

Note that $D_0 = P_0 = \mathbf{0}$. Finally, the matrix $M_{t|s^{t-1}}$ of dimension $|S| \times |S|$ is a diagonal matrix whose diagonal elements correspond to the scale of the arbitrage trade for each history s^t that shares a common

predecessor history s^{t-1} , as follows:

$$M_{t|s^{t-1}} = \begin{pmatrix} m_t(s^{t-1}, \underline{s}_t) & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & 0 & & \vdots \\ 0 & 0 & m_t(s^{t-1}, s_t) & 0 & 0 \\ \vdots & & 0 & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & m_t(s^{t-1}, \bar{s}_t) \end{pmatrix}.$$

Here we define the partition of $x_{t-1}^\alpha(s^{t-1})$:

$$x_{t-1}^\alpha(s^{t-1}) = \begin{pmatrix} x_{t-1}^{\alpha 1}(s^{t-1}) \\ \tilde{x}_{t-1}^\alpha(s^{t-1}) \end{pmatrix} = \begin{pmatrix} 1 \\ \tilde{x}_{t-1}^\alpha(s^{t-1}) \end{pmatrix}. \quad (\text{A.3})$$

We analogously define partitions of the matrices $D_{t|s^{t-1}}$ and $P_{t|s^{t-1}}$ as follows:

$$\begin{aligned} D_{t|s^{t-1}} &= \begin{pmatrix} D_{t|s^{t-1}}^1 & \cdots & D_{t|s^{t-1}}^j & \cdots & D_{t|s^{t-1}}^{|J|} \end{pmatrix} = \begin{pmatrix} D_{t|s^{t-1}}^1 & \tilde{D}_{t|s^{t-1}} \end{pmatrix} \\ P_{t|s^{t-1}} &= \begin{pmatrix} P_{t|s^{t-1}}^1 & \cdots & P_{t|s^{t-1}}^j & \cdots & P_{t|s^{t-1}}^{|J|} \end{pmatrix} = \begin{pmatrix} P_{t|s^{t-1}}^1 & \tilde{P}_{t|s^{t-1}} \end{pmatrix}, \end{aligned}$$

where $D_{t|s^{t-1}}^1$ and $P_{t|s^{t-1}}^1$ denote the first columns of the matrices $D_{t|s^{t-1}}$ and $P_{t|s^{t-1}}$ and $\tilde{D}_{t|s^{t-1}}$ and $\tilde{P}_{t|s^{t-1}}$ denote the matrix formed by the second to last columns of the matrices $D_{t|s^{t-1}}$ and $P_{t|s^{t-1}}$. Hence, we can explicitly write

$$D_{t|s^{t-1}}^1 = \begin{pmatrix} d_{t|s^{t-1}}^1(\underline{s}^t) \\ \vdots \\ d_{t|s^{t-1}}^1(s^t) \\ \vdots \\ d_{t|s^{t-1}}^1(\bar{s}^t) \end{pmatrix} \quad \text{and} \quad \tilde{D}_{t|s^{t-1}} = \begin{pmatrix} d_{t|s^{t-1}}^2(\underline{s}^t) & \cdots & d_{t|s^{t-1}}^j(\underline{s}^t) & \cdots & d_{t|s^{t-1}}^{|J|}(\underline{s}^t) \\ \vdots & \ddots & & & \\ d_{t|s^{t-1}}^2(s^t) & & d_{t|s^{t-1}}^j(s^t) & & d_{t|s^{t-1}}^{|J|}(s^t) \\ \vdots & & & \ddots & \\ d_{t|s^{t-1}}^2(\bar{s}^t) & \cdots & d_{t|s^{t-1}}^j(\bar{s}^t) & \cdots & d_{t|s^{t-1}}^{|J|}(\bar{s}^t) \end{pmatrix}$$

and

$$P_{t|s^{t-1}}^1 = \begin{pmatrix} p_{t|s^{t-1}}^1(\underline{s}^t) \\ \vdots \\ p_{t|s^{t-1}}^1(s^t) \\ \vdots \\ p_{t|s^{t-1}}^1(\bar{s}^t) \end{pmatrix} \quad \text{and} \quad \tilde{P}_{t|s^{t-1}} = \begin{pmatrix} p_{t|s^{t-1}}^2(\underline{s}^t) & \cdots & p_{t|s^{t-1}}^j(\underline{s}^t) & \cdots & p_{t|s^{t-1}}^{|J|}(\underline{s}^t) \\ \vdots & \ddots & & & \\ p_{t|s^{t-1}}^2(s^t) & & p_{t|s^{t-1}}^j(s^t) & & p_{t|s^{t-1}}^{|J|}(s^t) \\ \vdots & & & \ddots & \\ p_{t|s^{t-1}}^2(\bar{s}^t) & \cdots & p_{t|s^{t-1}}^j(\bar{s}^t) & \cdots & p_{t|s^{t-1}}^{|J|}(\bar{s}^t) \end{pmatrix}.$$

Finally, we partition the matrix $X_{t|s^{t-1}}^\alpha$ as

$$X_{t|s^{t-1}}^\alpha = \begin{pmatrix} x_t^{\alpha 1}(s^{t-1}, \underline{s}_t) & \cdots & x_t^{\alpha j}(s^{t-1}, \underline{s}_t) & \cdots & x_t^{\alpha |J|}(s^{t-1}, \underline{s}_t) \\ \vdots & \ddots & & & \vdots \\ x_t^{\alpha 1}(s^{t-1}, s_t) & & x_t^{\alpha j}(s^{t-1}, s_t) & & x_t^{\alpha |J|}(s^{t-1}, s_t) \\ \vdots & & & \ddots & \vdots \\ x_t^{\alpha 1}(s^{t-1}, \bar{s}_t) & \cdots & x_t^{\alpha j}(s^{t-1}, \bar{s}_t) & \cdots & x_t^{\alpha |J|}(s^{t-1}, \bar{s}_t) \end{pmatrix} = \begin{pmatrix} X_{t|s^{t-1}}^{\alpha, 1} & \tilde{X}_{t|s^{t-1}}^\alpha \end{pmatrix}.$$

Therefore, $\tilde{x}_{t-1}^\alpha (s^{t-1})$ is a vector of dimension $(|J| - 1) \times 1$, while $\tilde{D}_{t|s^{t-1}}$, $\tilde{P}_{t|s^{t-1}}$ and $\tilde{X}_{t|s^{t-1}}^\alpha$ are matrices of dimension $|S| \times (|J| - 1)$.

Derivative matrix definitions

We also define the following vectors and matrices of derivatives:

$$\begin{aligned} \frac{d\tilde{x}_t^\alpha}{dm_\tau} &= \left(\frac{dx_t^{\alpha 2}(s^t)}{dm_\tau}, \dots, \frac{dx_t^{\alpha j}(s^t)}{dm_\tau}, \dots, \frac{dx_t^{\alpha |J|}(s^t)}{dm_\tau} \right)', \\ \frac{\partial \tilde{P}_{t+1|s^t}^\alpha}{\partial m_\tau} &= \left(\frac{\partial p_{t+1|s^t}^{\alpha 2}(s^{t+1})}{\partial m_\tau}, \dots, \frac{\partial p_{t+1|s^t}^{\alpha j}(s^{t+1})}{\partial m_\tau}, \dots, \frac{\partial p_{t+1|s^t}^{\alpha |J|}(s^{t+1})}{\partial m_\tau} \right)', \\ \frac{d\tilde{X}_{t+1|s^t}^{\alpha j}}{dm_\tau} &= \left(\frac{d\tilde{x}_{t+1|s^t}^{\alpha j}(s^{t+1})}{dm_\tau}, \dots, \frac{d\tilde{x}_{t+1|s^t}^{\alpha j}(s^{t+1})}{dm_\tau}, \dots, \frac{d\tilde{x}_{t+1|s^t}^{\alpha j}(s^{t+1})}{dm_\tau} \right)', \end{aligned}$$

$$\frac{\partial \tilde{x}_t^\alpha}{\partial P_{t+1|s^t}^j} = \begin{pmatrix} \frac{\partial x_t^{\alpha 2}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha 2}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha 2}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial x_t^{\alpha j}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & & \frac{\partial x_t^{\alpha j}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & & \frac{\partial x_t^{\alpha j}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} \\ \vdots & & \vdots & \ddots & \vdots \\ \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial p_{t+1|s^t}^j(s^{t+1})} \end{pmatrix},$$

and

$$\frac{\partial \tilde{x}_t^\alpha}{\partial X_{t+1|s^t}^j} = \begin{pmatrix} \frac{\partial x_t^{\alpha 2}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha 2}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha 2}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial x_t^{\alpha j}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & & \frac{\partial x_t^{\alpha j}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & & \frac{\partial x_t^{\alpha j}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} \\ \vdots & & \vdots & \ddots & \vdots \\ \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} & \dots & \frac{\partial x_t^{\alpha |J|}(s^t)}{\partial x_{t+1|s^t}^j(s^{t+1})} \end{pmatrix}.$$

Dynamic effects of arbitrage

As in the static model, a change in the amount of arbitrage at date 0 affects the equilibrium prices at 0. However, when the economy is dynamic, changes in the amount arbitrated at date 0 also affect all future prices.

First, note that the market clearing conditions in the general model can be written as

$$\sum_i q_t^{ij}(s^t) + x_t^{\alpha j}(s^t) m_t(s^t) - x_{t-1}^{\alpha j} m_{t-1}(s^{t-1}) = 0, \quad \forall j, \forall t \in \{0, \dots, T-1\}, \forall s^t,$$

where $x_t^{\alpha 1}(s^t) = 1, \forall t, \forall s^t$. Therefore, the equilibrium price $p_t^j(s^t)$ depends on $m_t(s^t)$ directly and indirectly through the effect $m_t(s^t)$ may have on $\Delta q_t^{ij}(s^t)$ and on the arbitrage directions $x_{t-1}^{\alpha j}$ and $x_t^{\alpha j}$.

The optimality conditions of investors determine the demand function $\Delta q_t^{ij}(s^t)$. These optimality

conditions are defined by standard Euler equations:

$$p_t^j(s^t) = \sum_{s^{t+1}|s^t} \pi(s^{t+1}|s^t) \frac{\beta_i u_{t+1}^{i'}(c_{t+1}^i(s^{t+1}))}{u_t^{i'}(c_t^i(s^t))} \left(d_{t+1}^j(s^{t+1}) + p_{t+1}^j(s^{t+1}) \right) + \eta^i \frac{\partial \Phi^i}{\partial q_t^{ij}(s^t)} = 0, \forall j, \forall t \in \{0, \dots, T-1\}, \forall s^t, \quad (\text{A.4})$$

where $\pi(s^{t+1}|s^t)$ denotes a conditional probability, η denotes a vector of Lagrange multipliers, and where $p_T^j(s^T) = 0, \forall j$. Iterating the Euler equations forward it is easy to see that $\Delta q_t^{ij}(s^t)$ is a function of all future prices conditional on history s^t . Moreover, the Euler equation depends on the current marginal utility $u_t^{i'}(c_t^i(s^t))$ which implies that $\Delta q_t^{ij}(s^t)$ also depends on the asset holdings of investor i at date t , which depend on all past asset prices along history s^t , including p_0 . Since p_0 depends on m_0 at least directly through the market clearing conditions at date 0, all of the investor's demand functions, and hence the equilibrium prices, depend on m_0 .

The amount of arbitrage at date 0, m_0 , has an additional effect through the changes in the directions of future arbitrage trades. From the budget constraint of the arbitrageurs we have that

$$(D_{t|s^{t-1}} + P_{t|s^{t-1}}) x_{t-1}^\alpha(s^{t-1}) m_{t-1}(s^{t-1}) = M_{t|s^{t-1}} \text{diag} \left(P_{t|s^{t-1}} \left(X_{t|s^{t-1}}^\alpha \right)' \right). \quad (\text{A.5})$$

Note that arbitrage direction at date $t-1$, $x_{t-1}^\alpha(s^{t-1})$, depends on the arbitrage directions and prices at date t and on the current arbitrage trade scale m_{t-1} . Iterating this equation forward it is easy to see that the future arbitrage directions in $X_{t|s^{t-1}}^\alpha$ depend on m_0 through all future equilibrium prices following history s^{t-1} . The following lemma characterizes the effect of m_0 on arbitrage directions.

Moreover, under the presumption that $\tilde{D}_{t+1|s^t} + \tilde{P}_{t+1|s^t}$ is invertible, we can write \tilde{x}_t^α as

$$\tilde{x}_t^\alpha(s^t) = - \underbrace{(\tilde{D}_{t+1|s^t} + \tilde{P}_{t+1|s^t})^{-1}}_{\text{myopic}} \left(D_{t+1|s^t}^1 + P_{t+1|s^t}^1 \right) + \underbrace{(\tilde{D}_{t+1|s^t} + \tilde{P}_{t+1|s^t})^{-1} \frac{M_{t+1|s^t}}{m_t(s^t)} \left(P_{t+1|s^t}^1 + \text{diag} \left(\tilde{P}_{t+1|s^t} \tilde{X}_{t+1|s^t}^\alpha \right)' \right)}_{\text{continuation}} \forall s^t, \forall t \in \{0, \dots, T-1\}, \quad (\text{A.6})$$

since $x_t^{\alpha 1} = 1$ for all t . Therefore, $\tilde{x}_t^\alpha(s^t)$ is a function of $\{m_t(s^t), M_{t+1|s^t}, P_{t+1|s^t}, \tilde{X}_{t+1|s^t}\}$ and

$$\frac{d\tilde{x}_t^\alpha(s^t)}{dm_\tau} = \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial m_\tau} + \sum_j \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial P_{t+1|s^t}^j} \frac{dP_{t+1|s^t}^j}{dm_\tau} + \sum_j \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial \tilde{X}_{t+1|s^t}^{\alpha j}} \frac{d\tilde{X}_{t+1|s^t}^{\alpha j}}{dm_\tau}.$$

Lemma 1. (Effects on arbitrage directions) *The marginal effect of the scale of the arbitrage trade at date 0 m_0 on the arbitrage directions is given by*

$$\frac{d\tilde{x}_t^\alpha(s^t)}{dm_0} = \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial m_0} + \sum_j \sum_{h=0}^{T-t-1} \sum_{s^{t+h}|s^t} \delta_{t,t+h}(s^{t+h}) \frac{\partial \tilde{x}_{t+h+1}^\alpha(s^{t+h+1})}{\partial P_{t+h|s^{t+h}}^j} \frac{dP_{t+h+1|s^{t+h}}^j}{dm_0},$$

where

$$\delta_{t,t+h}(s^{t+h}) = \begin{cases} 1 & \text{if } h = 0 \\ \prod_{n=0}^{h-1} \frac{\partial \tilde{x}_{t+n}^\alpha(s^{t+n})}{\partial \tilde{x}_{t+n+1}^\alpha(s^{t+n+1})} & \text{if } h \neq 0 \end{cases}.$$

Proof. First, note that our normalization implies that $x_t^{\alpha,1} = 1$ for all t . The only relevant effect for our purposes is the effect of m_0 on \tilde{x}_t^α . From the budget constraint of the arbitrageur in Equation (A.5) we have that the total derivative of \tilde{x}_t^α with respect to m_0 , which corresponds to a vector of dimension $(|J| - 1) \times 1$, can be expressed as follows:

$$\frac{d\tilde{x}_t^\alpha(s^t)}{dm_0} = \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0} + \sum_j \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial P_{t+1|s^t}^j} \frac{dP_{t+1|s^t}^j}{dm_0} + \sum_j \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial \tilde{X}_{t+1|s^t}^{\alpha j}} \frac{d\tilde{X}_{t+1|s^t}^{\alpha j}}{dm_0}, \quad (\text{A.7})$$

where $\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0}$ is a vector of dimension $(|J| - 1) \times 1$, $\frac{dP_{t+1|s^t}^j}{dm_0}$ and $\frac{d\tilde{X}_{t+1|s^t}^{\alpha j}}{dm_0}$ are vectors of dimension $|S| \times 1$, and $\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial P_{t+1|s^t}^j}$ and $\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial \tilde{X}_{t+1|s^t}^{\alpha j}}$ are matrices of dimension $(|J| - 1) \times |S|$ as explicitly characterized in the previous section.

Note that

$$\sum_j \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial \tilde{X}_{t+1|s^t}^{\alpha j}} \frac{d\tilde{X}_{t+1|s^t}^{\alpha j}}{dm_0} = \sum_{s^{t+1}|s^t} \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial \tilde{x}_{t+1}^\alpha(s^{t+1})} \frac{d\tilde{x}_{t+1}^\alpha(s^{t+1})}{dm_0}.$$

Moreover, we characterize

$$\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0} = \begin{cases} 0 & \forall t > 0 \\ -\frac{1}{m_0} (\tilde{D}_1 + \tilde{P}_1)^{-1} \frac{M_1}{m_0} (P_1^1 + \text{diag}(\tilde{P}_1 \tilde{X}_1^{\alpha'}) & \text{for } t = 0 \end{cases} \quad (\text{A.8})$$

Then, using that $\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0} = 0$ for all $t > 0$, in Equation (A.7) we have

$$\frac{d\tilde{x}_{t+1}^\alpha(s^{t+1})}{dm_0} = \sum_j \frac{\partial\tilde{x}_{t+1}^\alpha(s^{t+1})}{\partial P_{t+2|s^{t+1}}^j} \frac{dP_{t+2|s^{t+1}}^j}{dm_0} + \sum_{s^{t+2}|s^{t+1}} \frac{\partial\tilde{x}_{t+1}^\alpha(s^{t+1})}{\partial \tilde{x}_{t+2}^\alpha(s^{t+2})} \frac{d\tilde{x}_{t+2}^\alpha(s^{t+2})}{dm_0} \quad \forall t.$$

Iterating this expression forward and using that

$$\frac{d\tilde{x}_T^\alpha(s^T)}{dm_0} = \frac{\partial\tilde{x}_T^\alpha(s^T)}{\partial m_0} = 0$$

we get

$$\frac{d\tilde{x}_t^\alpha(s^t)}{dm_0} = \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0} + \sum_j \left(\frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial P_{t+1|s^t}^j} \frac{dP_{t+1|s^t}^j}{dm_0} + \sum_{h=1}^{T-t-1} \sum_{s^{t+h}|s^t} \left(\prod_{n=0}^{h-1} \frac{\partial\tilde{x}_{t+n}^\alpha(s^{t+n})}{\partial \tilde{x}_{t+n+1}^\alpha(s^{t+n+1})} \right) \frac{\partial\tilde{x}_{t+h}^\alpha(s^{t+h})}{\partial P_{t+h+1|s^{t+h}}^j} \frac{dP_{t+h+1|s^{t+h}}^j}{dm_0} \right)$$

or alternatively

$$\frac{d\tilde{x}_t^\alpha(s^t)}{dm_0} = \frac{\partial\tilde{x}_t^\alpha(s^t)}{\partial m_0} + \sum_j \sum_{h=0}^{T-1} \sum_{s^h} \delta_{0,h}(s^h) \frac{\partial\tilde{x}_h^\alpha(s^h)}{\partial P_{h+1|s^h}^j} \frac{dP_{h+1|s^h}^j}{dm_0},$$

where

$$\delta_{t,t+h}(s^{t+h}) = \begin{cases} 1 & \text{if } h = 0 \\ \prod_{n=0}^{h-1} \frac{\partial\tilde{x}_{t+n}^\alpha(s^{t+n})}{\partial \tilde{x}_{t+n+1}^\alpha(s^{t+n+1})} & \text{if } h \neq 0 \end{cases}.$$

□

An investor i 's welfare in an arbitrage equilibrium indexed by $\{m_t\}_{t \in \{0, \dots, T-1\}}$ is given by

$$V^i \left(\{m_t\}_{t \in \{0, \dots, T-1\}} \right) = u_0^i(c_0^i) + \sum_{t=1}^T \beta_t^i \sum_{s^t} \pi_t(s^t) u_t^i(c_t^i(s^t)),$$

where

$$c_t^i(s^t) = n_t^i(s^t) + \sum_j \left(d_t^j(s^t) + p_t^j(s^t) \right) q_{t-1}^{ij}(s^{t-1}) - p_t^j(s^t) q_t^{ij}(s^t), \quad \forall s^t, \forall t.$$

Then,

$$\frac{dV^i}{d\lambda_0^i} = -\frac{dp_0}{dm_\tau} \cdot \Delta q_0^i + \sum_{t=1}^T \sum_{s^t} \pi_t(s^t) \frac{\beta_t^i u_t^{i'}(c_t^i(s^t))}{u_0^{i'}(c_0^i)} \sum_j \frac{dp_t^j(s^t)}{dm_\tau} \Delta q_t^{ij}(s^t)$$

where $\lambda_0^i \equiv u_0^{i'}(c_0^i)$ and $\Delta q_t^{ij}(s^{t-1}, s_t) \equiv q_t^{ij}(s^{t-1}, s_t) - q_{t-1}^{ij}(s^{t-1})$.

Aggregate investor welfare is given by

$$\int \frac{dV^i}{d\lambda_0^i} di = \frac{dp_0}{dm_\tau} \cdot m_0 x_0^\alpha + \int \sum_{t=1}^T \sum_{s^t} \pi_t(s^t) \frac{\beta_t^i u_t^{i'}(c_t^i(s^t))}{u_0^{i'}(c_0^i)} \sum_j \frac{dp_t^j(s^t)}{dm_\tau} \Delta q_t^{ij}(s^t) di, \quad (\text{A.9})$$

where we used the market clearing condition $\int \Delta q_0^i di = m_0 x_0^\alpha$.

Welfare arbitrageurs

Arbitrageurs welfare is given by the arbitrage profit at date 0, i.e.,

$$V^\alpha = -m_0 p_0 \cdot x_0^\alpha.$$

Then,

$$\frac{dV^\alpha}{dm_0} = - \left(p_0 \cdot x_0^\alpha + m_0 \left(x_0^\alpha \cdot \frac{dp_0}{dm_0} + \tilde{p}_0 \cdot \frac{d\tilde{x}_0^\alpha}{dm_0} \right) \right).$$

Using Lemma 1 we can write

$$\frac{dV^\alpha}{dm_0} = - \left(p_0 \cdot x_0^\alpha + m_0 \left(x_0^\alpha \cdot \frac{dp_0}{dm_0} + \tilde{p}_0 \cdot \frac{\partial \tilde{x}_0^\alpha}{\partial m_0} + \tilde{p}_0 \cdot \sum_j \sum_{h=0}^{T-1} \sum_{s^h} \delta_{0,h}(s^h) \frac{\partial \tilde{x}_{h+1}^\alpha(s^{h+1})}{\partial P_{h+1|s^h}^j} \frac{dP_{h+1|s^h}^j}{dm_0} \right) \right).$$

Note that using the expression for \tilde{x}_0^α in Equation (A.6) and Equation (A.8) we have

$$\tilde{x}_0^\alpha + m_0 \frac{\partial \tilde{x}_0^\alpha}{\partial m_0} = - (\tilde{D}_1 + \tilde{P}_1)^{-1} (D_1^1 + P_1^1).$$

Therefore,

$$\frac{dV^\alpha}{dm_0} = -p_0^1 + \tilde{p}_0 \cdot (\tilde{D}_1 + \tilde{P}_1)^{-1} (D_1^1 + P_1^1) - m_0 x_0^\alpha \cdot \frac{dp_0}{dm_0} - m_0 \tilde{p}_0 \cdot \sum_j \sum_{h=0}^{T-1} \sum_{s^h} \delta_{0,h}(s^h) \frac{\partial \tilde{x}_{h+1}^\alpha(s^{h+1})}{\partial P_{h+1|s^h}^j} \frac{dP_{h+1|s^h}^j}{dm_0}. \quad (\text{A.10})$$

Proof of Proposition 4 (Marginal social value of arbitrage: general model)

Using the expressions for investor and arbitrageur welfare in Equations (A.9) and (A.10), respectively, we have that the marginal social welfare gains of arbitrage are

$$\begin{aligned}
\frac{dW}{dm_0} &= \frac{dV^\alpha}{dm_0} + \sum_i \frac{\frac{dV^i}{dm_0}}{\lambda_0^i} & (A.11) \\
&= -p_0^1 + \tilde{p}_0 \cdot (\tilde{D}_1 + \tilde{P}_1)^{-1} (D_1^1 + P_1^1) \\
&\quad - \sum_j \sum_{t=0}^{T-1} \sum_{s^{t+1}} \left(\tilde{p}_0 \cdot \delta_{0,t}(s^t) \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial p_{t+1}^j(s^{t+1})} m_0 + \int \pi_{t+1}(s^{t+1}) \frac{\beta_i^{t+1} u_{t+1}^{i'}(c_{t+1}^i(s^{t+1}))}{u_0^{i'}(c_0^i)} \Delta q_{t+1}^{ij}(s^{t+1}) di \right) \frac{dp_{t+1}^j(s^{t+1})}{dm_0},
\end{aligned}$$

where we define

$$\Xi_{t+1}^j(s^{t+1}) \equiv \tilde{p}_0 \cdot \delta_{0,t}(s^t) \frac{\partial \tilde{x}_t^\alpha(s^t)}{\partial p_{t+1}^j(s^{t+1})} m_0 + \int_i \pi_{t+1}(s^{t+1}) \frac{\beta_i^{t+1} u_{t+1}^{i'}(c_{t+1}^i(s^{t+1}))}{u_0^{i'}(c_0^i)} \Delta q_{t+1}^{ij}(s^{t+1}) di.$$

ONLINE APPENDIX

Section C of this Online Appendix provides explicit microfoundations for a variety of frictions that may limit the size of the arbitrage trade chosen by arbitrageurs in practice. Section D of this Online Appendix includes several special cases of the general model that may be helpful to understand how the results of the baseline model extend to more general scenarios. Sections E through G of this Online Appendix include additional information related to the empirical applications in Sections 5 and 6.

C Limits to Arbitrage: Explicit Microfoundations

We have shown in this paper that it is possible to measure the value of closing arbitrage gaps without having to fully specify the frictions that prevent such gaps from being closed. However, it is useful to understand why arbitrage gaps may emerge in practice. In this section, in the context of the baseline model introduced in Section 2, we formally describe how a variety of frictions can determine the positions that arbitrageurs would choose within our framework if they had the ability to do so.

First, as a benchmark, we describe the case in which arbitrageurs can frictionlessly choose the size of the arbitrage trade m . In that case, the competitive equilibrium features equal prices in both markets and the size of the arbitrage trade is m^* . Next we sequentially consider environments that feature i) trading costs, ii) market power in the arbitrageur sector, iii) short-selling and borrowing constraints, and iv) price-dependent collateral constraints. These frictions limit the size/scale of the arbitrage trade chosen by arbitrageurs and prevent price equalization in equilibrium.

The same set of frictions considered here in the context of the baseline model can also shape the size of future arbitrage gaps, $m_1(s)$, in the general dynamic model in Section 4. Across all of these applications, we use the following notion of competitive equilibrium, which is standard.

Definition. (*Competitive equilibrium*) A competitive equilibrium is defined as a set of consumption allocations, asset holdings, and prices p^A and p^B such that i) investors and arbitrageurs maximize utility subject to their budget constraints, and ii) the asset markets A and B clear.

C.1 Frictionless benchmark

It is clear from $V^\alpha(m)$, defined in Equation (6), that the problem of choosing m by the arbitrageur sector does not have an interior solution whenever there is a price differential between markets A and B . Formally, the optimal size of the arbitrage trade m^α takes in this case the form

$$m^\alpha = \begin{cases} -\infty, & \text{if } p^A > p^B \\ \infty, & \text{if } p^A < p^B. \end{cases}$$

Therefore, as long as arbitrageurs can trade across both markets frictionlessly, the arbitrage equilibrium is identical to an integrated equilibrium in which $p^A(m^*) = p^B(m^*)$ and, to support market clearing, $m^\alpha = m^*$.

C.2 Trading costs

In this case, we consider a scenario in which arbitrageurs face linear trading costs τ^A and τ^B in markets A and B , respectively. That is, when trading in market i , an arbitrageur that trades a quantity q faces a cost

$\tau^i |q|$. In this case, the profit of an arbitrageur who buys m units of asset A and sells m units of asset B is

$$V^\alpha(m) = (p^B(m) - p^A(m))m - (\tau^A + \tau^B)|m|.$$

If the arbitrageur sector is competitive, the total amount arbitrated between markets A and B is such that the profits from the arbitrage activity are zero. In this case, the arbitrage gap satisfies

$$p^B(m) - p^A(m) = \tau^A + \tau^B.$$

C.3 Strategic arbitrageurs

Here we consider a scenario with a finite number of arbitrageurs, indexed by $h = \{1, \dots, H\}$, who take into account their price impact when trading in each market. Then, an arbitrageur h who buys m^h units of asset A and sells m^{-h} units of asset B , when the total amount arbitrated by the other arbitrageurs is m^{-h} , solves the following problem:

$$V^\alpha(m^h, m^{-h}) = (p^B(m^h + m^{-h}) - p^A(m^h + m^{-h}))m^h.$$

In a symmetric equilibrium with H strategic arbitrageurs, the optimal amount arbitrated m^α is such that

$$p^B(m^\alpha) - p^A(m^\alpha) = \left(\frac{\partial p^A(m^\alpha)}{\partial m} - \frac{\partial p^B(m^\alpha)}{\partial m} \right) \frac{m^\alpha}{H}.$$

Note that when the arbitrageur sector is competitive, i.e., if $H \rightarrow \infty$, $m^\alpha \Rightarrow m^*$, and the arbitrage gap is 0.

Finally, note that in the particular case in which there is a single monopolistic arbitrageur, using the notation of Section 3, the optimal choice of m is given by

$$m^\alpha = \frac{p^B(m^\alpha) - p^A(m^\alpha)}{-\left(\frac{dp^B(m^\alpha)}{dm} - \frac{dp^A(m^\alpha)}{dm} \right)}.$$

C.4 Short sales/Borrowing constraints

In this case, suppose that the arbitrageurs face a short selling constraint \underline{m} such that $0 < m \leq \underline{m}$. Then, if the arbitrageur sector is competitive, the total amount arbitrated is $\min\{\underline{m}, m^*\}$, where m^* is the gap-closing arbitrage trade. Then, the arbitrage gap is

$$p^B(\min\{\underline{m}, m^*\}) - p^A(\min\{\underline{m}, m^*\}) \geq 0.$$

If the short selling constraint is binding, i.e., if $\underline{m} < m^*$, then the arbitrage gap is positive.

C.5 Price-dependent collateral constraints

Finally, we study the simplest environment that yields price-dependent collateral constraints. We assume that the arbitrage trade takes place in two stages. In the first stage, the arbitrageur gets paid $p^B m$ for selling m units of asset B and receives the m units purchased of asset A . In the second stage, the arbitrageur delivers the units of asset B and pays $p^A m$ for the amount bought in market A . We assume that between stages 1 and 2, the arbitrageur can hide a fraction $1 - \theta$ of the proceeds from the short sale of asset B . Hence, the arbitrageur can commit i) the remaining fraction from the short sale proceeds, $\theta p^B m$, and ii) his

endowment, e_1 , to repaying the purchase of asset A . The collateral constraint is then given by

$$p^A m \leq \theta p^B m + e_1.$$

If the collateral constraint binds and the arbitrageur sector is competitive, the total amount arbitrated is

$$m^\alpha = \frac{e_1}{p^A (m m^\alpha) - \theta p^B (m^\alpha)}.$$

As one would expect, the larger the commitment to repaying, i.e., the larger θ , the higher the amount that can be arbitrated and the lower the arbitrage gap. In this case, the arbitrage gap satisfies

$$p^A (m^\alpha) - p^B (m^\alpha) = \frac{e_1}{m^\alpha} - (1 - \theta) p^B (m^\alpha).$$

Note that if $p^A \leq \theta p^B$ the collateral constraint never binds. Moreover, if the endowment of the arbitrageur is large enough, i.e., if

$$e_1 \geq (1 - \theta) p^B (m^*) m^*$$

the collateral constraint does not bind.

D Special Cases and Extensions

In this section, we describe several special cases of the general model studied in Section 4 and two additional extensions of the baseline model. In Section D.1 we consider a model with multiple assets and uncertainty. In Section D.2 and Section D.3, we present the results in a three-date model without and with uncertainty, respectively. In Section D.4 we show how our results can be extended to economies with multiple goods/currencies. Finally, in Section D.5 we extend the results to an economy with production.

D.1 Uncertainty and Multiple Assets

There are two dates $t \in \{0, 1\}$ and a single consumption good (dollar), which serves as numeraire. There is a finite number of investor types (investors, for short), indexed by $i \in I$, where $I = \{1, \dots, |I|\}$. At date 1, a state $s \in S$, where $S = \{1, \dots, |S|\}$, is realized. There is a finite number of assets, indexed by $j \in J$, where $j \in \{1, \dots, |J|\}$. We denote the payoff of asset j in state s by $d^j(s)$. First, we describe the problem that each type of investor faces. Subsequently, we describe the problem of arbitrageurs. We then present the counterpart of Proposition 1.

Investors The demand vector of type i investors is given by the solution to

$$\max_{\{q_0^{ij}\}_j} u_i(c_0^i) + \beta_i \sum_s \pi(s) u_i(c_1^i(s)),$$

where $\pi(s)$ represent probabilities, subject to the $|S| + 1$ budget constraints:

$$\begin{aligned} c_0^i &= n_0^i - \sum_j p^j \Delta q_0^{ij} \\ c_1^i(s) &= n_1^i(s) + \sum_j d^j(s) q_0^{ij}, \end{aligned}$$

where n_0^i and $n_1^i(s)$ denote endowments of consumption good, q_{-1}^{ij} and q_t^{ij} are investor i 's endowment of asset j at date 0 and his asset holdings of asset j at the end of date 0, respectively, and $d^j(s)$ denote asset payoffs. We denote by $\Delta q_0^{ij} \equiv q_0^{ij} - q_{-1}^{ij}$ the net demand of asset j of investor i . The first order conditions for the investor imply

$$p^j = \sum_s \beta_i \pi(s) \frac{u'_i(c_1^i(s))}{u'_i(c_0^i)} d^j(s).$$

Arbitrageurs The insights from Section 2 extend naturally to the case in which multiple assets need to be combined to form a replicating portfolio. In this case, m can be defined as the scale of the arbitrage trade and the vector $x = (x^1, \dots, x^{|J|})$ determines the direction of the arbitrage. Formally, the date 0 revenue of arbitrageurs is given by

$$V^\alpha(m) = - \sum_j p^j x^j m, \tag{A.12}$$

where $m = q^{\alpha 1}$ and $x^j = \frac{q^{\alpha j}}{q^{\alpha 1}}$, subject to zero-cash-flow constraints at date 1

$$c_1^\alpha(s) = \sum_j d^j(s) x^j m = 0.$$

Under the assumption that $|J| = |S| + 1$, this system has a solution for x^j 's in terms of $d^j(s)$. In the case of Section 2, $x = \begin{pmatrix} +1 \\ -1 \end{pmatrix}$.

Equilibrium

The zero-cash-flow constraints for the arbitrageur determine the direction of the arbitrage trade. The size/scale of the trade m is a parameter in our model and it indexes the equilibrium in this economy.

Definition. (*General arbitrage equilibrium*) An arbitrage equilibrium, parameterized by the scale of the arbitrage trade m , is defined as a set of consumption allocations, asset holdings, and prices $p^j(m)$ such that i) investors maximize utility subject to their budget constraints, and ii) the asset markets clear, that is,

$$\sum_i \Delta q_0^{ij} + x^j m = 0, \quad \forall j.$$

Given this equilibrium definition, we can define the payoffs for the investors and the arbitrageur as functions of m .

Welfare

Proposition 6 directly generalizes Proposition 1. The counterpart of Proposition 2 is straightforward and we omit it.

Proposition 6. (Marginal value of arbitrage: uncertainty, multiple assets)

a) (*Marginal individual value of arbitrage*) The marginal value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade m , measured in date 0 dollars, for type i investors and arbitrageurs,

is respectively given by

$$\begin{aligned}\frac{dV^i}{\lambda_0^i} &= -\sum_j \frac{dp^j}{dm} \Delta q_0^{ij}, \quad \forall i \\ \frac{dV^\alpha}{\lambda_0^\alpha} &= -\sum_j \left(\frac{dp^j}{dm} x^j m + p^j x^j \right),\end{aligned}$$

where λ_0^i and λ_0^α represent the marginal value of consumption at date 0 for agent $i \in I$ and the arbitrageur, respectively.

b) (Marginal social value of arbitrage) The marginal social value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade m , aggregated and measured in date 0 dollars, is given by

$$\frac{dW}{dm} = -\sum_j p^j x^j.$$

Proof. a) Differentiating V^i with respect to m we have that the value of an additional unit of traded in the arbitrage portfolio for investor i is

$$\frac{dV^i}{dm} = -u'_i(c_0^i) \sum_j \frac{dp^j}{dm} \Delta q_0^{ij},$$

where we use the optimality conditions of the investor's problem. Dividing by $\lambda_0^i = u'_i(c_0^i)$ expresses the marginal value in date 0 dollars and gives the first result above.

Differentiating Equation (A.12) with respect to m , we have that the marginal value of arbitrage for the arbitrageur is

$$\frac{dV^\alpha}{dm} = -\sum_j \left(\frac{dp^j}{dm} x^j m + p^j x^j \right).$$

Since $\lambda_0^\alpha = 1$, this shows the second part of a) above.

b) The marginal social value of m is given by

$$\frac{dW}{dm} = \frac{dV^\alpha}{\lambda_0^\alpha} + \sum_i \frac{dV^i}{\lambda_0^i}.$$

Using the results from part a), we have

$$\frac{dW}{dm} = -\sum_j \left(\frac{dp^j}{dm} x^j m + p^j x^j \right) - \sum_j \frac{dp^j}{dm} \Delta q_0^{ij} = -\sum_j p^j x^j.$$

□

D.2 Dynamic Model: No Uncertainty

In this section, we extend the model introduced in Section 2 to a dynamic setting with three dates. All the results extend naturally.

There are three dates $t \in \{0, 1, 2\}$ and a single consumption good (dollar), which serves as numeraire. There are two markets, indexed by A and B , at each date t . The economy is populated by type A investors, type B investors, and arbitrageurs. At each date, type A investors exclusively trade in market A while type B investors exclusively trade in market B . In each market $i \in \{A, B\}$ a riskless asset i is traded. Asset i

pays a dividend d_t^i at date t .

First, we describe the problem that both types of investors face. Subsequently, we describe the arbitrageur's profits. We then present the counterpart of Proposition 1.

Investors In each market $i \in \{A, B\}$, the representative investor's problem is

$$V^i(p_0^i, p_1^i) \equiv \max_{q_0^i, q_1^i} u_i(c_0^i) + \beta_i u_i(c_1^i) + \beta_i^2 u_i(c_2^i),$$

subject to

$$\begin{aligned} c_0^i &= n_0^i - p_0^i \Delta q_0^i \\ c_1^i &= n_1^i + (d_1^i + p_1^i) q_0^i - p_1^i q_1^i \\ c_2^i &= n_2^i + d_2^i q_1^i. \end{aligned}$$

The first order conditions for the investor imply

$$\begin{aligned} p_0^i &= \frac{\beta_i u_i'(c_1^i)}{u_i'(c_0^i)} (d_1^i + p_1^i) \\ p_1^i &= \frac{\beta_i u_i'(c_2^i)}{u_i'(c_1^i)} d_2^i. \end{aligned}$$

Arbitrageurs The arbitrageur follows an arbitrage strategy given by $(q_0^{\alpha A}, q_0^{\alpha B}, q_1^{\alpha A}, q_1^{\alpha B})$ where $q_t^{\alpha i}$ is the arbitrageur's position in market i at the end of date t . The arbitrage strategy has to satisfy the following period-by-period zero-cash-flow constraints:

$$0 = c_1^\alpha = (d_1^A + p_1^A) q_0^{\alpha A} + (d_1^B + p_1^B) q_0^{\alpha B} - p_1^A q_1^{\alpha A} - p_1^B q_1^{\alpha B} \quad (\text{A.13})$$

$$0 = c_2^\alpha = d_2^A q_1^{\alpha A} + d_2^B q_1^{\alpha B}. \quad (\text{A.14})$$

We define trading directions as

$$x_0 \equiv \frac{q_0^{\alpha B}}{q_0^{\alpha A}} \quad \text{and} \quad x_1 \equiv \frac{q_1^{\alpha B}}{q_1^{\alpha A}},$$

and the scales of the arbitrage trades by

$$m_0 \equiv q_0^{\alpha A} \quad \text{and} \quad m_1 \equiv q_1^{\alpha A}.$$

This allows us to rewrite the budget constraints as

$$0 = c_1^\alpha = (d_1^A + p_1^A) m_0 + (d_1^B + p_1^B) x_0 m_0 - p_1^A m_1 - p_1^B x_1 m_1 \quad (\text{A.15})$$

$$0 = c_2^\alpha = (d_2^A + d_2^B x_1) m_1. \quad (\text{A.16})$$

The zero-cash-flow constraints determine the direction of the arbitrage trade for given scales (m_0, m_1) . Formally, from Equations (A.13) and (A.14) it follows, whenever $m_1 \neq 0$, that

$$x_0 = -\frac{d_1^A + p_1^A}{d_1^B + p_1^B} + \frac{(p_1^A + p_1^B x_1) \frac{m_1}{m_0}}{d_1^B + p_1^B} \quad (\text{A.17})$$

$$x_1 = -\frac{d_2^A}{d_2^B}. \quad (\text{A.18})$$

The direction of the arbitrage trade at date 1 depends only on the relative payoff of the assets at date 2, just as in the model introduced in Section 2. Interestingly, the direction of the arbitrage trade at date 0 depends on the relative payoff of the assets at date 1 and on an additional term that incorporates the arbitrageur's profits from the arbitrage trade at date 1.

Equilibrium

For any given scales of the arbitrage trades m_0 and m_1 , the market clearing conditions in markets A and B in periods $t \in \{0, 1\}$ together with the four first order conditions of the investors, the investors' budget constraints, and the zero-cash-flow constraints in Equations (A.17) and (A.18) determine the equilibrium allocations for the investors, equilibrium prices, and trading directions for the arbitrageurs.

Definition. (*Dynamic arbitrage equilibrium*) A dynamic arbitrage equilibrium, parameterized by the scales of the arbitrage trades m_0 and m_1 , is defined as a set of consumption allocations, asset holdings, and prices $p_0^A(m_0, m_1)$, $p_0^B(m_0, m_1)$, $p_1^A(m_0, m_1)$, and $p_1^B(m_0, m_1)$ such that i) investors maximize utility subject to their budget constraints, and ii) the asset markets A and B clear at dates $t \in \{0, 1\}$, that is,

$$\begin{aligned} \Delta q_0^A + m_0 &= 0 \\ \Delta q_0^B + x_0 m_0 &= 0 \\ \Delta q_1^A + m_1 - m_0 &= 0 \\ \Delta q_1^B + x_1 m_1 - x_0 m_0 &= 0. \end{aligned}$$

The direction of the arbitrage trade at date 0 in Equation (A.17) can be written as function of the equilibrium prices at $t = 1$ and the scales of the arbitrage trades, i.e.,

$$x_0(p_1^A, p_1^B; m_0, m_1),$$

which is helpful in decomposing the direct effect and the pecuniary effects from changes in m_0 and m_1 . Differentiating x_0 with respect to m_t we have

$$\frac{dx_0}{dm_t} = \frac{\partial x_0}{\partial m_t} + \frac{\partial x_0}{\partial p_1^A} \frac{dp_1^A}{dm_t} + \frac{\partial x_0}{\partial p_1^B} \frac{dp_1^B}{dm_t},$$

where the first term represents the direct effect of a change in m_t on the direction of the arbitrage trade at date 0 and the last two terms represent the pecuniary effects from changing m_t . Lemma 2 below characterizes $\frac{dx_0}{dm_t}$.

Lemma 2. (*Effects on arbitrage direction at date 0*) The total derivatives of x_0 with respect to m_0 and m_1 are given by

$$\frac{dx_0}{dm_0} = \frac{1}{m_0} \left(-x_0 - \frac{d_1^A + p_1^A}{d_1^B + p_1^B} + \frac{m_1 - m_0}{d_1^B + p_1^B} \frac{dp_1^A}{dm_0} + \frac{x_1 m_1 - x_0 m_0}{d_1^B + p_1^B} \frac{dp_1^B}{dm_0} \right)$$

and

$$\frac{dx_0}{dm_1} = \frac{1}{m_0} \left(\frac{p_1^A + p_1^B x_1}{d_1^B + p_1^B} + \frac{m_1 - m_0}{d_1^B + p_1^B} \frac{dp_1^A}{dm_1} + \frac{x_1 m_1 - x_0 m_0}{d_1^B + p_1^B} \frac{dp_1^B}{dm_1} \right).$$

Proof. Partially differentiating Equation (A.17) with respect to m_0 and m_1 gives

$$\begin{aligned} \frac{\partial x_0}{\partial m_0} &= -\frac{p_1^A + p_1^B x_1}{d_1^B + p_1^B} \frac{m_1}{m_0} \frac{1}{m_0} = -\frac{1}{m_0} \left(x_0 + \frac{d_1^A + p_1^A}{d_1^B + p_1^B} \right) \\ \frac{\partial x_0}{\partial m_1} &= \frac{p_1^A + p_1^B x_1}{d_1^B + p_1^B} \frac{1}{m_0} = \frac{1}{m_1} \left(x_0 + \frac{d_1^A + p_1^A}{d_1^B + p_1^B} \right) \end{aligned}$$

and with respect to p_1^A and p_1^B gives

$$\begin{aligned} \frac{\partial x_0}{\partial p_1^A} &= \frac{1}{m_0} \frac{m_1 - m_0}{d_1^B + p_1^B} \\ \frac{\partial x_0}{\partial p_1^B} &= \frac{1}{m_0} \frac{x_1 m_1 - x_0 m_0}{d_1^B + p_1^B}. \end{aligned}$$

Using that

$$\frac{dx_0}{dm_t} = \frac{\partial x_0}{\partial m_t} + \frac{\partial x_0}{\partial p_1^A} \frac{dp_1^A}{dm_t} + \frac{\partial x_0}{\partial p_1^B} \frac{dp_1^B}{dm_t}$$

gives the results. \square

Welfare

The arbitrageur's utility from following a trading strategy with scales m_0 and m_1 is

$$V^\alpha(m_0, m_1) = -\left(p_0^B q_0^{\alpha B} + p_0^A q_0^{\alpha A}\right) = -\left(p_0^B x_0 m_0 + p_0^A m_0\right), \quad (\text{A.19})$$

where x_0 depends on m_0 and m_1 and is given by Equation (A.17). Moreover, the utility of an investor i depends on m_0 and m_1 only through the equilibrium prices p_0^i and p_1^i . The following proposition characterizes the individual (for investors and arbitrageurs) and marginal social values of arbitrage.

Proposition 7. (Marginal value of arbitrage: dynamic model, no uncertainty)

a) *(Marginal individual value of arbitrage)* The marginal value of arbitrage, that is, the marginal values of increasing the scale of the arbitrage trades m_0 and m_1 , measured in date 0 dollars, for type i investors and arbitrageurs, are respectively given by

$$\begin{aligned} \frac{\frac{dV^i}{dm_t}}{\lambda_0^i} &= -\frac{dp_0^i}{dm_t} \Delta q_0^i - \frac{\beta_i u'_i(c_1^i)}{u'_i(c_0^i)} \frac{dp_1^i}{dm_t} \Delta q_1^i \\ \frac{\frac{dV^\alpha}{dm_0}}{\lambda_0^\alpha} &= \left(\frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A \right) - \left(\frac{m_1 - m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^A}{dm_0} + \frac{x_1 m_1 - x_0 m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^B}{dm_0} \right) - \left(\frac{dp_0^B}{dm_0} x_0 + \frac{dp_0^A}{dm_0} \right) m_0 \\ \frac{\frac{dV^\alpha}{dm_1}}{\lambda_0^\alpha} &= -\frac{p_1^A + p_1^B x_1}{\frac{d_1^B + p_1^B}{p_0^B}} - \left(\frac{m_1 - m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^A}{dm_1} + \frac{x_1 m_1 - x_0 m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^B}{dm_1} \right) - \left(\frac{dp_0^B}{dm_1} x_0 + \frac{dp_0^A}{dm_1} \right) m_0, \end{aligned}$$

where λ_0^j represents the marginal value of consumption at date 0 for agent $j \in \{A, B, \alpha\}$.

b) *(Marginal social value of arbitrage)* The marginal social value of arbitrage, that is, the marginal values of increasing the scale of the arbitrage trades m_0 and m_1 , aggregated and measured in date 0 dollars, are given

by

$$\begin{aligned}\frac{dW}{dm_0} &= \left(\frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A \right) - \left(\frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \right) \frac{dp_1^A}{dm_0} (m_1 - m_0) \\ \frac{dW}{dm_1} &= -\frac{p_1^A + p_1^B x_1}{\frac{d_1^B + p_1^B}{p_0^B}} - \left(\frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \right) \frac{dp_1^A}{dm_1} (m_1 - m_0),\end{aligned}$$

where $x_1 = -\frac{d_2^A}{d_2^B}$.

Proof. a) The value of an additional unit of arbitrage at date 0 for investor i is

$$\frac{dV^i}{dm_t} = -u'_i(c_0^i) \frac{dp_0^i}{dm_t} \Delta q_0^i - \beta_i u'_i(c_1^i) \frac{dp_1^i}{dm_t} \Delta q_1^i,$$

where we use the optimality conditions of the investor's problem. Dividing by $u'_i(c_0^i)$ expresses the marginal value in date 0 dollars and gives the first result above.

Differentiating the expression in Equation (A.19) we have that the marginal value of m_0 for the arbitrageur is

$$\frac{dV^\alpha}{dm_0} = -\left(p_0^B \frac{d(x_0 m_0)}{dm_0} + p_0^A \right) - \left(\frac{dp_0^B}{dm_0} x_0 + \frac{dp_0^A}{dm_0} \right) m_0. \quad (\text{A.20})$$

Using Lemma 2 we have

$$\begin{aligned}\frac{d(x_0 m_0)}{dm_0} &= -\underbrace{\frac{d_1^A + p_1^A}{d_1^B + p_1^B}}_{=x_0 + \frac{\partial x_0}{\partial m_0} m_0} + \frac{m_1 - m_0}{d_1^B + p_1^B} \frac{dp_1^A}{dm_0} + \frac{x_1 m_1 - x_0 m_0}{d_1^B + p_1^B} \frac{dp_1^B}{dm_0},\end{aligned}$$

and using this expression in Equation (A.20) gives the second result in a) above, where we used that $\lambda_\alpha = 1$.

Differentiating the expression in Equation (A.19) with respect to m_1 we have

$$\frac{dV^\alpha}{dm_1} = -p_0^B \frac{dx_0}{dm_1} m_0 - \left(\frac{dp_0^B}{dm_1} x_0 + \frac{dp_0^A}{dm_1} \right) m_0.$$

Using Lemma 2 gives

$$\frac{dV^\alpha}{dm_1} = -\frac{p_1^A + p_1^B x_1}{\frac{d_1^B + p_1^B}{p_0^B}} - \left(\frac{m_1 - m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^A}{dm_1} + \frac{x_1 m_1 - x_0 m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^B}{dm_1} \right) - \left(\frac{dp_0^B}{dm_1} x_0 + \frac{dp_0^A}{dm_1} \right) m_0,$$

which concludes the proof of part a).

b) The marginal social value of m_t is given by

$$\frac{dW}{dm_t} = \frac{dV^\alpha}{\lambda_0^\alpha} + \frac{dV^A}{\lambda_0^A} + \frac{dV^B}{\lambda_0^B}.$$

Using the results from part a) for m_0 this implies

$$\begin{aligned} \frac{dW}{dm_0} &= \left(\frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A \right) - \left(\frac{m_1 - m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^A}{dm_0} + \frac{x_1 m_1 - x_0 m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^B}{dm_0} \right) - \left(\frac{dp_0^B}{dm_0} x_0 + \frac{dp_0^A}{dm_0} \right) m_0 \\ &\quad - \frac{dp_0^A}{dm_0} \Delta q_0^A - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \frac{dp_1^A}{dm_0} \Delta q_1^A - \frac{dp_0^B}{dm_0} \Delta q_0^B - \frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} \frac{dp_1^B}{dm_0} \Delta q_1^B, \end{aligned}$$

and using the market clearing conditions and the investors' first order conditions gives

$$\frac{dW}{dm_0} = \left(\frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A \right) + \left(\frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \right) \frac{dp_1^A}{dm_0} (m_0 - m_1),$$

which proves the first part of b).

Analogously, using the results from part a) for m_1 we have

$$\begin{aligned} \frac{dW}{dm_1} &= -\frac{p_1^A + p_1^B x_1}{\frac{d_1^B + p_1^B}{p_0^B}} - \left(\frac{m_1 - m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^A}{dm_1} + \frac{x_1 m_1 - x_0 m_0}{\frac{d_1^B + p_1^B}{p_0^B}} \frac{dp_1^B}{dm_1} \right) - \left(\frac{dp_0^B}{dm_1} x_0 + \frac{dp_0^A}{dm_1} \right) m_0 \\ &\quad - \frac{dp_0^A}{dm_1} \Delta q_0^A - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \frac{dp_1^A}{dm_1} \Delta q_1^A - \frac{dp_0^B}{dm_1} \Delta q_0^B - \frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} \frac{dp_1^B}{dm_1} \Delta q_1^B. \end{aligned}$$

Using market clearing and the investors' first order conditions gives

$$\frac{dW}{dm_1} = -\frac{p_1^A + p_1^B x_1}{\frac{d_1^B + p_1^B}{p_0^B}} - \left(\frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)} \right) \frac{dp_1^A}{dm_1} (m_1 - m_0).$$

This proves the result since $\lambda_\alpha = 1$. □

Note that for both $\frac{dW}{dm_0}$ and $\frac{dW}{dm_1}$ the marginal social value of arbitrage is given by i) a direct effect, given by the arbitrage gap at either date 0 or date 1 (appropriately discounted) and ii) a set of distributive pecuniary externalities, using the language of [Dávila and Korinek \(2018\)](#), who show that such externalities are always given by a) differences in marginal rates of substitution, in this case $\frac{\beta_B u'_B(c_1^B)}{u'_B(c_0^B)} - \frac{\beta_A u'_A(c_1^A)}{u'_A(c_0^A)}$; b) price sensitivities, in this case $\frac{dp_1^A}{dm_t}$; and c) net trading positions, in this case $m_1 - m_0$. As long as markets are incomplete, these externalities have a first-order effect on welfare. However, these are zero-sum when measured in terms of date 1 dollars.

Even when markets are incomplete, there are situations in which such distributive pecuniary externalities do not enter welfare calculations. For instance, when arbitrageurs follow a buy-and-hold strategy, $m_1 = m_0$ and the distributive externalities vanish. This is natural, since there are no net trades taking place at date 1 in that case.

Corollary 2. *The marginal social value of a buy-and-hold strategy is equal to the adjusted price gap. Formally,*

$$\left. \frac{dW}{dm_0} \right|_{m_1=m_0} = \frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A.$$

D.3 Dynamic Model: Uncertainty

In this section, we extend the model introduced in the previous subsection to a dynamic setting with uncertainty.

There are three dates $t \in \{0, 1, 2\}$ and a single consumption good (dollar), which serves as numeraire. At date 1 the state $s \in S$, where $S = \{1, \dots, |S|\}$, is realized. There are two markets, indexed by A and B , at each date t . The economy is populated by type A investors, type B investors, and arbitrageurs. At each date, type A investors exclusively trade in market A while type B investors exclusively trade in market B . In each market $i \in \{A, B\}$ a potentially risky asset i is traded. Asset i pays a dividend $d_t^i(s)$ at date $t \in \{1, 2\}$ — note that all uncertainty is realized with s at date 1, so $d_t^i(s)$ is a deterministic function of s realized at date 1.

Investors In each market i , the representative investor's problem is

$$V^i(p_0^i, p_1^i) \equiv \max_{q_0^i, \{q_1^i(s)\}} u_i(c_0^i) + \beta_i \sum_s \pi(s) u_i(c_1^i(s)) + \beta_i^2 \sum_s \pi(s) u_i(c_2^i(s))$$

subject to

$$\begin{aligned} c_0^i &= n_0^i - p_0^i \Delta q_0^i \\ c_1^i(s) &= n_1^i(s) + (d_1^i(s) + p_1^i(s)) q_0^i - p_1^i q_1^i(s) \\ c_2^i(s) &= n_2^i(s) + d_2^i(s) q_1^i(s). \end{aligned}$$

The first order conditions for the investor imply

$$\begin{aligned} p_0^i &= \sum_s \pi(s) \frac{\beta_i u_i'(c_1^i(s))}{u_i'(c_0^i)} (d_1^i(s) + p_1^i(s)) \\ p_1^i &= \frac{\beta_i u_i'(c_2^i(s))}{u_i'(c_1^i(s))} d_2^i(s), \quad \forall s. \end{aligned}$$

Arbitrageurs The arbitrageur follows an arbitrage strategy given by $(q_0^{\alpha A}, q_0^{\alpha B}, q_1^{\alpha A}(s), q_1^{\alpha B}(s))$ where $q_t^{\alpha i}$ is the arbitrageur's position in market i at the end of date t . Since we allow potentially for $|S| > |J|$, we consider arbitrage strategies of the form $c_1^\alpha(s) \geq 0$. The arbitrage strategy has to satisfy the following period by period cash-flow constraints:

$$0 \leq c_1^\alpha(s) = (d_1^A(s) + p_1^A(s)) q_0^{\alpha A} + (d_1^B(s) + p_1^B(s)) q_0^{\alpha B} - p_1^A(s) q_1^{\alpha A}(s) - p_1^B(s) q_1^{\alpha B}(s), \quad \forall s, \quad (\text{A.21})$$

$$0 = c_2^\alpha(s) = d_2^A(s) q_1^{\alpha A}(s) + d_2^B(s) q_1^{\alpha B}(s), \quad \forall s, \quad (\text{A.22})$$

Analogous to the case without uncertainty we define trading directions as

$$x_0 \equiv \frac{q_0^{\alpha B}}{q_0^{\alpha A}} \quad \text{and} \quad x_1(s) \equiv \frac{q_1^{\alpha B}(s)}{q_1^{\alpha A}(s)}, \quad \forall s.$$

and the scales of the arbitrage trades by

$$q_0^{\alpha A} \equiv m_0 \quad \text{and} \quad q_1^{\alpha A}(s) \equiv m_1(s), \quad \forall s.$$

Note that the scale and direction of the arbitrage trade are contingent on the state s . The cash-flow constraints determine the direction of the arbitrage trade for given scales $(m_0, m_1(s))$. Formally, from

Equations (A.21) and (A.22) it follows that

$$x_0 m_0 \geq -\frac{d_1^A(s) + p_1^A(s)}{d_1^B(s) + p_1^B(s)} m_0 + \frac{p_1^A(s) + p_1^B(s) x_1(s)}{d_1^B(s) + p_1^B(s)} m_1(s), \quad \forall s, \quad (\text{A.23})$$

$$x_1(s) = -\frac{d_2^A(s)}{d_2^B(s)}, \quad \forall s. \quad (\text{A.24})$$

The direction of the arbitrage trade at date 1 depends only on the relative payoffs of the assets at date 2 in state s , just as in the model introduced in Section 2 and the dynamic model without uncertainty. However, while in the model without uncertainty one could freely choose the scales of the arbitrage trades m_0 and m_1 , this is no longer the case when there is uncertainty. The $|S|$ equations in (A.23) impose $|S|$ restrictions that need to be satisfied by the $|S| + 1$ scales of the arbitrage trades, $m_0, \{m_1(s)\}_{s \in S}$ and the direction of the arbitrage trade at $t = 0$, x_0 . In particular, there may be scales $\{m_1(s)\}_{s \in S}$ such that it is impossible to find an arbitrage trade from the perspective of date 0. More specifically, there may be some arbitrage trades at $\{m_1(s)\}_{s \in S}$ for which negative consumption at date 1 by the arbitrageur may be unavoidable: this is exactly the scenario considered in Shleifer and Vishny (1997).

When the assets are riskless, i.e., $d_t^i(s) = d_t^i$, an arbitrage buy-and-hold strategy in which $m_1(s) = m_0$ always exists as long as

$$x_0 \geq \frac{-d_1^A + p_1^B(s) \left(-\frac{d_2^A}{d_2^B}\right)}{d_1^B + p_1^B(s)}, \quad \forall s.$$

Note that in this case we are outside of the Shleifer and Vishny (1997) scenario, because there are no interim portfolio adjustments for the arbitrageur. Finally, if $d_t^i = d_t$, the $|S|$ constraints on x_0 collapse to $x_0 \geq -1$, and we are back to the model studied in Section D.1.

D.4 CIP Model

Here we briefly describe how to extend Proposition 2 to the CIP case, in which there are multiple currencies.

Market A (USD) Investors in market A have time-separable utility, with discount factor β_A . They have euro endowments $n_0^{A\$}$ and $n_1^{A\$}$ and hold an initial position $q_{-1}^{A\$}$ in the traded asset. Hence, type A investors choose $q_0^{A\$}$ as the solution to

$$\max_{q_0^{A\$}} u_A(c_0^{A\$}) + \beta_A u_A(c_1^{A\$}),$$

subject to the budget constraints

$$\begin{aligned} p^{A\$} \Delta q_0^{A\$} + c_0^{A\$} &= n_0^{A\$} \\ c_1^{A\$} &= n_1^{A\$} + d_1^{A\$} q_0^{A\$}, \end{aligned}$$

where $\Delta q_0^{A\$} = q_0^{A\$} - q_{-1}^{A\$}$ and where $c_0^{A\$}$ and $c_1^{A\$}$ denote the consumption of type A investors at dates 0 and 1, respectively.

Market B (EUR) Investors in market B face the same problem as investors in market A . Investors in market B also have time-separable utility, with discount factor β_B . They have euro endowments $n_0^{B\text{€}}$ and $n_1^{B\text{€}}$ and hold an initial position $q_{-1}^{B\text{€}}$ in the traded asset. Hence, type B investors choose $q_0^{B\text{€}}$ as the solution

to

$$\max_{q_0^{B\epsilon}} u_B(c_0^{B\epsilon}) + \beta_B u_B(c_1^{B\epsilon}),$$

subject to the budget constraints

$$\begin{aligned} p^{B\epsilon} \Delta q_0^{B\epsilon} + c_0^{B\epsilon} &= n_0^{B\epsilon} \\ c_1^{B\epsilon} &= n_1^{B\epsilon} + d_1^{B\epsilon} q_0^{B\epsilon}, \end{aligned}$$

where $\Delta q_0^{B\epsilon} = q_0^{B\epsilon} - q_{-1}^{B\epsilon}$ and where $c_0^{B\epsilon}$ and $c_1^{B\epsilon}$ denote the consumption of type B investors at dates 0 and 1, respectively.

Arbitrageurs Arbitrageurs (indexed by α) are the only agents who can trade in both markets A and B . Arbitrageurs have no initial endowments of euros or dollars. Arbitrageurs can exchange date 0 euros for date 0 dollars at rate S_0 , and write contracts to exchange date 1 euros for date 1 dollars at rate F_0 .

Arbitrageurs implement a trading strategy with zero cash flows at date 1 in order to maximize the date 0 revenue raised by such a strategy. Formally, denoting by $q_0^{\alpha A\$}$ and $q_0^{\alpha B\epsilon}$ the respective asset purchases of arbitrageurs in markets A and B denominated in their domestic currencies, the objective function of an American arbitrageur (a symmetric setup obtains for a European arbitrageur) is given by

$$- \left(p^{A\$} q_0^{\alpha A\$} + \underbrace{p^{B\$}}_{S_0 p^{B\epsilon}} q_0^{\alpha B\epsilon} \right), \quad (\text{A.25})$$

subject to a zero-cash-flow constraint

$$d_1^{A\$} q_0^{\alpha A\$} + \underbrace{d_1^{B\$}}_{F_0 d_1^{B\epsilon}} q_0^{\alpha B\epsilon} = 0, \quad (\text{A.26})$$

where $p^{B\$} = S_0 p^{B\epsilon}$, and $d_1^{B\$} = F_0 d_1^{B\epsilon}$.

We can then exploit the zero-cash-flow constraint to rewrite,

$$q_0^{\alpha B\epsilon} = - \frac{d_1^{A\$}}{d_1^{B\epsilon}} \frac{1}{F_0} q_0^{\alpha A\$}.$$

Letting $m = q_0^{\alpha A\$}$, we can then write the arbitrageur's portfolio as $x = \begin{pmatrix} 1 \\ -\frac{d_1^{A\$}}{d_1^{B\epsilon}} \frac{1}{F_0} \end{pmatrix} m$.

Market clearing The market clearing conditions in this environment are given by

$$\begin{aligned} \Delta q_0^{A\$} + q_0^{\alpha A\$} &= 0 \\ \Delta q_0^{B\epsilon} + q_0^{\alpha B\epsilon} &= 0, \end{aligned}$$

which can be equivalently written as

$$\begin{aligned} \Delta q_0^{A\$} + m &= 0 \\ \Delta q_0^{B\epsilon} - \frac{d_1^{A\$}}{d_1^{B\epsilon}} \frac{1}{F_0} m &= 0. \end{aligned}$$

Without loss of generality, we assume that $d_1^{A\$} = 1$ and $d_1^{B\epsilon} = 1$.

Welfare Here we compute the marginal value of increasing m for both types of investors, measured in date 0 dollars:

$$\begin{aligned}\frac{\frac{dV^A}{dm}}{u'_A(c_0^A)} &= -\frac{dp^{A\$}}{dm}\Delta q_0^{A\$} \\ S_0\frac{\frac{dV^B}{dm}}{u'_B(c_0^B)} &= -S_0\frac{dp^{B\epsilon}}{dm}\Delta q_0^{B\epsilon}.\end{aligned}$$

We can then consider the indirect utility of the arbitrageurs, denoted by $V^\alpha(m)$, as

$$\frac{dV^\alpha(m)}{dm} = -\left(\frac{dp^{A\$}}{dm}q_0^{\alpha A\$} + p^{A\$}\frac{dq_0^{\alpha A\$}}{dm} + \frac{dp^{B\$}}{dm}q_0^{\alpha B\epsilon} + p^{B\$}\frac{dq_0^{\alpha B\epsilon}}{dm}\right).$$

Leaving aside pecuniary effects, we can write the marginal social value of arbitrage as

$$\begin{aligned}\frac{dW}{dm} &= \frac{d_1^{A\$}}{d_1^{B\epsilon}}\frac{S_0}{F_0}p_B^\epsilon - p_A^\$ \\ &= \frac{S_0}{F_0}p_B^\epsilon - p_A^\$ \\ &= \frac{S_0}{F_0}\frac{1}{1+r^\epsilon} - \frac{1}{1+r^\$},\end{aligned}$$

where we set $d_1^{A\$} = d_1^{B\epsilon}$ in both countries without loss of generality.

D.5 Production

It is straightforward to introduce production in our framework.

Investors As in the baseline model, we assume that the problem of investors in market i is given by

$$\max_{q_0^i, k^i} u_i(c_0^i) + \beta_i u_i(c_1^i),$$

subject to

$$\begin{aligned}p^i\Delta q_0^i + c_0^i + \Upsilon_i(k^i) &= n_0^i \\ c_1^i &= n_1^i + d_1 q_0^i + f_i(k^i).\end{aligned}$$

In this case $\Upsilon_i(k^i)$ denotes a cost of production/adjustment cost, and $f_i(k^i)$ denotes the output of the production process, which materializes at date 1. Note that in this formulation, the amount of output produced depends on the equilibrium p^i , so different arbitrage gaps are associated with different production levels.

The first order conditions for q_0^i and k^i are

$$p_0^i = \frac{\beta_i u_i'(c_1^{i*})}{u_i'(c_0^{i*})} d_1$$

$$\Upsilon_i'(k^{i*}) = \frac{\beta_i u_i'(c_1^{i*})}{u_i'(c_0^{i*})} f_i'(k^{i*}).$$

The indirect utility of an investor in market i is given by

$$V^i(p^i) = u_i(c_0^{i*}) + \beta_i u_i(c_1^{i*}).$$

We omit the description of the problem of arbitrageurs, since it is the same as the one they face in the baseline model.

Equilibrium

We extend the definition of equilibrium to account for production.

Definition. (*Arbitrage equilibrium with production*) An arbitrage equilibrium, parameterized by the scale of the arbitrage trade m , is defined as a set of consumption and capital allocations, asset holdings, and prices $p^i(m)$, such that i) investors maximize utility subject to their budget constraints, and ii) the asset markets clear, that is,

$$\Delta q_0^A + m = 0$$

$$\Delta q_0^B + m = 0.$$

Proposition 8. (Marginal value of arbitrage with production)

a) (*Marginal individual value of arbitrage*) The marginal value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade $m \in [0, m^*]$, measured in date 0 dollars, for type A investors, type B investors, and arbitrageurs, is respectively given by

$$\frac{\frac{dV^A(m)}{dm}}{\lambda_0^A} = \frac{dp^A(m)}{dm} m > 0$$

$$\frac{\frac{dV^B(m)}{dm}}{\lambda_0^B} = \frac{dp^B(m)}{dm} (-m) > 0$$

$$\frac{\frac{dV^\alpha(m)}{dm}}{\lambda_0^\alpha} = \left(\frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm} \right) m + p^B(m) - p^A(m) \geq 0.$$

b) (*Marginal social value of arbitrage*) The marginal social value of arbitrage, that is, the marginal value of increasing the scale of the arbitrage trade m , aggregated and measured in date 0 dollars, is given by

$$\frac{dW(m)}{dm} = p^B(m) - p^A(m) > 0. \tag{A.27}$$

Proof. a) Differentiating the indirect utility of investors and applying the envelope theorem implies

$$\begin{aligned} \frac{dV^i}{dm} &= (-u'_i(c_0^{i*}) p^i + \beta_i u'_i(c_1^{i*}) d_1) \frac{dq_0^i}{dm} + (-u'_i(c_0^{i*}) \Upsilon'(k^{i*}) + \beta_i u'_i(c_1^{i*}) f'_i(k^{i*})) \frac{dk^{i*}}{dm} - u'_i(c_0^{i*}) \frac{dp^i}{dm} \Delta q_0^i \\ &= -u'_i(c_0^{i*}) \frac{dp^i}{dm} \Delta q_0^i. \end{aligned}$$

Defining $\lambda_0^i = u'_i(c_0^{i*})$ and using the market clearing condition in market i yield the first two results in a). The marginal value for the arbitrageurs is as in the baseline model.

b) Combining the result in a) with the characterization of $\frac{dV^\alpha}{dm}$ in the baseline model, we find that

$$\begin{aligned} \frac{dW(m)}{dm} &= \frac{dV^A(m)}{\lambda_0^A} + \frac{dV^B(m)}{\lambda_0^B} + \frac{dV^\alpha(m)}{\lambda_0^\alpha} \\ &= p^B(m) - p^A(m), \end{aligned}$$

exactly as in the baseline model. □

E Detailed Data Description/Institutional Background

This Appendix describes in detail the sources of data used in the CIP application in Section 5 and the Dual-Listed Companies application in Section 6.

E.1 CIP Application: FX Futures Data

Our price impact estimates rely on bid, offer, and trade data from FX markets traded at the Chicago Mercantile Exchange (CME). Direct feed “L1” (top of book and trade) data from the CME was recorded on a server in real time and stored in a database. Every transaction, bid change, and offer change was recorded with a millisecond timestamp corresponding to the time that the CME’s matching engine processed the order book change or trade. Bid and offer updates consist of a price and size that reflect the highest bid and lowest offer at each point in time. Transaction updates reflect all executed trades in the studied markets, with a transaction price and size.

Table OA-1: CME Data Summary Statistics (1)

Market	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
Transactions	7,951,935	8,656,364	6,150,197	17,607,450	10,088,058
Volume	27,221,569	27,314,833	18,659,840	51,922,628	28,683,280
Trading Days	305	305	305	305	305
Minimum Fluctuation (Tick)	0.00005	0.0001	0.00005	0.00005	0.0000005
Contract Multiplier	100,000	62,500	100,000	125,000	12,500,000

Note: Table OA-1 provides summary statistics of the various contract markets at the Chicago Mercantile Exchange. Aggregate figures in the top section of the table cover all trading days, including holidays and roll periods, while trading day totals exclude holidays and the days immediately preceding liquidity migration from the front-month to the subsequent contract. Contract specifications are included along with statistics specific to the period 12/15/2019 – 02/26/2021.

The dataset covers the period between 12/15/2019 and 02/26/2021, and corresponds to five contract

months for each of five different futures contracts. Specifically, the March 2020, June 2020, September 2020, December 2020, and March 2021 futures contracts are covered for the AUD/USD, GBP/USD, CAD/USD, EUR/USD, and JPY/USD markets.

All studied CME Group FX markets have similar contract specifications. They feature a quantity of currency, a pricing convention, and a minimum fluctuation. For the EUR/USD, for instance, the quantity is 125,000 since the seller of a single contract (the minimum allowed size in any futures market) is promising to deliver 125,000 euros in exchange for dollars; the pricing convention is dollars per single Euro since in “EUR/USD” the “EUR” is listed first. The minimum fluctuation is 0.00005 ($\$6.25 = 125,000 \times \0.00005 in trading gains), meaning that prices must move in integer multiples of 0.00005. The CME Group determines minimum fluctuations and contract sizing based on liquidity and microstructure considerations. Less liquid markets generally feature larger minimum fluctuations whereas highly liquid markets feature smaller minimum fluctuations (also known as “tick sizes”).

The CME’s FX markets are open from 6pm Eastern Time on Sunday to 5pm Eastern Time on Friday with a daily maintenance window from 5pm Eastern Time to 6pm Eastern Time. All trading takes place via a continuous limit order book, except for at the market open during each 24 hour cycle, in which case a “Pre-Open” period allows an opening price auction to take place, during which time the market’s opening price and quantity are determined by orders in the book. This paper uses all transactions from the 24 hour trading cycle except for trades between 3pm and 5pm Eastern Time to avoid issues relating to diminished post-settlement time liquidity.

During the continuous trading period, a single large order that is matched against multiple smaller orders appears in the data feed as a string of consecutive individual prices and sizes. This paper’s CIP price impact analysis aggregates these strings of transactions (which originate from a single “aggressive” order that has crossed the spread to take liquidity from resting orders on the opposite size) into an average fill price and total trade size. By way of illustration, if an aggressive 1,000 lot buy order was filled 500@1.2000 and 500@1.2001, we would study it as a 1,000 lot order filled at an average price of 1.20005.

Table OA-2: CME Data Summary Statistics (2)

	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
Min Price	0.55	1.14	0.68	1.07	0.0089
Mean Price	0.71	1.30	0.75	1.16	0.0094
Max Price	0.80	1.41	0.80	1.23	0.0097
Avg Daily Dollar Volume (\$B)	6.11	7.14	4.49	23.93	10.8
P10 Daily Volume	57,541	57,865	40,040	113,076	53,588
Median Daily Volume	82,833	83,979	56,416	159,742	78,199
P90 Daily Volume	118,345	124,929	81,095	225,866	147,717
# Transactions/Day	25,287	27,824	19,630	56,458	32,475

Note: Table OA-2 contains price, daily volume, and daily transaction summary statistics for the different contract markets’ datasets. All FX futures contracts are quoted in USD per single unit of foreign currency: this results in values in the range of 0.5 to 1.5 for most markets but values in the neighborhood of 0.01 for the yen. The average daily dollar volume is generally increasing in the size of the corresponding economy, with the smallest transaction volume in the AUD/USD and CAD/USD markets and the largest volumes in the EUR/USD (the most actively traded pair).

Table OA-3: CIP Summary Statistics

	Mean	St. Dev.	P10	Median	P90
AUD/USD	-0.0007	0.0011	-0.0013	-0.0004	-0.0001
GBP/USD	-0.0014	0.0014	-0.0024	-0.0009	-0.0005
CAD/USD	-0.0015	0.0010	-0.0021	-0.0012	-0.0009
EUR/USD	-0.0005	0.0007	-0.0011	-0.0004	0.0000
JPY/USD	-0.0017	0.0012	-0.0029	-0.0013	-0.0008

Note: This table provides statistics for the three-month cross-currency bases by currency for the period 02/29/2008 through 02/26/2021. We use three-month secured government paper for the cross-currency basis, whenever possible, and three-month LIBOR (unsecured) rates otherwise.

Table OA-4: Dual-Listed Companies Summary Statistics

	Mean	St. Dev.	P10	Median	P90
Royal Dutch/Shell	1.07510	0.05599	0.99436	1.07869	1.14640
Smithkline/Beecham	1.07968	0.04993	0.99439	1.09231	1.13279
Rio Tinto PLC/Ltd	0.98152	0.04666	0.91744	0.98540	1.03904

Note: Table OA-4 provides summary statistics of the relative pricing of various twin shares over the sample periods. A value of 1 represents parity (no arbitrage opportunity exists), while 1.1 or 0.9, for example, represent 10% deviations. Royal Dutch/Shell and Smithkline/Beechman in particular feature substantial and persistent fluctuations away from parity, with median deviations of around 8% and 9% respectively.

E.2 CIP Application: Price/Rates Data

We use Bloomberg to obtain data for spot currencies, forward points, and interest rates (secured and unsecured three-month lending rates). The exact data series used can be found by inputting the following symbols into a Bloomberg terminal: GB03 Govt, GTDEM3MO Corp, ADBB3M CMPN Curncy, CDOR03 Index, BP0003M, JY0003M Index, EUR BGN Curncy, EUR3M BGN Curncy, AUD BGN Curncy, AUD3M BGN Curncy, CAD BGN Curncy, CAD3M BGN Curncy, GBP BGN Curncy, GBP3M BGN Curncy, JPY BGN Curncy, JPY3M BGN Curncy; these represent (respectively) 3M T-Bills, 3M EUR German Debt, 3M Australian Bank Bills, 3M Canadian Bankers Acceptances, 3M GBP LIBOR, 3M JPY LIBOR, EUR/USD Spot, EUR 3M Forward Points, AUD/USD Spot, AUD 3M Forward Points, CAD/USD Spot, CAD 3M Forward Points, GBP/USD Spot, GBP 3M Forward Points, JPY/USD Spot, and JPY 3M Forward Points.

E.3 Dual-Listed Companies Application: Price Data

We use publicly available data on dual-listed companies provided by Mathijs A. Van Dijk. This data is based on [De Jong, Rosenthal and Van Dijk \(2009\)](#) and can be found on the website: <http://www.mathijsavandijk.com/dual-listed-companies>. Van Dijk's data includes share prices, currency conversions, and volumes, which are used in our paper to compute the deviations of twin share prices from parity as well as the annual average daily dollar volume required to use the price impact specification in [Frazzini, Israel and Moskowitz \(2018\)](#).

E.4 Foreign Exchange Market: Institutional Background

Foreign Exchange (“FX” or “Currency”) trading takes place through bilateral agreements between market participants as well as across a wide variety of trading platforms. While some platforms facilitate anonymous trading, other platforms are relationship-based. In the context of our CIP application, it is helpful to understand the institutional details of the spot FX markets and the futures FX markets.

The spot FX market is a fragmented market where institutions or individuals can exchange currencies. The most commonly traded and liquid currencies are frequently referred to as the “majors,” though the definition can differ across papers and platforms. The USD (US Dollar), EUR (Euro Currency), JPY (Japanese Yen), GBP (British Pound), CHF (Swiss Franc), AUD (Australian Dollar), and CAD (Canadian Dollar) are the group of major currency pairs components, with the most commonly traded pairs being EUR/USD, JPY/USD, GBP/USD, USD/CHF, USD/CAD, and AUD/USD. When quoting a currency pair, the convention is that the price is quoted in units of the second currency per single unit of the first. By way of illustration, a price of 1.20 for the EUR/USD implies that 1 Euro = 1.20 USD. “Cross currencies” are combinations of liquid currencies that are less commonly traded than their dollar counterparts: these include markets like the EUR/GBP or EUR/JPY. As this paper considers deviations from covered interest parity between US and foreign rates, we consider only majors.

Spot FX is largely transacted through major dealing platforms, the two most well-known institutional platforms being EBS (owned by CME Group) and Refinitiv (owned by the London Stock Exchange Group). The platforms are offered in various flavors, some of which involve relationship-based dealing where liquidity takers (hedgers, hedge funds, and other buy-side institutions) and liquidity providers (market-makers, banks, and other sell-side institutions) interact subject to greater rules and longer processing delays (EBS Direct, for instance). Other platforms involve an anonymous centralized limit order book that would be familiar to participants in other centralized markets like futures. Swaps and forwards are also traded through dealing platforms and interbank relationships.

In parallel with the spot FX markets, futures FX markets provide trading opportunities in currencies. The primary venue for such trading is the CME Group, though other exchanges also offer such contracts throughout the world. The CME Group’s FX trading takes place on the Chicago Mercantile Exchange (CME) Designated Contract Market (DCM). The CME Group also offers trading in other sectors beyond FX, with trading taking place across the CME DCM as well as its other DCMs, the NYMEX (New York Mercantile Exchange, mostly energies), the COMEX (Commodity Exchange, mostly metals), and the CBOT (Chicago Board of Trade). Trading in FX futures takes place across individual contracts, which are specified in a quarterly cycle each year (the coding scheme used in this paper’s raw data files as well as by all futures exchanges is H = March, M = June, U = September, Z = December). Cessation of trading in quarterly FX contracts typically takes place two days before the third Wednesday of the delivery month, after which point the contract is physically settled via the contract seller’s delivery on the third Wednesday to the contract buyer via the exchange’s clearinghouse. All trades formally take place between market participants and the exchange clearinghouse so that the clearinghouse is the counterparty to all transactions.

Liquidity is fragmented in FX markets between the spot FX and futures markets, both of which can be thought of as approximately equal sources of direct, platform-based trading volumes. Although futures involve delivery of a specified amount of foreign exchange at a specific date in the future just like an OTC (over-the-counter) forward contract, institutional traders commonly express directional opinions about currencies through an almost interchangeable combination of futures and spot FX due to the similar price impact and the fact that price impact in one of the two markets spills over almost instantaneously into the other market. For example, if a large trader wished to purchase \$500M worth of Euros, the trader could

either purchase that amount in the spot FX market, purchase that amount via an appropriately sized futures trade (taking into account the fact that a single EUR/USD contract is for the purchase of 125,000 euros in exchange for dollars), or simultaneously purchase $\$Y$ in the spot market and $\$500M - Y$ in the futures market. All of these trades push the price of EUR/USD up by essentially the same amount because they are simply different ways of expressing the same trade.

Currencies are typically quoted to the fifth decimal place with bid-ask spreads commonly quoted as 1.21005/1.21010, for example, with a spread of half a “pip” (a “pip” is a ten thousandth, the fourth decimal place). If a participant sends an order to purchase simultaneously at two different venues (say, at the CME Group via both their futures market for EUR/USD as well as the EBS dealing platform), the participant can take advantage of any resting liquidity in both markets, but because of how closely spot foreign exchange markets and futures markets are tied together (by CIP), any large orders in one market have significant spillover effects in the other market if trades are executed sequentially. For example, a large buy order (relative to resting order liquidity) submitted in the futures generally push up the price in the spot market by approximately the same amount. Arbitrageurs can take advantage of any fleeting differences in price induced by the continuous limit order book structure, a form of “stale quote sniping”.

Spot FX markets and futures markets differ in their margin requirements as well. Spot FX markets generally allow very high levels of leverage, with the CFTC (Commodity Futures Trading Commission), which regulates currency and futures trading in the US, imposing restrictions that cap leverage at 50:1 for retail traders (who are naturally subject to more restrictive regulation than institutional traders). Futures markets feature margin requirements determined by the exchanges, which at the CME Group consist of the SPAN margining system. SPAN allows participants to see reductions in margin requirements based on relationships across contracts. For example, a position long 100 E-mini S&P 500 contracts and short 100 E-mini Nasdaq 100 contracts would require less margin than a position long 100 in both markets, as they typically have a very high positive correlation. Futures Commission Merchants (FCMs), which are the clearing member organizations that process trades from clients and clear them through the exchange clearinghouse, might require higher margin requirements as a buffer based on client credit risk.

F CIP Application: Additional Results

F.1 Price Impact Estimation

F.1.1 Summary of Estimates

Table OA-5 reports the average estimates of the price impact coefficients over our full sample.

Table OA-5: Price Impact Estimates

Market	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
α	0.1333	0.1289	0.1215	0.1225	0.1160
$SE(\alpha)$	0.00014	0.00012	0.00012	0.00008	0.00019

Note: Table OA-5 presents the price impact coefficients α and their standard errors when $\beta = \frac{1}{2}$ estimated over our full sample: December 15, 2019 to February 26, 2021.

Note that we estimate both α , the coefficient on the square root of the quantity traded, and, θ , the intercept. As expected, the daily estimates of θ hover around zero through the whole sample, with estimates

that are generally two orders of magnitude smaller than those generated by the α coefficients, so we set $\theta = 0$ for our computations.

F.1.2 Assumptions on Cross-Price Impact

The ideal experiment to measure the price impact function for the currency legs of a CIP arbitrage trade would involve the simultaneous placement of opposite signed orders in the spot and three-month forward markets followed by observations of changes in the price over the immediate period after the trades. By randomly varying the trade sizes and direction over time, one could estimate the total price impact of a marginal CIP trade on the spot and forward markets.

Instead, we parsimoniously model the cross-price impact for the CIP arbitrage spot and forward legs as follows. Let α be the estimated futures price impact coefficient for the square root functional form of Equation (23), which we take to be equal to the spot and three-month forward price impact coefficients. Let α^O denote the impact of a trade on the “other” market, which is merely the three-month forward market in the case of a spot transaction or the spot market in the case of a forward market transaction. Spillover effects make a transaction of size Q in, for instance, the spot market have an impact of $\alpha|Q|^{0.5}$ in the spot market and $\alpha^O|Q|^{0.5}$ in the three-month forward market, where $\alpha^O < \alpha$. A simultaneous trade of $-Q$ in the three-month forward market then has an impact of $-\alpha|Q|^{-0.5}$ in the three-month forward market and $-\alpha^O|Q|^{0.5}$ in the spot market.

Summing up, we find the impact of a partially gap-closing trade of size $(Q, -Q)$ for the spot and forward markets to be equal to $(\alpha - \alpha^O)|Q|^{0.5}$ and $-(\alpha - \alpha^O)|Q|^{0.5}$ respectively. In all our computations, we assume that $\alpha^O = \frac{9}{10}\alpha$ so that the price impact is the same as if α is multiplied by a factor of $\frac{1}{10}$. We consider this to be a conservative assumption since it assumes that to move the spot price by the same amount as a directional \$1B trade in a currency market, one would have to simultaneously purchase \$100B of spot currency and sell \$100B in the three-month forward market to achieve the same effect under the square root functional form.

F.1.3 Estimates of Price Impact at Longer Intervals

The microstructure literature on price impact distinguishes between permanent price impact and temporary/transient price impact. The former represents the impact of a trade on the equilibrium price, generally taken to mean the impact either as time goes to infinity or a sufficiently large amount of time in the future as to render the notion of temporary impact implausible. Temporary price impact can be taken to be the transient fluctuations of price that occur subsequent to an order as supply and demand reach their equilibrium. A large order mechanically impacts price in a continuous limit order book market, but in a market with nearly perfectly elastic demand, the order book immediately replenishes and price fluctuations revert.

Our framework call for estimating the impact of a trade in the equilibrium price, hence permanent price impact. To check whether our price impact estimates revert back toward zero as we increase the transaction count interval for our dependent variable, we estimate the following equation:

$$F_{\tau+n} - F_{\tau} = \theta + \alpha \operatorname{sgn}(Q_{\tau}) |Q_{\tau}|^{\beta} + \varepsilon_{\tau}, \quad (\text{A.28})$$

allowing n to vary between $n = 1$ and $n = 5,000$.

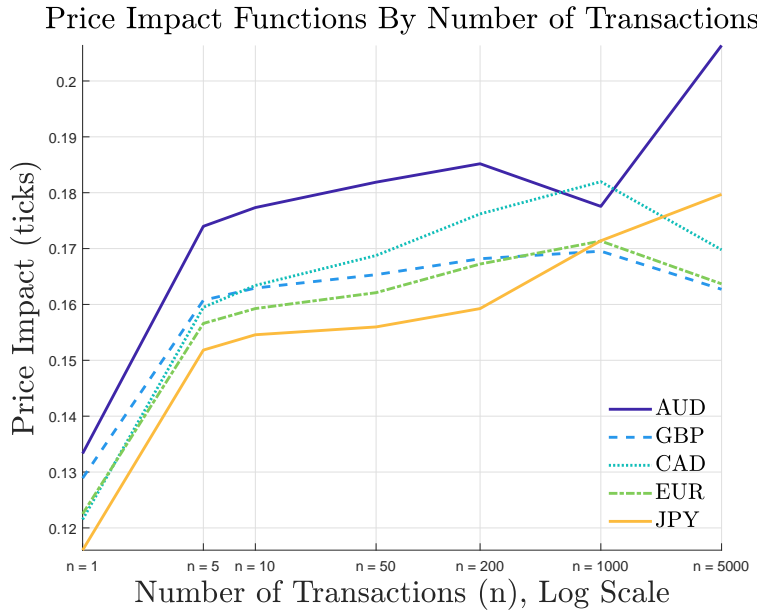


Figure OA-1: Price impact over long intervals

Note: Figure OA-1 shows the average estimate of the α coefficient from regressions of the form described in Equation A.28, when $\beta = 0.5$, where the dependent variable is the change in price over varying transaction time intervals. We show estimates for transaction intervals $n = \{1, 5, 10, 50, 200, 1000, 5000\}$. For visual clarity, we use a log scale when showing the horizontal axis. The baseline estimate in the text corresponds to $n = 1$. Estimated price impact increases when considering a few subsequent transactions, but then remains stable.

Figure OA-1 illustrates our findings when $\beta = 0.5$. Instead of finding a decreasing pattern, as would be expected with temporary price impact, we find an increasing pattern. This is unsurprising, as most of the price change occurs within several transactions of the initial transaction, likely because of autocorrelated order flows combined with slow diffusion of information. This finding is suggestive that our estimates capture permanent price impact, not microstructure noise.

F.2 CIP Deviations: Welfare Gains between 2010 and 2019

Figure OA-2 shows the welfare gains from closing CIP deviations from the beginning of 2010 through the end of 2019. To compute such welfare gains, we backward extrapolate the price impact estimates from the period 12/15/2019–02/26/2021, which we combine with the CIP deviations presented in Figure 5a. We recover estimates for welfare of similar order of magnitude as in the 2019–2021 period.




Figure OA-2: Welfare gains from CIP Arbitrage: 2010–2019

Note: Figure OA-2 shows the welfare gains from closing CIP deviations from the beginning of 2010 through the end of 2019. To compute such welfare gains, we backward extrapolate the price impact estimates from the period 12/15/2019–02/26/2021, shown in Figure 4, and combine them with the CIP deviations shown in Figure 5a. The Yen-Dollar cross-currency basis features deviations from CIP sufficient to cause a modest \$1.6B reduction in welfare at its extreme, while for other bases the welfare reductions remain small in magnitude.

F.3 CIP Deviations: Welfare Gains Isoquants

Figure OA-3 shows isoquants of CIP deviations and price impact estimates that yield the same level of welfare gains, corresponding to \$100M, \$1B, and \$10B in the EUR/USD case.




Figure OA-3: Isoquants: CIP Welfare Gains (EUR/USD)

Note: Figure OA-3 shows isoquants of CIP deviations and price impact estimates (the α coefficients) that yield the same level of welfare gains, corresponding to \$100M, \$1B, and \$10B in the EUR/USD case. A similar qualitative and quantitative relation applies to the other currency pairs considered in the paper. The solid blue dot represents the average estimates of CIP deviations and the average price impact coefficient α in our sample. The solid red dot represents the maximum measure of CIP deviations and the highest price impact estimate α in our sample. This figure illustrates that large welfare gains from closing CIP deviations can only emerge when CIP deviations are extremely large and price impact estimates are low.

F.4 CIP Deviations: LIBOR Estimates

We rerun the same analysis as the main body of the paper using exclusively the same LIBOR rates used in Du, Tepper and Verdelhan (2018) for the calculation of the various cross-currency bases. The results are detailed in Figure OA-4. The primary difference from the main results can be seen in the EUR/USD cross-currency basis, where a large spike toward the end of 2011 in welfare gains from closing arbitrage gaps is visible. Our main body’s results use secured three-month lending rates whenever such instruments are available, and do not find such a large cross-currency basis for T-Bills versus German three-month paper. One can then view the noticeable surge in 2011 as corresponding to differences between secured and unsecured lending in the midst of the European sovereign debt crisis.

As noted in Du, Tepper and Verdelhan (2018), the use of LIBOR to compute the cross-currency basis is imperfect because of the potential for some of the bases being accounted for by differential credit risk, which does not perfectly correspond to the pure arbitrage of identical legs described in our model. For completeness, however, we include it here and note that the switch to all unsecured lending rates does not change the order of magnitude of maximal welfare gains from closing arbitrage gaps and in fact, for the period where we are able to estimate daily price impact functions including the recent COVID-19 crisis, we find minimal gains from closing arbitrage gaps despite some of the most negative cross-currency bases in the larger dataset. We suspect that if we were able to obtain data for the EUR/USD futures in 2011, we would similarly find a larger price impact than our average estimates for December 2019–February 2021.

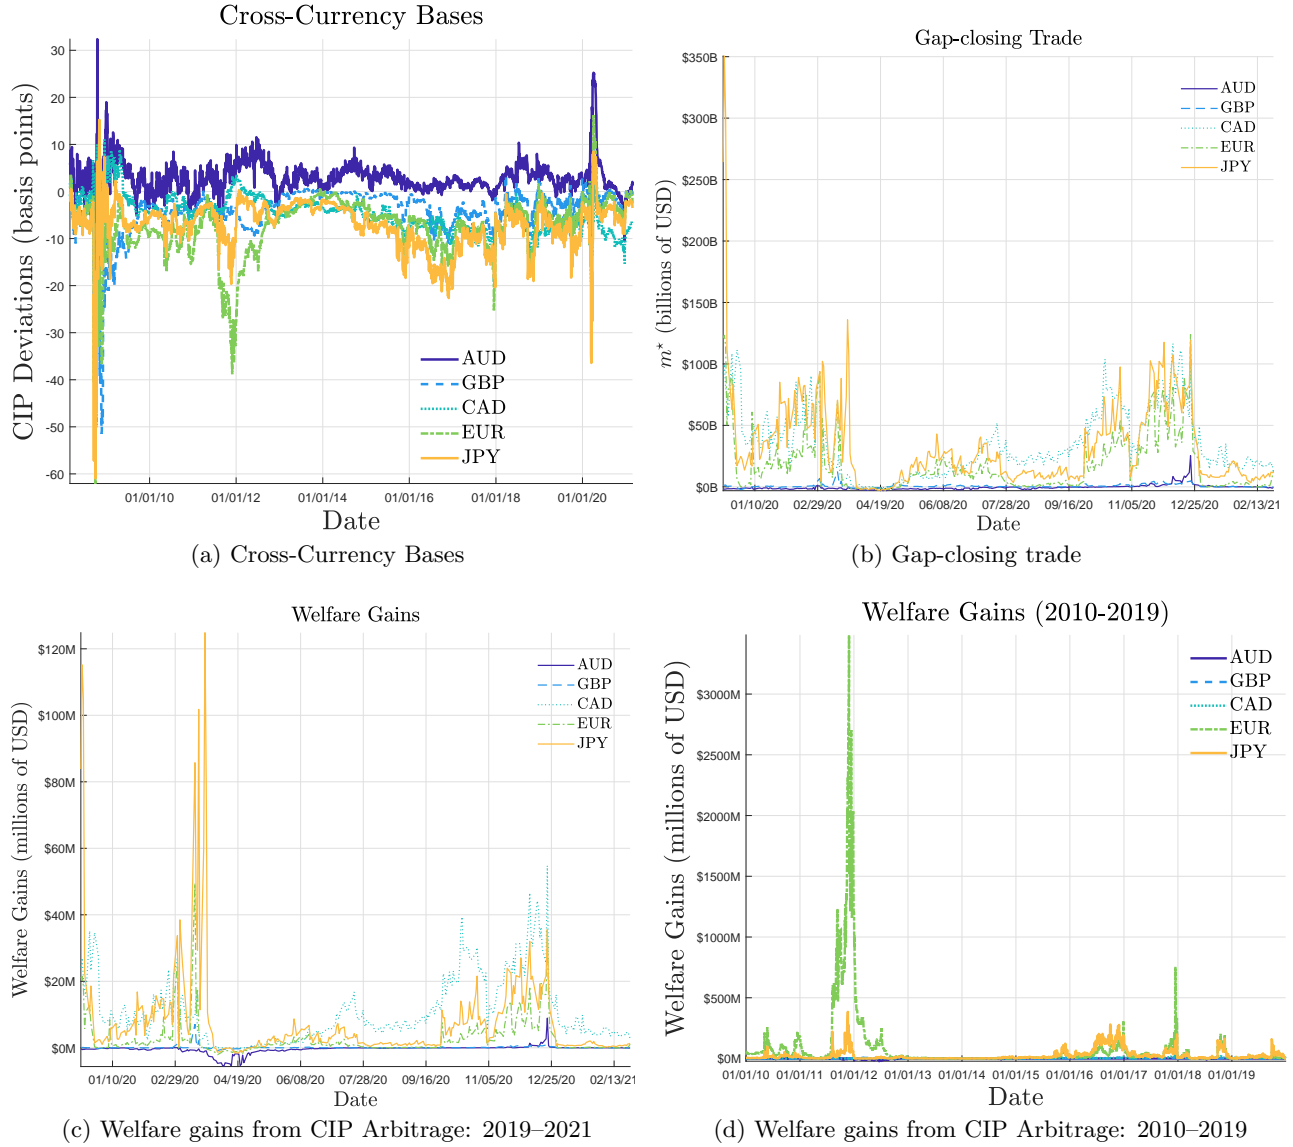


Figure OA-4: LIBOR Results

Note: Figure OA-4 includes the counterparts of Figures 5a, 6a, and 6b in the text, and Figure OA-2 of the Appendix, when using LIBOR as the riskless rate.

F.5 CIP Deviations: Estimates with Generalized Nonlinear Price Impact Estimates

In the text, we consistently make assumptions biased in favor of finding large welfare gains from CIP arbitrage: significant cross-price impact of simultaneous trades in spot and forward markets, spot and three-month forward market price impact consistent with estimated futures market price impact, and a square root functional form of price impact. Here, we relax the assumption of square root price impact to estimate a general nonlinear model of the form:

$$F_{\tau+1} - F_{\tau} = \theta + \alpha \operatorname{sgn}(Q_{\tau}) |Q_{\tau}|^{\beta} + \varepsilon_{\tau}, \quad (\text{A.29})$$

where β is not restricted to $\beta = \frac{1}{2}$. The theoretical literature generally specifies either $\beta = \frac{1}{2}$ or $\beta = 1$, while as previously discussed empirical work has largely found evidence consistent with an approximately square root price impact functional form.

For each of the five currency pairs, we estimate the general price impact function via Nonlinear Least Squares (NLLS) for each day of the sample. Finally, we estimate α and β using the entire sample and compute the usual heteroskedasticity-consistent standard errors: these can be found in Table OA-6. As is readily apparent from the table, the α coefficients are estimated more imprecisely than in the square root model because of the introduction of an additional free parameter, β , but the microstructure dataset is so vast that overprecision remains more concerning than imprecision; the β exponents are similarly found to have standard errors ranging from 0.0016 (EUR/USD) to 0.0099 (JPY/USD). Most interesting, however, are the estimates of β itself: instead of a square root functional form, we find evidence in favor of exponents between 0.56 and 0.64 across the currencies. We rerun our main welfare and gap-closing trade analysis, with results captured in Figure OA-5. Because of the extraordinarily large trade sizes, the difference between an exponent of 0.5 and 0.6 becomes substantial, resulting in minimal gains from closing arbitrage gaps and much smaller gap-closing trades. Looking at the crisis in March 2020 during which CIP deviations became substantial, we find a sudden increase in welfare gains from closing arbitrage gaps, with highs of around \$250M. This crisis-induced surge occurs despite the increase in the estimated α coefficients because the estimation finds a consistently smaller exponent β during the crisis stage.

Table OA-6: Price Impact Estimates: Generalized Price Impact Specification

Market	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
α	0.1123	0.1010	0.1087	0.1034	0.0937
$SE(\alpha)$	0.00043	0.00047	0.00034	0.00024	0.0015
β	0.5955	0.6441	0.5697	0.6049	0.6359
$SE(\beta)$	0.0022	0.0028	0.0022	0.0016	0.0099

Note: Table OA-6 presents the price impact parameters and their standard errors estimated over our full sample: December 15, 2019 to February 26, 2021.

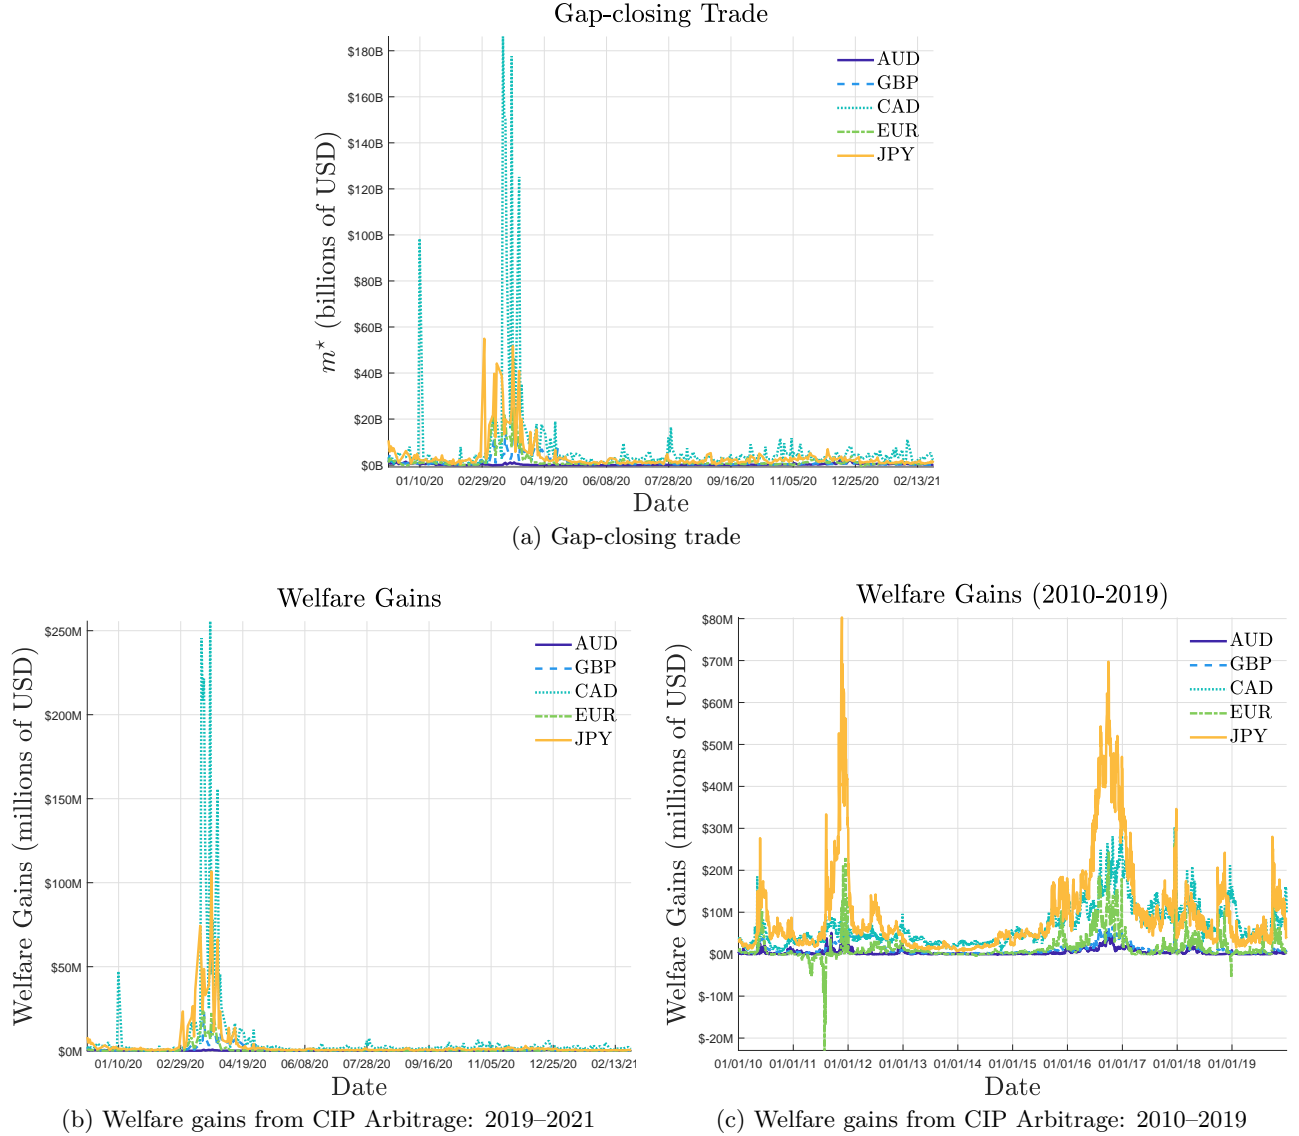


Figure OA-5: Generalized Price Impact Estimation Results

Note: Figure OA-5 includes the counterparts of Figures 6a and 6b in the text, and Figure OA-2 of the Appendix, when using estimates of a generalized nonlinear price impact function.

F.6 CIP Deviations: Estimates with Fixed Nonlinear Exponents

In the text, following existing literature, we adopt $\beta = 1/2$ as our baseline value for the power law parameter. In Section F.5, we freely estimated a nonlinear model via Equation A.29. Here, we explore how fixing the exponent in the equation impacts the estimated welfare gains from closing arbitrage gaps. Specifically, we use the same specification as before but we fix $\beta = \beta_0$, a constant, which we vary from 0.3 to 0.6, in 0.1 intervals:

$$F_{\tau+1} - F_{\tau} = \theta + \alpha \operatorname{sgn}(Q_{\tau}) |Q_{\tau}|^{\beta_0} + \varepsilon_{\tau}. \quad (\text{A.30})$$

Figure OA-6 illustrates the results. Imposing values for β lower than 0.5 results in larger gains from closing arbitrage gaps. Note that the baseline $\beta = 0.5$ assumption is already conservative in light of the previous

subsection’s general nonlinear model finding power laws with beta closer to 0.6 than 0.5; this means that the values of beta required to generate large estimates of gains from closing arbitrage gaps are implausible. For instance, in the case of $\beta = 0.3$ for the JPY/USD pair, the maximal gap-closing trade equals nearly a millennium’s worth of GDP.

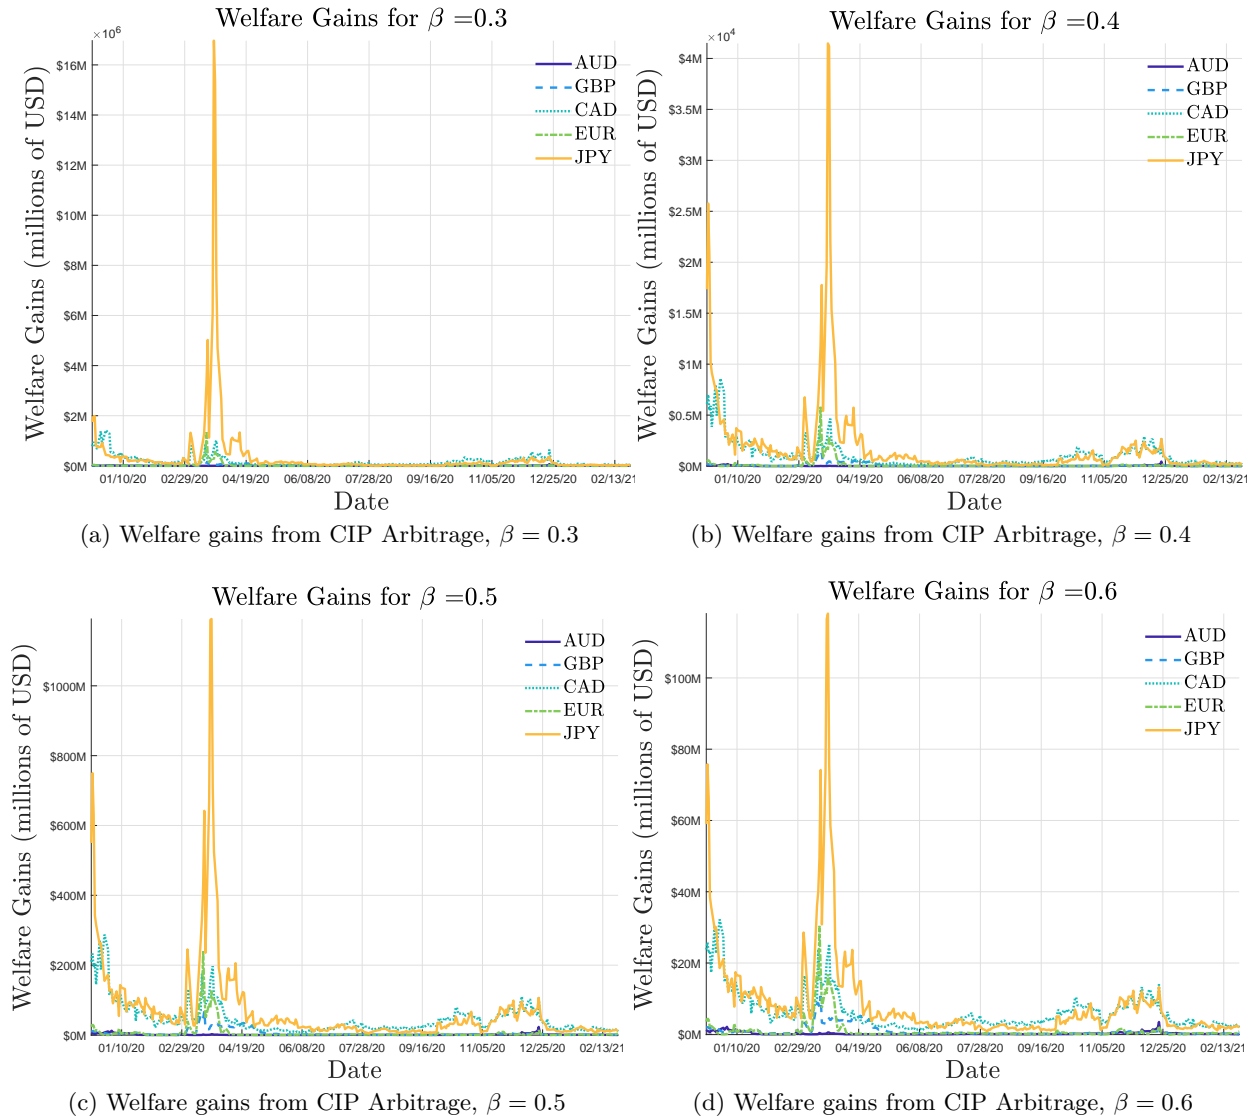


Figure OA-6: Price Impact Estimation Results with Varying Exponent

Note: Figure OA-6 shows the counterparts of Figure 6b in the text when imposing power law parameters $\beta_0 = \{0.3, 0.4, 0.5, 0.6\}$ when estimating Equation A.30.

G Price Impact Estimates: Gabaix and Koijen (2021) Multipliers

Gabaix and Koijen (2021) highlight the importance of stock market order flows in causing fluctuations in stock market prices. Using a granular instrument variables (GIV) approach, which uses the granular character of idiosyncratic trading by large institutions, they show that flows into US equities generally have a “multiplier” of approximately 5. This means that, according to their estimates, one dollar of flows into the stock market

increases the market capitalization of the stock market by around five dollars. This large multiplier implies that demand is more inelastic than standard frictionless models predict. [Gabaix and Kojien \(2021\)](#) refer to this insensitivity to price as the “inelastic markets hypothesis.” As a consequence of our model’s need for price impact estimates to compute welfare gains from closing arbitrage gaps, we have tested this hypothesis in currency markets and find strong evidence that even modest flows of funds in currency markets can have a material impact on prices.

Table OA-7: Price Impact Estimates: [Gabaix and Kojien \(2021\)](#) Style Multipliers

	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
\$50M Trade	7.042	13.869	7.210	36.087	12.875
\$100M Trade	4.979	9.807	5.098	25.517	9.104
\$1B Trade	1.575	3.101	1.612	8.069	2.879

Note: Table [OA-7](#) reports multipliers in the style of [Gabaix and Kojien \(2021\)](#) — expressed as the percentage movement in the currency price versus the percent of foreign GDP that the order represents — from our estimation of price impact in the FX futures markets. Since our estimates of price impact are nonlinear, we report estimates of the multipliers for different trade sizes. Multipliers are decreasing in the size of the trade, since we estimate a concave price impact function. We find a high multiplier for the EUR, JPY, and GBP because a given trade size represents a very small amount as a fraction of GDP for those countries/regions. See [Table OA-8](#) for exact computations.

Table [OA-7](#) shows our estimates of the multipliers for different trade sizes: the percentage movement in the currency price versus the percent of foreign GDP that the order represents. For example, if we had a row for \$180B it would represent around 1% of Eurozone GDP, and so the multiplier would be equal to the percent price change that we estimate would be induced by the order divided by 1%. We must present different multipliers depending on the size of the trade because we find strong support in currencies for the approximately square root price impact specification of [Gabaix et al. \(2006\)](#); small trades then carry large multipliers while the largest trades carry much smaller multipliers. The biggest trades listed in this table correspond to over 10% of a typical day’s dollar transaction volume in outright futures contracts at the CME and can be considered comparable to transactions from the largest institutions. These values represent significant responses of price to institutional trade sizes and underline the central message of [Gabaix and Kojien \(2021\)](#) that order flows by large institutions and investors create much larger market fluctuations than captured by most current models.

Table OA-8: Percent (in bps) of foreign GDP of a given order

	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
\$50M Trade	0.35714	0.17668	0.28736	0.02732	0.09843
\$100M Trade	0.71429	0.35336	0.57471	0.05464	0.19685
\$1B Trade	7.14286	3.53357	5.74713	0.54645	1.96850

Note: Table [OA-8](#) reports the size of the trade for each currency normalized by the GDP of the country. That is, a \$100M trade of AUD represents 0.71bps of Australia’s GDP.

The diminishing multiplier we find versus the linear GIV approach used in [Gabaix and Kojien \(2021\)](#) on equity market data is why welfare gains from closing arbitrage gaps are a convex function of the price wedge. Linear price impact results in the estimation of a triangular wedge that can be substantial even for

moderate gap-closing trades. Instead, concave price impact implies that small wedges will be closed with small orders, so the importance of price wedges to welfare is an increasing function of the absolute size of the wedge.