The Value of Arbitrage*

Eduardo Dávila† Daniel Graves‡ Cecilia Parlatore§

July 2021

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 social marginal value of executing an arbitrage trade. We further show that arbitrage gaps and measures of price impact are sufficient to compute the aggregate total gain from closing an arbitrage gap. Theoretically, we show that, for a given arbitrage gap, the total 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 range of estimates of the value of closing arbitrage gaps varies across applications.

JEL Codes: G12, G18, D61

Keywords: arbitrage, welfare, price impact, covered interest parity, dual-listed stocks/Siamese twins

*We would like to thank Markus Brunnermeier, Ricardo Caballero, John Campbell, John Cochrane, Wenxin Du, Darrel Duffie, Emmanuel Farhi, Xavier Gabaix, Joel Hasbrouck, Ralph Koijen, Guido Lorenzoni, Urban Jermann, Martin Oehmke, Stavros Panageas, Adriano Rampini, Alexi Savov, Jesse Schreger, Jeremy Stein, Dimitri Vayanos, and Adrien Verdelhan for helpful comments and discussions. Jalen Li, Jack Kelly, and Flint O'Neil provided excellent research assistance. Daniel Graves gratefully acknowledges support of the NSF GRFP program; this material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1752134.

†Yale University and NBER. Email: eduardo.davila@yale.edu

‡Yale University. Email: d.graves@yale.edu

§New York University, Stern School of Business, NBER and CEPR. Email: cparlato@stern.nyu.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 price 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 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 to move from an autarky equilibrium to a fully integrated equilibrium.

Our first main result shows that the social marginal 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 social marginal 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 social marginal value of arbitrage. Intuitively, the difference in prices fully captures the difference in willingness to pay between investors in different markets.

While the arbitrage gap measures the social marginal value of arbitrage, an arbitrage trade creates distributive pecuniary externalities — using the terminology of Dávila and Korinek (2018) — which cancel out in the aggregate. Accounting for these externalities, we

[1] In Section B in 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.

can also characterize the individual marginal 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 b) the welfare of the arbitrageur sector — given by the arbitrage profits — initially increases with the scale of the arbitrage trade, reaches a maximum, and then decreases, becoming zero once the arbitrage gap is closed.

Our second main result shows, exploiting the fundamental theorem of calculus, that the aggregate *total* value of arbitrage can be recovered by measuring the arbitrage gap for different levels of the arbitrage trade m. Intuitively, while prices are sufficient 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. Therefore, 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 our results do not rely on specific functional form assumptions and apply widely to a large class of environments.

Our third main result shows that the value of arbitrage depends on the degree of market illiquidity. For a given arbitrage gap, the aggregate total 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 welfare losses are small. Alternatively, finding an arbitrage gap in a market in which prices are not very sensitive to the quantities traded implies that the value of arbitrage is potentially large.

In the Appendix of our paper, we generalize our results to dynamic environments and environments with uncertainty, multiple assets, and arbitrary cross-holdings by investors. These extensions allow us to refine the interpretation of the results of the baseline model and give us the opportunity to apply the insights from our approach in complex real-world scenarios.

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 compute the magnitude of the gap-closing arbitrage trades and the value of arbitrage for different currency pairs. We find that all welfare gains from closing CIP arbitrage, even during the March 2020 start of the COVID-19 crisis in the US, are less than $1B in magnitude, which can be considered small in relation to the size of FX markets. The largest estimated gap-closing arbitrage trades in the sample period do not exceed $1T and never exceed $250B outside the yen-dollar basis. Our findings can be explained by the fact that precisely when CIP violations reach their peak in March 2020, price impact also surges, resulting in only a moderate increase in the gap-closing trade size. More broadly, scenarios of market distress

are commonly associated with illiquidity, which can both contribute to divergence from covered interest parity as well as increases in the price impact of 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"). 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 to dual-listed companies, 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 $1.7 trillion in stock transactions spanning two decades. We find that the welfare gains from arbitrage in the Royal Dutch/Shell case peak at approximately $2B in 1996 when converted to USD. This price divergence appears to cause a significant welfare loss, in contrast to our CIP results. However, the magnitude of the arbitrage gap for Royal Dutch/Shell is extreme – the price deviation at the time of maximum welfare gains is 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 20B pound sterling is required to close the arbitrage gap. Unlike with CIP deviations, the persistence of extreme deviations from parity over many years results in an extended period over which the welfare gains from closing price deviations 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 PLC/Ltd, which experiences the smallest mean absolute divergence of all the studied twin shares, we find minimal welfare gains from closing price deviations 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, consistent with our theory.

Related Literature The absence of arbitrage opportunities is considered by many to be the fundamental theorem of finance, 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 funding constraints in Xiong (2001); of collateral constraints in Gromb and Vayanos (2002),

Liu and Longstaff (2004), 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 agents 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 (2014). 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 are the potential gains from reducing business cycles, which corresponds to 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.[2]

Our empirical exercises are motivated by the recent evidence documenting and rationalizing the systematic breakdown of covered interest parity as well as by the persistence of deviations from arbitrage relations in dual-listed companies. The covered interest parity (CIP) literature includes Du, Tepper and Verdelhan (2018), who find persistent deviations from CIP starting with the Great Financial Crisis. They demonstrate that these violations occur for many currency pairs in the post-crisis period 2010 to 2016, and 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. Meanwhile, Amador, Bianchi, Bocola and Perri (2020) develop a model of how CIP violations can arise as a direct result of the zero lower bound constraints on monetary policy

[2]Conceptually, 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.

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 work of Froot and Dabora (1999) and De Jong, Rosenthal and Van Dijk (2009). 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 motivated by extensive theoretical and empirical literatures that stress the square root 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, Gopikrishnan, Plerou and Stanley (2003) 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 data sets and varied empirical approaches. In particular, Gabaix, Gopikrishnan, Plerou and Stanley (2006) find a square root power law using large-cap French stock data from the 1990s, 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 asset prices. Section G in 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. Sections 4 and 5 develop the empirical applications and Section 6 concludes. All proofs and derivations are in the Appendix. The Appendix also include extensions of our baseline model and robustness results for our applications.

2 Baseline Model

In this section, 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 the approach that we develop in this paper. In the Appendix, we generalize our results to an environment with uncertainty and multiple assets.

2.1 Environment

There are two dates $t = \{0, 1\}$ and a single consumption good (dollar), which serves as numeraire. There is no uncertainty. There are 2 markets, indexed by A and B. The economy is populated by type A investors, type B investors, and arbitrageurs. Type A investors exclusively trade in market A, while type B investors exclusively trade in market B.

In market A, type A investors (in unit measure) trade a risk-free asset (asset A) with payoff $d_1 > 0$ at date 1. The price of this asset is p^A. In market B, type B investors (also in unit measure) trade a different risk-free asset (asset B) with an identical payoff d_1 at date 1. The price of this asset is p^B.

First, we describe the problem that both types of investors face. Subsequently, we describe the problem of arbitrageurs.

Market A: Investors' Problem Investors in market A have time-separable utility, with a flow utility of consumption $u_A(c)$, which satisfies standard regularity conditions, and a discount factor β_A. They have dollar endowments n_0^A and n_1^A and hold an initial position q_{-1}^A in the traded asset. Hence, the demand of type A investors for asset A is given by the solution to

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

subject to the budget constraints

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

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: Investors' Problem Investors in market B face the same problem as investors in market A. Investors in market B also have time-separable utility, with a flow utility of consumption $u_B(c)$, which satisfies standard regularity conditions, and a discount factor β_B. They have dollar endowments n_0^B and n_1^B and hold an initial position q_{-1}^B in the traded asset. Hence, type B investors choose q_0^B as the solution to

$$\max_{q_0^B} u_B\left(c_0^B\right) + \beta_B u_B\left(c_1^B\right)$$

subject to the budget constraints

$$p^B \Delta q_0^B + c_0^B = n_0^B$$
$$c_1^B = n_1^B + d_1 q_0^B,$$

where $\Delta q_0^B = q_0^B - q_{-1}^B$ and where c_0^B and c_1^B 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. For simplicity, we assume that arbitrageurs have no initial endowments of dollars or assets, and that their flow utility is linear. Arbitrageurs implement a trading strategy with zero cash flows at date 1 and receive the date 0 revenue generated by such a strategy. Formally, denoting the asset purchases of arbitrageurs in markets A and B by $q_0^{\alpha A}$ and $q_0^{\alpha B}$, respectively, the date 0 revenue of arbitrageurs is given by

$$-\left(p^A q_0^{\alpha A} + p^B q_0^{\alpha B}\right), \tag{1}$$

which is subject to the following zero-cash-flow constraint at date 1:

$$d_1 \left(q_0^{\alpha A} + q_0^{\alpha B}\right) = 0. \tag{2}$$

From Equation (2), it follows that arbitrageurs always follow a one-for-one long/short strategy that includes assets A and B, that is, $q_0^{\alpha A} = -q_0^{\alpha B}$. By defining

$$m \equiv q_0^{\alpha A} = -q_0^{\alpha B}$$

as the size/scale of the arbitrage trade, we can express the utility of the arbitrageurs as a function of m, denoted by $V^\alpha\left(m, p^A, p^B\right)$, after normalizing their date 1 utility to 0, as

$$V^\alpha\left(m, p^A, p^B\right) = \left(p^B - p^A\right) m. \tag{3}$$

Whenever $m > 0$, arbitrageurs are buyers in market A and sellers in market B, and vice versa when $m < 0$. As we describe next, instead of making assumptions on the behavior of arbitrageurs (i.e., whether they are competitive, strategic, face different types of financing constraints, etc.), we take the scale of the arbitrage trade m as a primitive of the model and focus on the impact of varying m on social welfare.[3]

[3]If arbitrageurs had concave utility, with endowments n_0^α and n_1^α, their utility would take the form:

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

2.2 Equilibrium

We now define and then characterize a notion of competitive equilibrium — an arbitrage equilibrium — in this environment.

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

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

Going forward, we proceed under the two following 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 ($m = 0$). Formally,

$$p^B(0) - p^A(0) > 0.$$

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

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

This 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 4 and 5 are consistent with this assumption. It is straightforward to provide conditions on primitives that guarantee that this assumption is satisfied.

An important object in our main results is the price differential between the two identical assets, in particular 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*.[4] Formally, $\mathcal{G}_{BA}(m)$ is given by

$$\mathcal{G}_{BA}(m) \equiv p^B(m) - p^A(m). \qquad \text{(arbitrage gap)} \tag{4}$$

In terms of the newly defined arbitrage gap, $\mathcal{G}_{BA}(m)$, our second assumption implies that $\mathcal{G}'_{BA}(m) = \frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm} < 0$. That is, the arbitrage gap narrows as the size of the arbitrage trade m increases. We denote by $m^\star$ the scale of the arbitrage trade that closes

[4]Sometimes $\mathcal{G}_{BA}(m)$ is called an arbitrage basis.

the gap, such that $p^A(m^\star) = p^B(m^\star)$. 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 $[0, m^\star]$.

Figure 1 illustrates the behavior of the allocations and prices of an arbitrage equilibrium as a function of m. Its left panel shows the equilibrium allocations of assets A and B for type A and type B investors as a function of the scale of the arbitrage trade m. Its right panel shows the equilibrium prices in market A and market B as a function of the scale of the arbitrage trade m. The right panel illustrates how the arbitrage gap converges to 0 as m approaches the gap-closing trade $m^\star$.

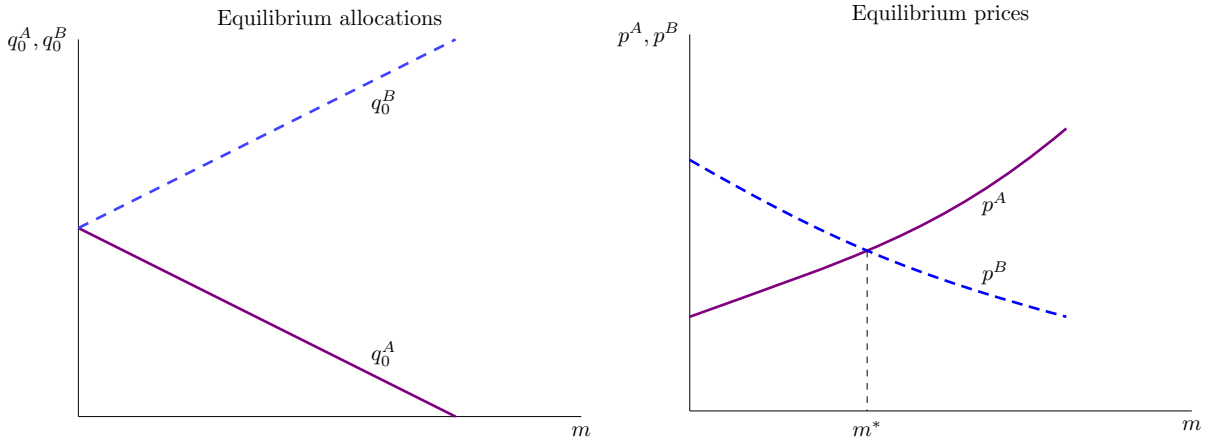


Figure 1: Equilibrium allocations and prices

Note: The left panel of Figure 1 shows the equilibrium allocation of the type A and B investors as a function of the scale of the arbitrage trade m. The right panel of Figure 1 shows the equilibrium price in market A and market B as a function of the size of the arbitrage trade m. It includes the gap-closing arbitrage trade, $m^\star$.

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

Remark 1. *(Unspecified limits to arbitrage)* It is obvious that a fully unconstrained perfectly competitive sector would compete away any arbitrage profits until $p^B(m^\star) = p^A(m^\star)$. Therefore, when indexing the equilibrium by the scale of the arbitrage trade m, we implicitly assume that there is some 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 independent of the specific underlying frictions that are relevant for a particular scenario, we purposefully avoid modeling any friction explicitly. However, we know that there is a one-to-one mapping between varying the scale of arbitrage m and loosening/tightening the friction that prevents complete price equalization.[5] In Section B

[5]If the scale of arbitrage m were chosen by a mass of unconstrained competitive arbitrageurs, the solution

in the 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 in this section.

Remark 2. *(Autarky equilibrium and integrated equilibrium as special cases of arbitrage equilibrium)* In the Appendix, we define two related notions of equilibrium: 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 and arbitrageurs are redundant. In an integrated equilibrium, $p^A = p^B$ by assumption. In general, it is hard to compare the outcomes of an autarky equilibrium and an integrated equilibrium. The approach we develop in this paper can be interpreted as a smooth way to connect both equilibria. In the Appendix, we formally show that increasing the scale of the arbitrage trade m is a form of smoothly converging to the integrated equilibrium from the autarky equilibrium.

3 The Value of Arbitrage

After characterizing the form taken by an arbitrage equilibrium, we now study its welfare properties. We denote the indirect utilities of type A and B investors as a function of the relevant asset price by $V^A\left(p^A(m)\right)$ and $V^B\left(p^B(m)\right)$, respectively. As described in Equation (3), $V^\alpha\left(m, p^A(m), p^B(m)\right)$ denotes the utility of arbitrageurs.

3.1 Marginal Value of Arbitrage

First, we characterize how individual and aggregate welfare change when we vary the scale of the arbitrage trade m. We refer to this as the marginal value of arbitrage. We express all individual welfare assessments in money-metric terms, measured in date 0 dollars, i.e., normalizing by an agent's date 0 marginal utility of consumption, denoted by λ_0^A, λ_0^B, and λ_0^α, respectively, and aggregate welfare in terms of this money-metric representation.[6] Formally, the change in aggregate welfare induced by increasing the scale of the arbitrage

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

[6]This approach can be interpreted as selecting a set of uniform "generalized social marginal welfare weights" — described in Saez and Stantcheva (2016) — for all agents.

trade m, which we denote by $\frac{dW}{dm}$, takes the form

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

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

$$\begin{aligned}
\frac{\frac{dV^A}{dm}(m)}{\lambda_0^A} &= \frac{dp^A(m)}{dm} m > 0 \\
\frac{\frac{dV^B}{dm}(m)}{\lambda_0^B} &= \frac{dp^B(m)}{dm}(-m) > 0 \\
\frac{\frac{dV^\alpha}{dm}(m)}{\lambda_0^\alpha} &= \left(\frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm}\right) m + p^B(m) - p^A(m) \gtreqless 0.
\end{aligned}$$

b) (Social marginal value of arbitrage) The social marginal 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

$$\frac{dW(m)}{dm} = p^B(m) - p^A(m) > 0. \tag{5}$$

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, a unit increase in the scale of the arbitrage trade yields a profit of $p^B(m) - p^A(m)$ dollars to the arbitrageurs. The pecuniary effects can be interpreted as the "distributive" pecuniary externalities of increasing the scale of the arbitrage trade, using the terminology in Dávila and Korinek (2018). These externalities, which include the terms in which $\frac{dp^A(m)}{dm}$ and $\frac{dp^B(m)}{dm}$ appear, are well understood.[7] In Walrasian environments, the dollar value of changes in equilibrium prices in any given market is zero-sum, due to market clearing. Note that Proposition 1a) implies that the impact of pecuniary effects is zero when $m \to 0$. This

[7] Gromb and Vayanos (2002) is the first paper that identifies this type of pecuniary externality in models of arbitrage. See Geanakoplos and Polemarchakis (1986), Lorenzoni (2008), or Dávila and Korinek (2018), for work that studies this type of externality in different contexts.

implies that only the direct effects that accrue to arbitrageurs matter when the amount arbitraged is small.

Proposition 1b) shows that the pecuniary welfare effects cancel out in the aggregate in dollar terms for any value of the arbitrage trade m — not only when $m \to 0$ — leaving the direct effects as the single source of marginal aggregate welfare gains. While intuitive, we are to our knowledge the first to identify the arbitrage gap as the measure of the social marginal value of arbitrage in a general equilibrium environment. 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 social marginal value of arbitrage.

While most of the paper is focused on characterizing and measuring aggregate welfare effects, Proposition 1a) has clear 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 size 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.*

Figure 2 illustrates the individual and social marginal value of arbitrage for each of the agents. While the welfare of investors in both markets increases with m, the welfare of the arbitrageur sector follows a "Laffer" curve. First, the structure of the arbitrage problem is such that the distributive pecuniary effects always 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.

Second, an increase in m always nets arbitrageurs the arbitrage gap, but the distributive pecuniary effects are always negative for them, for the opposite reasons that distributive effects make investors better off: an increase in m increases the price in market A, in which arbitrageurs buy, and lowers the price in market B, in which arbitrageurs sell. If we had a monopolist arbitrageur, m would be given by the solution to $\frac{dV^\alpha}{dm} = 0$.[8] Interestingly, introducing a monopolist arbitrageur generates a Pareto improvement in the

[8]Formally, the optimality condition for m of a monopolist arbitrageur is $m = \frac{p^B(m) - p^A(m)}{-\left(\frac{dp^B(m)}{dm} - \frac{dp^A(m)}{dm}\right)}$, while the condition that defines the equilibrium when the arbitrageur sector is competitive is given by $p^B - p^A$.

model considered here. As discussed above, competitive arbitrageurs would trade until $p^B = p^A$, in turn eliminating all arbitrage profits.

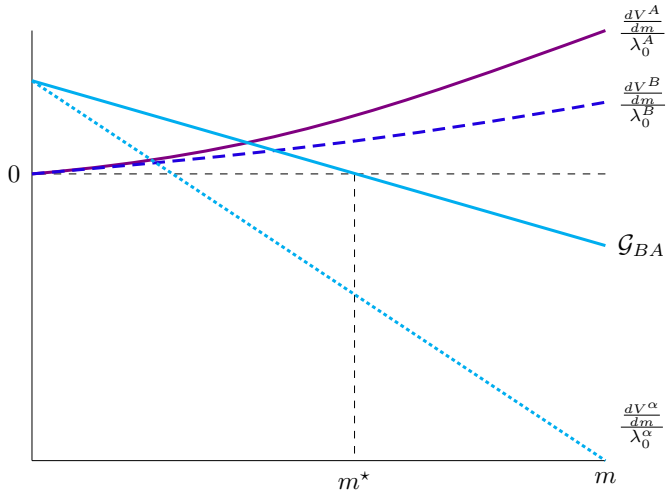


Figure 2: Marginal Value of Arbitrage

Note: Figure 2 shows the marginal value of arbitrage for type A investors, for type B investors, and for arbitrageurs, all as a function of the scale of the arbitrage trade m, as characterized in Proposition 1. For reference, it also shows the arbitrage gap.

While the distributional consequences of closing arbitrage gaps are interesting and worth studying further, our paper is focused on aggregate welfare implication for two reasons. First, in more complex environments in which a given investor trades in many asset markets — like the one considered in Section C — 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 that same individual. Second, in the context of the empirical applications, one would need i) portfolio level information (net trade positions), ii) investors' valuations (differences in MRS) for different investors, and iii) the pecuniary impact of arbitrage trades to be able to measure individual welfare effects.[9] Measuring these objects lies outside of the scope of this paper.

In the following remark, we highlight that our model and, consequently, our welfare results, do not account for additional externalities.

Remark 3. *(Absence of additional externalities)* Note that the economy presented 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 social marginal value of arbitrage would have to account for how changing m affects those.

[9] Dávila and Korinek (2018) identify these three elements as the key determinants of distributive pecuniary externalities. Despite its importance, there has been no theory-based empirical work of any of these elements.

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. Importantly, it is not the presence of additional choices which calls for augmenting the characterization of marginal social welfare gains, but whether these choices are associated with externalities, i.e., are constrained inefficient, or not. In any case, the direct marginal welfare effect of closing gap, in Equation (5), never vanishes. For this reason, it is natural to interpret our welfare calculations as capturing the direct welfare gains from closing arbitrage gaps.

3.2 Total Value of Arbitrage

Next, by relying on the fundamental theorem of calculus, we characterize the total value of arbitrage and study its properties. Proposition 2 introduces the second main result of the paper.

Proposition 2. (Social value of arbitrage)

a) (Social value of arbitrage) The (total) social value of arbitrage, that is, the aggregate value associated with a change in the size of the arbitrage trade from m_0 (any given arbitrage trade) to $m^\star$ (the gap-closing arbitrage trade), measured in date 0 dollars, is given by

$$W(m^\star) - W(m_0) = \int_{m_0}^{m^\star} W'(m)\, dm = \int_{m_0}^{m^\star} \mathcal{G}_{BA}(m)\, dm. \tag{6}$$

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, $\frac{dp^A}{dm}$ and $\frac{dp^B}{dm}$, to exactly compute the social value of arbitrage, since

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

Proposition 2a) combines the fundamental theorem of calculus and our characterization in Proposition 1b). Intuitively, if the social marginal value of arbitrage is given by $\mathcal{G}_{BA}(m)$, by adding over the 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 gaps evolves with m (through price impact) is sufficient to be able to compute the social value of arbitrage.

It is worth highlighting that the exact characterization of the 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 the size of the gap-closing arbitrage trade, 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 4 and 5.

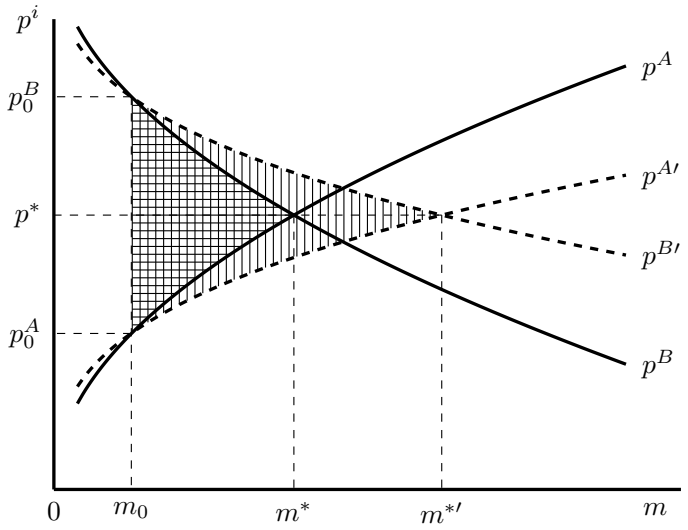


Figure 3: The value of arbitrage and market liquidity

Note: Figure 3 illustrates the social value of closing an arbitrage gap for two different economies that start with the same arbitrage gap. The area with horizontal lines defines the total social value of arbitrage for the economy in which markets are illiquid (price impact is high, so a given arbitrage trade moves prices substantially). The area with vertical lines defines the social value of arbitrage for the economy in which markets are liquid (price impact is low, so an identical arbitrage trade moves prices substantially).

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:

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

However, since finding the gap-closing arbitrage trade $m^\star$ implicitly involves making assumptions about price impact in both markets, there is little to gain by using this approximation instead of the exact characterization in Equation (6).

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, which can be seen as our main theoretical result, shows that the value of arbitrage depends on the degree of market illiquidity 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. *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 the market is illiquid implies that welfare losses are small. Alternatively, finding an arbitrage gap in a market in which prices are not very sensitive to the quantities traded implies that the value of arbitrage is potentially large.

Figure 3 illustrates Proposition 3 by comparing two economies that when $m = m_0$ have the same arbitrage gap, which implies that the social marginal value of arbitrage is the same, but that feature different levels of price impact. Given that, from Proposition 2, we can express $W(m^\star) - W(m_0)$ as $\int_{m_0}^{m^\star} \left(p^B(m) - p^A(m)\right) dm$, the triangle between the pricing functions defines the social value of arbitrage. It therefore becomes evident that steeper pricing functions (associated with more illiquid markets) generate lower welfare costs and vice versa. In practical terms, Proposition 3 opens the door for small arbitrage gaps in highly liquid markets to be associated with potentially large welfare losses. In Section 4 onwards, we discuss whether this is empirically the case.

4 Application 1: Covered Interest Parity

In our first empirical application, we use of our theoretical results to provide estimates of the welfare losses associated with violations of Covered Interest Parity (CIP). 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 2008 financial crisis, CIP went from consistently holding within a narrow band to being systematically violated for many currencies, 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 significant CIP deviations as a starting point and, building on the theoretical results derived in the previous two sections, explore instead the welfare losses associated with such deviations.

4.1 Measurement Approach

To simplify the exposition, in Sections 2 and 3 we presented our theoretical results in a single-currency environment. In Proposition 4, we characterize the social marginal value of arbitrage in the context of a multi-currency model — fully developed in the Appendix — that nests our baseline model and allows for deviations from CIP.

Proposition 4. *(CIP: Social marginal value of arbitrage) The social marginal 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*

$$\frac{dW(m)}{dm} = \frac{S(m)}{F(m)} \frac{1}{1+r^f(m)} - \frac{1}{1+r^d(m)}, \tag{7}$$

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 7, is positive, so that the return on the foreign leg is lower than the return on the domestic leg, then there is scope to sell the foreign leg and buy the domestic leg to close the arbitrage gap.[10] 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 value of arbitrage in Proposition 4, we can use Proposition 2 to find the total value of arbitrage. Therefore, conceptually, to find empirical counterparts of the social total value of arbitrage, we must i) measure the cross-currency basis, which is easily observable, and ii) estimate price impact measures for each of the markets involved in the arbitrage trade, which is substantially harder. In practice, we will use a high-frequency dataset of the FX futures market, whose centralized order book and

[10] Note that in a single currency economy, $S(m) = F(m) = 1$, and Equation (7) collapses to Equation (5).

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 are forced to base our welfare calculations on the following three assumptions.

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 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 maximal size of the 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 on the true available liquidity, the impact of executing the gap-closing trades on interest rates will be extremely small, justifying our assumption. Note that this assumption biases our estimate of welfare gains upwards. In other words, if we allowed for 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.[11] 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 each are at some point in time equivalent to the futures).

Finally, we must make an assumption on cross-price impact. Because the spot and forward prices are almost perfectly 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%. With a square root functional form, this is the same as assuming that $100B of transactions has 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 us against finding limited welfare gains. We run robustness checks under alternative assumptions in the Appendix.

Overall, these three assumptions are conservative in the sense that they bias our

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

results upward and against finding limited welfare gains. Therefore, our results should be interpreted as upper bounds for the direct welfare gains from closing arbitrage gaps.

4.2 Estimating Price Impact

We briefly describe the data used to estimate price impact in the FX futures market. In Section D in the Appendix, we include summary statistics and a more detailed description of the data sources. In Section E in the Appendix, we provide useful institutional background on the FX futures and spot markets. Next, we describe our econometric 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 Australian Dollar/USD, British Pound/USD, Canadian Dollar/USD, Euro/USD, and Yen/USD.

Estimation Procedure As shown in Section 3, 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, Bouchaud (2010), Almgren, Thum, Hauptmann and Li (2005), Gabaix, Gopikrishnan, Plerou and Stanley (2003), 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.[12]

Therefore, we use non-linear least squares to estimate the following functional form for price impact in the FX futures markets:

$$F_{\tau+1} - F_\tau = \theta + \alpha \operatorname{sgn}(Q_\tau) |Q_\tau|^\beta + \varepsilon_\tau, \tag{8}$$

where F_τ denotes the price of a futures contract right before a given transaction takes place

[12]For theoretical microstructure models that justify the power law functional form, see Gabaix, Gopikrishnan, Plerou and Stanley (2006) and Donier, Bonart, Mastromatteo and Bouchaud (2015).

— its imputation is described in Equation (9) below — and Q_τ denotes the actual size of the transaction.[13] Three features of our estimation procedure are worth highlighting.

First, note that the unit of observation in our non-linear regression is a transaction. Hence, every time a transaction takes place, we record the difference in prices right before and right after the transaction, as well as the size of the transaction to generate an observation that will inform the estimation of price impact.

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 following weighted average of the bid and ask prices:

$$F_\tau = \omega_\tau F_\tau^{\mathcal{B}} + (1-\omega_\tau)\, F_\tau^{\mathcal{A}}, \quad \text{and} \quad \omega_\tau = \frac{M_\tau^{\mathcal{A}}}{M_\tau^{\mathcal{B}} + M_\tau^{\mathcal{A}}}, \tag{9}$$

where $F_\tau^{\mathcal{B}}$ and $F_\tau^{\mathcal{A}}$ respectively denote the bid and ask prices and $M_\tau^{\mathcal{B}}$ and $M_\tau^{\mathcal{A}}$ respectively denote the size of the bid and the ask right before a transaction takes place.[14]

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

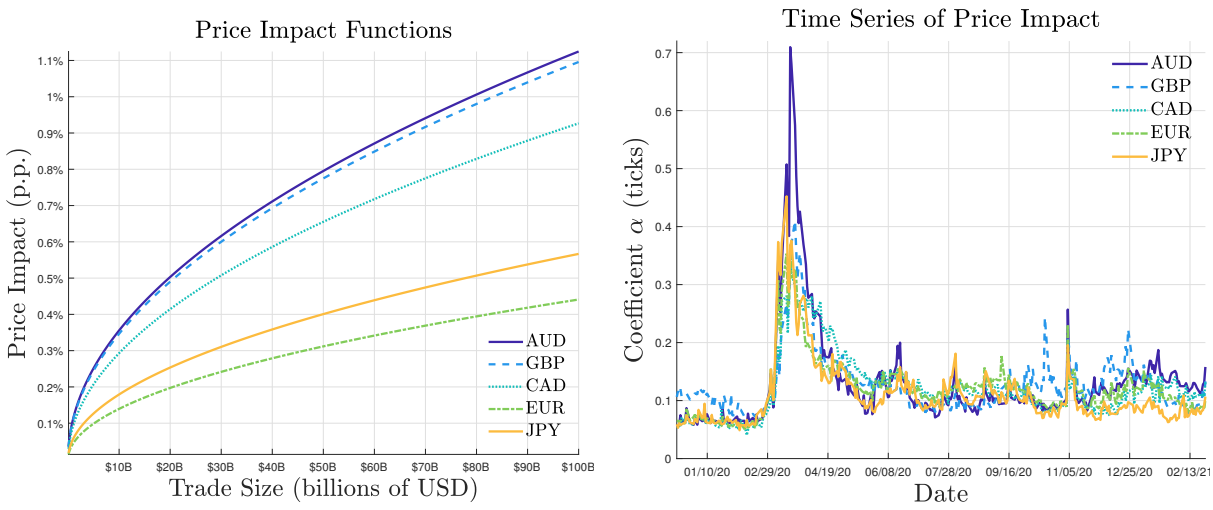
Price Impact Estimates Given the high-frequency nature of our data, we are able to compute precise estimates of price impact on a daily basis between December 15, 2019 and February 26, 2021, which allows us to capture the time-varying nature of price impact. We present average and daily price impact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

Figure 4a shows the estimated price impact functions of the form described in Equation (8) for the five currency pairs that we study using average estimates over the sample. Since our measure of price impact is non-linear, it is more informative to present the whole estimated function, rather than simply reporting the estimated α coefficients.[15] In terms of magnitudes, our estimates imply that a trade of \$10B will move the FX futures rate by roughly 0.1% in the case of the Euro and to roughly 0.35% in the case of the Australian Dollar, with the other currency pairs in between. The non-linear (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 will move the FX futures rate by roughly 0.45% in the case of the Euro and 1.1% in the case of the Australian Dollar,

[13] In our dataset, transactions may occur within microseconds of each other.

[14] See Graves (2021) for a detailed discussion of these and other (e.g., midpoint) imputation methods.

[15] In Table 5 in the Appendix, we include summary statistics of the estimated daily price impact coefficients.



(a) FX futures: Price Impact Functions

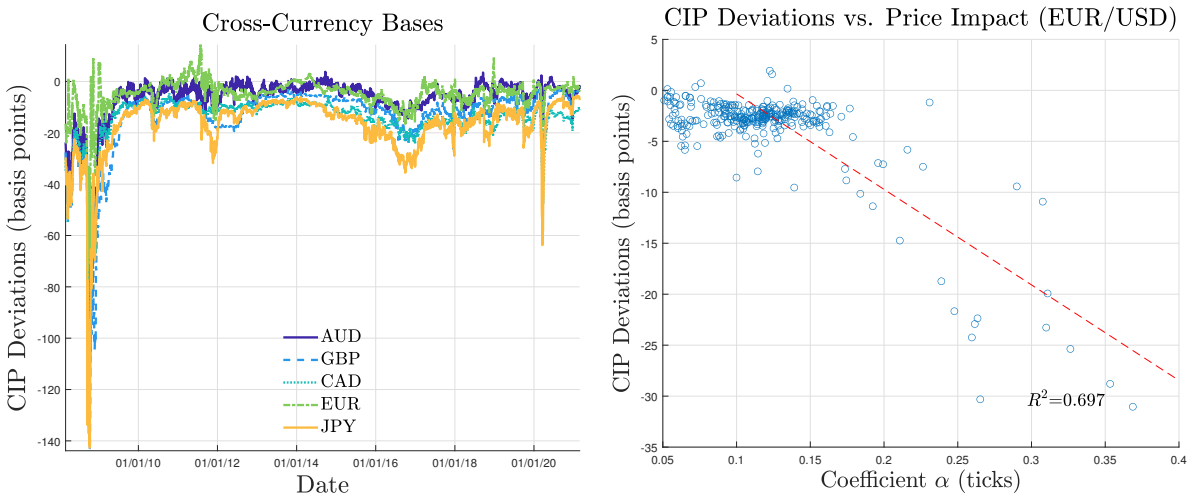
(b) FX futures: Time Series of Price Impact

Figure 4: FX futures: Price Impact Estimation

Note: Figure 4a shows the estimated price impact functions of the form introduced in Equation (8) for the five currency pairs that we study using average estimates over our full sample: December 15, 2019 to February 26, 2021. Figure 4b shows daily estimates of the α coefficients from the regression introduced in Equation (8). 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 1 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, the order of magnitude of the estimates is stable, and precision is obtained even for single days via the high-frequency data set, which generally contains of the order of tens of thousands of observations per market per day.

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 was associated with high uncertainty, limited liquidity and significant CIP violations — as we will show next. This is an important observation for our welfare calculations, since the period of elevated price impacts corresponds precisely to the greatest violations of covered interest parity since the Financial Crisis in 2008. 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 the pre-pandemic levels.



(a) Cross-Currency Bases

(b) CIP Deviations and Price Impact

Figure 5: CIP Deviations/Correlation with Price Impact Estimates

Note: Figure 5a shows the three-month CIP deviations for the five currencies we consider. The basis is calculated as $\frac{S}{F}\frac{1}{\left(1+r^f\right)^{1/4}} - \frac{1}{\left(1+r^d(m)\right)^{1/4}}$, 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 Great Financial Crisis, 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 value of the bases reported by Du, Tepper and Verdelhan (2018), which are annualized. In Section F.4 in the Appendix we include results using LIBOR. Figure 5b shows that there exists a negative relation between our daily price impact estimates (the α coefficients) and the observed CIP deviations. Figure 5b shows the relation for the case of the Euro for the period between December 15, 2019 and February 26, 2021. The $R^2 = 0.697$ and the regression line 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 the other four currencies: the analogous R^2s are 0.289 (AUD), 0.683 (GBP), 0.792 (CAD), and 0.637 (JPY).

4.3 Measuring CIP Deviations

The next input 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 in 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 02, 2008 and February 26, 2021. Table 3 in 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 relevant five 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 basis for the USD versus the five currencies that we consider, which are some of the world's most liquid currencies. As discussed extensively in Du, Tepper and Verdelhan (2018), the post-financial crisis period contains substantially and persistently negative cross-currency bases. Two major events are immediately visible in the figure: the 2008 financial crisis and the COVID-19 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 combines the daily measures of cross-currency bases (CIP deviations) with our daily price impact estimates introduced in Figure 4. This Figure shows that there exists a clear negative relation between the daily price impact estimates (the α coefficients) and the 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. The negative comovement between price impact measures and CIP deviations will be relevant for our our welfare conclusions, as described next.

4.4 Welfare Estimates and Gap-closing Trades

Armed with both price impact estimates and cross-currency bases, we are ready to leverage the theoretical framework introduced in Sections 2 and 3 to compute gap-closing arbitrage trade sizes and social welfare gains from closing the arbitrage gap. First, given our specification of price impact for a given currency pair market, we compute the gap-closing arbitrage trade $m^\star$ as the size of the arbitrage that would close the arbitrage gap, that is,

$$\frac{S\left(m^\star\right)}{F\left(m^\star\right)}\frac{1}{1+r^f\left(m^\star\right)}=\frac{1}{1+r^d\left(m^\star\right)}. \tag{10}$$

Second, given that we have estimates of $\frac{dW(m)}{dm}$, as defined in Equation (7), 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,

$$\int_0^{m^\star}\frac{dW\left(\tilde{m}\right)}{dm}d\tilde{m}. \tag{11}$$

Since the cross-currency bases are generally negative, this number will be negative but its absolute value reflects the welfare gains associated with closing the gap. As discussed in the

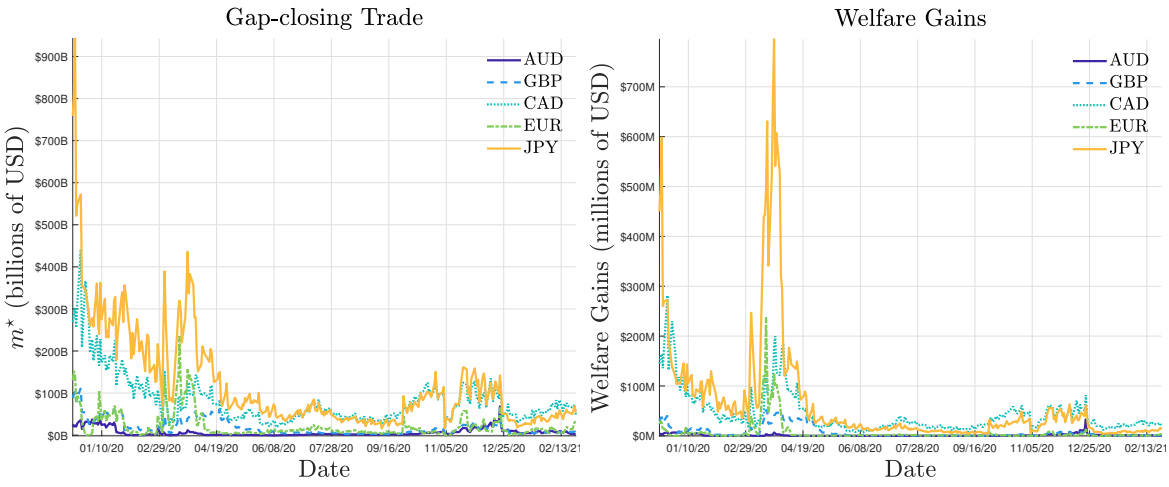
context of our dynamic model in Section B in the Appendix, the computed welfare gains have the exact interpretation of closing the arbitrage gap today while holding the position until the strategy unwinds.[16]

Figure 6 presents daily estimates of i) gap-closing arbitrages 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 \$1T and never exceeds \$450B outside the yen-dollar basis. One can readily see why gap-closing trades do not spike in periods with high CIP deviations: consistently 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 the gap-closing trade size. 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. Note that the gap-closing quantities necessary are within what a market participant could execute.

Figure 6b shows that the estimated welfare gains from closing CIP deviations never exceed \$1B for the five studied pairs, and never exceed \$250M outside the yen-dollar basis. To put these magnitudes in perspective, the total daily volumes in the spot FX markets for the EUR/USD and the JPY/USD pars are respectively \$416B and \$260B (BIS, 2019). Therefore, the estimated welfare gains from closing arbitrage gaps are always lower than 0.5% of daily dollar volume in the spot FX markets, which one could argue makes the gains from closing gaps small, despite taking place in some of the most highly liquid markets that exist. There are two features of our analysis that explain these results. First, the CIP deviations are significant, but not that large in magnitude — this is in contrast to Section 5, in which we show that our approach yields large welfare gains when arbitrage gaps are large, as in the case of dual-listed companies. Second, as described above, the fact that CIP deviations and price impact estimates comove negatively makes the size of the gap-closing trade small, lowering the scope for welfare gains.

In the Appendix, we present additional exercises. First, by backward extrapolating our price impact estimates from 2019–2021, we can find 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 2019–21 is immediately apparent in Figure A.1 in 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 in markets outside of the yen-dollar pair. This is true despite the persistent, significant

[16]Note that our approach abstracts from cross-currency spillovers, for which it is very hard to find credible estimates. Hence, since we do not account for the impact that closing a given CIP deviation has on other CIP deviations, our approach is once again likely to overestimate the welfare gains of closing all arbitrage gaps.



(a) Gap-closing trade

(b) Welfare Gains from closing CIP deviations

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

Note: Figure 6a shows the gap-closing arbitrage trade $m^\star$ such that the Covered Interest Parity condition holds. 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 11, 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 the yen-dollar basis.

deviations from CIP documented in the aforementioned literature.

Next, we characterize the combinations of price impacts and CIP violations that must exist to generate large welfare gains from closing CIP deviations. In Figure A.2 in the Appendix, we show isoquants of CIP deviations and price impact estimates (α coefficients) that yield the same level of welfare gains, corresponding to \$100M, \$1B, and \$10B in the EUR/USD case. This figure 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, we also show 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 A.2 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 although three-month CIP deviations do not result in substantial welfare gains, five-year CIP deviations might. Du, Tepper and Verdelhan (2018) and others study the magnitude of such CIP violations, which are found to be greater than their three-month counterparts but have the same order of magnitude. The five-year T-Note and corresponding foreign interest rate markets, however, are not nearly as liquid as the three-month T-Bill

and Eurodollar markets. More importantly, however, the five-year forward market with its distant delivery date is far less liquid than the three-month forward market. Quoted spreads in Bloomberg are much wider, and if comprehensive five-year forward quote and trade data were available, the estimated price impact coefficients would almost certainly be far larger than those estimated from front-month futures, which are among the most liquid markets in the world. Although the welfare gains from longer timer horizons might seem larger in magnitude due to the large CIP deviations documented in the literature, the lack of liquidity (large price impact) would make the gap-closing trade size small. Furthermore, it is straightforward that the cross-price impact of currency contracts — from which we conservatively abstract in our calculations — with distant delivery dates must be significantly less than for those with nearby delivery dates.

Finally, note that our conclusions regarding the magnitude of gap-closing trade sizes as well as our assumptions on cross-price impact seem to be supported by the Federal Reserve's quarterly reports and press announcements. On September 18, 2008, right as the CIP violations peaked, the FOMC 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 to reduce 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 we cannot infer causality, 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, as well as the underlying cross-price impact assumptions we use.

5 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 opportunities in the case of 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.[17]

Under a dual-listing arrangement, two share classes trade on different exchanges but represent fixed claims to cash flows of the same company. 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 (Froot and Dabora, 1999). Given that the dividend streams were fixed at a $1.5:1$ ratio, the law of one price/absence of arbitrage implies that the market capitalizations of both companies should always respect 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 20–30%. Like CIP, this represents a textbook arbitrage opportunity.[18][19]

5.1 Measurement Approach

To determine the welfare gains from closing arbitrage gaps in the case of dual-listed companies, the main obstacle is determining the price impact function of each involved stock, since the only readily available information (from the online De Jong, Rosenthal and Van Dijk (2009) dataset or from publicly available sources) are daily stock prices, volumes, exchange rates, shares outstanding, and exchange rates. Since we do not have access to detailed transaction data — as we do in the case of the Future FX market — we adopt the price impact estimates for global equities of Frazzini, Israel and Moskowitz (2018).

Using a global database of a hedge fund's $1.7 trillion stock transactions spanning two decades, Frazzini, Israel and Moskowitz (2018) estimate the price impact of trades as a function of trade size and stock characteristics. Consistent with earlier literature, they find that the price impact of quantity traded satisfies a square root functional form. Their main specification of price impact for a given stock j takes the form:

$$\frac{p_t^j(m) - p_t^j(0)}{p_t^j(0)} = \alpha_t^j \operatorname{sgn}(m) (m)^{\frac{1}{2}},$$

where $p_t^j(m)$ denotes the price of stock j at time t after a trade of size m — expressed as the

[17] According to Loewenstein (2000), more than half of the losses incurred by LTCM in equity pairs trading (of $286 million) was due to the Royal Dutch/Shell trade.

[18] As originally noted by Shleifer and Vishny (1997), the need to unwind arbitrage trades before the strategies converge can make this type of arbitrage risky. We discuss how this possibility can be incorporated in our framework and how this affects the interpretation of our results in Section C in the Appendix.

[19] Throughout this section, we use the data from Mathijs van Dijk's 2009 paper and related website, http://www.mathijsavandijk.com/dual-listed-companies

fraction of daily volume, which is the trade size as a percentage of one-year average daily volume — takes place.[20] The price before the trade takes place is denoted by $p_t^j(0)$ and α_t^j is the key estimate that captures potentially time-varying liquidity conditions. Over their full sample, they find a price impact coefficient of $\alpha = 0.000889$. We use this specification as a conservative estimate, in particular in the Royal Dutch/Shell scenario. Indeed, Frazzini, Israel and Moskowitz (2018) unsurprisingly find a lower price impact for large-cap stocks, and Royal Dutch/Shell is and was one of the largest companies in the world, with an average market capitalization of 70B sterling in the 1990s and a market cap of around $160B today. Rio Tinto and Smithkline Beecham are similarly large companies.

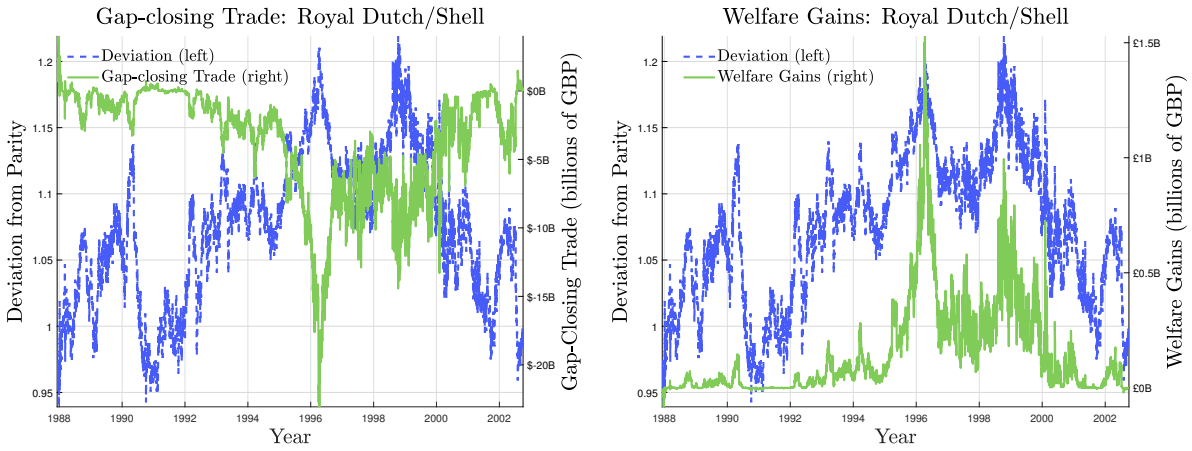
For each trading day between October 27, 1986 and October 3, 2002, we combine the data on dual-listed stocks from De Jong, Rosenthal and Van Dijk (2009) with the price impact specification of Frazzini, Israel and Moskowitz (2018) to compute the social welfare gains from closing the arbitrage gaps following our theoretical results in Section 3 and in the same manner as for Covered Interest Parity. Specifically, given a share price of p_t^D for Royal Dutch and p_t^S for Shell, one can purchase m shares of Royal Dutch and short sell $\frac{3}{2}m$ shares of Shell, which generates a marginal welfare gain of

$$\frac{dW(m)}{dm} = p_t^D(m) - \frac{3}{2}p_t^S(m).$$

As in the model and the previous application, it is straightforward to find the gap-closing arbitrage trade for each day in the sample as well as the welfare gains from closing arbitrage gaps. We follow an analogous procedure in the case of Smithkline Beecham and Rio Tinto.

Before presenting our welfare estimates, it is worth discussing the external validity of the estimates by Frazzini, Israel and Moskowitz (2018), since the magnitude of the gap-closing trades will be substantial in relation to the shares outstanding. The Frazzini, Israel and Moskowitz (2018) estimates cover transactions that amount to 0–13.1% of the daily volume of a stock. The trade sizes required to close these arbitrage gaps are so large as to account for over 13.1% of the entire float of the company in the case of Royal Dutch/Shell. At the same time, Frazzini, Israel and Moskowitz (2018) provide reason to suspect that their methodology might overestimate price impact for large trades (and thus underestimate welfare implications): they find a power law of 0.35 in the data but estimate a 0.50 (square root) power law functional form instead so as to be consistent with past theory and empirical papers. Additionally, the time periods during which extreme Royal Dutch/Shell parity deviations persist last years, implying that if the transaction were spread over time, the percentage of daily volume would be sufficiently small so as to allow for execution of the

[20] Because the price impact specification of Frazzini, Israel and Moskowitz (2018) relies on a normalization of trading volume that employs one-year volume averages, it may be the case that we find different welfare estimates in periods in which the arbitrage gaps are identical.



(a) Gap-Closing Trade: Royal Dutch/Shell

(b) Welfare Gains: Royal Dutch/Shell

Figure 7: Royal Dutch/Shell

Note: Figure 7a shows the gap-closing arbitrage trade for Royal Dutch/Shell. It also shows the price deviation relative to parity (normalized to 1). At its peak, the gap-closing trade reaches just over £20B (around $30B). This corresponds to an over 20% deviation from parity in the dual-listed shares, with an underlying market capitalization of over $100B. Figure 7b shows the welfare gains from closing the arbitrage gap for Royal Dutch/Shell. It also shows the price deviation relative to parity. The welfare gains peak in 1996 corresponding to over $2B USD. The welfare calculations use the price impact specification of Frazzini, Israel and Moskowitz (2018). Peak deviations here are associated with significant welfare losses, with much of the 1996–2000 period seeing welfare gains from closing the gap of over $1B USD.

trade.

5.2 Welfare Estimates and Gap-closing Trades

5.2.1 Royal Dutch/Shell

Figure 7b shows the welfare gains from arbitrage in GBP over the nearly sixteen-year period, with a peak of approximately $2B in 1996 when converted to USD. This twin share divergence appears to generate significant potential welfare gains, which is at odds with our CIP results, for which we find lower welfare gains. Yet the magnitude of the deviations from parity in the Royal Dutch/Shell is extreme — as detailed in Figure 7b, 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 1.5:1 parity as well as square root price impact, a trade size equal to around 20B pound sterling is required to close the arbitrage gap (see Figure 7a below). Unlike the case of CIP deviations, the persistence of extreme deviations from parity over many years results in an extended period over which the welfare gains from closing price deviations in Royal Dutch/Shell are over $1B.

The existence of arbitrage opportunities in easy to trade dual-listed companies (DLC)

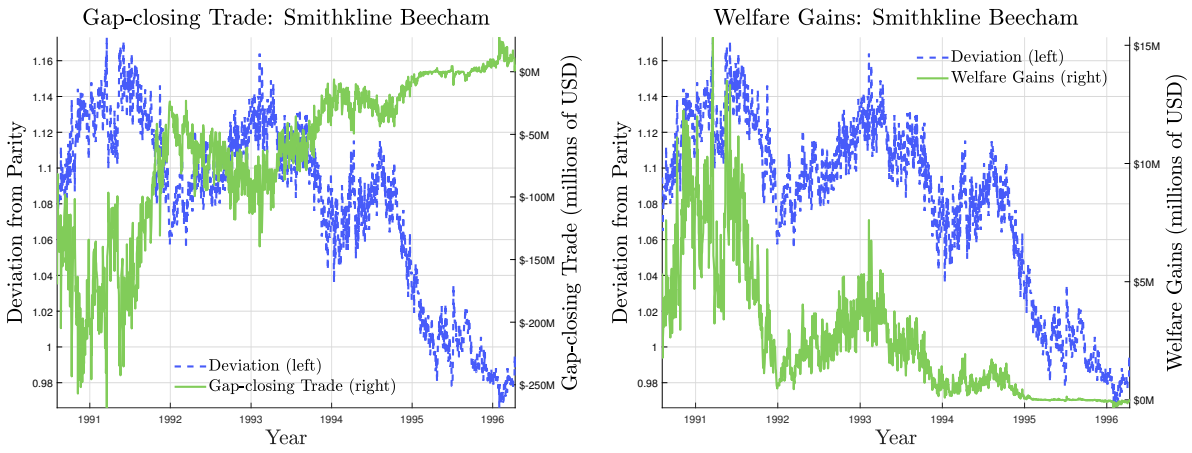
may be hard to explain. Yet the arbitrage opportunity is not quite as promising as it may appear: in a real world analogue of Shleifer and Vishny (1997), an arbitrageur in 1907 would have had to wait nearly a century to see unification (or less time but still an average of decades for a parity crossing point), during which time the marked-to-market value of the positions would have oscillated wildly. Furthermore, unwinding such a large position would simply have resulted in the price wedge being driven back to where it started – unification of the DLC is the vital ingredient for a trade of this magnitude to succeed. In Appendix C, we describe how to think about effects of the type described by Shleifer and Vishny (1997) within our framework. There we show that, leaving aside distributive pecuniary externalities, our measures capture the welfare gains from closing the arbitrage gap at a specific point in time using a buy-and-hold strategy.

5.2.2 Smithkline Beecham

GlaxoSmithKline was created in 2000 as a combination of Glaxo Wellcome and Smithkline Beecham. Smithkline Beecham was one of the world's major pharmaceutical companies until the merger, and 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), as described in De Jong, Rosenthal and Van Dijk (2009). 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 8b, which shows the welfare implications of the Smithkline share class divergence, illustrates how the arbitrage opportunity is small in magnitude and is several orders of magnitude smaller than the maximal gains from arbitrage seen with Royal Dutch/Shell. This can be seen as a direct result of the limited liquidity in the UK-traded H-share class.

5.2.3 Rio Tinto PLC/Ltd

Using the dataset from De Jong, Rosenthal and Van Dijk (2009), we explore the welfare implications of arbitrage and gap-closing trade sizes in Rio Tinto, a dual-listed company that represents one of the largest mining corporations in the world. Dual listings in Australia and the United Kingdom resulted in share price divergences that cannot be explained by a frictionless rational framework. As is readily apparent from Figure 9b, the share price ratio of Rio Tinto's dual listings is far less stable than the ratio of Royal Dutch/Shell and diverges from unity less dramatically. As noted in De Jong, Rosenthal and Van Dijk (2009), Rio Tinto experiences the smallest mean absolute divergence of all the studied twin shares in their paper. Largely as a result of this fact, we find minimal welfare implications to arbitrage throughout most of the price history, except when the ratio briefly diverges to



(a) Gap-Closing Trade: Smithkline Beecham

(b) Welfare Losses: Smithkline Beecham

Figure 8: Smithkline Beecham

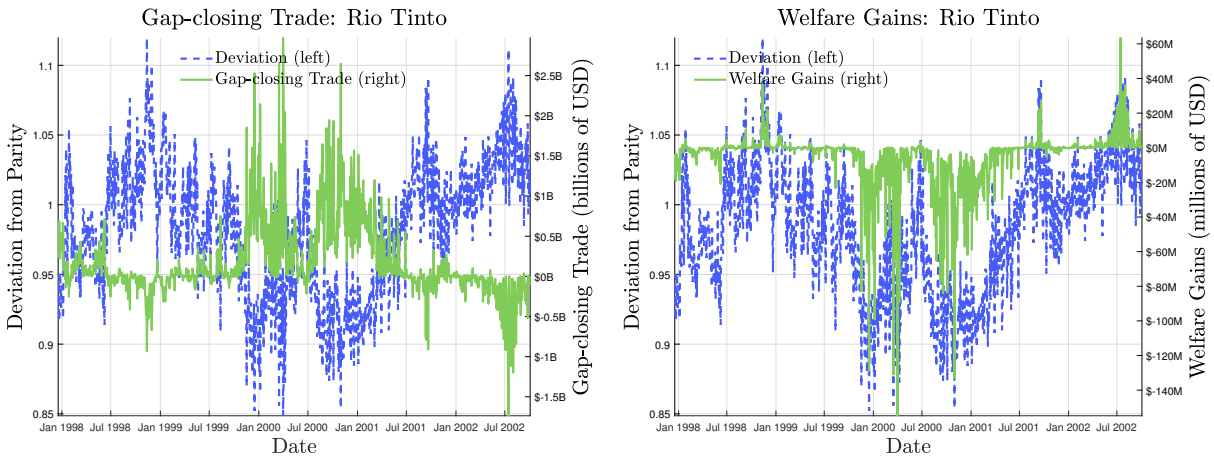
Note: Figure 9a shows the gap-closing arbitrage trade for Smithkline Beecham. The maximum trade size reaches around $250M and displays a clear time trend along with the price ratio. Figure 9b shows the welfare gains from closing price deviations for Smithkline Beecham. The gains never exceed $16M due to the limited liquidity of one of the two listings. Both sub-figures include the price deviation relative to parity.

its most extreme points with peak welfare gains from closing arbitrage deviations reaching around $150M USD and maximal gap-closing trade sizes of approximately $3B USD (see Figure 9a). The combination of lower liquidity than Royal Dutch/Shell as measured by dollar trade volume and limited divergence from parity results in far less striking welfare implications.

5.2.4 Final Remarks

As illustrated by this paper's three DLC examples, the welfare implications of arbitrage for dual-listed companies can vary by several orders of magnitude. These empirical findings can be traced directly to the theoretical connection between liquidity and welfare: arbitrage gaps that are resistant to closure because of minimal price impact from trades imply significant welfare losses from the price wedge, whereas large price impact implies that the gap can be closed easily and with minimal social gains from arbitrage. Given that virtually every empirical and theoretical exploration of price impact finds a direct connection between dollar liquidity of an instrument and price impact, the welfare gains from arbitrage in dual-listed companies are directly tied to the liquidity of the dual-listed shares.

An intuitive prerequisite for finding large welfare implications of DLC divergence is a balanced distribution of trading volume across the two venues. 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



(a) Gap-Closing Trade: Rio Tinto

(b) Welfare Gains: Rio Tinto

Figure 9: Rio Tinto PLC/LTD

Note: Figure 9a shows the gap-closing arbitrage trade for Rio Tinto PLC/Ltd. The maximum trade size reaches around $3B, corresponding to the period of maximal deviation from parity. The visible high-frequency oscillations are related to the unstable oscillations in the share price ratio, possibly due to the non-overlapping time zones of the listings. Figure 9b shows the welfare gains from closing the arbitrage gap for Rio Tinto PLC/Ltd. The gains never exceed $150M and display rapid fluctuations daily, consistent with the fluctuations in the share price ratio. Both sub-figures include the price deviation relative to parity.

only a small transaction because the leg of the trade transacting in the illiquid stock will trigger a sudden jump up price. Alternatively, if one venue represents the primary trading venue, while the other venue is a foreign listing with limited trading volumes, then even extreme divergences 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, particularly if the company possesses an outsized market capitalization.

6 Conclusion

We show that arbitrage gaps and measures of price impact are sufficient to compute the welfare gains associated with closing arbitrage gaps. The approach introduced in this paper can be applied to any environment in which arbitrage conditions are violated. Our results imply that tracing how arbitrage gaps react to changes in the size of arbitrage trades is sufficient to be able to characterize the value of arbitrage. Therefore, measures of quantities (price impact) are needed in addition to arbitrage gaps to understand the value of arbitrage. While our welfare assessments are exact in models in which there are strict arbitrage opportunities, one may expect 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.

BIS. 2019. "Global foreign exchange market turnover in 2019." *TriennialCentral 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*.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics*, 130(4): 1547–1621.

Campbell, John Y. 2017. *Financial Decisions and Markets: A Course in Asset Pricing.* Princeton University Press.

Cochrane, John. 2005. *Asset Pricing: (Revised).* . Revised ed., Princeton University Press.

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.

Donier, Jonathan, Julius Bonart, Iacopo Mastromatteo, and J-P 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. 2014. "A theory of macroprudential policies in the presence of nominal rigidities." NBER Working Paper.

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.

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.

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: Section 3

Proposition 1. (Marginal value of arbitrage)

a) This result follows from differentiating the indirect utility of investors (applying the envelope theorem) and the utility of arbitrageurs.

b) This result follows immediately after aggregating the three objects characterized in part a).

Proposition 2. (Social value of arbitrage)

a) This result follows from the fundamental theorem of calculus and Proposition 1b).

b) This result follows from part a) and the fact that

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

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

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

$$W(m^\star) - W(m_0) = \int_{m_0}^{m^\star} \mathcal{G}_{BA}(m)\, dm, \tag{A.1}$$

where $\mathcal{G}_{BA}(m) \equiv p^B(m) - p^A(m) = \left(p^B(0) - p^A(0)\right) + \int_0^m \left(\frac{dp^B(\tilde{m})}{dm} - \frac{dp^A(\tilde{m})}{dm} \right) d\tilde{m}$. Price impact is higher when the arbitrageur moves the price against him more, i.e., when $\left| \frac{dp_i(s)}{dm} \right|$ is higher. Therefore, since

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

more liquid markets will 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.1).

B Explicit Models of Limits to Arbitrage

There are alternative frictions that can lead to arbitrage gaps. In this section, we formally describe how several leading explanations can be mapped to our framework. In particular, we consider trading costs, market power in the arbitrageur sector, and short selling, funding or collateral constraints. All these features limit the size of the arbitrage trade implemented by arbitrageurs and prevent asset prices from equalizing in equilibrium. In other words, these frictions prevent the actual arbitrage trade m from differing from the gap-closing trade $m^\star$, which would be chosen by competitive unconstrained arbitrageurs. In our dynamic extension in Section C, the same frictions can be used to explain why a given $m_1(s)$ is not chosen by individual agents to close an arbitrage gap.

B.1 Trading costs

Suppose that there are linear trading costs τ^A and τ^B in markets A and B, respectively. That is, when trading in market i, an investor or arbitrageur trading 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

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

If the arbitrageur sector is competitive, the total amount arbitraged between markets A and B will be m_0 such that the profits of arbitraging are zero. In this case, the arbitrage gap is given by

$$p^B(m_0) - p^A(m_0) = \tau^A + \tau^B.$$

B.2 Strategic arbitrageurs

Suppose that there is a finite number of arbitrageurs indexed by $h = 1, ..., H$ 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 arbitraged by the other $l \neq h$ arbitrageurs is m^{-h}, solves the following problem

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

In a symmetric equilibrium with H strategic arbitrageurs, the optimal amount arbitraged m_0 will be such that

$$p^B(m_0) - p^A(m_0) = \left(\frac{\partial p^A(m_0)}{\partial m} - \frac{\partial p^B(m_0)}{\partial m}\right) \frac{m_0}{H}.$$

Note that when the arbitrageur sector is competitive, i.e., if $H \to \infty$, the arbitrage gap is 0 and $m_0 = m^\star$.

B.3 Short sales/Borrowing constraints

Suppose that the arbitrageurs face a short selling constraint $\underline{m}$ such that $m \leq \underline{m}$. Then, if the arbitrageur sector is competitive, the total amount arbitraged will be $\min\{\underline{m}, m^\star\}$, where $m^\star$ is the gap-closing arbitrage trade. Then, the arbitrage gap will be

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

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

B.4 Price-dependent collateral constraint

We develop 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. Then the total amount of resources that the arbitrageur can commit to repaying the purchase of asset A are given by his endowment e_1 and the remaining fraction from the short sale proceeds. The collateral constraint then is 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 arbitraged will be

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

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

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

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

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

the collateral constraint does not bind.

B.5 No Frictions

It is clear from $V^\alpha(m)$ that the problem of 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 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 will be identical to the integrated equilibrium. We will index our arbitrage equilibrium by the size of the arbitrage trade m. This is the simplest way of considering frictions that prevent arbitrageurs from fully equalizing prices. We can think that some frictions, which we left unmodeled, restrict the ability of arbitrageurs to close the arbitrage gap, so by moving m we are effectively relaxing those frictions. Alternatively, market power in the arbitrage market would imply that the arbitrage gap does not close.

Note that there exists a level of m, which we denote by $m^\star$, such that $p^A = p^B$. We refer to this level of trade as the optimal arbitrage trade. When $m = 0$, the arbitrage equilibrium is identical to the autarky equilibrium. When $m = m^\star$, the prices in the autarky equilibrium correspond to the integrated economy.

C Model Extensions

C.1 Static Framework: Uncertainty and Multiple Assets

In this section, we extend the model introduced in Section 2 to allow for uncertainty and an arbitrary asset structure. All the results extend naturally. In Section C.4 we show how our results can be further extended to multi-good economies.

There are two dates $t = \{0, 1\}$ and a single consumption good (dollar), which serves as numeraire. At date 1, a state $s = 1, ..., S$ is realized. There are J tradeable assets indexed by $j = 1, \ldots, J$. Investors are indexed by $i = 1, \ldots, I$.

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

$$\max_{\left\{q^i_{0,j}\right\}_j} u_i\left(c^i_0\right) + \beta_i \sum_s \pi(s)\, u_i\left(c^i_1(s)\right),$$

subject to the $S+1$ budget constraints:

$$c^i_0 = n^i_0 - \sum_j p_j \Delta q^i_j$$
$$c^i_1(s) = n^i_1 + \sum_j d_j(s)\, q^i_j.$$

The first order conditions for the investor imply

$$p_j = \sum_s \beta_i \pi(s) \frac{u_i'\left(c^i_1(s)\right)}{u_i'\left(c^i_0\right)} d_j(s).$$

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 $x = (x_1, ..., x_J)$ determines the direction of the arbitrage. Formally, the date 0 revenue of arbitrageurs is given by

$$V^\alpha(m) = -\sum_j p_j x_j m, \tag{A.2}$$

where $m = q^\alpha_1$ and $x_j = \frac{q^\alpha_j}{q^\alpha_1}$, subject to zero-cash-flow constraints at date 1

$$c^\alpha_1(s) = \sum_j d_j(s)\, x_j m = 0.$$

Under the assumption that $J \geq S$, 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 of the trade, m, is a parameter in our model and it will index the equilibrium in this economy.

Definition. *(General arbitrage equilibrium)* An arbitrage equilibrium, parameterized by the scale of the arbitrage trades 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,

$$\sum_i \Delta q_j^i + x_j m = 0 \quad \forall j.$$

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

Welfare

Proposition 5 directly generalizes Proposition 1. The counterpart of Proposition 2 is straightforward and we omit it.

Proposition 5. (Marginal value of arbitrage: general environment)

a) (Individual marginal 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

$$\frac{\frac{dV^i}{dm}}{\lambda_0^i} = -\sum_j \frac{dp_j}{dm} \cdot \Delta q_j^i \quad \forall i$$
$$\frac{\frac{dV^\alpha}{dm}}{\lambda_0^\alpha} = -\sum_j \left(\frac{dp_j}{dm} x_j m + p_j x_j \right),$$

where λ_0^i and λ_0^α respectively represent the marginal value of consumption at date 0 for agent $i = 1, ..., I$ and the arbitrageur, respectively.

b) (Aggregate marginal value of arbitrage) The social marginal 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

$$\frac{dV^i}{dm} = -u_i'(c_0^i) \sum_j \frac{dp_j^i}{dm} \Delta q_j^i,$$

where we used the optimality conditions of the investor's problem. Dividing by $u_i'\left(c_0^i\right)$ expresses the marginal value in date 0 dollars and gives the first result above.

Differentiating Equation A.2 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 social marginal value of m is given by

$$\frac{dW}{dm} = \frac{\frac{dV^\alpha}{dm}}{\lambda_0^\alpha} + \sum_j \frac{\frac{dV^j}{dm}}{\lambda_0^j}.$$

Using the results from part a) we have

$$\frac{dW}{dm} = -\sum_j \left(\frac{dp_j}{dm} x_j m + p_j x_j\right) - \sum_j \frac{dp_j^i}{dm} \Delta q =_j^i - \sum_j p_j x_j.$$

□

C.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 $t = \{0, 1, 2\}$ 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 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' problem In each market i, the representative investor's problem is

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

subject to

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

The first order conditions for the investor imply

$$p_0^i = \frac{\beta_i u_i'\left(c_1^i\right)}{u_i'\left(c_0^i\right)}\left(d_1^i + p_1^i\right)$$
$$p_1^i = \frac{\beta_i u_i'\left(c_2^i\right)}{u_i'\left(c_1^i\right)} d_2^i.$$

Arbitrageurs The arbitrageur follows an arbitrage strategy given by $\left(q_0^{\alpha A}, q_0^{\alpha B}, q_1^{\alpha A}, q_1^{\alpha B}\right)$ 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.

$$0 = c_1^\alpha = \left(d_1^A + p_1^A\right) q_0^{\alpha A} + \left(d_1^B + p_1^B\right) q_0^{\alpha B} - p_1^A q_1^{\alpha A} - p_1^B q_1^{\alpha B} \tag{A.3}$$
$$0 = c_2^\alpha = d_2^A q_1^{\alpha A} + d_2^B q_1^{\alpha B}. \tag{A.4}$$

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

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

The zero-cash-flow constraints determine the direction of the arbitrage trade for given scales (m_0, m_1). Formally, from Equations (A.3) and (A.4) it follows that

$$x_0 = -\frac{d_1^A + p_1^A}{d_1^B + p_1^B} + \frac{\left(p_1^A + p_1^B x_1\right)\frac{m_1}{m_0}}{d_1^B + p_1^B} \tag{A.5}$$
$$x_1 = -\frac{d_2^A}{d_2^B}. \tag{A.6}$$

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 static 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 $t = 0, 1$ together with the four first order conditions of the investors, the investors' budget constraints, and the zero-cash-flow constraints in Equations (A.5) and (A.6) 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 date $t = 0, 1$, 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.5) can be written as function of the equilibrium prices at $t = 1$ and the scales of the arbitrage trades, i.e.,

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

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. The Lemma below characterizes $\frac{dx_0}{dm_t}$.

Lemma 1. *(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.5) with respect to m_0 and m_1 gives

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

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. □

Welfare

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

$$V^\alpha(m_0, m_1) = -\left(p_0^B q_0^{\alpha B} + p_0^A q_0^{\alpha A}\right) = -\left(p_0^B x_0 m_0 + p_0^A m_0\right), \tag{A.7}$$

where x_0 depends on m_0 and m_1 and is given by Equation (A.5). 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 social marginal values of arbitrage.

Proposition 6. (Marginal value of arbitrage: dynamic no uncertainty)

a) (Individual marginal 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

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

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

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

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

where $x_1 = -\frac{d_1^B}{d_1^A}$.

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

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

where we used the optimality conditions of the investor's problem. Dividing by $u_i'\left(c_0^i\right)$ expresses the marginal value in date 0 dollars and gives the first result above.

Differentiating the expression in Equation (A.7) we have that the marginal value of m_0 for the arbitrageur is

$$\frac{\partial V^\alpha}{\partial m_0} = -\left(p_0^B \frac{d\left(x_0 m_0\right)}{dm_0} + p_0^A\right) - \left(\frac{dp_0^B}{dm_0} x_0 + \frac{dp_0^A}{dm_0}\right) m_0. \tag{A.8}$$

Using Lemma 1 we have

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

and using this expression in Equation A.8 gives the second result in a) above, where we used that $\lambda_\alpha = 1$.

Differentiating the expression in Equation (A.7) with respect to m_1 we have

$$\frac{\partial V^\alpha}{\partial m_1} = -p_0^B \frac{dx_0}{dm_1} m_0 - \left(\frac{dp_0^B}{dm_1} x_0 + \frac{dp_0^A}{dm_1}\right) m_0.$$

Using Lemma 1 gives

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

which concludes the proof of part a).

b) The marginal social value of m_t is given by

$$\frac{\partial W}{\partial m_t} = \frac{\frac{\partial V^\alpha}{\partial m_t}}{\lambda_0^\alpha} + \frac{\frac{\partial V^A}{\partial m_t}}{\lambda_0^A} + \frac{\frac{\partial V^B}{\partial m_t}}{\lambda_0^B}.$$

Using the results from part a) for m_0 this implies

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

and using the market clearing conditions and the FOCs for the investors gives

$$\frac{dW}{dm_0} = \left(\frac{d_1^A + p_1^A}{d_1^B + p_1^B} p_0^B - p_0^A\right) + \left(\frac{\beta_B u_B'\left(c_1^B\right)}{u_B'\left(c_0^B\right)} - \frac{\beta_A u_A'\left(c_1^A\right)}{u_A'\left(c_0^A\right)}\right) \frac{dp_1^A}{dm_0} \Delta q_1^A,$$

which proves the first part of b).

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

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

Using market clearing and the FOCs for the investors gives

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

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

Note that for both $\frac{dW}{dm_0}$ and $\frac{dW}{dm_1}$ the social marginal value of arbitrage is given by i) a direct effect, given by the arbitrage gap at either date 0 or at date 1 (appropriate discounted) and ii) a set of distributive pecuniary externalities (using the language of Dávila and Korinek (2018)). Dávila and Korinek (2018) show that distributive pecuniary externalities are always given by i) differences in MRS's, in this case $\frac{\beta_B u'_B\left(c_1^B\right)}{u'_B\left(c_0^B\right)} - \frac{\beta_A u'_A\left(c_1^A\right)}{u'_A\left(c_0^A\right)}$, ii) price sensitivities, in this case $\frac{dp_1^A}{dm_1}$, and iii) net trading positions, in this case As long as markets are incomplete, these externalities will 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 not net trades taking place at date 1 in that case.

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

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

C.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 $t = \{0, 1, 2\}$ and a single consumption good (dollar), which serves as numeraire. At date 1 the state $s = 1, ..., S$ 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 a potentially risky asset i is traded. Asset i pays a contingent dividend $d_t^i(s)$ at date $t = 1, 2$.

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

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

subject to

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

The first order conditions for the investor imply

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

Arbitrageurs The arbitrageur follows an arbitrage strategy given by $\left(q_0^{\alpha A}, q_0^{\alpha B}, q_1^{\alpha A}\left(s\right), q_1^{\alpha B}\left(s\right)\right)$ 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 cash-flow constraints.

$$0 \leq c_1^\alpha\left(s\right) = \left(d_1^A\left(s\right) + p_1^A\left(s\right)\right) q_0^{\alpha A} + \left(d_1^B\left(s\right) + p_1^B\left(s\right)\right) q_0^{\alpha B} - p_1^A\left(s\right) q_1^{\alpha A}\left(s\right) - p_1^B\left(s\right) q_1^{\alpha B}\left(s\right), \quad \forall s, \tag{A.9}$$

$$0 = c_2^\alpha\left(s\right) = d_2^A\left(s\right) q_1^{\alpha A}\left(s\right) + d_2^B q_1^{\alpha B}\left(s\right), \quad \forall s. \tag{A.10}$$

Analogous to the case without uncertainty we define trading directions as

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

and the scales of the arbitrage trades by

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

Note that the scale and direction of the arbitrage trade are contingent on the state s. The zero-cash-flow constraints determine the direction of the arbitrage trade for given scales $\left(m_0, m_1\left(s\right)\right)$. Formally, from Equations (A.9) and (A.10) it follows that

$$x_0 m_0 \geq -\frac{d_1^A\left(s\right) + p_1^A\left(s\right)}{d_1^B\left(s\right) + p_1^B\left(s\right)} m_0 + \frac{\left(p_1^A\left(s\right) + p_1^B\left(s\right) x_1\left(s\right)\right)}{d_1^B\left(s\right) + p_1^B\left(s\right)} m_1\left(s\right), \quad \forall s, \tag{A.11}$$

$$x_1\left(s\right) = -\frac{d_2^A\left(s\right)}{d_2^B\left(s\right)}, \quad \forall s. \tag{A.12}$$

The direction of the arbitrage trade at date 1 depends only on the relative payoff of the assets at date 2 in state s, just as in the static model introduced in Section 2 and dynamic model without

uncertainty. However, while in the model without uncertainty one could freely choose two 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.11) impose S restrictions that need to be satisfied by the $S+1$ scales of the arbitrage trades, $m_0, \{m_1(s)\}_{s=1,...,S}$ and the direction of the arbitrage trade at $t=0$, x_0. In particular, there may be scales $\{m_1(s)\}_{s=1,...,S}$ such that an arbitrage trade does not exist. More specifically, there may be some arbitrage trades at $\{m_1(s)\}_{s=1,...,S}$ for which negative consumption at date 1 by the arbitrageur may be unavoidable: this is exactly the scenario consider 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

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

Note that in this case we are outside the Shleifer and Vishny (1997) scenario, because there are not 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 static model case.

C.4 CIP Model

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

Market A (USD) The problem of type A investors is given by

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

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 $\Delta q_0^B = q_0^A - q_{-1}^A$ 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 flow utility of consumption $u_B(c) = \frac{c^{1-\gamma_B}}{1-\gamma_B}$ and discount factor β_B. They have dollar endowments $n_0^{B\epsilon}$ and $n_1^{B\epsilon}$ and hold an initial position $q_{-1}^{B\epsilon}$ in the traded asset. Hence, type B investors choose q_0^B as the solution to

$$\max_{q_0^B} u_B\left(c_0^B\right) + \beta_B u_B\left(c_1^B\right)$$

subject to the budget constraints

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

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

$$-\left(p_A^{\$} q_0^{\alpha A\$} + \underbrace{p_B^{\$}}_{S_0 p_B^{\epsilon}} q_0^{\alpha B\epsilon} \right), \tag{A.13}$$

subject to a zero-cash-flow constraint

$$d_1^{A\$} q_0^{\alpha A\$} + \underbrace{d_1^{B\$}}_{F_0 d_1^{B\epsilon}} q_0^{\alpha B\epsilon} = 0, \tag{A.14}$$

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

$$\Delta q_0^{A\$} + q_0^{\alpha A\$} = 0$$
$$\Delta q_0^{B\epsilon} + q_0^{\alpha B\epsilon} = 0,$$

which can be equivalently written as

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

Therefore

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

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

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

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

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

$$\frac{dV_\alpha(m)}{dm} = \frac{dp_A^\$}{dm}q_0^{\alpha A\$} + p_A^\$\frac{dq_0^{\alpha A\$}}{dm} + \frac{dp_B^\$}{dm}q_0^{\alpha B€} + p_B^\$\frac{dq_0^{\alpha B€}}{dm}.$$

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

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

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

D Detailed Data Description

This Appendix describes in detail the sources of data used in the CIP application in Section 4 and the Dual-Listed Companies application in Section 5.

D.1 CIP Application: FX futures Data

Price impact estimation was conducting using bid, offer, and trade data from FX markets traded at the CME Group. Direct feed "L1" (top of book and trade) data from the Chicago Mercantile

Exchange (CME) were recorded on a server in real time and stored in a database. Every transaction, bid change, and offer change were 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 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 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 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 12/15/2019 – 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.

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) so 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 size").

The CME's FX markets are open from 6pm EST on Sunday to 5pm EST on Friday with a daily maintenance window from 5pm EST to 6pm EST. 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 for an opening price auction to take place during which time the market's opening price and quantity are determined by orders in the book.

During the continuous trading period, a single large order that is matched against multiple smaller orders will appear 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 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 1000 lot buy order was filled 500@1.2000 and 500@1.2001, we would study it as a 1000 lot order filled at an average price of 1.20005.

Table 2: CME Data Summary Statistics (2)

	AUD/USD	GBP/USD	CAD/USD	EUR/USD	JPY/USD
Min Price	0.5512	1.1439	0.682	1.0671	0.0089
Mean Price	0.7056	1.3033	0.7548	1.1557	0.0094
Max Price	0.7956	1.4142	0.799	1.233	0.0097
Avg Daily Dollar Volume	6114468402	7140772800	4497738880	23939036247	10805096107
10P Daily Volume	26567	16253	13153	44751	19237
Median Daily Volume	86988	87992	59771	166217	92155
90P Daily Volume	212480	205379	185292	370850	407432
# Transactions/Day	25287	27824	19630	56458	32475

Note: Table 2 contains price, daily volume, and daily transaction summary statistics for the different contract markets' data sets. 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. For parsimony, all price values are converted into ticks for the empirical analysis: a tick is the minimal allowed price fluctuation in the contract market. The average daily dollar volume is generally increasing in the size of the corresponding economy, with smallest transaction volume in the AUD/USD and CAD/USD markets and the largest volumes in the EUR/USD (the most actively traded pair).

D.2 CIP Application: Price/Rates Data

We used 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 be 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 .

D.3 Dual-Listed Companies Application: Price Data

We use the 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

Table 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. Wherever possible, three-month secured government paper was used for the cross currency basis, and three-month LIBOR (unsecured) rates were used whenever necessary.

Table 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 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), 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.

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 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 workings 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, USD/JPY, 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 the US and foreign bonds, we consider only majors.

Spot FX are largely transacted through major dealing platforms, the two most well-known institutional platforms being EBS (owned by CME Group) and Retfinitiv (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 2 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 will 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 will spill over almost instantaneously into the other market. For example, if a large trader wished to purchase $500M USD 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 will 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 will 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 will 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; this kind of "stale quote sniping" is extensively reviewed in Budish, Cramton and Shim (2015).

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 consists of the SPAN margining system. SPAN allows for 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 Additional Results

F.1 Price Impact Estimates

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

Table 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 5 presents the average of the daily price impact estimates over our full sample: December 15, 2019 to February 26, 2021.

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

Figure A.1 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.

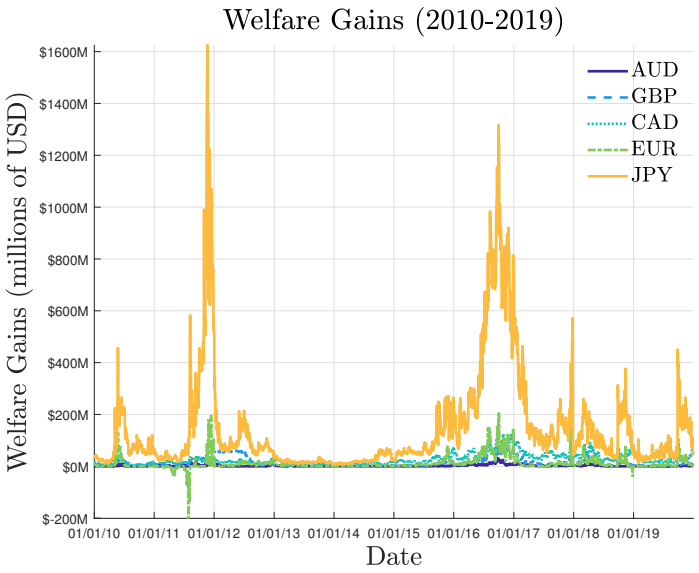


Figure A.1: Welfare gains from CIP Arbitrage: 2010-2019

Note: Figure A.1 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 A.2 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.

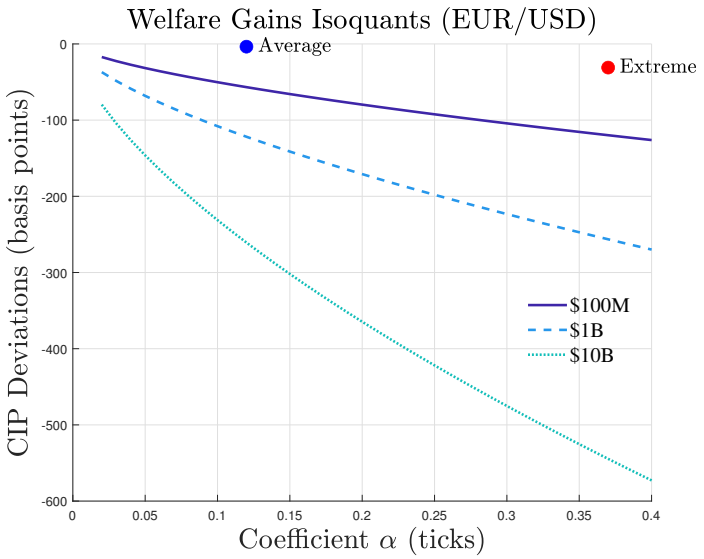


Figure A.2: Isoquants: CIP Welfare Gains (EUR/USD)

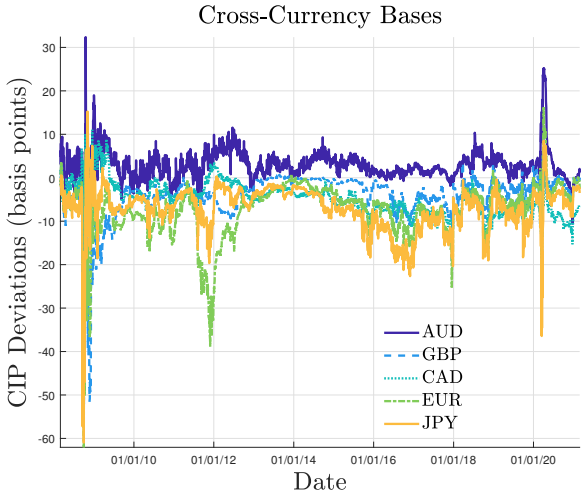
Note: Figure A.2 shows isoquants of CIP deviations and price impact estimates (α coefficients) that yield the same level of welfare gains, corresponding to $\$100M$, $\$1B$, and $\$10B$ in the EUR/USD case. A similar qualitative and quantitative relation applies to the other currencies considered in the paper. The solid red 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

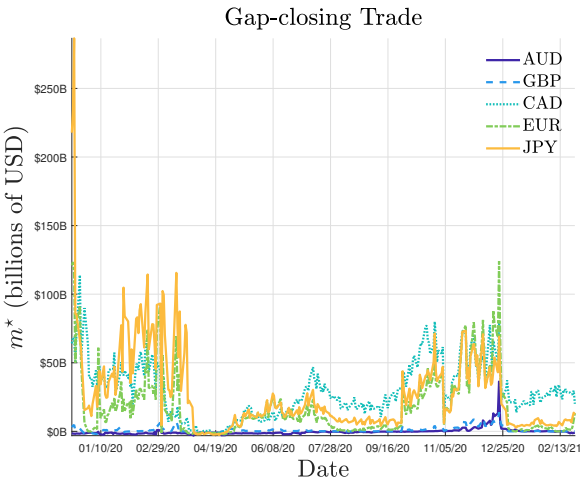
We rerun the same analysis as the main body of the paper using exclusively the same LIBOR rates used in Du, Tepper and Verdelhan (2018) for the calculation of the various cross-currency bases. The results are detailed in Figure A.3. The primary difference versus the main results can be seen in the EUR/USD cross-currency basis, where a large spike toward the end of 2011 in welfare gains from arbitrage 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), use of LIBOR to compute the cross-currency basis is imperfect because of the potential for some of the basis 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 arbitrage 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 arbitrage 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 larger price impact than our average estimates

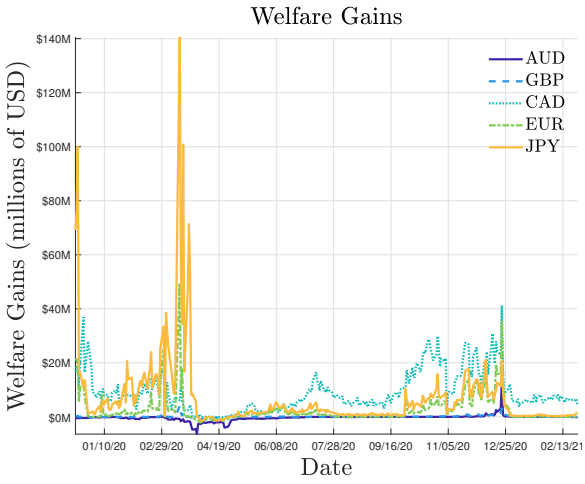
for December 2019-February 2021.



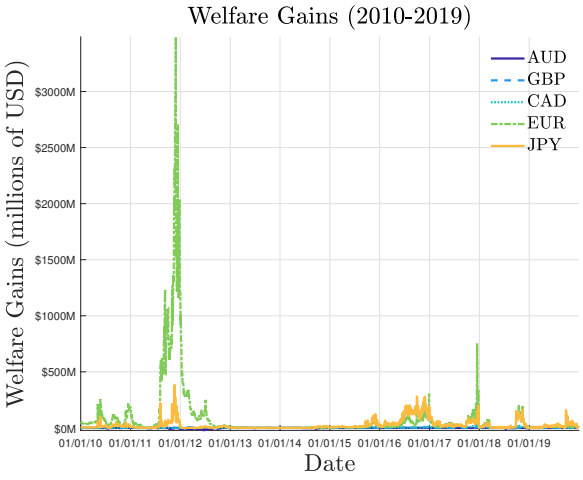
(a) Cross-Currency Bases



(b) Gap-closing trade



(c) Welfare gains from from CIP Arbitrage: 2019-2021



(d) Welfare gains from CIP Arbitrage: 2010-2019

Figure A.3: LIBOR Results

Note: Figure A.3 includes the counterparts of Figures 5a, 6a, and 6b in the text, and Figure A.1 in the Appendix when using LIBOR as riskless rate.

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 Koijen (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 arbitrage, 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 6: Price Impact Estimates: Gabaix and Koijen (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 6 reports multipliers in the style of Gabaix and Koijen (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 non-linear, 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 7 for exact computations.

Table 6 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 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, Gopikrishnan, Plerou and Stanley (2006); small trades then carry a large multiplier while the largest trades carry a much smaller multiplier. 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 so 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 Koijen (2021) that order flows by large institutions and investors create much larger market fluctuations than captured by most current models.

Table 7: Percent 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 7 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 Koijen (2021) on equity market data is why welfare gains from arbitrage 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 trade sizes, whereas concave price impact implies that small wedges will be closed with small orders and the importance of price wedges to welfare is an increasing function of the absolute size of the wedge.