

Externalities as Arbitrage*

Benjamin Hébert[†]

February 16, 2023

Abstract

How can we assess whether macro-prudential regulations are having their intended effects? If these regulations are optimal, their marginal benefit of addressing externalities should equal their marginal cost of distorting risk-sharing. These risk-sharing distortions will manifest as trading opportunities that intermediaries are unable to exploit. Focusing in particular on arbitrage opportunities, I construct an “externality-mimicking portfolio” whose returns track the externalities that would rationalize existing regulations as optimal. I conduct a revealed-preference exercise using this portfolio and test whether the recovered externalities are sensible. I find that the signs of existing CIP violations are inconsistent with optimal policy.

*Acknowledgements: This paper builds on the work of and is dedicated to the memory of Emmanuel Farhi. I would also like to thank Jonathan Berk, Eduardo Dávila, Sebastian Di Tella, Wenxin Du, Darrell Duffie, Gary Gorton, Denis Gromb, Valentin Haddad, Barney Hartman-Glaser, Piero Gottardi, Bob Hodrick, Jens Jackwerth, Anton Korinek, Arvind Krishnamurthy, Matteo Maggiori, Semyon Malamud, Ian Martin, Tyler Muir, Jesse Schreger, Amit Seru, Ludwig Straub, Jeremy Stein, Dimitri Vayanos, and seminar participants at the MIT Sloan Juniors Conference, Yale SOM, EPFL Lausanne, Stanford, UCLA, UC Irvine, Columbia Juniors Conference, Duke Macro Jamboree, Gerzensee ESSFM, Barcelona FIRS, LSE, Wharton, HSE Moscow, and the Virtual Finance Seminar for helpful comments, and in particular Hanno Lustig for many useful comments and Peter Kondor for a helpful discussion. I would like to thank Michael Catalano-Johnson of Susquehanna Investment Group for help with bounds on American option prices, Colin Eyre and Alex Reinstein of DRW for help understanding SPX and SPY options, and Jonathan Wallen and Christian Gonzalez-Rojas for outstanding research assistance. All remaining errors are my own.

[†]Stanford University and NBER. bhebert@stanford.edu

Following the great financial crisis (GFC), apparent arbitrage opportunities emerged in financial markets. These arbitrage opportunities, such as the gap between the federal funds rate and the interest on excess reserves (IOER) rate, or violations of covered interest rate parity (CIP), are notable in part because they have persisted for years after the peak of the financial crisis. Authors such as Du et al. (2018) have argued that macro-prudential regulations put in place after the GFC have enabled these arbitrages to exist and persist. If these apparent arbitrages opportunities are made possible by macro-prudential regulations, does that imply that there is something wrong with the regulations? More generally, what can be learned from the patterns of arbitrage across assets induced by regulation? Can we use these patterns to assess whether macro-prudential regulations are having their intended effects?

Suppose we observe the price of an asset a that is traded by both intermediaries and households, and the price of another asset a' with identical payoffs that is traded by intermediaries only. If the prices of a and a' are not equal, an apparent arbitrage exists. Such an opportunity can exist in equilibrium if intermediaries are bound by a constraint that treats a and a' differently (say, because it treats derivatives and cash assets differently). Let us further suppose that this constraint is regulatory— it is a leverage constraint, capital control, or some other kind of policy.

Now imagine that a household approaches an intermediary and wishes to buy asset a . The intermediary contemplates selling asset a to the household and hedging by purchasing asset a' . If the price of a is higher than the price of a' , this arbitrage trade is profitable; the intermediary would execute the trade if unconstrained. From this we can infer, as outside observers, that regulation is constraining the extent to which the intermediary can sell the asset a to the household. That is, the regulation is keeping the risks associated with asset a within the intermediary sector instead of allowing those risks to be transferred to the household sector. We can then adopt a macro-prudential viewpoint and ask whether or not it is desirable, from a social perspective, to shift the holdings of asset a from households to intermediaries.

In this paper I develop a formal, multi-asset version of the argument above. The procedure I develop takes the form of a revealed preference exercise. I ask: what externalities would justify the distortions in risk-sharing I infer from various arbitrage opportunities? Equivalently, what risks are pushed by regulation into the intermediary sector and what risks are pushed by regulation into the household sector? The key step in the analysis is the construction of what I call the “externality-mimicking portfolio.” This portfolio’s returns are an estimate of the externalities that would justify the observed patterns of arbitrage, or equivalently an estimate of the effect of regulation on risk-sharing between households and intermediaries.

I construct this portfolio using CIP violations. In the context of CIP violations, the assets a and a' are combinations of bond and exchange rate transactions. Empirically, the sign of the observed arbitrage, as documented by Du et al. (2018), is related to the direction of the carry trade— it

is more expensive for the household to execute the carry trade than it is for the intermediary to replicate the carry trade using currency forwards. Consequently, I infer that the relevant regulatory constraint(s) have the effect of keeping carry trade risk within the intermediary sector, instead of allowing intermediaries to transfer those risks to households. This result is surprising given the risky nature of the carry trade.

The basic issue is that some CIP violations (e.g. AUD-USD and JPY-USD) have the wrong sign. That is, because JPY appreciates and AUD depreciates vs. USD in bad times, optimal policy should encourage intermediaries to be long JPY and short AUD (i.e. short the carry trade). But the signs of the CIP violations are such that they encourage intermediaries to be long the carry trade, taking on more macro-economic risk. I speculate that this issue arises from an interaction between leverage constraints (which do not consider the “sign” of a trade) and demand from customers, as suggested by Du et al. (2018).

This conclusion does not rely on specific assumptions about which regulations are binding in the case of CIP violations. This is desirable in light of the “alphabet soup” of potentially relevant post-crisis regulations.¹ That is, the methodology developed in this paper is relatively model-agnostic and does not require strong assumptions about the nature of the relevant regulation; in contrast, the traditional approach of using a fully specified model necessarily involves such assumptions.

The procedure I develop has two key limitations. First, it assumes that the arbitrages under consideration are “enabled” by regulation. I will say an arbitrage is enabled by regulation if, in the absence of regulation, intermediaries would take advantage of the arbitrage until it closed (i.e. became unprofitable after accounting for bid-offer spreads and other transactions costs). I use the term “enabled,” as opposed to the term “caused,” to emphasize that arbitrage violations are rarely caused solely by regulation; typically, it is the combination of a demand imbalance amongst non-intermediaries and quantity-based regulatory constraints on intermediaries that generates arbitrage.²

Absent regulation, banks and other intermediaries would likely be willing to engage in large quantities of profitable, risk-free, short-maturity borrowing and lending, and indeed did exactly this prior to the GFC. For this reason, in my empirical analysis, I will restrict attention to short-dated arbitrages (with a one-month or shorter horizon) that were insignificant prior to the GFC. My empirical analysis focuses on two particular kinds of arbitrage, the fed funds vs. IOER arbitrage and CIP arbitrages, but the methods I develop are applicable to all such arbitrages. Examples

¹In the context of CIP violations, Du et al. (2018) discuss leverage ratios, risk-weighted asset calculations, stress tests, value-at-risk, the Volcker rule, liquidity coverage ratios, and the recent money market reform; most of these are implemented differently across jurisdictions, and some are currency- or country-specific for multi-national banks.

²This description is applicable to both theoretical models (Gromb and Vayanos (2002)) and to, for example, the explanation of CIP arbitrage offered by Du et al. (2018).

include the cash-future arbitrages discussed in Siriwardane et al. (2021), mortgage dollar rolls (Song and Zhu (2019)), spreads between tri-party and general collateral repo rates (Duffie and Krishnamurthy (2016)), the “box trade” involving equity options (Van Binsbergen et al. (2022)), and put-call parity arbitrages on stocks or stock indices (one such example is discussed in the Appendix, Section B.3).

It is less clear whether longer-dated arbitrages are enabled by regulation, and for this reason I exclude them from my analysis. For example, the “anomalies” described by Lamont and Thaler (2003) are violations of the law of one price between long-lived assets, and can become larger before converging. Both regulatory and non-regulatory limits to arbitrage (as in Shleifer and Vishny (1997)) might prevent intermediaries from eliminating these anomalies.³

Second, the externality mimicking portfolio estimates the overall macro-prudential effect of existing regulation, without identifying which regulations are relevant. Some regulations, such as a leverage constraint, might be justified on micro-prudential grounds, and yet have macro-prudential effects; likewise, redistributive policies can have macro-prudential effects. When I use the returns of the externality-mimicking portfolio to study the optimality of policy, I am implicitly assuming that policymakers have a rich enough set of macro-prudential tools to achieve their desired outcomes, even if that requires offsetting the macro-prudential effects of non-macro-prudential policies.

I start from the general equilibrium with incomplete markets (GEI) framework of Farhi and Werning (2016), which encompasses many existing studies of macro-prudential policy.⁴ I augment this framework by distinguishing between two classes of agents, “households” and “intermediaries,” as is standard when studying arbitrage (as in, e.g., Gromb and Vayanos (2002)). In this framework, optimal policy equates the marginal benefits of addressing externalities with the marginal cost of distorting risk-sharing. Under some additional assumptions about how policy is implemented, these risk-sharing distortions will manifest themselves as arbitrage opportunities. To clarify the underlying mechanism and provide a concrete example, I develop the example of capital controls, building on Fanelli and Straub (2021).

The central contribution of the paper uses the relationship between arbitrages and externalities to construct what I call the “externality-mimicking portfolio.” The returns of this portfolio are the projection of the externalities that would rationalize existing regulations onto the space of returns. Equivalently, they describe the direction of the marginal risk-transfer between households and

³In particular, short-sale costs or constraints may play a key role in several of the anomalies discussed in Lamont and Thaler (2003), but it is unclear whether these costs or constraints are the result of regulations or other frictions.

⁴I simplify the Farhi and Werning (2016) in several respects to clarify the exposition, but the procedure I develop is applicable to the general framework. This framework is quite general, but does not span the set of all possible models. Appendix D briefly discusses how hidden trade, private information, moral hazard, and other sources of inefficiency would affect the main results of the present paper.

intermediaries under existing regulation. The portfolio can also be thought of as representing the minimum difference between the household and intermediary SDFs necessary to explain observed arbitrages (an analog of Hansen and Richard (1987)), or as the portfolio that maximizes what I call the “Sharpe ratio due to arbitrage” (an analog of Hansen and Jagannathan (1991)).

Using data on interest rates, foreign exchange spot and forward rates, and foreign exchange options, I construct an externality-mimicking portfolio. The weights in this portfolio are entirely a function of asset prices; no estimation is required. If policy is optimal, this portfolio’s returns track the externalities the social planner perceives when considering transfers of wealth between the households and intermediaries in various states of the world. Irrespective of whether policy is optimal, the returns describe the direction of risk transfer between household and intermediaries induced by policy. When its returns are positive (negative), the planner must perceive positive (negative) externalities when transferring wealth from intermediaries to households to rationalize existing policy. In “bad times,” we would expect this portfolio to have negative returns, consistent with the idea that the planner would like to encourage intermediaries to hold more wealth in these states.

I consider two definitions of “bad times.” First, intuitively, bad times can be defined as times in which the intermediaries have a high marginal utility of wealth. Using this definition, I show that it is sufficient to study the expected returns of the externality-mimicking portfolio, and test if they are positive. Second, I define “bad times” using the stress test scenarios developed by the Federal Reserve. I argue that these tests are statements about when the Fed would like intermediaries to have more wealth, and as a result the returns of the externality-mimicking portfolio should be negative in the stress test scenarios.

However, I find that the expected return of the externality-mimicking portfolio is generally negative, and that its returns in the stress tests are often positive. This implies that the externalities that would justify current regulation are positive in bad times; put another way, existing regulations are having the effect of retaining risk within the intermediary sector. This is inconsistent with intuition and suggests that regulations are not having their desired effect.

This test is weak, in that it asks only whether existing regulations have the effect of moving risk out of or into the intermediary sector, and not whether the magnitude of this risk transfer is optimal. For a variety of reasons, we should not expect policy to be optimal at all times and in all states of the world. However, the tests I conduct do not require this; they instead ask only if policy on average moves risk in the right direction as opposed to the wrong direction. The failure of policy to pass such a weak test is striking.

My procedure does not directly identify why policy is sub-optimal. After presenting the main results, I discuss several possible explanations for why policy makers might inadvertently implement policies that fail my tests.

My theoretical framework builds on the GEI framework of Geanakoplos and Polemarchakis (1986) and Farhi and Werning (2016). My example of capital controls resembles both Fanelli and Straub (2021) and example 5.4 of Farhi and Werning (2016). In most theoretical models of macro-prudential policy (such as Farhi and Werning (2016) and Dávila and Korinek (2017)), a planner can implement policy using quantity constraints, agent-state-good-specific taxes, or some combination thereof. In this paper, I focus on quantity constraints, which is both realistic, in the sense that most regulation of banks takes this form, and enables the empirical exercise that follows. This paper is also related to Davila et al. (2012); both papers attempt to measure how close existing allocations are to constrained efficiency.

My empirical work considers short-term arbitrages such as the fed funds/IOER spread (Bech and Klee, 2011) and CIP violations (Du et al., 2018), and hence this paper lies at the intersection of the theoretical literature mentioned above and the empirical literature on arbitrage opportunities. The central and most surprising result of the paper is, in effect, that this intersection exists. The techniques I use to characterize the externality-mimicking portfolio that links the theory with the data build on Hansen and Richard (1987) and Hansen and Jagannathan (1991). There is also a significant literature that studies CIP violations in the context of particular models (as opposed to the general GEI framework). Examples include Amador et al. (2017); Andersen et al. (2019); Du et al. (Forthcoming); Gabaix and Maggiori (2015); Ivashina et al. (2015). The framework I develop allows for the empirical analysis of CIP violations within a relatively minimal theoretical structure, and hence enables more general conclusions about the optimality or sub-optimality of policy. My analysis focuses on the sign of arbitrage violations and asks whether macro-prudential policy is moving risk in a sensible direction; in complementary analysis, Dávila et al. (2022) attempt to quantify the welfare losses of the associated risk-sharing distortions.

1 Externalities and Arbitrage

I begin by outlining an endowment economy version of the economy studied by Farhi and Werning (2016),⁵ with a specific asset structure that I will introduce below. I will describe only the key aspects of the GEI framework; a more formal presentation of the model can be found in appendix section C.

Let S_1 be the set of future states, let s_0 be the initial state, and define $S = S_1 \cup \{s_0\}$. Let J_s be the set of goods in each state $s \in S$. The model will feature two sets of agents, \mathcal{I} (intermediaries) and \mathcal{H} (households). An intermediary $i \in \mathcal{I}$ will have income I_s^i in state $s \in S$, and face prices $\{P_{j,s}\}_{j \in J_s}$, resulting in an indirect utility in state s of $V^i(I_s^i, \{P_{j,s}\}_{j \in J_s}; s)$. Likewise, a household $h \in \mathcal{H}$ will have income I_s^h and indirect utility $V^h(I_s^h, \{P_{j,s}\}_{j \in J_s}; s)$. Each agent evaluates her

⁵For simplicity, I also assume flexible prices and simpler constraints on transfers.

expected utility using a full-support “reference” probability measure $\{\pi_s^r\}_{s \in S_1}$.⁶

The simplest interpretation of this setup is as a two-date model with multiple goods in the second date, in which case these indirect utility functions are the result of a standard consumer problem. The model can also be interpreted as having more than two dates and a single consumption good at each date, in which case each $j \in J_s$ corresponds to consumption at some date.⁷ The key assumption is that there is at least one relative price in the future states $s \in S_1$.

The incomes I_s^i and I_s^h are determined by the agents’ endowments and the allocation of a set of assets. Let A be the set of assets in the economy, and let $Z_{a,s}(\{P_{j,s}\}_{j \in J_s})$ denote the payoff of asset $a \in A$ in state $s \in S_1$, given the goods prices $\{P_{j,s}\}_{j \in J_s}$.⁸ Let Q_a be the price of asset $a \in A$.

Intermediaries and households differ with respect to their ability to trade assets in two ways. First, intermediaries can trade certain assets (the set $A^I \subset A$) that households cannot. Second, households cannot trade directly with each other, only via intermediaries.⁹ Both features are standard in models of financial intermediation such as Gromb and Vayanos (2002).

Consider a planner who can regulate asset allocations for each agent (denoted D_a^i and D_a^h) and transfer wealth between agents in the initial state s_0 (denoted T^i and T^h).¹⁰ By regulating the asset allocations of and transfers to each agent, the planner can influence the incomes I_s^h and I_s^i . Define the income of household h in state s as the sum of her endowment income, asset payoffs, and transfers,

$$I_s^h = \sum_{j \in J_s} P_{j,s} Y_{j,s}^h + \sum_{a \in A} D_a^h Z_{a,s}(\{P_{j,s}\}_{j \in J_s}) + T^h \mathbf{1}\{s = s_0\},$$

where $Y_{j,s}^h$ is the endowment of good j for household h in state s . Intermediary income I_s^i is defined in similar fashion.

Regulations and Arbitrage. Suppose that the planner regulates the trades of intermediaries but not households, and moreover does not constrain trade in the intermediary-only assets A^I for some intermediary $i^* \in \mathcal{I}$. I will show below that this is without loss of generality, in the sense that any optimal policy can be implemented this fashion. For now, however, let us consider any regulations of this form, optimal or sub-optimal.

⁶This probability measure is arbitrary; any differences between an agent’s subjective beliefs and these probabilities can be incorporated into the state-dependent indirect utility function V .

⁷See Farhi and Werning (2016) for several examples along these lines.

⁸I adopt the convention that all assets are ‘ex-dividend’ (have zero payoff in state s_0), and that all asset payoffs are homogenous of degree one in the goods prices.

⁹Formally, $A \setminus A^I$ is the union of a set of disjoint sets $\{A^h\}_{h \in \mathcal{H}}$, each containing assets tradable by the household h and the intermediaries (but not other households). Note that multiple households can trade “the same” asset in the sense of payoffs; this formalism is simply a way of preventing households from trading directly with each other.

¹⁰The ability of the planner to transfer income in state s_0 isolates the macro-prudential motives of the planner (correcting externalities) from redistributive motives. This captures the idea that governments can use taxes and transfers for redistributive purposes while at the same time using macro-prudential regulations to correct externalities. See Appendix D for a discussion of the interpretation of the results in the absence of this assumption.

Consider an asset tradable by households, $a \in A \setminus A^I$, and suppose there is an asset tradable only by intermediaries, $a' \in A^I$, with identical cashflows ($Z_{a,s}(\cdot) = Z_{a',s}(\cdot)$). The household's SDF will price the assets $a \in A \setminus A^I$. That is, for any $a \in A \setminus A^I$, there exists an $h \in \mathcal{H}$ such that

$$Q_a = \sum_{s \in S_1} \pi_s^r M_s^{h,r} Z_{a,s}(\{P_{j,s}\}_{j \in J_s}), \quad (1)$$

where $M_s^{h,r}$ denotes the household's SDF.¹¹ In contrast, regulations might constrain intermediaries' trade in a , in which case the intermediaries' SDF will not price a .

The asset $a' \in A^I$, however, is priced by i^* ,

$$Q_{a'} = \sum_{s \in S_1} \pi_s^r M_s^{i^*,r} Z_{a',s}(\{P_{j,s}\}_{j \in J_s}), \quad (2)$$

where $M_s^{i^*,r}$ is the SDF of i^* , and not necessarily by any household's SDF.

If the regulations on intermediaries do not bind with respect to asset a , the intermediaries' SDF will price both assets, and the prices of a and a' will be identical. If regulations are binding with respect to asset a , meaning (by the definition of a constraint 'binding') that the intermediaries' SDF does not price asset a , then the prices of the assets a and a' will not be identical. That is, an apparent arbitrage opportunity (specifically, a law-of-one-price violation) will exist.

Apparent arbitrage opportunities of this form are driven by the difference between the household and intermediary SDFs,

$$Q_{a'} - Q_a = \sum_{s \in S_1} \pi_s^r (M_s^{i^*,r} - M_s^{h,r}) Z_{a,s}(\{P_{j,s}\}_{j \in J_s}). \quad (3)$$

This equation can be used in two ways. First, given a set of arbitrages (pairs $a \in A \setminus A^I$ and $a' \in A^I$ with observable prices Q_a and $Q_{a'}$), it is possible to construct an estimate of $M_s^{i^*,r} - M_s^{h,r}$.¹² This is one interpretation of what I call the "externality-mimicking portfolio," described in section 3 below. Second, given an estimate of $M_s^{i^*,r} - M_s^{h,r}$ and an observable asset price Q_a , this equation can be used to predict $Q_{a'}$, even for assets without an intermediary-only replicating asset.¹³ That is, this equation makes it possible to extrapolate the effects of regulation from assets for which it is easily observed (e.g. CIP violations) to other traded assets. The difference between the SDFs,

¹¹ $\pi_s^r M_s^{h,r}$ is the ratio of h 's marginal utility of income in state $s \in S_1$, $\frac{\partial}{\partial I} V^h(I, \{P_j\}_{j \in J_s}; s)|_{I=I_s^h}$, relative to the marginal utility of income in state $s \in S_0$. I have adopted the convention that the subjective discount factor and subjective state probabilities are incorporated into the indirect utility function.

¹² I use the term "estimate" here to indicate that, except in special cases, we cannot recover a unique $M_s^{i^*,r} - M_s^{h,r}$ from a given set of arbitrages.

¹³ I use the terms replicating asset and replicating portfolio to describe an asset or portfolio with the same payoffs, which may or may not have the same price.

$M_s^{*,r} - M_s^{h,r}$, summarizes the direction in which regulation is pushing trade at the margin; if $Q_{a'} < Q_a$, intermediaries are being discouraged by regulation from selling asset a to households.

Before proceeding, I will note that these conclusions do not depend on the optimality of policy, only that policy is implemented by regulating intermediaries as opposed to households, and that at least one intermediary is unconstrained with respect to the intermediary-only assets.. I view both of these assumptions as approximations of real world policy. First, macro-prudential regulations are usually implemented by regulating banks, as opposed to attempting to directly regulate the portfolios of households. Second, banks' trades in products such as derivatives are far less regulated than trades with households and firms.¹⁴ Interpreting the intermediary-only assets as derivatives, the arbitrage arises due to the household's inability to trade derivatives and the regulatory constraints facing intermediaries, exactly as suggested by e.g. Du et al. (2018).

This discussion has assumed that in the absence of regulation, intermediaries would be free to trade both assets until their prices equalized. In the real world, it is possible that non-regulatory constraints prevent arbitrage. In the fully specified model of appendix section C, I consider this possibility, and conclude that the argument above holds for assets for which non-regulatory constraints do not prevent the planner from reallocating the asset between the intermediary and household sectors. In my empirical exercise, I address this issue by focusing on arbitrages that are short maturity (circumventing limits-to-arbitrage issues) and absent prior to the GFC.

Lastly, observe that the logic above extends without modification from assets to portfolios of assets. I will call an asset $a \in A \setminus A^I$ "arbitrage-able" if there exists a portfolio of assets in A^I that replicate its payoff, regardless of the goods prices that occur in equilibrium. That is, for any arbitrage-able asset, there exists portfolio weights $w_{a'}(a)$ such that, for all states $s \in S_1$ and all price levels $\{P_{j,s}\}$,

$$Z_{a,s}(\{P_{j,s}\}) = \sum_{a' \in A^I} w_{a'}(a) Z_{a',s}(\{P_{j,s}\}).$$

Let A^* denote the set of arbitrage-able securities.

Optimal Policy. Let us now consider what pattern of arbitrage should exist under an optimal policy. Suppose that the planner can choose any asset allocations, transfers, and goods prices for each state, subject to market clearing constraints for goods and assets, a constraint that transfers sum to zero, and the constraints on feasible asset allocations. Under these conditions and subject to a participation constraint for intermediaries, suppose the planner maximizes household welfare.

¹⁴Specifically, as noted by Du et al. (2022), bonds have an impact on leverage constraints (which are the relevant constraints post-GFC) that is an order of magnitude greater than the impact of swaps (i.e. derivatives) of the same notional. Moreover, all of the arbitrages documented in e.g. Siriwardane et al. (2021) involve the use of balance sheet, and I am not aware of any derivative-only arbitrages, consistent with the idea that at least one intermediary is unconstrained with respect to derivatives trading.

For a formal definition of the problem, see definition 3 in appendix section C.

Let us consider in particular the market clearing constraint for good j in state s . Let $X_{j,s}^h(I_s^h, \{P_{j',s}\}_{j' \in J_s})$ denote the demand function for good j by household h in state s , and define $X_{j,s}^i(\cdot)$ as the same for intermediaries. Market clearing requires that for each state $s \in S$ and good $j \in J_s$,

$$\sum_{h \in H} (X_{j,s}^h(I_s^h, \{P_{j',s}\}_{j' \in J_s}) - Y_{j,s}^h) = \sum_{i \in I} (X_{j,s}^i(I_s^i, \{P_{j',s}\}_{j' \in J_s}) - Y_{j,s}^i).$$

Let $\mu_{j,s}$ be the multiplier on this constraint, scaled to the units of prices,¹⁵ and let $P_{j,s}^*$ be the price chosen by the planner for good j in state s . The multiplier $\mu_{j,s}$ can be interpreted as the additional social cost of good j in state s above the price $P_{j,s}^*$.

If the prices that clear markets are also the prices that maximize social welfare (e.g. if the classic welfare theorems apply), the ratio of the social cost $P_{j,s}^* + \mu_{j,s}$ to the private cost $P_{j,s}^*$ will be the same for all goods $j \in J_s$ within each state $s \in S$.¹⁶ In this case, the agents' and the planner's preferences are perfectly aligned, and the solution to the constrained planner's problem is also a competitive equilibrium.

However, if markets are incomplete, then generically, the solution to the constrained planner's problem will not coincide with a competitive equilibrium (Geanakoplos and Polemarchakis (1986)). If prices are rigid, or if there are constraints on agents' goods allocations that depend on prices, pecuniary externalities will lead to generic constrained inefficiency regardless of whether markets are complete or incomplete (Farhi and Werning (2016)). In these cases, the multipliers $\mu_{j,s}$ are non-zero, and the ratio of $P_{j,s}^* + \mu_{j,s}$ to $P_{j,s}^*$ is not the same for all goods within each state. In what follows, I will focus on the incomplete markets case, but the results I derive will hold regardless of whether the underlying source of inefficiency is incomplete markets, nominal rigidities, prices in constraints, or some combination thereof.

To quantify these inefficiencies, I define “wedges” (following Farhi and Werning (2016)). The wedge $\tau_{j,s}^r$ is the difference between the social/private cost ratio for the good $j \in J_s$ in state $s \in S_1$ and the average ratio for all goods in that state,¹⁷

$$\pi_s^r \tau_{j,s}^r = -\frac{P_{j,s}^* + \mu_{j,s}}{P_{j,s}^*} + \frac{1}{|J_s|} \sum_{j' \in J_s} \frac{P_{j',s}^* + \mu_{j',s}}{P_{j',s}^*}.$$

¹⁵Converting from social marginal utility to price units requires scaling by a Pareto weight and an agent's marginal utility; see section 2 below for an example.

¹⁶Only relative prices matter. If j_0 is the numeraire (implying $\mu_{j_0,s} = 0$ and $P_{j_0,s}^* = 1$), then the ratio of social to private cost will be the same for all goods if and only if $\mu_{j,s} = 0$ for all goods.

¹⁷This definition of wedges is essentially the same as the one employed by Farhi and Werning (2016), adjusted for the difference between production and endowment economies and scaled by the reference measure π_s^r . It is not necessary in what follows to define wedges for the state s_0 .

The wedge $\tau_{j,s}^r$ is scaled by the probability measure $\pi_s^r > 0$. In the applications I consider, this measure is a risk-neutral measure or the physical probability measure. The wedges $\tau_{j,s}^r$ are defined in the context of this reference measure, and if defined instead under an alternative reference measure $\pi_s^{r'}$ would be rescaled, $\tau_{j,s}^{r'} = \frac{\pi_s^r}{\pi_s^{r'}} \tau_{j,s}^r$.

The wedge is the difference between the first-order conditions of the planner and of the agents—the latter do not account for effects of their demands on goods prices, and these pecuniary externalities, due to market incompleteness, affect welfare. It is positive if the social cost of a good is low relative to its price.

The wedges can be compensated for by transferring income in state s between agents. Let h and i be two agents in the economy. Let $X_{I,j,s}^h = \frac{\partial}{\partial I} X_{j,s}^h(I, \{P_{j',s}^*\}_{j' \in J_s})|_{I=I_s^{h*}}$ be the change in h 's consumption of good j in state s if given a marginal unit of income, holding prices constant, evaluated at the income I_s^{h*} and prices $P_{j',s}^*$ that solve the constrained planner's problem, and let $X_{I,j,s}^i$ be the same income effect for i . If the wedge-weighted difference of these income effects,

$$\Delta_s^{h,i,r} = \sum_{j \in J_s} P_{j,s}^* \tau_{j,s}^r (X_{I,j,s}^h - X_{I,j,s}^i),$$

is positive, transferring income from i to h in state s has a benefit, from the planner's perspective, because it alleviates externalities. I will call $\Delta_s^{h,i,r}$ the “externalities” because they summarize this benefit.¹⁸

For goods $j \in J_s$ with positive wedges $\tau_{j,s}^r$, the social cost of the good is lower than the price, and it is desirable to increase demand for the good. If $X_{I,j,s}^h > X_{I,j,s}^i$, then transferring income from i to h will indeed increase demand for the good. Summing these effects across goods determines the marginal benefit $\Delta_s^{h,i,r}$ of a transfer of income from i to h in state s .

Under optimal policy, the marginal benefit of reallocating income between i and h across the states $s \in S_1$ is offset by a marginal cost—distorting risk-sharing. Absent regulation, agents share risks by trading assets, ignoring externalities. The planner, in contrast, distorts risk-sharing to address these externalities.

Consider an asset $a \in A$ that can be freely traded by both i and h in the solution to the constrained planner's problem. The planner can reallocate the asset between these agents; as a result, the marginal benefit of such a reallocation must equal the marginal cost under optimal policies. Reallocating the asset between h and i has a cost if it prevents those agents from equating their valuations of the asset. The following proposition shows how the planner equates the marginal benefit of reducing externalities and the marginal cost of distorting risk-sharing. Let $Z_{a,s}^* = Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s})$ be the asset payoffs and let $M_s^{h,r,*}$ and $M_s^{i,r,*}$ be the agents' SDFs in the solution to the planner's problem.

¹⁸Farhi and Werning (2016) define an object $\tau_{D,s}^i$, which is closely related, $\pi_s^r \Delta_s^{h,i,r} = \tau_{D,s}^h - \tau_{D,s}^i$.

Proposition 1. *In the solution to the planner's problem, for any agents h and i , and any asset $a \in A$, if the planner is free to reallocate a between h and i , then*

$$\sum_{s \in \mathcal{S}_1} \pi_s^r \Delta_s^{h,i,r} Z_{a,s}^* = \sum_{s \in \mathcal{S}_1} \pi_s^r (M_s^{i,r,*} - M_s^{h,r,*}) Z_{a,s}^*. \quad (4)$$

These results holds for both the endowment economy of appendix section C and the production economy of section 4 of Farhi and Werning (2016).

Proof. See the appendix, section F.1, or Farhi and Werning (2016). □

If there is a complete market of securities that can be freely reallocated by the planner between h and i , then the externalities $\Delta_s^{h,i,r}$ must exactly equal the difference of the agents' SDFs, $M_s^{i,r,*} - M_s^{h,r,*}$. In this case, the externalities can be non-zero if there are other agents who cannot trade in the complete securities market (as shown in the example of section 2 below).¹⁹ In the incomplete markets case, the externalities $\Delta_s^{h,i,r}$ are equal to the difference of the agents' SDFs within the span of the payoff space of the assets that can be reallocated between the agents. That is, because the planner can move these assets between the agents, the planner must equate the marginal benefit of doing so (alleviating the externalities) with the marginal cost (distorting risk-sharing).

Implementation. The marginal cost vs. marginal benefit tradeoff just described arises from a planning problem in which the planner allocates assets for each of the agents. Because the planner chooses each agent's asset allocation, asset prices do not enter the constrained planner's problem. I next describe how the planner can implement optimal policy using asset markets. In this implementation, there will be a single price for each asset, and hence I will be able to discuss asset prices.

There is tension between assuming that each asset has a single price and the results of Proposition 1. In the presence of externalities, Proposition 1 requires that the willingness to pay for asset a of h be different from that of i , $\sum_{s \in \mathcal{S}_1} \pi_s^r M_s^{h,r,*} Z_{a,s}^* \neq \sum_{s \in \mathcal{S}_1} \pi_s^r M_s^{i,r,*} Z_{a,s}^*$. Consequently, h and i cannot both be free to trade the asset at the price Q_a . To implement optimal policy, the planner must place constraints on one or both of the agents' ability to trade the asset.²⁰ These constraints are what I will call macro-prudential policy; leverage restrictions are an example.

The planner has a great deal of latitude about the form of these constraints. The agents, when trading the asset, consider both the asset price and the shadow cost of the constraints (as in, e.g.,

¹⁹The externalities can also be non-zero if prices are rigid or if prices enter constraints on agents' goods allocations, as discussed above.

²⁰The planner could also use agent-specific taxes on asset holdings, so that the post-tax asset price faced by the agents is different even if the pre-tax price is the same. The FDIC fees charged to US banks are an example along these lines. I focus on quantity constraints because these appear more common in practice.

Garleanu and Pedersen (2011) or Du et al. (Forthcoming)). As long as these prices and shadow costs are consistent with Proposition 1, then the resulting equilibrium will be constrained efficient. That is, the functional form of the constraints does not matter, as long as the constraints generate the appropriate shadow costs. In particular, the constraints could be a function of both asset prices and portfolio choices (like a capital requirement), but this is not required.²¹

Under the financial intermediation structure I have assumed, proposition 2 below shows it is without loss of generality for the planner to implement optimal policy via portfolio constraints on intermediaries only. Because households trade through intermediaries, by regulating the trade of intermediaries with each household and with each other, the planner can dictate the asset allocation for all agents.²²

Now consider an arbitrage-able security $a \in A^*$, tradable by some household h . Applying (1) to this asset and (2) to its replicating portfolio illustrates the relationship between the arbitrage on asset a and the externalities.

Proposition 2. *The planner can implement the solution to the constrained planning problem using portfolio constraints on intermediaries only, and without constraining the trades of one intermediary, $i^* \in \mathcal{I}$, in the intermediary-only assets A^I .*

In this implementation, for any arbitrage-able asset $a \in A^$ that is tradable by household h and that the planner is free to reallocate between h and i^* ,*

$$\underbrace{-Q_a + \sum_{a' \in A^I} w_{a'}(a) Q_{a'}}_{\text{arbitrage violation}} = \underbrace{\sum_{s \in \mathcal{S}_1} \pi_s^r \Delta_s^{h, i^*, r} Z_{a, s}^*}_{\text{expected externality-weighted payoffs}}. \quad (5)$$

Proof. See the appendix, section F.2. □

This equation demonstrates the tight connection under optimal policy between arbitrage and the externalities the planner attempts to correct. To correct pecuniary externalities, the planner must distort risk-sharing. Under the assumed structure of financial intermediation, the planner can implement the optimal risk-sharing distortions by regulating intermediaries. In this implementation, certain assets will be priced by households (because households are not directly regulated) while others will be priced by intermediaries (because these assets are not tradable by households). For the subset of assets that are arbitrage-able (tradable by households with an intermediary-only

²¹For a formal definition of these portfolio constraints, see the appendix, Section §C.

²²This point also illustrates the limits of the Farhi and Werning (2016) framework, which does not allow for either private information or hidden trade, both of which would limit the set of implementable allocations. I discuss in appendix section D why my basic conclusions would continue to hold in the presence of private information or hidden trade.

replicating portfolio), this implementation of optimal policy will lead to an apparent arbitrage opportunity.

Strikingly, arbitrage is a generic feature of constrained efficient allocations (if the planner implements the constrained efficient allocation in the manner described by Proposition 2).²³ Thus, the absence of arbitrage is not a sign of efficiency, but rather a sign of inefficiency in the presence of incomplete markets (or other frictions, more generally). That said, not all patterns of arbitrage can be justified. An arbitrage-able asset should be cheap relative to its replicating portfolio if and only if its payoffs occur mainly in states in which the planner would like to transfer wealth from intermediaries to households.

The implementation described by Proposition 2 assumes that there exists an intermediary who is completely unconstrained with respect to trade in the replicating portfolio. This does not mean the intermediary is unregulated; it means only that regulatory constraints do not bind for that intermediary with respect to intermediary-only assets. Interpreting intermediary-only assets as derivatives, this implementation assumes only that derivatives are priced by some intermediary.²⁴

Lastly, note that these results inherit the generality of the Farhi and Werning (2016) framework. In particular, they allow for arbitrary heterogeneity across agents in both preferences and endowments. Depending on the nature of the assumed heterogeneity, optimal policies may or may not need to discriminate between different groups of households or intermediaries.²⁵ I next provide a concrete example to further illustrate the connection between externalities and arbitrage.

2 An Example of Externalities as Arbitrage

This section describes a version of Fanelli and Straub (2021) (see also example 5.4 of Farhi and Werning (2016)) to illustrate the meaning of Proposition 2. In this example, a planner limits foreign-currency lending by intermediaries (i.e. uses capital controls) to stabilize the real exchange rate. This creates a CIP violation whose size and direction are determined by the externalities as in (5).

This example connects to the empirical exercise that follows in that it illustrates how CIP violations can arise from optimal macro-prudential policy. However, this example focuses on

²³Generically, externalities are non-zero and are in the span of the payoff space of A^* .

²⁴It is likely more realistic to suppose regulations impact derivatives to a small degree; in this case, the result should be viewed as an approximation.

²⁵Specifically, if it is optimal for some particular subset of households or intermediaries to bear a certain risk, optimal macro-prudential policy must ensure that those agents end up bearing that risk, which may require implementing asymmetric regulations across different groups households /intermediaries. Such precise targeting may or may not be feasible; to simplify the analysis, I do not consider “equal-treatment” constraints. Examples of asymmetric regulations include the special regulations for “globally systemically important banks” and borrowing limits that vary across geographic areas.

capital controls, and hence is more naturally interpreted as concerning developing economies, whereas my empirical application focuses on developed markets and bank regulation. I present this example because it more transparently illustrates the principles behind the exercise.

There are two types of domestic households, Ricardians and non-participants ($\mathcal{H} = \{r, n\}$). In each state, there are two goods, tradable and non-tradable, $J_s = \{T, NT\}$. Both households have log utility preferences over a Cobb-Douglas aggregate of tradables and non-tradables, with share parameter α on tradables.

The future state can be either good (g) or bad (b), $S_1 = \{g, b\}$. Non-participants are endowed with non-tradables Y_{NT} and tradables $Y_{T,s}^n$, with $Y_{T,g}^n > Y_{T,b}^n$. Ricardian households are endowed only with tradables Y_T^r . Only $Y_{T,s}^n$ varies across the states; otherwise, the states are identical. Foreign intermediaries are risk-neutral, consume tradables only, and have a large endowment of tradables in all states. The discount factor for both households is $\beta < 1$, and is one for the intermediaries.

The tradables price is stable in the foreign currency, and the domestic price index is stable in the domestic currency, and I will use the tradable good as the numeraire ($P_{T,s} = 1$). The exchange rate is therefore $e_s = (P_{NT,s})^{1-\alpha}$, Ricardians can trade both a foreign-currency risk-free bond a_{fc} , with $Z_{a_{fc},s}(P_{NT,s}) = 1$, and a domestic risk-free bond a_{dc} , with $Z_{a_{dc},s}(P_{NT,s}) = (P_{NT,s})^{1-\alpha}$. Intermediaries can trade both of these bonds, an intermediary-only foreign currency risk-free bond a_I , and a currency forward a_F at exchange rate F , $Z_{a_F,s}(P_{NT,s}) = (P_{NT,s})^{1-\alpha} - F$. The bonds a_{fc} and a_{dc} are both arbitrage-able: $Z_{a_I,s}(\cdot) = Z_{a_{fc},s}(\cdot)$ and $F \times Z_{a_I,s}(\cdot) + Z_{a_F,s}(\cdot) = Z_{a_{dc},s}(\cdot)$. Non-participants cannot trade any assets.

The planner's problem is to maximize a weighted sum of household utility, subject to a participation constraint for the foreign intermediaries. Note that there is a complete market traded between intermediaries and Ricardians, as the exchange rate will vary across $S_1 = \{g, b\}$ and the domestic and foreign currency bonds will have different returns. As a result, the participation constraint of the intermediaries creates a unified budget constraint for the Ricardians.

Let π_s^p be the physical measure,²⁶ with $\pi_{s_0}^p = 1$, let T_{s_0} be the transfer in s_0 to non-participants, and let $\lambda^p > 0$ and $\lambda^r > 0$ be Pareto weights. The planner solves

$$\max_{\{I_s^r \geq 0, I_s^n \geq 0\}_{s \in S}, \{P_{NT,s} \geq 0\}_{s \in S}, T_{s_0}} \sum_{h \in \{r, n\}} \lambda^h \sum_{s \in S} \pi_s^p V_s^h(I_s^h, P_{NT,s}),$$

$$\text{subject to non-tradable market clearing, } Y_{NT} = \sum_{h \in \{r, p\}} X_{NT,s}^h(I_s^h, P_{NT,s}), \forall s \in S,$$

$$\text{the non-participants budget constraints, } I_s^n = P_{NT,s} Y_{NT} + Y_{T,s}^n + \mathbf{1}\{s = s_0\} T_{s_0}, \forall s \in S,$$

$$\text{and the Ricardian budget constraint, } I_{s_0}^r + \pi_g^p I_g^r + \pi_b^p I_b^r \leq 2Y_T^r - T_{s_0}.$$

²⁶Because intermediaries are risk-neutral, this is also the intermediaries' risk-neutral measure.

The functional forms in this example lead to

$$V_s^h(I_s^h, P_{NT,s}) = \begin{cases} \beta[\ln(I_s^h) - \ln(P_{NT,s}^{1-\alpha}) + (1-\alpha)\ln(1-\alpha)] & s \in \{g, b\}, \\ \ln(I_s^h) - \ln(P_{NT,s}^{1-\alpha}) + (1-\alpha)\ln(1-\alpha) & s = s_0, \end{cases}$$

$$X_{NT,s}^h(I_s^h, P_{NT,s}) = (1-\alpha)\frac{I_s^h}{P_{NT,s}}.$$

The market-clearing condition highlights the pecuniary externality present in the model.²⁷ If the Ricardian households sell a bond, reallocating income from the states in S_1 to the state s_0 , this will increase the price of the non-tradable good in s_0 and reduce the price in the states $\{g, b\}$. These price changes have an effect on welfare because the poor households face incomplete (non-existent) markets. The additional social cost of the non-tradable good, $\mu_{NT,s}$, is determined by the planner's first-order condition with respect to $P_{NT,s}$, and can be written for $s \in S_1$ as

$$\frac{\mu_{NT,s}}{P_{NT,s}^*} = \beta \frac{I_{s_0}^n}{\lambda^n} \pi_s^p \left(\frac{\lambda^n + \lambda^r}{I_s^n + I_s^r} - \frac{\lambda^n}{I_s^n} \right),$$

where $\frac{\lambda^n}{I_{s_0}^n} \mu_{NT,s}$ is the multiplier on the goods market clearing constraint. In states in which the income share of non-participants, $\frac{I_s^n}{I_s^n + I_s^r}$, is lower than the relative welfare weight $\frac{\lambda^n}{\lambda^n + \lambda^r}$, the planner would like to increase non-participant incomes. Because non-participants are net sellers of non-tradables, it is desirable in this case to increase $P_{NT,s}$ and the social cost of non-tradables is less than the private cost.²⁸

The wedges are $\pi_s^p \tau_{NT,s}^p = -\pi_s^p \tau_{T,s}^p = -\frac{1}{2} \frac{\mu_{NT,s}}{P_{NT,s}^*}$, and the externalities simplify to

$$\Delta_s^{r,i,p} = -(1-\alpha) \frac{\mu_{NT,s}}{P_{NT,s}^* \pi_s^p} = (1-\alpha) \beta \frac{I_{s_0}^n}{I_s^n} \left(1 - \frac{\lambda^n + \lambda^r}{\lambda^n} \frac{I_s^n}{I_s^n + I_s^r} \right).$$

Let us now consider the first-order conditions of the planner's problem with respect to Ricardian households' income and with respect to the transfer T_{s_0} ,

$$\beta \pi_s^p \frac{\lambda^r}{I_s^r} - (1-\alpha) \frac{\mu_{NT,s}}{P_{NT,s}^*} \frac{\lambda^n}{I_{s_0}^n} = \pi_s^p \frac{\lambda^r}{I_{s_0}^r} - \pi_s^p (1-\alpha) \frac{\mu_{NT,s_0}}{P_{NT,s_0}^*} \frac{\lambda^n}{I_{s_0}^n} = \pi_s^p \frac{\lambda^n}{I_{s_0}^n} - \pi_s^p (1-\alpha) \frac{\mu_{NT,s_0}}{P_{NT,s_0}^*} \frac{\lambda^n}{I_{s_0}^n}.$$

The transfer ensures that the goods market at date zero is efficient ($\mu_{NT,s_0} = 0$). Combining these

²⁷Note, in this model, that the pecuniary externality arises from a failure of risk-sharing due to incomplete markets; the planner could be said to trade-off risk-sharing between Ricardians and non-participants against risk-sharing between Ricardians and foreigners. This perspective is specific to examples in which the externalities arise from market incompleteness.

²⁸Note that I have ignored tradable goods market clearing (Walras' law), and without loss of generality $\mu_{T,s} = 0$.

equations produces the complete markets analog of (4),

$$\pi_s^p (M_s^{i,p} - M_s^{r,p}) = \pi_s^p \Delta_s^{r,i,p}, \quad (6)$$

where $M_s^{i,p} = 1$ and $M_s^{r,p} = \beta \frac{I_{s0}^r}{I_s^r}$ are the SDFs of the intermediaries and Ricardians, respectively. Summing this equation weighted by the payoffs $Z_{a,s}(P_{NT,s}^*)$ gives a version of (4) for log-utility households and risk-neutral intermediaries.

Moreover, the externalities must be non-zero in the solution to the planner's problem (and hence that competitive equilibria are inefficient). If the externalities were zero, the income shares $\frac{I_g^m}{I_g^m + I_g^r}$ and $\frac{I_b^m}{I_b^m + I_b^r}$ would both equal $\frac{\lambda^n}{\lambda^n + \lambda^r}$, by (6) the Ricardian incomes would be equal, $I_g^r = I_b^r$, and therefore the non-participant incomes would be equal, $I_g^m = I_b^m$. Market clearing in non-tradables requires that if incomes are identical across $\{g, b\}$, so are prices. But if the non-tradable price is the same in g and b , non-participant income cannot be equal in those states by the assumption that $Y_{T,g}^n > Y_{T,b}^n$, and therefore the externalities must be non-zero.

More specifically, in the solution to the planner's problem, the non-tradable price will be lower in b than in g . Consequently, the domestic bond has a lower return in b than in g . In the absence of regulation, in the competitive equilibrium of this example the Ricardians will borrow from intermediaries using the foreign currency bond (because $\beta < 1$ and the Ricardians have no income risk). The planner, to increase the price of non-tradables in state b relative to state g , would instead prefer that Ricardians borrow from intermediaries using the domestic bond, thereby increasing Ricardian income in b relative to g . A macro-prudential regulation limiting the quantity of foreign-currency lending by intermediaries is one way of implementing this outcome. Depending on parameters, the planner might also limit the total amount of lending by intermediaries.

When the planner implements optimal policies as described in proposition 2, the externalities $\Delta_s^{r,i,p}$ will manifest as arbitrages. The Ricardians will price the foreign-currency and domestic-currency bonds a_{fc} and a_{dc} . The foreign intermediaries will price the intermediary-only (i.e. off-shore) foreign-currency risk-free bond a_I and the currency forward a_F . The resulting arbitrages are

$$Q_{a_I} - Q_{a_{fc}} = \pi_g^p \Delta_g^{r,i,p} + \pi_b^p \Delta_b^{r,i,p}, \quad (7)$$

$$F \times Q_{a_I} + Q_{a_F} - Q_{a_{dc}} = \pi_g^p \Delta_g^{r,i,p} (P_{NT,g}^*)^{1-\alpha} + \pi_b^p \Delta_b^{r,i,p} (P_{NT,b}^*)^{1-\alpha}. \quad (8)$$

The first is a difference between the price intermediaries use when borrowing or lending with each other and the price they use when borrowing or lending to the Ricardians. The second is a CIP violation that involves the domestic currency bonds (the asset Ricardians can trade) and a replicating portfolio only intermediaries can trade (the currency forward and the intermediary-only bond). These two arbitrages closely resemble the arbitrages I study in the empirical exercise that

follows.

3 The Externality-Mimicking Portfolio

Let us now adopt the perspective of a financial economist who observes asset prices, and wants to know what externalities would justify the patterns of arbitrage in those asset prices. Suppose the financial economist observes prices for a set of arbitrage-able assets A^* tradable by some household h , along with the prices of the corresponding replicating portfolios of intermediary-only assets (e.g. derivatives). Further suppose that the financial economist believes these arbitrages are enabled by regulation. In the context of the preceding example, $A^* = \{a_{fc}, a_{dc}\}$ (the foreign and domestic currency bonds), and h is a Ricardian household.

Suppose regulatory policy is implemented as described in section 1; then (3) will hold. In this section, I show how (3) can be “inverted” to recover the difference of the household and intermediary SDFs from asset prices. When the assets in A^* form a complete market (i.e. in the example of the previous section), we can perfectly recover the difference of the SDFs from asset prices. When A^* does not form a complete market, we will instead recover the projection of the difference of the SDFs onto the space of returns. In both cases, we will recover the (projected) difference of the SDFs by constructing a portfolio, the “externality-mimicking portfolio.” The externality-mimicking portfolio’s returns track the externalities that would justify existing regulation as optimal (by Proposition 1).

For simplicity, I use the space of returns, $R_{a,s} = \frac{Z_{a,s}^*}{Q_a}$, as opposed to the space of payoffs, and therefore assume that every arbitrage-able asset and its replicating portfolio have strictly positive prices.²⁹ The return of an arbitrage-able asset $a \in A^*$, $R_{a,s}$, and the return of its replicating portfolio, $R_{a,s}^I$, are linked by the relationship $R_{a,s}^I = (1 - \chi_a)R_{a,s}$, where

$$\chi_a = \frac{-Q_a + \sum_{a' \in A^I} w_{a'}(a) Q_{a'}}{\sum_{a' \in A^I} w_{a'}(a) Q_{a'}} \quad (9)$$

is a scale-free measure of arbitrage. Intuitively, when the asset is cheaper than its replicating portfolio, its returns are higher. Using this notation, we can rewrite (7) and (8) from the example in the previous section as

$$\begin{bmatrix} \chi_{a_{fc}} \\ \chi_{a_{dc}} \end{bmatrix} = \begin{bmatrix} \pi_g^p R_{a_{fc},g}^I & \pi_b^p R_{a_{fc},b}^I \\ \pi_g^p R_{a_{dc},g}^I & \pi_b^p R_{a_{dc},b}^I \end{bmatrix} \cdot \begin{bmatrix} \Delta_g^{r,i,p} \\ \Delta_b^{r,i,p} \end{bmatrix}.$$

²⁹This is without loss of generality if there is a risk-free arbitrage-able security, as one could always add some amount of the risk-free security to another other security to ensure that its price is positive, while still ensuring that a replicating portfolio exists.

Now consider the portfolio of replicating portfolios (in the example, a portfolio of the currency forward and intermediary-only bond, expressed as weights on the replicating portfolios of the bonds traded by Ricardian households) defined by

$$\begin{bmatrix} \theta_{a_{fc}}^* \\ \theta_{a_{dc}}^* \end{bmatrix} = \underbrace{\begin{pmatrix} \begin{bmatrix} R_{a_{fc},g}^I & R_{a_{dc},g}^I \\ R_{a_{fc},b}^I & R_{a_{dc},b}^I \end{bmatrix} \\ \begin{bmatrix} \pi_g^p R_{a_{fc},g}^I & \pi_b^p R_{a_{fc},b}^I \\ \pi_g^p R_{a_{dc},g}^I & \pi_b^p R_{a_{dc},b}^I \end{bmatrix} \end{pmatrix}^{-1}}_{\text{inverse second moment matrix of } R_{a,s}^I} \cdot \begin{bmatrix} \chi_{a_{fc}} \\ \chi_{a_{dc}} \end{bmatrix}.$$

This portfolio has returns that are equal to the externalities in each state,

$$\begin{bmatrix} \Delta_g^{r,i,p} \\ \Delta_b^{r,i,p} \end{bmatrix} = \begin{bmatrix} R_{a_{fc},g}^I & R_{a_{dc},g}^I \\ R_{a_{fc},b}^I & R_{a_{dc},b}^I \end{bmatrix} \cdot \begin{bmatrix} \theta_{a_{fc}}^* \\ \theta_{a_{dc}}^* \end{bmatrix}.$$

This portfolio, which is externality-mimicking portfolio in the context of the capital controls example, is the projection of the externalities onto space of returns. In this example, which features complete markets, the portfolio's returns are the unique set of externalities that would justify the observed pattern of arbitrage. More generally, in the incomplete markets case, the externality-mimicking portfolio's returns are a (not unique) set of externalities that would justify the observed pattern of arbitrage.

These results are analogous to (and build on) the results of Hansen and Richard (1987), who study the projection of an SDF onto the space of returns. Those authors also show that their projection is equivalent to minimizing the variance of an SDF, subject to the constraint that the SDF price a set of assets.³⁰ Hansen and Jagannathan (1991) then show that the portfolio whose returns are the projection of the SDF is also the portfolio with the maximum available Sharpe ratio. Below, I develop analogous interpretations of the externality-mimicking portfolio.

Constructing the externality-mimicking portfolio requires three ingredients:

1. A set of arbitrage-able assets A^* ,
2. Prices for the arbitrage-able assets in A^* and their replicating portfolios, and
3. Expected returns and a variance-covariance matrix for the assets in A^* .

I assume, to simplify the exposition, that A^* includes a risk-free asset, whose return is R_f . Let $R_f^I = (1 - \chi_f)R_f$ be the return on the replicating portfolio of the risk-free asset, and let χ^{A^*} be the vector of scaled arbitrages χ_a for the risky assets in A^* . Let $v^{A^*,r}$ and $\Sigma^{A^*,r}$ be the mean vector and variance-covariance matrix of the returns $R_{a,s}$ for each risky arbitrage-able asset $a \in A^*$, under some measure π_s^r . Given $v^{A^*,r}$ and $\Sigma^{A^*,r}$, the mean returns and variance-covariance matrix for the returns $R_{a,s}^I$, $v^{A^*,I,r}$ and $\Sigma^{A^*,I,r}$, are defined by the relationship $R_{a,s}^I = (1 - \chi_a)R_{a,s}$. I assume there

³⁰This equivalence holds if the constraint that the SDF be positive does not bind.

are no redundant risky arbitrage-able assets ($\Sigma^{A^*,r}$ has full rank). From these objects, I define the externality-mimicking portfolio.

Note, by definition, that the space of returns of the arbitrage-able assets is identical to the space of returns of the replicating portfolios. As a result, the externality-mimicking portfolio can be defined as either a portfolio of the arbitrage-able assets or as a portfolio of replicating portfolios. It is convenient for what follows to define it as a portfolio of replicating portfolios; for the alternative, see appendix section E.

Definition 1. The externality-mimicking portfolio is a portfolio of the replicating portfolios of A^* , with weights on the risky replicating portfolios equal to

$$\theta^{A^*,r} = (\Sigma^{A^*,I,r})^{-1} (\chi^{A^*} - \chi_f \frac{v^{A^*,I,r}}{R_f^I}), \quad (10)$$

and a weight on the risk-free replicating portfolio equal to

$$\theta_f^{A^*,r} = -(\theta^{A^*,r})^T \frac{v^{A^*,I,r}}{R_f^I} + \frac{1}{(R_f^I)^2} \chi_f. \quad (11)$$

Proposition 3 below demonstrates four facts about this portfolio. The first two facts show that this portfolio is the projection of the externalities (under the assumptions of proposition 2) onto the space of replicating portfolio returns (i.e. the space of returns for assets in A^*). The third fact shows that the return of the portfolio is the difference between the household and intermediary SDFs, projected onto the space of returns. The fourth fact shows that the portfolio maximizes the “Sharpe ratio due to arbitrage,” which I define next.

Given an arbitrary portfolio θ of risky assets in A^* , consider both its Sharpe ratio, $S^{A^*,r}(\theta)$, and the Sharpe ratio of the replicating portfolio,³¹

$$S^{A^*,r}(\theta) = \frac{\frac{\theta^T v^{A^*,I,r}}{R_f^I} - \sum_{a \in A^*} \theta_a}{(\theta^T \Sigma^{A^*,I,r} \theta)^{\frac{1}{2}}}.$$

$S^{A^*,r}(\theta)$ is defined similarly, with $(v^{A^*,r}, \Sigma^{A^*,r}, R_f)$ in the place of $(v^{A^*,I,r}, \Sigma^{A^*,I,r}, R_f^I)$.

Because the prices of the replicating portfolios are not the same as the prices of the arbitrage-able assets, an allocation in dollars to arbitrage-able assets and the same dollar asset allocation to the replicating portfolios are in fact claims to different cashflows.³² I instead compare portfo-

³¹The definition of the Sharpe ratio given here is signed, and might be scaled by the inverse of R_f^I when compared to other definitions of the Sharpe ratio. Note also that the portfolio weights θ do not need sum to one (the units of θ are “dollars”, not percentages), and that the Sharpe ratio is homogenous of degree zero.

³²For example, if both intermediaries and households can buy stocks at \$1/share but households pay \$2/bond

lios that are claims to the same cashflows but perhaps have different prices. Define the portfolio transformation $\tilde{\theta}(\theta)$ by

$$\tilde{\theta}_a(\theta) = (1 - \chi_a)\theta_a.$$

This transformation converts an allocation in dollars at the replicating portfolio prices to an allocation in dollars at the arbitrage-able asset prices.

I define the ‘‘Sharpe ratio due to arbitrage’’ as the difference between the Sharpe ratio on a set of claims and the replicating Sharpe ratio of those same claims,

$$\hat{S}^{A^*,I,r}(\theta) = S^{A^*,r}(\tilde{\theta}(\theta)) - S^{A^*,I,r}(\theta).$$

A little algebra shows that the Sharpe ratio due to arbitrage is the ratio of the excess arbitrage, $\chi^{A^*} - \chi_f \frac{v^{A^*,I,r}}{R_f^I}$, to the volatility of the portfolio,

$$\hat{S}^{A^*,I,r}(\theta) = \frac{\theta^T \cdot (\chi^{A^*} - \chi_f \frac{v^{A^*,I,r}}{R_f^I})}{(\theta^T \Sigma^{A^*,I,r} \theta)^{\frac{1}{2}}}.$$

I next summarize the properties of the externality-mimicking portfolio.

Proposition 3. *Under the assumptions of proposition 2, the externality-mimicking portfolio has the following properties:*

1. *The externalities are the return on the portfolio plus a zero-mean residual, uncorrelated with the returns of all arbitrage-able assets $a \in A^*$: for all $i \in I$,*

$$\Delta_s^{h,i^*,r} = \sum_{a \in A^*} R_{a,s}^I \theta_a^{A^*,r} + \varepsilon_s^{A^*,r},$$

$$\sum_{s \in \mathcal{S}_1} \pi_s^r R_{a,s}^I \varepsilon_s^{A^*,r} = 0 \quad \forall a \in A^*.$$

2. *The variance of the externalities, $\sum_{s \in \mathcal{S}_1} \pi_s^r (\Delta_s^{h,i^*,r} - \frac{\chi_f}{R_f^I})^2$, is weakly greater than the variance of the externality-mimicking portfolio’s return, $(\theta^{A^*,r})^T \Sigma^{A^*,I,r} \theta^{A^*,r}$.*
3. *Let $m_s^{I,r}$ be any SDF that prices the replicating portfolios under the measure π_s^r . Then $m_s^r =$*

whereas intermediaries pay \$1/bond, an allocation of \$4 split equally between stocks and bonds means two shares and two bonds for the intermediaries, but two shares and one bond for the households.

$m_s^{I,r} + \sum_{a \in A^*} R_{a,s}^I \theta_a^{A^*,r}$ is the solution to the problem:

$$\begin{aligned} \min_{m \in \mathbb{R}^{|\mathcal{S}_1|}} \sum_{s \in \mathcal{S}_1} \pi_s^r (m_s - m_s^{I,r})^2 \text{ subject to} \\ \sum_{s \in \mathcal{S}_1} \pi_s^r m_s R_{a,s} = 1 \quad \forall a \in A^*. \end{aligned}$$

4. The Sharpe ratio due to arbitrage of the externality-mimicking portfolio, $\hat{S}^{A^*,I,r}(\theta^{A^*,r})$, is weakly greater than the Sharpe ratio due to arbitrage of any other portfolio of replicating portfolios of the assets in A^* .

Proof. See the appendix, section F.3. □

The first two claims follow from the least-squares projection. The third is the analog of Hansen and Richard (1987), and shows that the estimated household SDF $m_s^r = m_s^{I,r} + \sum_{a \in A^*} R_{a,s}^I \theta_a^{A^*,r}$ is the one that makes the household's and intermediary's SDFs as close as possible, subject to the constraint that it price the arbitrage-able assets A^* (and therefore differ sufficiently from $m_s^{I,r}$, which prices the replicating portfolios).³³ The fourth claim is the analog of Hansen and Jagannathan (1991), and shows that the externality-mimicking portfolio is also the one that maximizes the Sharpe ratio due to arbitrage.

The externality-mimicking portfolio is a reflection of what regulation is actually accomplishing. Consider a state s in which the externality-mimicking portfolio has a negative 10% return. If policy is optimal, the best linear prediction of the externalities in this state is negative 10%. That is, the planner would be indifferent between being able to transfer ex-post one extra dollar from households to intermediaries in state s and receiving an additional $10\% \times \pi_s^r$ dollars in the initial state s_0 .

The externality-mimicking portfolio is defined in the context of the reference measure π^r . In my empirical exercises, I focus on the intermediaries' risk-neutral measure, $\pi_s^{i^*} = \pi_s^p R_f^I M_s^{i^*,r}$, and consider the physical (or actual) probability measure, π^p , in robustness exercises. The two corresponding externalities are the "risk-neutral externalities" Δ_s^{h,i^*,i^*} and the "physical externalities" $\Delta_s^{h,i^*,p}$, and are linked, $\pi_s^p \Delta_s^{h,i^*,p} = \pi_s^{i^*} \Delta_s^{h,i^*,i^*}$. This connection reflects the usual equivalence in asset pricing between state-dependent preferences and beliefs. Using the risk-neutral externality-mimicking portfolio has a particular advantage, which is that all expected returns are equal to the risk-free rate, and hence observable. Moreover, options and quanto option³⁴ prices (which I presume are traded only by intermediaries) can reveal risk-neutral variances and covariances (Kremens and Martin (2019)). If we consider only arbitrage-able assets for which options and

³³That is, m_s^r maximizes a "market integration" measure along the lines of Chen and Knez (1995).

³⁴A quanto option is an option that involves both an exchange rate and an asset (such as the S&P 500). See appendix section B.4 or Kremens and Martin (2019) for details.

quantos are available (currencies and the S&P 500), no estimation is required when constructing the risk-neutral externality-mimicking portfolio. Using the physical measure, in contrast, requires estimating both expected returns and a variance-covariance matrix.

The externality-mimicking portfolio reveals what externalities would justify observed patterns of arbitrage under an optimal policy. The next step in our revealed preference exercise is to ask whether the recovered externalities make sense. We generally expect that externalities (and hence the mimicking portfolio returns) are negative in “bad” states of the world. That is, governments seem tempted to bailout intermediaries in bad states, not in good states. To test whether regulations are consistent with this intuition, we need to define what we mean by “bad” states. I consider two definitions, which result in two different tests. The first definition is to define bad times as being bad for the intermediaries, which is to say that the intermediaries’ SDF is high. The second definition involves studying a particular situation– the “stress tests” conducted by the Federal Reserve– in which the regulator is concerned about externalities, and would like intermediaries to have more wealth. Presumably, the idea behind the stress tests is to ensure that intermediaries have sufficient wealth in the stress scenario so as to avoid a bailout ex-post.

The first approach yields a simple test. The covariance of the intermediary SDF and the risk-neutral externality-mimicking portfolio, under the physical measure, is

$$\begin{aligned} Cov^p(M_s^{i^*,r}, \Delta_s^{h,i^*,i^*}) &= \frac{1}{R_f^I} \left(\frac{\chi_f}{R_f^I} - \sum_{s \in S_1} \pi_s^p \Delta_s^{h,i^*,i^*} \right) \\ &= -\frac{1}{R_f^I} (\theta^{A^*,i^*})^T \cdot (v^{A^*,I,p} - R_f^I) + Cov^p(M_s^{i^*,r}, \varepsilon_s^{A^*,i^*}). \end{aligned}$$

If the externalities are negatively correlated with the intermediaries’ SDF, the expected excess return of the risk-neutral externality-mimicking portfolio under the physical measure should be positive. Therefore, after constructing the externality-mimicking portfolio, I will estimate its expected returns and ask whether or not they are positive. The covariance term between the projection error and the intermediary SDF shows that this test would be biased if there are components of the SDF that are not spanned by the space of arbitrage-able asset returns, and which are correlated with components of the externalities that are also unspanned.

The second approach uses stress tests to identify a particular state (the stress test scenario) in which externalities should be negative. The purpose of the stress test is to verify that intermediaries have sufficient wealth in the stress scenario. To the extent that regulations achieve this goal, they must operate by inducing the intermediaries to hold different assets and issue different liabilities than they otherwise would have. Consequently, the intermediaries’ counterparties (the households) must also hold different assets and issue different liabilities than they otherwise would have. In other words, if the regulations act to raise intermediaries’ wealth in certain scenarios, they must

lower the wealth of households in those scenarios (at least in an endowment economy). That is, stress test scenarios are a statement when the regulator perceives negative externalities associated with transferring wealth from intermediaries to households (negative $\Delta_s^{h,i,r}$). Consequently, if the regulations are having the desired effect, returns on the externality-mimicking portfolio in the stress test scenario should be negative. This test is biased if the unspanned component of the externalities is large in absolute value in the stress scenario.

Note, with regards to both of these tests, that the component of the externalities that is orthogonal to all asset returns is irrelevant; there is nothing the planner can do at the margin with regards to this component. Rather, the question is whether there are assets that are not arbitrage-able (say, because there is only a cash price and no derivative price available, or vice-versa) and whose returns are correlated both with the unspanned component of the externalities and some arbitrage-able asset's return. The presence of such an asset will cause these tests to be biased (in a manner akin to the usual omitted-variable bias), and this bias can be mitigated by considering a richer set of arbitrage-able assets (which I do in Appendix Section A).

4 Data

In this section, I describe the arbitrages, data sources, and additional assumptions I use to conduct the tests described in the previous section.

The set of arbitrage-able securities A^* should be limited to arbitrages enabled by regulation. Two classes of arbitrage that fit the criteria discussed in the introduction are the difference between the federal funds rate and the IOER rate and CIP violations (Bech and Klee (2011); Du et al. (2018)).

Constructing the externality-mimicking portfolio requires determining what is the “asset” $a \in A^*$ and what is the “replicating portfolio.” My framework offers two ways of making this distinction: assets $a \in A^*$ are tradable by households, whereas replicating portfolios are not, and intermediaries' trades in assets $a \in A^*$ are regulated, whereas trades in the replicating portfolio are not. To a first approximation, the difference between “cash assets” and “derivatives” lines up with both of these distinctions: derivatives are both less accessible to households and less regulated under various leverage and capital requirements. This also helps clarify how we should classify agents between the household and intermediary sectors: if an agent can relatively easily trade all legs of the arbitrages in question, we should think of them as an intermediary; if they face significant costs or barriers, we should think of them as a household.

Consider first the fed funds/IOER arbitrage, which is the difference of two risk-free rates. A bank can earn interest on excess reserves held at the Fed, whereas a household cannot. If there is no meeting of the FOMC within the next month, the bank is essentially guaranteed to earn one month's

worth of interest at the current overnight rate.³⁵ A household could instead purchase treasury bills, highly-rated commercial paper, repo agreements (via money market funds), bank deposits, or the like. I use 1-month OIS swap rates, which closely track the yields of one-month maturity highly-rated commercial paper in the US, as a proxy for a risk-free rate available to households that provides no liquidity benefits. These rates tend to be higher than the rates on one-month constant maturity treasuries, but lower than LIBOR rates (which may include credit risk).³⁶ In the notation of the model, R_f is the one-month OIS swap rate, and R_f^l is the interest rate on excess reserves.

Next, consider CIP violations. Here, guided by the “cash vs. derivatives” heuristic, I assume that households can purchase foreign-currency bonds, but cannot trade derivatives.³⁷ The asset $a \in A^*$ is a claim to one euro in one month, and households can purchase this asset by spot exchanging dollars for euros and then purchasing a risk-free euro-denominated bond. Following Du et al. (Forthcoming), I use OIS rates as proxies for the risk-free rates available to households in various currencies. The replicating portfolio involves an intermediary earning the dollar IOER rate for one month and using a one-month FX forward to lock in the dollar/euro exchange rate.³⁸

My data sample begins on January 4, 2011, and runs through March 12, 2018. My sample is restricted to include only days on which the settlement of the one-month currency forward occurs before the next FOMC meeting, excluding days with an FOMC meeting. Because the FOMC holds eight scheduled meetings each year, roughly one quarter of all non-weekend days are included in the dataset. My data on spot and forward exchange rates, FX options, and OIS rates are from Bloomberg. I use the London closing time for all of these instruments, following Du et al. (2018). I focus on the euro, yen, and pound, because these currencies are major currencies that are modeled explicitly in the Federal Reserve’s stress test scenarios, and the Australian dollar, which plays a role in the “carry trade.” For details of the data construction, see appendix section B.

Information on the “stress test” scenarios comes from the Federal Reserve’s website.³⁹ The “severely adverse” scenario described in the tests shows, among other variables, the level of euro, yen, and pound, as well as the Dow Jones Industrial Average, at a quarterly frequency. I collect both the one and four-quarter percentage changes for each of the assets I study, and in my analysis will pretend that these are returns that occur over a one-month horizon. For AUD, which is not explicitly

³⁵In rare circumstances, the Fed might change the IOER rate between meetings, but such changes have low ex-ante likelihood and are unlikely to materially alter the expected interest rate.

³⁶For example, on August 19th, 2016, the one-month constant-maturity treasury rate was 27bps, the AA non-financial one-month commercial paper rates was 37bps, the one-month OIS rate was 40bps, the IOER rate was 50bps, and one-month LIBOR was 52bps.

³⁷That is, either the household cannot trade derivatives, or the transactions costs on derivatives trades for households are high enough that households cannot profitably execute the arbitrage.

³⁸Implicitly, I am assuming that the default risk on the forward contract is negligible (or, to be more precise, that the pricing data reflects forward rates available to a risk-free counterparty). Du et al. (2018) argue, persuasively in my view, that this risk is negligible.

³⁹<https://www.federalreserve.gov/supervisionreg/dfa-stress-tests.htm>

modeled in the stress test scenarios, I impute the returns in each stress scenario by running a daily regression predicting AUD returns using the contemporaneous GBP, EUR, JPY, and stock returns over the preceding 720 days, and then use these regression coefficients along with the one or four-quarter stress returns.

To conduct the tests described in the previous section, several additional assumptions are required. To construct the risk-neutral externality-mimicking portfolio, I require a full variance-covariance matrix under the risk-neutral measure. I construct such a matrix from currency options on each currency possible currency pair. For details, see appendix section §B. To estimate that portfolio’s expected excess returns under the physical measure, an estimate of expected excess returns is required. Motivated by Meese and Rogoff (1983) and the related literature, I assume that exchange rates are random walks over my one-month horizon.

Table 1 presents the sample means and standard deviations of the arbitrage associated with each currency and the risk-free arbitrage. Conceptually, these statistics correspond to the term χ_a defined in (9). For example, for euros, it represents the percentage difference in price, in dollars today, of purchasing a single euro one month in the future by buying the euro at spot today and saving at OIS (the asset $a \in A^*$), and obtaining the same euro one month in the future by savings at the IOER rate and using a currency forward (the replicating portfolio). I also present the difference between the dollar OIS rate and IOER (R_f vs. R_f^I). The arbitrages have a one month horizon, but are scaled to annualized values.

This table also shows the option-implied volatility and correlations of each currency (with respect to the US dollar).⁴⁰ Finally, the table reports the empirical correlations between the currencies and both the SPDR ETF and the daily He et al. (2017) (HKM) intermediary capital factor, and the quanto-implied correlations between the currencies and the S&P 500 (as in Kremens and Martin (2019), see appendix section B for details). A positive correlation means appreciation relative to the dollar when the S&P 500 has positive returns.

From Table 1, we can observe several notable features of the data. First, intermediaries are able to earn a higher rate of interest than households (IOER vs. OIS). However, the positive sign on the euro, pound, and yen arbitrages implies that it is more expensive for intermediaries to use derivatives to purchase e.g. a euro one month in the future in exchange for a dollar today than it is for intermediaries to use products also available to households. The opposite pattern holds for the Australian dollar. Note also that the euro and Australian dollar are positively correlated with the S&P 500 and the HKM factor, while the yen is negatively correlated with the S&P 500. Immediately, by Proposition 2, we can observe that the AUD and JPY CIP violations do not have the expected sign, if either the S&P 500 or the HKM factor is a reasonable proxy for the externali-

⁴⁰A version of the table with physical measure estimated volatilities and correlations is in appendix section A. The average volatilities and correlations are strikingly similar to their risk-neutral counterparts.

ties. Figure 1 shows the time series of the “risk-neutral” excess arbitrages, $\chi_a - \chi_f$, for Australian dollar, euro, and yen. Using the risk-neutral excess arbitrage, as opposed to the physical measure excess arbitrage, eliminates the dependence on an estimate of expected returns.

5 Results

I first construct the risk-neutral externality-mimicking portfolio, using the euro, Australian dollar, and yen assets, and a risk-free asset. This portfolio can be constructed at daily frequency using Definition 1 and data on the arbitrages and the risk-neutral variance-covariance matrix implied by FX options prices. Figure 2 displays the time series of the portfolio weights on the risky assets (EUR, AUD, JPY).

A few patterns in the data are apparent. First, the portfolio is generally long yen and euro and short AUD, and long currencies overall. That is, the portfolio is short US dollars and short the carry trade.⁴¹ The “short US dollars” part is likely to generate positive expected returns, whereas the “short the carry trade” generates negative expected returns, and this latter effect will dominate (see Table 2 below). Interpreted through the lens of the model, this portfolio implies that a strong desire to transfer wealth from households to intermediaries (negative externalities) coincides with an appreciation of the US dollar and high returns for the carry trade. If the planner would like to transfer wealth to intermediaries in “bad times,” the first part seems sensible, in light of the safe haven role of the US dollar (see, e.g., Maggiori (2017)), but the second is surprising. Lustig and Verdelhan (2007) show that negative carry trade returns are associated with falls in consumption, and we would generally presume that these times are times when the planner would like intermediaries to have relatively more wealth. Second, the noticeable spikes in the euro and yen CIP deviations that occur around quarter- and year-end result in large changes to the portfolio weight. This is not surprising, as there is no corresponding large change in implied volatilities that would offset the effect. Interpreted through the lens of the model, suddenly binding constraints could only be justified by large changes in externalities, and hence in the externality-mimicking portfolio.

I next consider the predictions that this portfolio has about other arbitrages. I deliberately excluded GBP from the set of currencies used to form the externality-mimicking portfolio. This allows me to test whether the arbitrage predicted using the externality-mimicking portfolio is consistent with the arbitrage actually observed for the dollar-pound currency pair. Formally, I compute

$$\chi_{GBP} - \chi_f = \Sigma_{GBP}^{A^*, I^*} \theta^{A^*, I^*}, \quad (12)$$

⁴¹In the terminology of Lustig et al. (2011), the portfolio is long the “level” factor and short the “slope” factor (the slope is with respect to interest rates).

where θ^{A^*,i^*} is the externality-mimicking portfolio in (10) and $\Sigma_{GBP}^{A^*,I,i^*}$ is the covariance, under the intermediaries' risk-neutral measure, between the dollar-pound exchange rate and the risky assets used to form the externality-mimicking portfolio (EUR, AUD, JPY). This is equivalent to computing the excess arbitrage under the projected externalities, which coincides with the arbitrage under the true externalities if there is no covariance between the pound and the error of the projection.

Figure 3 displays the results graphically. The actual excess arbitrage in pounds is constructed from OIS rates in dollars and pounds, and the spot and forward dollar-pound exchange rates (using excess arbitrage eliminates the dependence on the IOER rate). The predicted excess arbitrage is constructed entirely from those same variables in euros, AUD, and yen, along with options prices on all possible currency pairs, which are used to both construct the externality-mimicking portfolio (in the matrix Σ^{A^*,I,i^*}) and to construct the covariances $\Sigma_{GBP}^{A^*,I,i^*}$. Note that the sets of financial instruments used to construct the actual and predicted excess arbitrages do not overlap. Nevertheless, the predicted and actual excess arbitrages track each other, except near the end of 2011. The R^2 of a regression of the actual arbitrage on the predicted arbitrage, with no constant, is 83%. For predictions involving other currencies, see appendix section A.1.

I next consider the expected return of this portfolio (the first test described in the previous section). Intuitively, because the portfolio is generally short the carry trade, the expected return on the portfolio is negative. This contradicts the intuition that the externalities should be negatively correlated with the SDF. Figure 4 presents the time series of expected excess returns on the portfolio, and Table 2 formally tests whether the average expected return over my sample is greater than or equal to zero (a one-sided test). I show results for the full sample, only for dates for which the trade crosses a quarter-end, and only for dates for which the trade crosses a year-end.⁴² I also formally test whether the quarter-end dates are different from other dates, and whether the year-end dates are different from other quarter-end dates. Both Bech and Klee (2011), for fed funds vs. IOER, and Du et al. (2018), for CIP, have documented that the arbitrage spikes near quarter-ends. As these results demonstrate, the problem of negative expected returns documented above is particularly acute at quarter and year-ends.⁴³

I now turn to the second test, using the stress tests. Once per year, the Federal Reserve describes a “severely adverse” scenario and requires banks to maintain various leverage and capital ratios in this scenario. In Table 3, I present the returns of the yen, euro, and stocks in the stress test scenarios, at both the one quarter and four quarter horizons, for each stress test conducted. A general pattern emerges: recent stress tests have involved sizable euro depreciations relative to the dollar, and

⁴²The trade crosses quarter/year end if the settlement dates of the spot and forward FX trades are before and after the end of some quarter/year, respectively.

⁴³The fact that the arbitrages spike at quarter- and year-ends is itself suggestive of inefficiency (why would such a policy be optimal?). However, this is a distinct issue from the point highlighted by my analysis: that the signs of the arbitrages are inconsistent with the expected direction of risk-transfer.

sizable yen appreciations. This pattern is consistent with the observation that, during my sample, stock market declines tend to coincide with euro depreciation and yen appreciation relative to the dollar, and that these sorts of correlations might influence how the Federal Reserve constructs the stress test scenarios. The stock return itself very negative in all of these scenarios.

The scenarios do not specify a return for the Australian dollar, presumably because it would be virtually impossible to specify the returns of every asset a bank might hold. I impute the return of the AUD using the stress tests returns on euro, yen, and pounds (not shown in Table 3) and the stock market. Because of the Australian dollar's positive correlation with the stock market and negative correlation with the yen, the imputed returns are quite negative. When banks calculate their stress scenario returns, they likely perform a similar kind of imputation.

Each of the stress test scenarios is associated with a particular date (listed in Table 3) which is the date at which the scenario starts. For each date in my sample that is also within 180 calendar days of the stress test date, I report the returns of the risk-neutral externality-mimicking portfolio under the associated stress test scenario. Requiring that the relevant financial market data come from a day that is within 180 days of the stress test date assigns almost all of the days in my sample to a single stress test per date, dropping only a handful of days.

We should expect that the stress test returns of the externality-mimicking portfolio are negative. However, this is not the case. At almost all points in time, the portfolio is long low-interest-rate currencies (EUR and/or JPY) and short high-interest-rate currencies (AUD), and as a result has positive returns in the stress scenario, because the carry trade performs poorly in the stress scenario.

Table 4 formally tests whether returns are negative, averaging across dates near a particular stress test. The p-values correspond to a one-sided test that the mean is less than or equal to zero. I am able to reject the hypothesis that returns are negative on average for all four stress test years beginning in 2014.

For robustness, appendix section A presents three sets of additional results. The first set uses the physical measure externality-mimicking portfolio instead of the risk-neutral externality-mimicking portfolio in the stress test exercise. The estimated physical and risk-neutral measure covariance matrices are similar, and the stress test results do not depend on the choice of reference measure. See appendix sections B.2 and A.2 for details on the construction of the portfolio and the results.

The second set of results incorporates an equity-based arbitrage between the SPDR ETF and SPY options into the risk-neutral externality-mimicking portfolio. The purpose of this robustness exercise is to demonstrate that the puzzling results of the main analysis are not driven by the choice of arbitrages to include in the portfolio. Including the SPDR-SPY arbitrage in the externality-mimicking portfolio increases both the complexity and data requirements of the exercise, and the noisiness of the results due to the imprecision in the measurement of the arbitrage. For these reasons, I do not include it in the main analysis. With this arbitrage, the results of the expected

return test are broadly unchanged— expected returns are robustly negative, contrary to expectations. The results for the stress test are “better,” in that more but not all of the stress returns are sharply negative.⁴⁴ This effect is driven by the combination of very negative equity returns in the stress scenario and a small positive SPDR-SPY arbitrage (on average). The data is described in appendix sections B.3 and B.4 and the results are presented in appendix section A.3.

The third set of results uses “carry” and “dollar” portfolios of currencies. These results support the interpretation that the externality-mimicking portfolio is short USD (which was expected) and short the carry trade (which was not expected). The negative expected returns and positive returns in the stress scenario that I document are due to the short carry aspect of the portfolio. See appendix sections B.5 and A.4 for details on the portfolios and the results.

Discussion. The risk-neutral externality-mimicking portfolio has negative expected returns and positive returns in the stress scenario, the opposite of what was expected. One possible explanation, of course, is that the goal of regulation is to encourage intermediaries to take more, not less, macro-prudential risk, and regulatory policy is accomplishing its goals. A more likely explanation, in my view, is that the current regulatory apparatus is not accomplishing its macro-prudential objectives.

There are several potential explanations for why regulation might move risk in the wrong direction. One possibility is a lack of coordination; there are many regulatory entities (both within and across countries), each with their own parochial interests. At a theoretical level, it is possible that each regulator is acting according to its own incentives, in a way that ends up retaining risk within the intermediary sector. This is not my preferred explanation, for the reason that regulatory coordination failures and jurisdiction shopping are generally thought to lead to a lack of binding regulation (and hence a lack of arbitrage), as opposed to strongly binding regulation.⁴⁵

My preferred explanation instead emphasizes the role of leverage constraints. Du et al. (2018) shows that the direction of the CIP arbitrage across currencies is predicted by the direction of the carry trade. A simple interpretation of this fact is that households or their proxies want to do the carry trade, and intermediaries are induced by the arbitrage to take the other side. Leverage constraints, such as the “supplementary leverage ratio,” prevent the intermediaries from fully satisfying households’ demands. If this story is correct, households are trying to take macro-economic risk and insure intermediaries from those risks, but regulation (the leverage ratio) is limiting this risk transfer, which is the exact opposite of what an optimal policy would do. Note that, under this interpretation, existing regulation does not by itself determine the sign of the arbitrage. Instead, the sign of the arbitrage under existing regulations arises from the interaction of customer demands and regulatory constraints.

⁴⁴These results taken together suggest that the stress returns might not be negative enough. Just because the returns in the stress scenario are negative does not mean macro-prudential regulation is working optimally.

⁴⁵That is, this kind of coordination failure seems more relevant for the pre-GFC period than the post-GFC period.

This, however, is a policy choice; instead of encouraging intermediaries to take more carry trade risk, regulation could be redesigned so as to discourage intermediaries from taking carry trade risk. One simple and specific proposal is to continue using stress test scenarios, but regulate intermediaries based on the return of their assets in those scenarios, as opposed to their leverage at the end of the stress test scenario. Under the former regime, a levered position in a “hedge” asset benefits the intermediary; in the latter (existing) regime, the effect of the leverage can end up penalizing the intermediary for hedging. This is precisely what is happening with respect to the USD-JPY and USD-AUD currency pairs in the current regime. Although the yen appreciates and the Australian dollar depreciates relative to the dollar in the stress scenarios, an intermediary holding a levered long yen or short AUD position is penalized for the leverage, instead of being rewarded for hedging.

This proposal also pushes against another potential explanation for the failure of the current regime, namely that un-modeled constraints prevent regulators from reaching better outcomes. It is certainly the case that regulators adopted leverage constraints in part because of the problems associated with risk-weighted constraints, an observation that emphasizes the difficulty of designing the optimal regulatory regime. That granted, emphasizing portfolio returns as opposed to leverage in the stress test scenario does not require solving for “optimal” risk weights and would facilitate a transfer of macro-prudential risk away from the intermediary sector.

6 Conclusion

There is a close connection between the externalities regulation attempts to address and the arbitrage that regulation creates. Using this connection, we can assess whether macro-prudential policies are achieving their objectives. I construct an externality-mimicking portfolio, which tracks the difference between intermediary and household SDFs, and hence the externalities that would rationalize existing policy. I argue that these externalities should negatively covary with the SDF, and be negative in “stress” scenarios, and develop tests demonstrating that current data is inconsistent with basic intuitions about the nature of optimal policies.

References

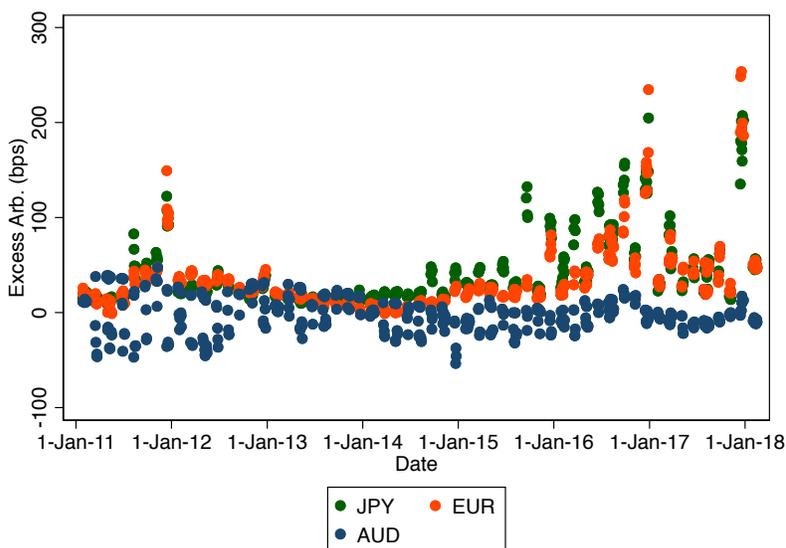
- Manuel Amador, Javier Bianchi, Luigi Bocola, and Fabrizio Perri. Exchange rate policies at the zero lower bound. *The Review of Economic Studies*, 2017.
- Leif Andersen, Darrell Duffie, and Yang Song. Funding value adjustments. *The Journal of Finance*, 74(1):145–192, 2019.
- Morten L Bech and Elizabeth Klee. The mechanics of a graceful exit: Interest on reserves and

- segmentation in the federal funds market. *Journal of Monetary Economics*, 58(5):415–431, 2011.
- Zhiwu Chen and Peter J Knez. Measurement of market integration and arbitrage. *The Review of Financial Studies*, 8(2):287–325, 1995.
- Eduardo Dávila and Anton Korinek. Pecuniary externalities in economies with financial frictions. *The Review of Economic Studies*, page rdx010, 2017.
- Eduardo Dávila, Daniel D Graves, and Cecilia Parlatore. The value of arbitrage. Technical report, National Bureau of Economic Research, 2022.
- Julio Davila, Jay H Hong, Per Krusell, and José-Víctor Ríos-Rull. Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica*, 80(6):2431–2467, 2012.
- Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018.
- Wenxin Du, Benjamin M Hébert, and Wenhao Li. Intermediary balance sheets and the treasury yield curve. Technical report, National Bureau of Economic Research, 2022.
- Wenxin Du, Benjamin M Hébert, and Amy Wang Huber. Are intermediary constraints priced? *The Review of Financial Studies*, Forthcoming.
- Darrell Duffie and Arvind Krishnamurthy. Passthrough efficiency in the fed’s new monetary policy setting. 2016.
- Sebastián Fanelli and Ludwig Straub. A theory of foreign exchange interventions. *The Review of Economic Studies*, 88(6):2857–2885, 2021.
- Emmanuel Farhi and Iván Werning. A theory of macroprudential policies in the presence of nominal rigidities. *Econometrica*, 84(5):1645–1704, 2016.
- Xavier Gabaix and Matteo Maggiori. International liquidity and exchange rate dynamics. *The Quarterly Journal of Economics*, 130(3):1369–1420, 2015.
- Nicolae Garleanu and Lasse Heje Pedersen. Margin-based asset pricing and deviations from the law of one price. *The Review of Financial Studies*, 24(6):1980–2022, 2011.
- John Geanakoplos and Heraklis M Polemarchakis. Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. *Uncertainty, information and communication: essays in honor of KJ Arrow*, 3:65–96, 1986.
- Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of financial Economics*, 66(2):361–407, 2002.
- Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of political economy*, 99(2):225–262, 1991.
- Lars Peter Hansen and Scott F Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica: Journal of the Econometric Society*, pages 587–613, 1987.
- Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017.
- Victoria Ivashina, David S Scharfstein, and Jeremy C Stein. Dollar funding and the lending behavior of global banks. *The Quarterly Journal of Economics*, 130(3):1241–1281, 2015.
- Lukas Kremens and Ian Martin. The quanto theory of exchange rates. *American Economic Review*, 109(3):810–43, 2019.

- Owen A Lamont and Richard H Thaler. Anomalies: The law of one price in financial markets. *The Journal of Economic Perspectives*, 17(4):191–202, 2003.
- Hanno Lustig and Adrien Verdelhan. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1):89–117, 2007.
- Hanno Lustig, Nikolai Roussanov, and Adrien Verdelhan. Common risk factors in currency markets. *The Review of Financial Studies*, 24(11):3731–3777, 2011.
- Matteo Maggiori. Financial intermediation, international risk sharing, and reserve currencies. *American Economic Review*, 107(10):3038–71, 2017.
- Richard A Meese and Kenneth Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics*, 14(1-2):3–24, 1983.
- Andrei Shleifer and Robert W Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1):35–55, 1997.
- Emil Siriwardane, Aditya Sunderam, and Jonathan Wallen. Segmented arbitrage. Available at SSRN 3960980, 2021.
- Zhaogang Song and Haoxiang Zhu. Mortgage dollar roll. *The Review of Financial Studies*, 32(8):2955–2996, 2019.
- Jules H Van Binsbergen, William F Diamond, and Marco Grotteria. Risk-free interest rates. *Journal of Financial Economics*, 143(1):1–29, 2022.

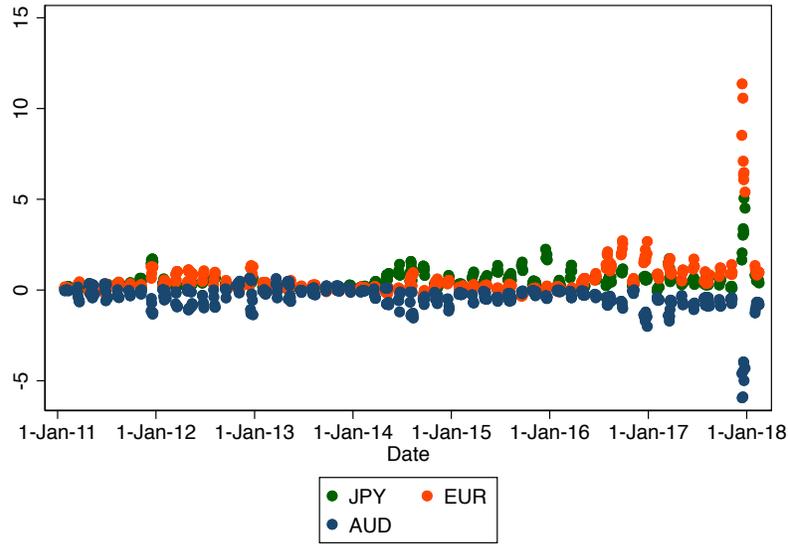
Tables and Figures

Figure 1: Time Series of Excess Arbitrage



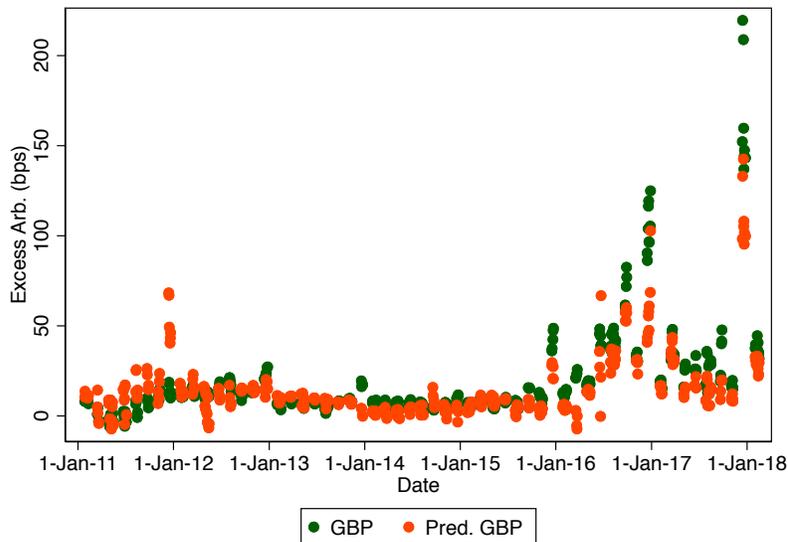
Notes: This figure plots the annualized excess arbitrage $\chi_a - \chi_f$, as defined as in (9), for the yen, euro, and Australian dollar. These excess arbitrages are approximately equal to the one month OIS-based CIP violation vs. USD for those currencies. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Figure 2: Externality-Mimicking Portfolio Weights



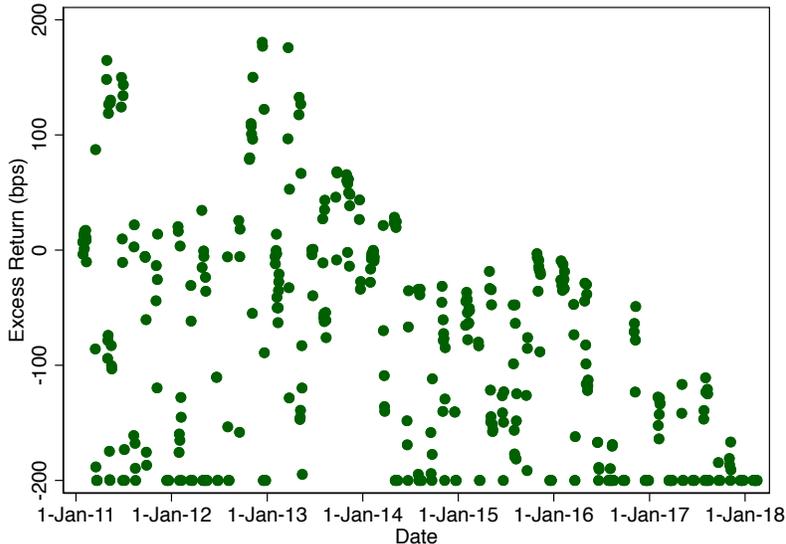
Notes: This figure plots the portfolio weights of the externality-mimicking portfolio (definition 1). The portfolio is constructed using a set of arbitrage-able assets A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The reference measure is the intermediaries' risk-neutral measure, meaning that expected returns are equal to the IOER rate and the variance-covariance matrix is inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Figure 3: Actual vs. Predicted Excess Arbitrage in Pounds



Notes: This figure plots excess the annualized pound excess arbitrage $\chi_{GBP} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based GBP-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Figure 4: Risk-Neutral EMP Expected Returns



Notes: This figure plots the excess expected return under the physical measure of the risk-neutral externality-mimicking portfolio (definition 1), under the assumption that currencies follow a random walk. The excess return is censored at +/- 200bps to enhance readability. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation is inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Table 1: Summary Statistics for Arbitrage

	Pounds	Euros	Yen	Aus. Dollar	OIS-IOER
Arbitrage Mean (bps/year)	6.7	22.4	28.3	-15.4	-12.5
Arbitrage SD (bps/year)	28.2	37.7	37.2	18.5	2.8
OI Vol. (bps/year)	859	950	977	1073	-
OI Corr. with Pound/USD	1.00	0.56	0.22	0.47	-
OI Corr. with Euro/USD	0.56	1.00	0.31	0.51	-
OI Corr. with Yen/USD	0.22	0.31	1.00	0.26	-
Empirical Corr. with SPDR	0.23	0.10	-0.34	0.37	-
Empirical Corr. with HKM	0.26	0.17	-0.31	0.31	-
Implied Corr. with S&P 500	0.28	0.11	-0.29	0.50	-
N	444	444	444	444	444

Notes: This table presents summary statistics for the sample of all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. Arbitrage mean χ_a is defined using (9) for a claim to e.g. one euro in one month, priced in dollars today. The OIS-IOER arbitrage is the risk-free arbitrage, based on a claim to one dollar in one month. Arbitrage SD is the daily standard deviation of χ_a . OI Vol. and OI Corr. variables for currencies are the time-series mean of a daily series extracted from variance-covariance matrices implied by currency options. Empirical Corr. with SPDR and Empirical Corr. with HKM are the time-series means of the correlations between the currency returns and the SPDR ETF (which tracks the S&P 500) and with the He et al. (2017) daily intermediary capital factor, as estimated on a rolling basis by the methodology described in appendix section B.2. Implied Corr. with S&P 500 is based on the time-series mean of the currency correlation with the S&P 500 extracted from quanto options and described in appendix section B.4.

Table 2: Risk-Neutral EMP Expected Returns

	N	Mean (bps)	Standard Deviation (bps)	Test	P-Value
Full Sample	444	-222	25.9	≥ 0	0.0000
Quarter-Ends	155	-431	69.7	≥ 0	0.0000
Year-Ends	46	-1005	209.2	≥ 0	0.0000
QE - NQE		-321		$= 0$	0.0000
YE - NYE QE		-816		$= 0$	0.0000

Notes: This table reports the excess expected return under the physical measure of the risk-neutral externality-mimicking portfolio (definition 1), under the assumption that currencies follow a random walk. The portfolio is constructed from an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation is inferred from currency options. The full sample is all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. The quarter-end and year-end sub-samples are restricted to days on which a quarter- or year-end occurs between the spot FX settlement date and the one-month FX settlement date. The QE-NQE and YE-NYE QE rows report the mean difference between quarter-end vs. non-quarter-end dates and year-end vs. non-year-end quarter-end. Test indicates the hypothesis about the mean being tested, and P-Value reports the associated p-value.

Table 3: Stress Test “Severely Adverse” Scenarios

Stress Test Date	Euro		Stocks		Yen		AUD*	*Imputed
	One-Quarter Return	Four-Quarter Return						
9/30/12	-7.7	-15.4	-19.3	-51.5	2.8	-1.0	-8.8	-24.5
9/30/13	-14.3	-21.4	-26.5	-49.5	3.1	-1.1	-17.2	-28.7
9/30/14	-12.0	-13.4	-16.3	-57.1	7.6	6.5	-5.1	-17.2
12/31/15	-7.7	-13.9	-20.2	-50.7	2.7	5.1	-7.7	-18.3
12/31/16	-9.1	-11.9	-34.0	-49.7	3.3	7.5	-17.0	-24.1
12/31/17	-6.6	-10.9	-51.3	-62.8	11.7	4.6	-23.7	-32.9

Notes: This table reports the percentage changes in the level of the euro, yen, and the Dow Jones Total Stock Market Index (“Stocks”) during the first one or four quarters of the associated “Severely Adverse Scenario” from that year’s stress test. These percentage changes are treated as returns in my analysis. Stress Test Date lists the date on which that year’s scenario begins. AUD shows the imputed return for the Australian dollar, using the imputation method described in the text.

Table 4: Risk-Neutral Returns in Stress Scenario

Stress Test Date	N	Mean (1Q,%)	S.D. (1Q,%)	P-value (1Q)	Mean (4Q,%)	S.D. (4Q,%)	P-value (4Q)
9/30/12	63	-0.8	0.3	0.9923	-1.8	1.0	0.9593
9/30/13	59	-0.9	0.4	0.9838	-2.1	0.7	0.9987
9/30/14	62	4.8	0.6	0.0000	9.9	1.0	0.0000
12/31/15	60	1.9	0.4	0.0000	4.5	0.8	0.0000
12/31/16	61	2.8	0.8	0.0000	7.0	1.2	0.0000
12/31/17	45	31.4	5.0	0.0000	28.6	4.3	0.0000

Notes: This table reports the mean and standard deviation of stress test scenario returns for the risk-neutral externality-mimicking portfolio. The risk-neutral EMP portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation is computed from currency options prices. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting and within 180 days of a stress test date. Each of these dates is assigned to the nearest stress test date. N reports the number of dates assigned to each stress test, and P-Value reports the p-value associated with a one-sided hypothesis test that the mean return is negative. Results are reported for both one-quarter (1Q) and four-quarter (4Q) returns from the stress test scenarios.

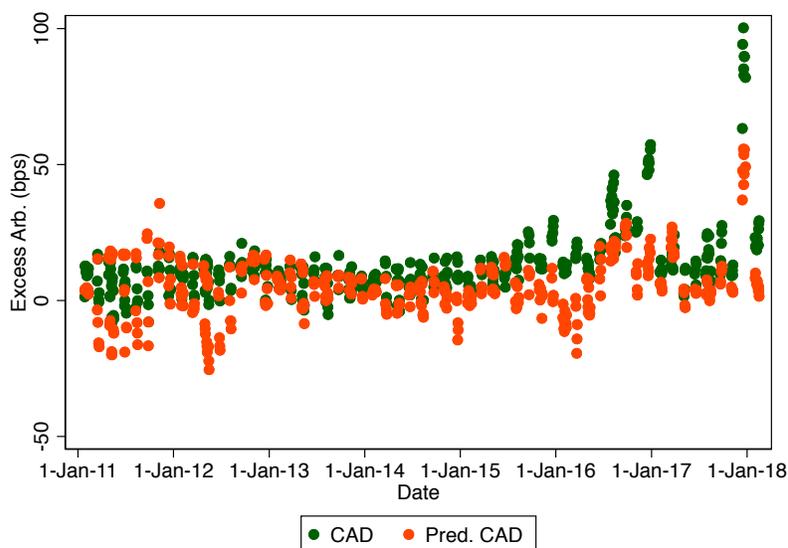
Appendix

A Additional Results

A.1 Predicted Arbitrage in Other Currencies

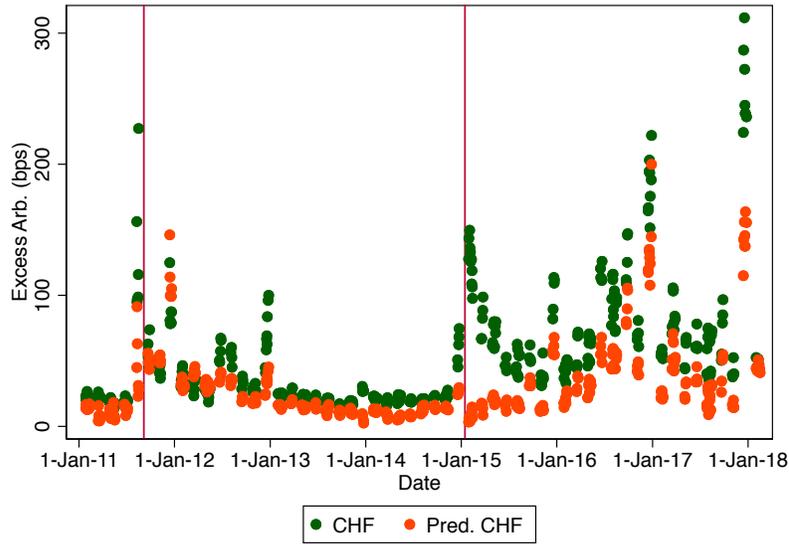
This sub-section presents the predicted vs. actual arbitrage using the risk-neutral externality-mimicking portfolio for three additional currencies: CAD, CHF, and SEK. These currencies have enough OIS swap and options data available to make these predictions, although for both CHF and SEK the sample size is reduced.

Figure 5: Actual vs. Predicted Excess Arbitrage in Canadian Dollar



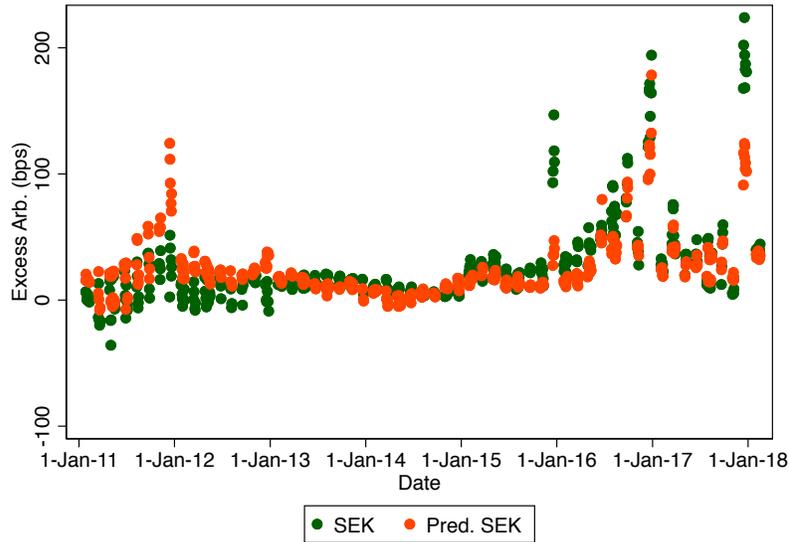
Notes: This figure plots excess the annualized CAD excess arbitrage $\chi_{CAD} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based CAD-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Figure 6: Actual vs. Predicted Excess Arbitrage in Swiss Franc



Notes: This figure plots excess the annualized CHF excess arbitrage $\chi_{CHF} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based CHF-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting. The two vertical lines indicate the beginning and end of a period during which CHF was pegged to EUR.

Figure 7: Actual vs. Predicted Excess Arbitrage in Swedish Krona



Notes: This figure plots excess the annualized SEK excess arbitrage $\chi_{SEK} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based SEK-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

A.2 Results using the Physical Measure

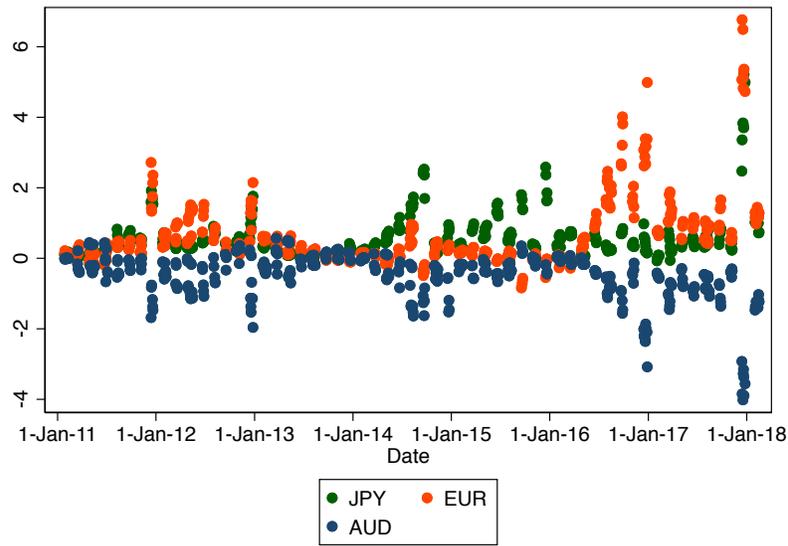
This sub-section presents results using the physical-measure externality-mimicking portfolio, as described in appendix section B.2. Table 5 presents a version of the summary statistics table, with an estimated variance-covariance matrix in the place of an option-implied variance-covariance matrix. I then present the portfolio weights, predicted vs. actual GBP arbitrage, and stress test returns for the physical measure. The results are similar to their counterparts from the main text.

Table 5: Summary Statistics for Physical-Measure Arbitrage

	Pounds	Euros	Yen	Aus. Dollar	OIS- IOER
Arbitrage Mean (bps/year)	6.7	22.4	28.3	-15.4	-12.5
Arbitrage SD (bps/year)	28.2	37.7	37.2	18.5	2.8
Empirical Vol. (bps/year)	812	852	922	1046	-
Empirical Corr. with Pound/USD	1.00	0.58	0.18	0.46	-
Empirical Corr. with Euro/USD	0.58	1.00	0.32	0.46	-
Empirical Corr. with Yen/USD	0.18	0.32	1.00	0.24	-
Empirical Corr. with SPDR	0.23	0.10	-0.34	0.37	-
Empirical Corr. with HKM	0.26	0.17	-0.31	0.31	-
N	444	444	444	444	444

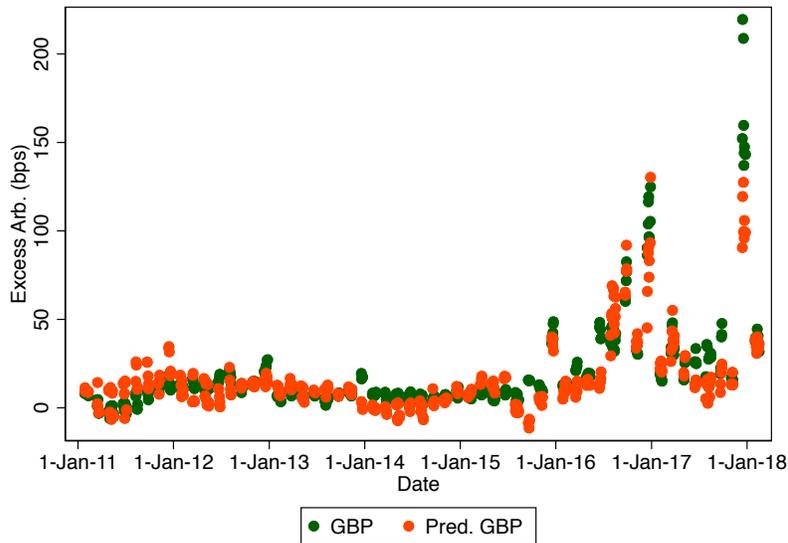
Notes: This table presents summary statistics for the sample of days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. Arbitrage mean χ_a is defined using (9) for a claim to e.g. one euro in one month, priced in dollars today. The OIS-IOER arbitrage is the risk-free arbitrage, based on a claim to one dollar in one month. Arbitrage SD is the daily standard deviation of χ_a . Empirical Vol. and Empirical Corr. variables for currencies are the time-series means of the correlations between the currency returns, as estimated on a rolling basis by the methodology described in appendix section B.2.. Empirical Corr. with SPDR and Empirical Corr. with HKM are the time-series means of the correlations between the currency returns and the SPDR ETF (which tracks the S&P 500) and with the He et al. (2017) daily intermediary capital factor, as estimated on a rolling basis by the methodology described in appendix section B.2.

Figure 8: Externality-Mimicking Portfolio Weights, Physical Measure



Notes: This figure plots the portfolio weights of the externality-mimicking portfolio (definition 1). The portfolio is constructed using a set of arbitrage-able assets A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The reference measure is the physical measure, meaning that expected returns calculated under the assumption that log exchange rates are random walks and the variance-covariance matrix is estimated using the RiskMetrics methodology. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Figure 9: Actual vs. Predicted Excess Arbitrage in Pounds, Physical Measure



Notes: This figure plots excess the annualized pound excess arbitrage $\chi_{GBP} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based GBP-USD CIP violation. The externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The reference measure is the physical measure, meaning that expected returns calculated under the assumption that log exchange rates are random walks and the variance-covariance matrix, as well as the covariances with the pound, is estimated using the RiskMetrics methodology. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Table 6: Physical Returns in Stress Scenario

Stress Test Date	N	Mean (1Q,%)	S.D. (1Q,%)	P-value (1Q)	Mean (4Q,%)	S.D. (4Q,%)	P-value (4Q)
9/30/12	63	-1.0	0.3	0.9998	-2.8	0.8	0.9995
9/30/13	59	-1.2	0.4	0.9963	-2.7	0.7	0.9999
9/30/14	62	6.2	0.9	0.0000	11.7	1.4	0.0000
12/31/15	60	1.9	0.5	0.0002	4.0	0.9	0.000
12/31/16	61	1.4	0.8	0.0384	4.9	1.1	0.0000
12/31/17	45	35.7	4.3	0.0000	31.8	3.3	0.0000

Notes: This table reports the mean and standard deviation of stress test scenario returns for the physical-measure externality-mimicking portfolio. The physical measure EMP portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset. The expected returns are calculated under the assumption that log exchange rates are random walks and the variance-covariance matrix is estimated using the RiskMetrics methodology. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting and within 180 days of a stress test date. Each of these dates is assigned to the nearest stress test date. N reports the number of dates assigned to each stress test, and P-Value reports the p-value associated with a one-sided hypothesis test that the mean return is negative. Results are reported for both one-quarter (1Q) and four-return (4Q) returns from the stress test scenarios.

A.3 Results including Equity Arbitrage

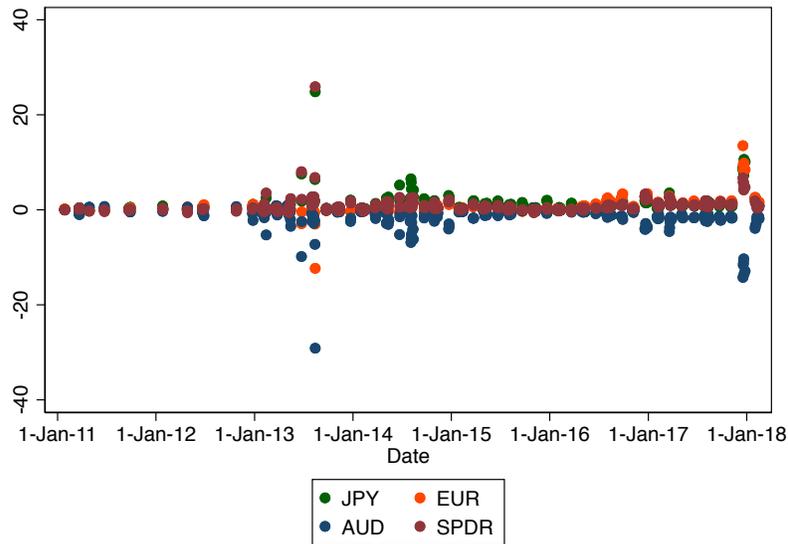
This sub-section presents results for a risk-neutral externality-mimicking portfolio that incorporates JPY, EUR, and AUD arbitrages as well as the SPY-based arbitrage described in appendix section B.3 and the risk-free rate arbitrage. The covariances and expected returns under the physical measure used in this section are described in appendix section B.4.

Table 7: Summary Statistics for Arbitrage including SPY

	Pounds	Euros	Yen	Aus. Dollar	SPY	OIS-IOER
Arbitrage Mean (bps/year)	6.7	22.4	28.3	-15.4	5.4	-12.5
Arbitrage SD (bps/year)	28.2	37.7	37.2	18.5	47.5	2.8
OI Vol. (bps/year)	859	950	977	1073	1566	-
OI Corr. with Pound/USD	1.00	0.56	0.22	0.47	0.28	-
OI Corr. with Euro/USD	0.56	1.00	0.31	0.51	0.11	-
OI Corr. with Yen/USD	0.22	0.31	1.00	0.26	-0.29	-
Empirical Corr. with SPDR	0.23	0.10	-0.34	0.37	1.00	-
Empirical Corr. with HKM	0.26	0.17	-0.31	0.31	0.66	-
Implied Corr. with S&P 500	0.28	0.11	-0.29	0.50	1.00	-
N	444	444	444	444	312	444

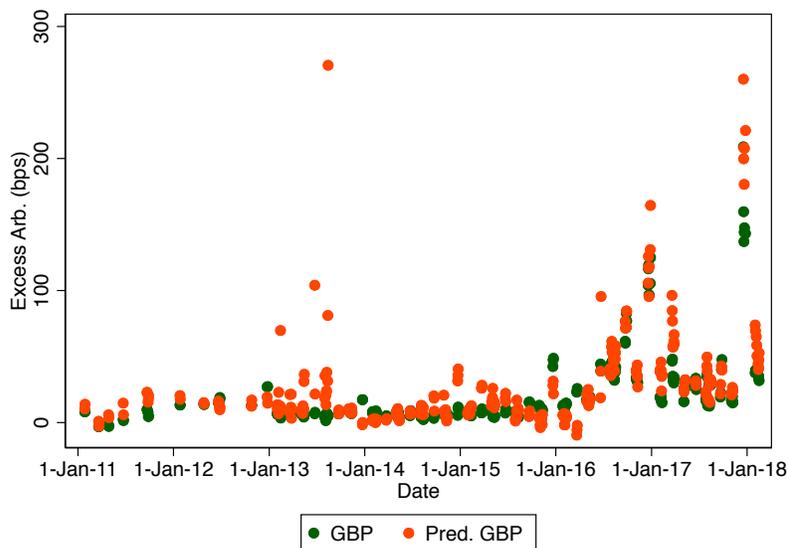
Notes: This table presents summary statistics for the sample of all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. The SPY statistics are restricted to dates at least one month before a SPDR ex-dividend date. Arbitrage mean χ_a is defined using (9) for a claim to e.g. one euro in one month, priced in dollars today. The OIS-IOER arbitrage is the risk-free arbitrage, based on a claim to one dollar in one month. Arbitrage SD is the daily standard deviation of χ_a . OI Vol. and OI Corr. variables for currencies are the time-series mean of a daily series extracted from variance-covariance matrices implied by currency options, SPY options, and quanto options. Empirical Corr. with SPDR and Empirical Corr. with HKM are the time-series means of the correlations between the currency returns and the SPDR ETF (which tracks the S&P 500) and with the He et al. (2017) daily intermediary capital factor, as estimated on a rolling basis by the methodology described in appendix section B.2. Implied Corr. with S&P 500 is based on the time-series mean of the currency correlation with the S&P 500 extracted from quanto options and described in appendix section B.4.

Figure 10: Risk-Neutral Externality-Mimicking Portfolio Weights with SPY



Notes: This figure plots the portfolio weights of the externality-mimicking portfolio (definition 1). The portfolio is constructed using a set of arbitrage-able assets A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset and the SPDR ETF. The reference measure is the intermediaries' risk-neutral measure, meaning that expected returns are equal to the IOER rate and the variance-covariance matrix is inferred from currency options, SPY options, and quanto options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting and one month before a SPDR ex-dividend date.

Figure 11: Actual vs. Predicted Excess Arbitrage in Pounds, Risk-Neutral Measure with SPY



Notes: This figure plots excess the annualized pound excess arbitrage $\chi_{GBP} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based GBP-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset and the SPDR ETF. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options, SPY options, and quanto options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting and one month before a SPDR ex-dividend date.

Table 8: Risk-Neutral EMP Expected Returns with SPY

	N	Mean (bps)	Standard Deviation (bps)	Test	P-Value
Full Sample	312	-111	43.3	≥ 0	0.0053
Quarter-Ends	93	-425	125.9	≥ 0	0.0005
Year-Ends	23	-1495	439.9	≥ 0	0.0013
QE - Full		-447.0		$= 0$	0.0000
YE - QE		-1421		$= 0$	0.0000

Notes: This table reports the excess expected return under the physical measure of the risk-neutral externality-mimicking portfolio (definition 1), under the assumption that currencies follow a random walk and with a 5% equity risk premium. The portfolio is constructed from an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset and the SPDR ETF. The variance-covariance matrix used in the computation is inferred from currency options, SPY options, and quanto options. The full sample is all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before a FOMC meeting and one month before a SPDR ex-dividend date. The quarter-end and year-end sub-samples are restricted to days on which a quarter- or year-end occurs between the spot FX settlement date and the one-month FX settlement date. The QE-NQE and YE-NYE QE report the mean difference between quarter-end vs. non-quarter-end dates and year-end vs. non-year-end quarter-end. Test indicates the hypothesis about the mean being tested, and P-Value reports the associated p-value.

Table 9: Risk-Neutral Portfolio with SPY, Returns in Stress Scenario

Stress Test Date	N	Mean (1Q,%)	S.D. (1Q,%)	P-value (1Q)	Mean (4Q,%)	S.D. (4Q,%)	P-value (4Q)
9/30/12	28	-1.9	0.6	0.9989	-6.0	2.1	0.9956
9/30/13	52	-0.7	1.5	0.6786	-17.7	4.3	0.9999
9/30/14	54	2.5	1.3	0.0304	-14.3	2.0	1.0000
12/31/15	53	0.1	0.7	0.4302	0.4	1.6	0.4142
12/31/16	55	-23.4	1.6	1.0000	-30.2	2.1	1.0000
12/31/17	41	0.0	5.4	0.5013	-16.7	4.8	0.9994

Notes: This table reports the mean and standard deviation of stress test scenario returns for the risk-neutral externality-mimicking portfolio. The risk-neutral EMP portfolio is constructed with an A^* that contains the yen, euro, and Australian dollar, as well as a risk-free asset and the SPDR ETF. The variance-covariance matrix used in the computation is computed from currency options prices, SPY options, and quanto options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before a FOMC meeting, one month before a SPDR ex-dividend date, and within 180 days of a stress test date. Each of these dates is assigned to the nearest stress test date. N reports the number of dates assigned to each stress test, and P-Value reports the p-value associated with a one-sided hypothesis test that the mean return is negative. Results are reported for both one-quarter (1Q) and four-return (4Q) returns from the stress test scenarios.

A.4 Results with Dollar and Carry Portfolios

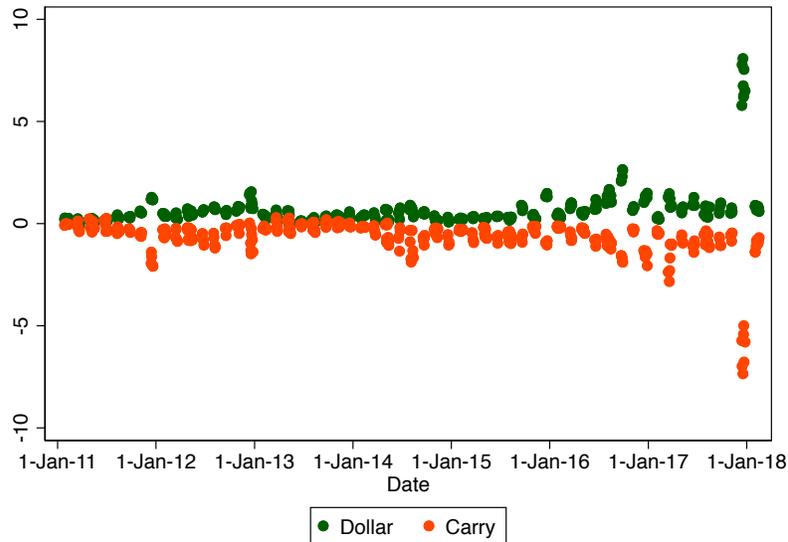
This sub-section presents for a risk-neutral externality-mimicking portfolio that incorporates “Carry” and “Dollar” arbitrages, as well as a risk-free rate arbitrage. Carry and Dollar (defined in appendix section B.5) are portfolios of currency trades. Carry is long two low-forward-premium currencies (AUD and CAD for most of the sample) and short two high-forward-premium currencies (JPY and EUR for most of the sample). Dollar is an equally weighted portfolio of currencies vs. USD. Note when interpreting the results below that Dollar is short USD, not long USD.

Table 10: Summary Statistics for Carry and Dollar Arbitrage

	Carry	Dollar	OIS-IOER
Arbitrage Mean (bps/year)	-44.1	8.6	-12.5
Arbitrage SD (bps/year)	31.4	23.3	2.8
OI Vol. (bps/year)	835	680	-
OI Corr. with Carry	1.00	0.15	-
OI Corr. with Dollar	0.15	1.00	-
Empirical Corr. with SPDR	0.56	0.21	-
Empirical Corr. with HKM	0.46	0.21	-
Implied Corr. with S&P 500	0.68	0.30	-
N	444	444	444

Notes: This table presents summary statistics for the sample of all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. Arbitrage mean χ_a is defined using (9) for a claim to e.g. one euro in one month, priced in dollars today. The OIS-IOER arbitrage is the risk-free arbitrage, based on a claim to one dollar in one month. Arbitrage SD is the daily standard deviation of χ_a . The OI Vol. and OI Corr. variables are the time-series mean of a daily series extracted from variance-covariance matrices implied by currency options. Empirical Corr. with SPDR and Empirical Corr. with HKM are the time-series means of the correlations between the currency returns and the SPDR ETF (which tracks the S&P 500) and with the He et al. (2017) daily intermediary capital factor, as estimated on a rolling basis by the methodology described in appendix section B.2. Implied Corr. with S&P 500 is based on the time-series mean of the currency correlation with the S&P 500 extracted from quanto options and described in appendix section B.4.

Figure 12: Risk-Neutral Externality-Mimicking Portfolio Weights with Carry and Dollar



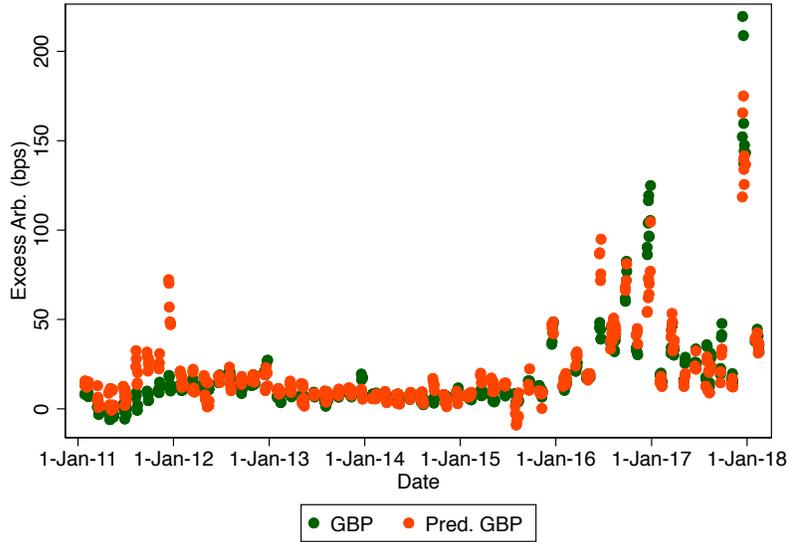
Notes: This figure plots the portfolio weights of the externality-mimicking portfolio (definition 1). The portfolio is constructed using a set of arbitrage-able assets A^* that contains the Carry and Dollar portfolios, as well as a risk-free asset. The reference measure is the intermediaries' risk-neutral measure, meaning that expected returns are equal to the IOER rate and the variance-covariance matrix is inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Table 11: Risk-Neutral EMP Expected Returns with Carry and Dollar

	N	Mean (bps)	Standard Deviation (bps)	Test	P-Value
Full Sample	444	-151	20.7	≥ 0	0.0000
Quarter-Ends	155	-315	56.1	≥ 0	0.0000
Year-Ends	46	-763	169.9	≥ 0	0.0000
QE - Full		-252		$= 0$	0.0000
YE - QE		-638		$= 0$	0.0000

Notes: This table reports the excess expected return under the physical measure of the risk-neutral externality-mimicking portfolio (definition 1), under the assumption that currencies follow a random walk. The portfolio is constructed from an A^* that contains the Carry and Dollar portfolios, as well as a risk-free asset. The variance-covariance matrix used in the computation is inferred from currency options. The full sample is all US trading days from Jan 4, 2011 to March 12, 2018 at least one month before an FOMC meeting. The quarter-end and year-end sub-samples are restricted to days on which a quarter- or year-end occurs between the spot FX settlement date and the one-month FX settlement date. The QE-NQE and YE-NYE QE report the mean difference between quarter-end vs. non-quarter-end dates and year-end vs. non-year-end quarter-end. Test indicates the hypothesis about the mean being tested, and P-Value reports the associated p-value.

Figure 13: Actual vs. Predicted Excess Arbitrage in Pounds, Risk-Neutral Measure with Carry and Dollar



Notes: This figure plots excess the annualized pound excess arbitrage $\chi_{GBP} - \chi_f$, as defined in (9), along with the predicted value defined as in (12). The excess arbitrage is approximately equal to the one month OIS-based GBP-USD CIP violation. The risk-neutral externality-mimicking portfolio is constructed with an A^* that contains the Carry and Dollar portfolios, as well as a risk-free asset. The variance-covariance matrix used in the computation and the covariances with the pound are inferred from currency options. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting.

Table 12: Stress Test “Severely Adverse” Scenarios

Stress Test Date	Carry*	*Imputed	Dollar*	*Imputed
	One-Quarter Return	Four-Quarter Return	One-Quarter Return	Four-Quarter Return
9/30/12	-7.4	-16.5	-4.1	-11.6
9/30/13	-8.9	-13.1	-10.5	-16.5
9/30/14	-2.1	-10.5	-3.2	-7.7
12/31/15	-3.8	-10.5	-4.1	-8.8
12/31/16	-11.3	-18.2	-7.9	-11.0
12/31/17	-23.6	-25.9	-8.0	-14.2

Notes: This table reports the imputed returns of the Carry and Dollar portfolios during the first one or four quarters of the associated “Severely Adverse Scenario” from that year’s stress test, using the imputation method described in the text. Stress Test Date lists the date on which that year’s scenario begins.

Table 13: Risk-Neutral Portfolio with Carry and Dollar, Returns in Stress Scenario

Stress Test Date	N	Mean (1Q,%)	S.D. (1Q,%)	P-value (1Q)	Mean (4Q,%)	S.D. (4Q,%)	P-value (4Q)
9/30/12	63	0.9	0.4	0.0144	0.4	0.9	0.3234
9/30/13	59	-1.8	0.4	1.0000	-2.9	0.6	1.0000
9/30/14	62	0.1	0.1	0.1707	3.9	0.6	0.0000
12/31/15	60	-0.5	0.1	0.9992	0.1	0.3	0.3965
12/31/16	61	3.7	0.6	0.0000	7.9	1.0	0.0000
12/31/17	45	27.5	4.4	0.0000	20.6	3.0	0.0000

Notes: This table reports the mean and standard deviation of stress test scenario returns for the risk-neutral externality-mimicking portfolio portfolio. The risk-neutral EMP portfolio is constructed with an A^* that contains the Carry and Dollar portfolios, as well as a risk-free asset. The variance-covariance matrix used in the computation is computed from currency options prices. The sample is all US trading days from Jan 4, 2011 to March 12, 2018 that are at least one month before an FOMC meeting and within 180 days of a stress test date. Each of these dates is assigned to the nearest stress test date. N reports the number of dates assigned to each stress test, and P-Value reports the p-value associated with a one-sided hypothesis test that the mean return is negative. Results are reported for both one-quarter (1Q) and four-return (4Q) returns from the stress test scenarios.

Internet Appendix (For Online Publication)

B Details on Data Construction

B.1 The Risk-Neutral Variance-Covariance Matrix

The risk-neutral variance-covariance matrix Σ^{A,I,i^*} is defined as variance-covariance matrix, under the intermediaries' risk-neutral measure, of the returns $R_{a,s}^I$. In the case of currencies, these are the returns of investing at the IOER rate for one month and using a currency forward to purchase, say, yen, and then immediately exchanging back to dollars the spot rate. Let $S_{j,t}$ be the spot exchange rate of currency j per dollar (e.g. euros per dollar, yen per dollar), and let $F_{j,t}$ be the one-month forward rate. The empirical analog of the return $R_{a,s}^I$ is

$$R_{j,t} = R_{f,t}^I \frac{F_{j,t}}{S_{j,t+1}},$$

where $R_{f,t}^I$ is the gross IOER rate accumulated over the one-month time horizon. Note that, because $R_{f,t}^I$ is the risk-free rate available to intermediaries, we must have $E_t^*[\frac{F_{j,t}}{S_{j,t+1}}] = 1$ where E_t^* denotes expectations taken under the intermediaries' risk-neutral measure.

To construct the risk-neutral variance-covariance matrix of currency returns, I use daily, London-closing at-the-money 1-month implied volatilities from Bloomberg for each currency pair. The volatilities are “percentage” volatilities from a log-normal Garman and Kohlhagen (1983) model. If $S_{j,t}$ and $S_{j',t}$ are two exchange rates vs. the US dollar at time t (e.g. euros per dollar and yen per dollar), then

$$Cov_t^*[\Delta s_{j,t+1}, \Delta s_{j',t+1}] = \frac{1}{2}(V_t^*[\Delta s_{j,t+1}] + V_t^*[\Delta s_{j',t+1}] - V_t^*[\Delta s_{j,t+1} - \Delta s_{j',t+1}])$$

where $\Delta s_{j,t} = \ln(\frac{S_{j,t+1}}{S_{j,t}})$ and Cov^* and V^* denote the risk-neutral variance and covariance, respectively. Under the assumption of log-normality,

$$\begin{aligned} Cov_t^*[\frac{S_{j,t}}{S_{j,t+1}}, \frac{S_{j',t}}{S_{j',t+1}}] &= E_t^*[\exp(-\Delta s_{j,t+1} - \Delta s_{j',t+1})] - E_t^*[\exp(-\Delta s_{j,t+1})]E_t^*[\exp(-\Delta s_{j',t+1})] \\ &= \frac{S_{j,t}S_{j',t}}{F_{j,t}F_{j',t}}(\exp(Cov_t^*[\Delta s_{j,t+1}, \Delta s_{j',t+1}]) - 1), \end{aligned}$$

where $F_{1,t}$ and $F_{2,t}$ are the forward rates. It follows that

$$\Sigma_{j,j',t}^{A,I,i} = (R_{f,t}^I)^2(\exp(\frac{1}{2}(V_t^*[\Delta s_{j,t+1}] + V_t^*[\Delta s_{j',t+1}] - V_t^*[\Delta s_{j,t+1} - \Delta s_{j',t+1}])) - 1).$$

In theory, I should make an adjustment to Bloomberg implied volatilities to use a discount rate associated with the IOER rate, instead of the more standard OIS rate. However, the difference amounts to about one basis point in the option price, and hence is negligible. I have also experimented with more sophisticated methods of computing the risk-neutral variance-covariance matrix (the “SVIX” method of Martin (2017)). Such methods avoid log-normality assumptions at the expense of additional data requirements and complexity, and have little impact on my results.

B.2 The Physical Expected Returns and Variance-Covariance Matrix

Expected returns under the physical measure are required both to conduct the expected returns test described in the text and to construct the physical measure externality-mimicking portfolio. The latter also requires an estimate of the physical measure variance-covariance matrix of returns.

As described in the text, I assume that currencies are random walks. Specifically, I assume log-normal exchange rates and that the log-exchange rate is a martingale. The expected excess return of using the IOER rate and a forward to purchase, say, one yen one month from now, is the determined by the difference between the forward and expected exchange rate. That is,

$$\mu_{j,t}^{A^*,I,P} = R_{f,t}^I E_t \left[\frac{F_{j,t}}{S_{j,t+1}} \right],$$

where S_{t+1} is the exchange rate in foreign currency per dollar, F_t is the one-month forward rate, $R_{f,t}^I$ is the IOER rate accumulated over the next month, and expectations are taken under the physical measure. Under the stated assumptions,

$$\mu_{j,t}^{A^*,I,P} = R_{f,t}^I \frac{F_{j,t}}{S_{j,t}} \exp\left(\frac{1}{2} V_t[\Delta s_{j,t+1}]\right)$$

where $V_t[\Delta s_{t+1}]$ is the conditional variance of the log change in the exchange rate. Consequently, armed with an estimate for $V_t[\Delta s_{t+1}]$, we can construct expected returns.

I estimate a daily physical-measure variance-covariance matrix using an exponentially weighted moving average of the daily series, with a decay factor of 0.97 (a procedure known as the “Risk-Metrics” methodology, see for example Alexander (2008)). This avoids using future information to estimate the variance-covariance matrix. I initialize my variance and covariance estimates at the beginning of 2011 with the realized variance/covariance for 2010. I then scale my daily estimated variance-covariance matrix to a one-month horizon.⁴⁶

I use this estimated variance-covariance matrix of log returns both to construct my estimates of mean returns, as above, and to construct a variance-covariance matrix for arithmetic returns, as

⁴⁶More sophisticated approaches that incorporate higher-frequency data might yield better results. See, for example, Ghysels et al. (2006).

described in appendix section B.1, under the assumption of log-normality. In this case,

$$\begin{aligned} Cov_t \left[\frac{S_{j,t}}{S_{j,t+1}}, \frac{S_{j',t}}{S_{j',t+1}} \right] &= E_t[\exp(-\Delta s_{j,t+1} - \Delta s_{j',t+1})] - E_t[\exp(-\Delta s_{j,t+1})]E_t[\exp(-\Delta s_{j',t+1})] \\ &= \frac{\mu_{j,t}^{A^*,I,P} \mu_{j',t}^{A^*,I,P}}{(R_{f,t}^I)^2} \frac{S_{j,t} S_{j',t}}{F_{j,t} F_{j',t}} (\exp(Cov_t[\Delta s_{j,t+1}, \Delta s_{j',t+1}]) - 1) \end{aligned}$$

and

$$\Sigma_{j,j',t}^{A^*,I,P} = (\mu_{j,t}^{A^*,I,P} \mu_{j',t}^{A^*,I,P}) (\exp(\frac{1}{2}(V_t[\Delta s_{j,t+1}] + V_t[\Delta s_{j',t+1}] - V_t[\Delta s_{j,t+1} - \Delta s_{j',t+1}])) - 1).$$

B.3 The SPDR ETF/SPY Option Arbitrage

In this section, I describe an equity-related arbitrage that I include in a robustness exercise. The arbitrage I consider is an arbitrage between the SPDR S&P 500 ETF and options on that ETF, which trade on the CBOE under the ticker SPY. This arbitrage is closely related to the classic index-future arbitrage involving S&P 500 futures (e.g. Chung (1991); MacKinlay and Ramaswamy (1988); Miller et al. (1994)).

The arbitrage I study considers the cost of purchasing a share of the SPDR ETF and holding it for one month as compared to the cost of purchasing that ETF via put-call parity (by buying a call and selling a put with a one month horizon). The ETF share itself is the arbitrage-able asset (both because it is easily purchased by households and because regulations affect intermediaries' trade in equity shares). The intermediary, to replicate the ETF, can buy a call on the ETF, sell a put on the ETF at the same strike, and invest enough cash at the IOER rate over the next month to cover the exercise price of the put/call. Regardless of whether the ETF ends up above or below the strike price, the intermediary will end up owning the ETF in one month.

The particular details of this arbitrage are complicated by the fact that the SPY options are "American" and not "European" options, meaning they can be exercised at any time. I deal with this issue by employing the Margrabe (1978) bound on American put prices, which in my setting (short time horizons and low interest rates) is reasonably tight. I discuss this issue in more detail below.

The ETF and the replicating portfolio will generate identical payoffs as long as there are no dividends over the course of the month (more precisely, that an ex-dividend date does not occur within the month). The ETF has ex-dividend dates quarterly, usually on the third Friday of March, June, September, and December.⁴⁷ I therefore limit my sample to avoid these dates. This illus-

⁴⁷The prospectus, available at https://us.spdrs.com/library-content/public/SPDR_500%20TRUST_PROSPECTUS.pdf, describes the details of how the ex-dividend dates are determined.

trates one of the two main advantages the ETF-based arbitrage has over the traditional S&P 500 cash-futures arbitrage. The stocks of the S&P 500 index pay dividends often, and hence most studies of index arbitrage assume either perfect foresight of dividends or use a dividend forecast, whereas no such assumptions are required for the ETF arbitrage. The second advantage relates to transactions costs and stale prices. The traditional index arbitrage involves buying and selling 500 stocks, generating substantial transactions costs and exacerbating the issue that prices might not be synchronized. Using the ETF, which is one of the most actively traded securities in the equity market and has a very small bid-offer, mitigates many of these issues. Of course, synchronizing the options prices and the ETF price is still critically important, as in Van Binsbergen et al. (2012).

For this arbitrage, I am assuming that the costs associated with posting margin on the options are negligible. That is, the margin is sufficiently small, and the interest rate the intermediary receives on the posted margin sufficiently close to the IOER rate, that these costs are negligible. This assumption is also, implicitly, being applied to the margin required by counterparties in the OTC market for FX swaps when studying CIP violations.

Because the SPY options are American, not European, I construct arbitrage bounds as opposed to a single arbitrage measure. It is straightforward to observe that an American call or put must be weakly more valuable than its European counterpart, but the possibility of early exercise implies that this weak inequality might be strict.

Let $p_a(K) \in A$ and $c_a(K) \in A$ denote the American put and American call of strike K . Following the argument of Margrabe (1978), let $\hat{p}_a(K)$ be an American put option to exchange the ETF for an amount that grows at the IOER rate, $K \exp(t \cdot \ln(R_f^I))$, where t is the time of the exchange. Because the option matures prior to the next FOMC meeting, the IOER rate is assumed to be constant. Because the intermediary is always indifferent between buying the ETF and the risk-free bond, there is never any advantage to early exercise. Consequently, the put option \hat{p}_a has the same value as a European style put option with the same expiry and a strike KR_f^I . Moreover, because $R_f^I \geq 1$, the option \hat{p}_a is more valuable than the American put p_a . Hence, by no-arbitrage, the following inequalities hold for the American put:

$$0 \leq Q_{p_e(K)} \leq Q_{p_a(K)} \leq Q_{\hat{p}_a(K)} = Q_{p_e(KR_f^I)},$$

where $p_e(K) \in A$ is the European put of strike K . Using essentially the same argument, let \hat{c}_a be an American call option to buy the ETF in exchange for $\frac{K}{R_f^I} \exp(t \cdot \ln(R_f^I))$ dollars at time t . Early exercise is again never optimal, and hence this call's value is equal to the European call with strike K and the same expiry. Moreover, this call dominates the American call with strike K and the same expiry, and hence early exercise is never optimal and the American and European calls have the

same value. That is,

$$Q_{c_a(K)} = Q_{c_e(K)},$$

where $c_e(K) \in A$ is the European call of strike K .

Let us now consider how to replicate the ETF. If I observed European options prices, I would calculate the arbitrage, for any strike K , as

$$\chi_e = \frac{Q_{c_e(K)} - Q_{p_e(K)} + \frac{K}{R_f^I} - Q_e}{Q_{c_e(K)} - Q_{p_e(K)} + \frac{K}{R_f^I}},$$

where Q_e is the ETF price. Using the put inequalities derived above,

$$\chi_e \geq \chi_{e,\min}(K) = \frac{Q_{c_a(K)} - Q_{p_a(K)} + \frac{K}{R_f^I} - Q_e}{Q_{c_a(K)} - Q_{p_a(K)} + \frac{K}{R_f^I}},$$

and, for $K' = \frac{K}{R_f^I}$,

$$\chi_e \leq \frac{Q_{c_a(K'R_f^I)} - Q_{p_a(K')} + K' - Q_e}{Q_{c_a(K'R_f^I)} - Q_{p_a(K')} + K'}. \quad (13)$$

If R_f^I is sufficiently close to one (and one month's worth of interest is indeed quite small), these bounds will be tight.

One implementation issue that arises from these inequalities is that the strike KR_f^I is unlikely to be traded. However, by the convexity of call prices (another cashflow dominance argument),

$$Q_{c_e(KR_f^I)} \leq \alpha Q_{c_e(K_1)} + (1 - \alpha) Q_{c_e(K_2)}$$

for any $K_1 \leq KR_f^I \leq K_2$ such that $\alpha K_1 + (1 - \alpha) K_2 = KR_f^I$. Choosing K_1 and K_2 to be as close as possible to KR_f^I generates the tightest bound. If KR_f^I is greater than the maximum traded strike K_{max} , then $Q_{c_e(KR_f^I)} \leq Q_{c_e(K_{max})}$. Using these bounds along with (13) generates an implementable upper-bound, which I will call $\chi_{e,\max}(K)$. Because these bounds must hold for all K , we are free to choose the greatest lower bound and least upper bound from the set of available strikes.

Another empirical issue to consider in the implementation of this trading strategy is whether to use bids and offers or mid-prices. Bid-offers are wide in options markets, and likely substantially overstate the bid-offer associated with "delta one" trades. That is, buying a call and selling a put together likely has a much smaller bid-offer than doing those trades separately. For this reason, authors such as Van Binsbergen et al. (2022) use mid-prices. However, mid-prices can exhibit strange behavior when bid-offers are particularly wide (which is why those authors use outlier-

robust methods of analysis). To deal with this issue, I consider only strikes K with sufficiently small bid-offers. In particular, I restrict attention to values of $\chi_{e,\min}(K)$ which the difference between the mid-price and the lowest bound that can be constructed from the various bids and offers is less than 0.05% of the spot price Q_e . Under this restriction, the lower bound $\chi_{e,\min}(K)$ constructed from mid prices is at most ~5bps too high in the worse-case scenario. Similarly, I require that the difference between $\chi_{e,\max}(K)$ and the highest bound constructed from bids and offers be less than 0.05% of Q_e . From these two sets of valid strikes (one for $\chi_{e,\min}$ and one for $\chi_{e,\max}$), I choose the strikes that generate the tightest possible bounds on χ_e .

After finding the maximum and minimum bounds, $\chi_{e,\min}$ and $\chi_{e,\max}$, I define the estimated arbitrage as

$$\chi_e = \begin{cases} \chi_{e,\min} & \max(\chi_{e,\min}, \chi_{e,\max}) > \chi_{RF} \\ \chi_{e,\max} & \min(\chi_{e,\min}, \chi_{e,\max}) < \chi_{RF} \\ \chi_{RF} & \text{otherwise.} \end{cases}$$

In other words, I will assume that there is zero risk-neutral excess arbitrage ($\chi_e - \chi_{RF}$) if this is possible, and assume the minimum amount, in absolute value terms, if it is not possible.⁴⁸ In practice, my final dataset never has $\chi_e = \chi_{RF}$, because the bounds are sufficiently tight.

The dataset is a high-frequency (minute-level) dataset of options quotes purchased from the CBOE DataShop. From this dataset (which contains quotes for all minutes the exchange is open), I have extracted the five minutes on each day immediately preceding the Bloomberg London closing time. On most days, this is 12:55pm-12:59pm EST, although the EST hour moves around due the asynchronous use of daylight savings time in the US and UK.

This dataset contains SPY options of many different expiries. Because I am interested in one-month options (where one-month is defined based on FX trading conventions) that do not cross an SPY dividend date, I restrict attention to expiries between 21 and 58 days in the future. These cutoffs ensure expiries are roughly one month and using these specific cutoffs simplifies the logic of determining whether an expiry occurs after the next ex-dividend date on the SPY. I also require that each expiry have at least eleven different strikes quoted to be included in the dataset.

The result of these restrictions and calculations is a dataset containing many estimates of χ_e on each day (five minutes times the number of valid expiries). From this set, for each day I search for minute/expiry pairs with non-missing data, expiries that cross neither the next SPY ex-dividend date nor the next FOMC meeting, and that have no arbitrage violations based on the bids and offers of options prices.⁴⁹ Among the surviving minute/expiry pairs, I choose first the expiries that are closest to the FX market definition of one month, and then among those choose the minute-expiry

⁴⁸Note that, because I am using mid prices, it is possible to have $\tilde{\chi}_{e,\min} > \tilde{\chi}_{e,\max}$.

⁴⁹This last filter eliminates a few days with bad options quotes.

pair with the narrowest bid-offer for the relevant arbitrage bounds. This procedure results in a unique value for χ_e on each day.

B.4 Expectations, Variance, and Covariance with the SPDR Arbitrage

To use the SPY arbitrage described in the preceding sub-section in my exercise, I require estimates of its variance and covariance with currency returns under both the physical and risk-neutral measures, as well as an estimate of its expected return under the physical measure.

I use, as an empirical analog of the one-month return $R_{a,s}^I = (1 - \chi_a)R_{a,s}$,

$$R_{e,t} = (1 - \chi_{e,t-1}) \frac{Q_{e,t}}{Q_{e,t-1}},$$

where $Q_{e,t}$ is the spot SPY price. Note that this definition does not include dividends, because I have restricted attention to dates on which dividends will not occur over the next month.

To compute expected returns under the physical measure, I assume an equity premium of 5% (roughly the average value of Martin (2017) in recent years). Although many predictors of time-varying equity returns have been documented in the literature, over a one-month horizon most of these predictors are quite weak, and it seems reasonable to use an estimate of the unconditional equity premium. Under this assumption,

$$\mu_t^{A^*,I,P} = (1 - \chi_{e,t})(R_{f,t} - 1 + 1.05^{\Delta t}),$$

where Δt is the time (in years) to the next month under the FX market convention and $R_{f,t}$ is the US OIS rate accumulated over that month.

To compute the physical-measure variance-covariance matrix, I use a daily series of surprise log-returns,

$$r_{e,t} = \ln(Q_{e,t}) - \ln(Q_{e,t-1}) - \ln(R_{f,t} - 1 + 1.05^{\Delta t}),$$

consistent with how I construct surprise currency returns. I then use the same ‘‘Risk-Metrics’’ methodology described in appendix section B.2.

I compute the risk-neutral variance-covariance matrix using the SVIX method of Martin (2017) and data on quanto options from Markit (as in Kremens and Martin (2019)). Applying the SVIX methodology of Martin (2017) (in particular, equation (11) of that paper) to the SPY options data used to construct the arbitrage series χ_e , I compute $V_{t-1}^* \left[\frac{Q_{e,t}}{Q_{e,t-1}} \right]$, and then scale by $(1 - \chi_{e,t-1})^2$ to compute the variance.

I extract covariances from data on quanto options on the S&P 500. A quanto call option is, for example, the right to buy the S&P 500 for a fixed amount of euros at a certain date. Such

options are traded in OTC markets, and Markit provides a pricing service to help dealers that trade these options mark their books. The prices represent the (trimmed, cleaned) averages of prices submitted by participating dealers. My data set includes prices for call and put options for all of currencies used in this paper. Unfortunately, these options have a twenty-four month expiry (this is essentially the only traded expiry), and the data is monthly rather than daily. I discuss how I deal with both of these issues below. My use of the quanto options is also complicated by the presence of arbitrage (CIP violations). I deal with this issue by pricing the options under the assumptions of the framework developed in this paper, and then extracting a risk-neutral covariance from those pricing formulas.

Let $S_{j,t}$ be the spot exchange rate (e.g. euros per dollar). The dollar price of a quanto call (qc) option, as a percentage of the spot price, with a strike equal to the current spot price is

$$qc_{j,t} = \frac{\frac{1}{R_{f,t,t+24}^l} E_t^* \left[\frac{X_{j,t}}{S_{j,t+24}} \max\{Q_{e,t+24} - Q_{e,t}, 0\} \right]}{Q_{e,t}},$$

where $X_{j,t}$ is the agreed-upon fixed exchange rate for the quanto option and $R_{f,t,t+24}^l$ is the intermediaries' cumulative discount factor over the next two years. The Markit data use the convention $X_{j,t} = S_{j,t}$.

Quanto-put (qp) prices follow an analogous formula, and by put-call parity, for the strikes in my data,

$$R_{f,t,t+24}^l (qc_{j,t} - qp_{j,t}) = E_t^* \left[\frac{S_{j,t} Q_{e,t+24}}{S_{j,t+24} Q_{e,t}} \right] - E_t^* \left[\frac{S_{j,t}}{S_{j,t+24}} \right].$$

My data also includes hypothetical prices for quanto call and put options under the assumption of zero correlation between the foreign exchange rate and the S&P 500. Under this assumption, the price of the quanto call with $X_{j,t} = S_{j,t}$ is

$$zc_{j,t} = E_t^* \left[\frac{S_{j,t}}{S_{j,t+24}} \right] \times \frac{\frac{1}{R_{f,t,t+24}^l} E_t^* [\max\{Q_{e,t+24} - Q_{e,t}, 0\}]}{Q_{e,t}}$$

and hence is equal to the (inverse) forward premium multiplied by the price of the vanilla (standard) call option on the S&P 500. That is, by asking for “zero-correlation” quanto call prices, Markit is not asking dealers to price a new exotic instrument but rather to report the levels of two standard contracts along with the the price of the quanto call option.

Again by put-call parity,

$$R_{f,t,t+24}^l (zc_{j,t} - zp_{j,t}) = E_t^* \left[\frac{S_{j,t}}{S_{j,t+24}} \right] E_t^* \left[\frac{Q_{e,t+24}}{Q_{e,t}} \right] - E_t^* \left[\frac{S_{j,t}}{S_{j,t+24}} \right],$$

where $zp_{j,t}$ is the zero-correlation quanto put price.

It follows that

$$E_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \frac{F_{j,t,t+24}}{S_{j,t+24}} \right] = \frac{1 + \frac{F_{j,t,t+24}}{S_{j,t}} R_{f,t,t+24}^I (qc_{j,t} - qp_{j,t})}{1 + \frac{F_{j,t,t+24}}{S_{j,t}} R_{f,t,t+24}^I (zc_{j,t} - zp_{j,t})}, \quad (14)$$

where $F_{j,t,t+24}$ is the two-year forward price.

As mentioned previously, the quanto dataset is a monthly dataset (with a few missing observations) of 24-month expiry options. The goal of this exercise is to extract one-month horizon covariances. To that end, I assume that correlations are constant over horizon, so that I can extract a 24-month risk-neutral correlation and then assume it is equal to the one-month correlation. I will use the most recent non-missing observation for each currency. In the data, the correlations I extract move slowly over time.

Ignoring these issues for a moment, the quantity of interest is

$$\Sigma_{j,e,t}^{A^*,I,i^*} = E_t^*[R_{e,t+1}R_{j,t+1}] - E_t^*[R_{e,t+1}]E_t^*[R_{j,t+1}],$$

where $R_{j,t}$ is the intermediary currency return defined in B.1. This is

$$\begin{aligned} \Sigma_{j,e,t}^{A^*,I,i^*} &= (1 - \chi_{e,t}) R_{f,t}^I \frac{E_t^*[Q_{e,t+1}]}{Q_t} \text{Cov}_t^* \left[\frac{Q_{e,t+1}}{E_t^*[Q_{e,t+1}]}, \frac{F_{j,t}}{S_{j,t+1}} \right] \\ &= (R_{f,t}^I)^2 \text{Corr}_t^* \left[\frac{Q_{e,t+1}}{E_t^*[Q_{e,t+1}]}, \frac{F_{j,t}}{S_{j,t+1}} \right] V_t^* \left[\frac{Q_{e,t+1}}{E_t^*[Q_{e,t+1}]} \right]^{\frac{1}{2}} V_t^* \left[\frac{F_{j,t}}{S_{j,t+1}} \right]^{\frac{1}{2}}. \end{aligned} \quad (15)$$

Under the assumption that correlations are constant across horizon,

$$\begin{aligned} \text{Corr}_t^* \left[\frac{Q_{e,t+1}}{E_t^*[Q_{e,t+1}]}, \frac{F_{j,t}}{S_{j,t+1}} \right] &= \frac{\text{Cov}_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]}, \frac{F_{j,t,t+24}}{S_{j,t+24}} \right]}{V_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \right]^{\frac{1}{2}} V_t^* \left[\frac{F_{j,t,t+24}}{S_{j,t+24}} \right]^{\frac{1}{2}}} \\ &= \frac{E_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \frac{F_{j,t,t+24}}{S_{j,t+24}} \right] - 1}{V_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \right]^{\frac{1}{2}} V_t^* \left[\frac{F_{j,t,t+24}}{S_{j,t+24}} \right]^{\frac{1}{2}}}. \end{aligned} \quad (16)$$

I compute risk-neutral variances $V_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \right]$ and $V_t^* \left[\frac{F_{j,t,t+24}}{S_{j,t+24}} \right]$ from Bloomberg at-the-money 2-year at-the-money SPX implied volatilities and 2-year FX volatilities (implicitly assuming log-normality). As discussed earlier, using an SVIX-based calculation would avoid log-normality assumptions at the expense of increased data requirements, computational complexity, and uncertainty related to bid-offers and illiquidity of out-of-the-money options. Under an assumption of log-normality,

$$V_t^* \left[\frac{Q_{e,t+24}}{E_t^*[Q_{e,t+24}]} \right] = \exp(V_t^*[\ln(Q_{e,t+24})]) - 1,$$

and likewise

$$V_t^* \left[\frac{F_{j,t,t+24}}{S_{j,t+24}} \right] = \exp(V_t^* [\ln(S_{jt+24})]) - 1.$$

One last implementation concerns the intermediaries’ two-year discount factor, $R_{f,t,t+24}^I$. In the main text, I study only dates at least one month before an FOMC meeting and use the IOER rate as the one-month rate. This approach does not allow me to construct a two-year rate. For simplicity, I use instead the two-year OIS rate and then add the sample mean spread between IOER and fed funds (see Table 1). This adjustment makes almost no difference to the estimated correlations. Van Binsbergen et al. (2022) offer a better approach: extracting a two-year intermediary discount factor using “box” trades (put-call parity for different strikes). As with the SVIX methodology, this approach is theoretically superior but increases data requirements and concerns about issues related to illiquidity, bid-offer spreads, and the like.

B.5 Construction of the Dollar and Carry Portfolios

As an additional robustness exercise, I construct an externality-mimicking portfolio under the assumption that A^* includes a risk-free asset and two portfolios of currency trades, which I will refer to as dollar and carry.

These portfolios are portfolios of five developed-market currencies vs. the US dollar. The five currencies (plus the US dollar) in the portfolio are: euro, yen, pound, Australian dollar, and Canadian dollar. To select these currencies, I started with the nine non-US-dollar G10 currencies. I removed the New Zealand dollar, Swedish krona, and Norwegian Krone due to limited data availability for FX options and OIS swaps. I removed the Swiss franc both because of problems with its OIS rate (discussed in Du et al. (Forthcoming)) and because of the pegging and de-pegging events that occur during the sample period.

From these five non-USD currencies, I define the dollar and carry portfolios, in the spirit of the factor approach of Lustig et al. (2011). Dollar is an equal-weighted basket of the five currencies vs. USD. Note that this portfolio is short USD vs. these other currencies, not long. This sign convention helps make the portfolio definition consistent with the exercise in the main text. Carry is a long-short portfolio that is long the two currencies with the smallest 1m forward premium and short the two countries with the largest forward premium. For almost all of the sample, this means long AUD and CAD and short JPY and EUR. To ensure a positive price, I add some of the risk-free USD investment to the Carry portfolio. This has no effect on the excess arbitrage $\chi_{carry} - \chi_{RF}$, and hence no effect on resulting externality-mimicking portfolio.

Expected returns and variance-covariance matrices for these portfolios can be constructed from the assumed expected returns and variance-covariance matrices of the individual currencies (as described in the previous parts of this appendix section). Because some of the currencies in these

portfolios are not explicitly modeled in the stress tests, I impute returns using the same procedure used in the main text for AUD.

C General Equilibrium with Intermediaries

This appendix section more formally describes the environment outlined in section 1.

The economy has a set of future states, S_1 , and an initial state, s_0 . Let $S = S_1 \cup \{s_0\}$ denote the set of all states, and let J_s be the set of goods in each state $s \in S$. The government can transfer income between agents in the initial state, s_0 , but not any other state, and these transfers must sum to zero. The goods available in each state are denoted by the set J_s .

Households $h \in \mathcal{H}$ maximize expected utility,

$$\sum_{s \in S} U^h(\{X_{j,s}^h\}_{j \in J_s}; s),$$

where $U^h(\{X_{j,s}^h\}_{j \in J_s}; s)$ is the utility of household h in state s , inclusive of the household's rate of time preference and the probability the household places on state s . I will assume non-satiation for at least one good in each state, implying that each household places non-zero probability on each state in S_1 , and that the utility functions are differentiable. Note that I will refer to each h as "a household," and assume price-taking; nothing would change if we thought of each h as representing a mass of identical households. In each state $s \in S$, household $h \in H$ has an endowment of good $j \in J_s$ equal to $Y_{j,s}^h$. In state s_0 , the household might also receive a transfer T^h .

The set of securities available in the economy, A , has securities which offer payoffs $Z_{a,s}(\{P_{j,s}\}_{j \in J_s})$ for security $a \in A$ in state $s \in S$. Note that the payoff may be a function of goods prices, which are endogenous, and I will assume that the payoffs are homogenous of degree one in prices, so that it is without loss of generality to fix the price for one good (the numeraire) in each state. Let D_a^h denote the quantity of security a purchased or sold by household h , and let Q_a be the "ex-dividend" price at time zero (i.e. under the convention that $Z_{a,s_0} = 0$).

In state s , the household's income used for consumption (i.e. consumption expenditure) is

$$I_s^h = \begin{cases} \sum_{j \in J_s} P_{j,s} Y_{j,s}^h + \sum_{a \in A} D_a^h Z_{a,s}(\{P_{j,s}\}_{j \in J_s}) & \text{if } s \neq s_0, \\ T^h + \sum_{j \in J_s} P_{j,s} Y_{j,s}^h - \sum_{a \in A} D_a^h Q_a & \text{if } s = s_0. \end{cases} \quad (17)$$

That is, for all states except the initial state, consumption expenditure is equal to income, and is the value of the household endowment plus the payoffs of the household's asset holdings. In the initial state, income used for consumption is the value of the endowment plus any transfers, less the purchase price of the household's asset holdings. In all states, there is a budget constraint for

consumption in the state,

$$\sum_{j \in J_s} P_{j,s} X_{j,s}^h \leq I_s^h.$$

Household asset allocations are constrained by limited participation constraints or other kinds of limits. The constraints on households' asset positions are summarized by

$$\Phi^h(\{D_a^h\}_{a \in A}) \leq \vec{0}, \quad (18)$$

where Φ^h is a vector-valued function, convex in D^h . Limited participation is the key form of constraint I am attempting to capture with these Φ functions, but it is not the only kind of constraint that fits into this framework.

Having defined expenditure and prices, I define the standard indirect utility function in each state,

$$V^h(I_s^h, \{P_{j,s}\}_{j \in J_s}; s) = \max_{\{X_{j,s}^h \in \mathbb{R}^+\}_{j \in J_s}} U^h(\{X_{j,s}^h\}_{j \in J_s}; s)$$

subject to

$$\sum_{j \in J_s} P_{j,s} X_{j,s}^h \leq I_s^h.$$

Using these indirect utility functions, we can write the portfolio choice problem as

$$\max_{\{D_a^h \in \mathbb{R}\}_{a \in A}, s \in S} \sum_{s \in S} V^h(I_s^h, \{P_{j,s}\}_{j \in J_s}; s)$$

subject to the budget constraints that define income (equation (17)) and the constraints on asset allocation (equation (18)). Note that I fold the household's discounting and subjective probability assessments into the state-dependent direct and indirect utility functions. In a competitive equilibrium (that is, taking asset prices Q and goods prices P as given, defined below), this is the problem households solve when choosing their asset allocation.

I will call the other type of agents in the economy intermediaries, and use $i \in \mathcal{I}$ to denote a particular intermediary. Intermediaries are like households (in the sense that all of the notation above applies, with some $i \in I$ in the place of an $h \in \mathcal{H}$), except that they face different constraints on their portfolio choices. In particular, households are constrained to trade only with intermediaries, but intermediaries can trade with both households and other intermediaries.

The constraint that households can trade only with intermediaries, but not each other, can be implemented using this notation in the following way. The set of assets, A , is a superset of the union of disjoint sets $\{A^h\}_{h \in \mathcal{H}}$, denoting trades with household h . For a given household h , the function $\Phi^{h,exog}$ implements the requirement that, for all $a \in A \setminus A^h$, $D_a^h = 0$. To be precise, if

$a \in A \setminus A^h$ and $D_a^h \neq 0$, then there exists an element of $\Phi^h(D^h)$ strictly greater than zero. The set of assets also includes assets that cannot be traded by any household. Define $A^I = A \setminus (\cup_{h \in H} A^h)$ as the set of securities tradable only by intermediaries. The exogenous constraints on intermediary trading are denoted $\Phi^{i,exog}$.

Regulation, in this framework, are additional convex functions $\Phi^{i,reg}(\{D_a^i\}_{a \in A})$ and $\Phi^{h,reg}(\{D_a^h\}_{a \in A})$ such that the intermediaries face the constraint $\Phi^i(\cdot) = \begin{bmatrix} \Phi^{i,reg} & \Phi^{i,exog} \end{bmatrix}$, and likewise for households. As discussed in the text, for notational simplicity I assume that these functions depend only on asset quantities and not on asset prices. Because the equilibrium gradient on these constraints is the only quantity that matters for equilibrium, this assumption is without loss of generality.

The notion of equilibrium, given the constraints Φ^i and Φ^h , is standard.

Definition 2. An equilibrium is a collection of consumptions $X_{j,s}^h$ and $X_{j,s}^i$, goods prices $P_{j,s}$, asset positions D_a^h and D_a^i , transfers T^h and T^i , and asset prices Q_a such that:

1. Households and intermediaries maximize their utility over consumption and asset positions, given goods prices and asset prices, respecting the constraints that consumption be weakly positive and the constraints on their asset positions,
2. Goods markets clear: for all $s \in S$ and $j \in J_s$,

$$\sum_{h \in \mathcal{H}} (X_{j,s}^h - Y_{j,s}^h) = \sum_{i \in \mathcal{I}} (X_{j,s}^i - Y_{j,s}^i),$$

3. Asset markets clear: for all $a \in A$,

$$\sum_{h \in \mathcal{H}} D_a^h + \sum_{i \in \mathcal{I}} D_a^i = 0, \tag{19}$$

4. The government's budget constraint balances,

$$\sum_{h \in \mathcal{H}} T^h + \sum_{i \in \mathcal{I}} T^i = 0. \tag{20}$$

The definition of equilibrium presumes price-taking by households and intermediaries. Absent government constraints, each household h can trade with every intermediary, and the price of asset $a \in A^h$ will be pinned down by competition between intermediaries. The equilibrium definition supposes that this will continue to be the case, even if the government places asymmetric constraints on intermediaries— for example, by granting a single intermediary a monopoly over trades with a particular household. In this case, it is as if the household had all of the bargaining power. Such a policy is unlikely to optimal, and will never be the unique optimum.

I next describe a planner's problem for this economy. I assume that the planner is unable to redistribute resources ex-post (doing so would allow the planner to circumvent limited participation).

Instead, in the spirit of Geanakoplos and Polemarchakis (1986), I will allow the planner to trade in asset markets on behalf of agents, trading for each agent only in markets the agent can participate in, to maximize a weighted sum of the household's indirect utility functions, subject to an ex-ante participation constraint for intermediaries. The planner is required respect the exogenous portfolio constraints, $\Phi^{h,exog}$ and $\Phi^{i,exog}$.

Definition 3. The constrained planner's problem is

$$\max_{\{D_a^h \in \mathbb{R}\}_{a \in A, h \in \mathcal{H}}, \{D_a^i \in \mathbb{R}\}_{a \in A, i \in \mathcal{I}}, \{P_{j,s} \in \mathbb{R}_+\}_{s \in S, j \in J_s}, \{T^i \in \mathbb{R}\}_{i \in \mathcal{I}}, \{T^h \in \mathbb{R}\}_{h \in \mathcal{H}}} \sum_{h \in \mathcal{H}} \lambda^h \sum_{s \in S} V^h(I_s^h, \{P_{j,s}\}_{j \in J_s}; s),$$

subject to the intermediaries' ex-ante participation constraint,

$$\sum_{s \in S} V^i(I_s^i, \{P_{j,s}\}_{j \in J_s}; s) \geq \bar{V}^i, \forall i \in \mathcal{I},$$

household's exogenous portfolio constraints,

$$\Phi^{h,exog}(\{D_a^h\}_{a \in A}) \leq \vec{0}, \forall h \in \mathcal{H},$$

intermediaries exogenous portfolio constraints,

$$\Phi^{i,exog}(\{D_a^i\}_{a \in A}) \leq \vec{0}, \forall i \in \mathcal{I},$$

the definition of incomes I_s^h and I_s^i (17), market clearing in assets (19), the government's budget constraint (20), and goods market clearing for each state $s \in S$ and good $j \in J_s$,

$$\sum_{h \in \mathcal{H}} (X_{j,s}^h(I_s^h, \{P_{j',s}\}_{j' \in J_s}) - Y_{j,s}^h) = \sum_{i \in \mathcal{I}} (X_{j,s}^i(I_s^i, \{P_{j',s}\}_{j' \in J_s}) - Y_{j,s}^i).$$

Here, $X_{j,s}^h(I_s^h, \{P_{j',s}\}_{j' \in J_s})$ denotes the demand function for good j by agent h in state s . I have chosen to write the planner's problem as maximizing household utility subject to an ex-ante participation constraint for intermediaries, because this fits best into the example from the text; nothing would change if instead the planner maximized a weighted combination of all agents' utilities. Lastly, I define constrained (in)efficiency:

Definition 4. A competitive equilibrium is constrained efficient if there exists Pareto weights λ^h and outside options \bar{V}^i such that the allocation of assets and goods in the competitive equilibrium

coincides with the solution to the planner's problem. Otherwise, the competitive equilibrium is constrained inefficient.

D Discussion of Alternative Sources of Inefficiency

In this appendix section, I will discuss informally how certain unmodeled sources of inefficiency might affect the interpretation of my results. I also discuss how the results should be interpreted if the purpose of macro-prudential policy is redistributive (i.e. in the absence of the assumption that transfers exist in state s_0).

Inefficiency in the states S_1 Suppose for the sake of arguments that agents have private information in some state $s \in S_1$, or that they can take hidden actions with externalities (i.e. a moral hazard problem exists) in $s \in S_1$, or even more generally that, for some reason, the welfare theorem fails to hold in $s \in S_1$. More formally, conditional on the incomes $\{I_s^i\}_{i \in \mathcal{I}}$, $\{I_s^h\}_{h \in \mathcal{H}}$, suppose any market-clearing prices $\{\bar{P}_{j,s}\}_{j \in J_s}$ are generically not the solution to an ex-post planner's problem. That, let $\{P_{j,s}^*\}_{j \in J_s}$ be a solution to

$$\max_{\{P_{j,s}\}_{j \in J_s}} \sum_{h \in \mathcal{H}} \lambda_s^h V^h(I_s^h, \{P_{j,s}\}; s) + \sum_{i \in \mathcal{I}} \lambda_s^i V^i(I_s^i, \{P_{j,s}\}; s)$$

subject to the constraint that the total endowment value equal total income,

$$\sum_{h \in \mathcal{H}} \sum_{j \in J_s} (P_{j,s} Y_{j,s}^h - I_s^h) = \sum_{i \in \mathcal{I}} \sum_{j \in J_s} (P_{j,s} Y_{j,s}^i - I_s^i).$$

If no $P_{j,s}^*$ satisfies market clearing, then a pecuniary externality will generically exist at the margin at any market-clearing prices $\{\bar{P}_{j,s}\}_{j \in J_s}$. That is, perturbations to prices in the neighborhood of $\{\bar{P}_{j,s}\}_{j \in J_s}$ that leave total income unchanged will generically affect welfare.⁵⁰ Consequently, this case is essentially isomorphic to the case of incomplete markets, a point due to Greenwald and Stiglitz (1986).

Summarizing, if private information, moral hazard, or any other form of inefficiency occurring the states $s \in S_1$ cannot be corrected by ex-post policies, it will give rise to pecuniary externalities of the form considered in the main body of the text. Such externalities can justify regulatory interventions, those interventions may create arbitrage opportunities, and the main part of my analysis follows unchanged. Of course, if the regulator can correct whatever problems arise from

⁵⁰This is true generically, and whenever the ex-post planner's problem is concave, but the possibility of local maxima, saddle points, or minima in prices means that it is not true for all possible parameter configurations.

private information or moral hazard ex-post, the regulator should do that; macro-prudential (ex-ante) policies are justified only when such ex-post corrections are imperfect.

Private Information or Hidden Trade in the state s_0 Private information or hidden trade in the initial state will limit the set of asset allocations the planner can implement.

Hidden trade in a particular asset is a straightforward case. Suppose an asset can be traded between intermediaries, but such trades cannot be observed by the planner. This in fact has no direct bearing on my results, which focus on trade between households and intermediaries. It can have an indirect (general equilibrium) effect, but the first-order conditions of the planner for the assets the planner can regulate must still hold. Consequently, all of the results in the main text continue to apply. Essentially the same point applies if intermediaries and households can secretly trade with each other; the first-order conditions of the planner with respect to the secretly-traded assets will not apply, but the first-order conditions of the planner with respect to other assets will still apply. Hidden trade can have consequences that affect equilibrium outcomes and the solutions to the planner's problem, but these will be captured in the endogenous variables (the externalities and asset prices) that enter the key formulas of the main text.

Private information can be accommodated within the framework using a virtual utility approach (see, e.g., Myerson (2007)). This involves, essentially, substituting the Lagrange multipliers associated with the incentive compatibility constraint into the objective. To keep things simple, I will consider private information for one type of household, $\hat{h} \in \mathcal{H}$; the extension to all households and intermediaries is straightforward. Suppose a continuum of these households draw private types $t \in \mathcal{T}$ from a distribution $\pi^{\mathcal{T}}$, and consider an incentive-compatible direct revelation mechanism version of the planner's problem in which this household reports a type $t' \in \mathcal{T}$. Let $\alpha(t'|t)$ be the multiplier on the incentive compatibility constraint.

The virtual utility version of the planner's objective is

$$\begin{aligned} & \sum_{h \in \mathcal{H}} \lambda^h \sum_{s \in \mathcal{S}} V^h(I_s^h, \{P_{j,s}\}_{j \in J_s}; s) + \\ & \sum_{t \in \mathcal{T}} (\pi^{\mathcal{T}}(t) \lambda^{\hat{h},t} - \sum_{t' \in \mathcal{T}} \alpha(t'|t)) \sum_{s \in \mathcal{S}} V^{\hat{h},t}(I_s^{\hat{h},t}, \{P_{j,s}\}_{j \in J_s}; s) + \\ & \sum_{t' \in \mathcal{T}} \alpha(t'|t) \sum_{s \in \mathcal{S}} V^{\hat{h},t'}(I_s^{\hat{h},t'}, \{P_{j,s}\}_{j \in J_s}; s). \end{aligned}$$

The “virtual SDF” that appears in the planner's first-order condition with respect to a transfer of an asset from i^* to (\hat{h}, t) (where t is the reported type) is now

$$\pi_s^t \tilde{M}_s^{\hat{h},t,r,*} = \frac{\pi^{\mathcal{T}}(t) \lambda^{\hat{h},t} V_I^{\hat{h},t}(I_s^{\hat{h},t,*}, \{P_{j,s}^*\}_{j \in J_s}; s) + \sum_{t' \in \mathcal{T} \setminus \{t\}} \alpha(t'|t) V_I^{\hat{h},t'}(I_s^{\hat{h},t,*}, \{P_{j,s}^*\}_{j \in J_s}; s)}{\pi^{\mathcal{T}}(t) \lambda^{\hat{h},t} V_I^{\hat{h},t}(I_{s_0}^{\hat{h},t,*}, \{P_{j,s_0}^*\}_{j \in J_{s_0}}; s_0) + \sum_{t' \in \mathcal{T} \setminus \{t\}} \alpha(t'|t) V_I^{\hat{h},t'}(I_{s_0}^{\hat{h},t,*}, \{P_{j,s_0}^*\}_{j \in J_{s_0}}; s_0)}.$$

That is, the case of private information in state s_0 is equivalent to the case of full information with a different collection of SDFs for the agents.

Define the “private information externalities” as

$$\tilde{\Delta}_s^{\hat{h},t,i,r} = \tilde{M}_s^{\hat{h},t,r,*} - \frac{V_I^{\hat{h},t}(I_s^{\hat{h},t,*}, \{P_{j,s}^*\}_{j \in J_s}; s)}{V_I^{\hat{h},t}(I_{s_0}^{\hat{h},t,*}, \{P_{j,s_0}^*\}_{j \in J_{s_0}}; s_0)},$$

observing that if the incentive compatibility constraints do not bind, these externalities are zero.

Regulation is more complex in this case; the planner must regulate the trades of households. However, as with intermediaries, it is without loss of generality to assume that one type t^* of the households in \hat{h} can freely trade all of the assets (to determine prices), as it is sufficient to constrain all of that household’s counterparties instead. Under this assumption, the externalities-as-arbitrage analysis of the main text applies to the (\hat{h}, t^*) household and i^* intermediary. In this case that the externalities recovered are the sum of the pecuniary externalities $\Delta_s^{\hat{h},t,i,r}$ and the private information externalities $\tilde{\Delta}_s^{\hat{h},t,i,r}$.

Regulation as Redistribution Lastly, I will discuss issues related to transfers and redistribution in the planner’s problem.

Because I have assumed all assets trade ex-dividend, in the absence of transfers, the planner has no ability at all to change incomes in state s_0 , putting the planner at a disadvantage relative to the agents (who can exchange assets for the state s_0 numeraire). In this case, the planner’s problem cannot serve as a definition of constrained efficiency.

To consider the planner’s problem in the absence of transfers, we must break the assumption that all assets trade ex-dividend. Let us instead suppose some asset trades both ex- and cum-dividend. This is equivalent to assuming an “asset” equivalent to an Arrow security in state s_0 exists. Since the planner controls the allocation of this asset, it is as-if the planner could make transfers. But if an Arrow security for the state s_0 cannot be synthesized, we are back to the assumption that the agents can accomplish something that the planner cannot. The broader point here is that it makes little sense to consider a planner who can is free to reallocate assets across agents but for some reason cannot transfer income at the same time.

Restrictions against income transfers by the planner make more sense in the context of a regulator who cannot freely assign assets to agents either. It is certainly plausible that regulators such as the Federal Reserve can implement various forms of capital and liquidity requirements but are not able to simply reallocate assets or income across agents. However, it is also not clear what the redistributive motives of the Fed are; in particular, the Fed may feel constrained to avoid “picking winners.” In this context, the transfers in state s_0 can be seen as a device to capture the idea that the

Fed wishes to move as close as possible to the Pareto frontier without regard for how this impacts the distribution of wealth across agents.

Let us set this issue aside, and suppose that the Fed is both limited in terms of what it can accomplish and animated by a desire to redistribute wealth across agents. In this case, the Fed might wish to use regulations to manipulate the terms of trade between agents. Put another way, regulatory policy must balance macro-prudential and redistributive objectives. The redistributive interpretation of the relationship between the sign of CIP violations and the carry trade is that the Fed is attempting to redistribute money from the agents seeking to do carry trade to banks. Whether this interpretation is more plausible than the “accidental by-product of leverage restrictions” hypothesis advanced in the main text is beyond the scope my analysis.

E Alternative Definition of the Externality-Mimicking Portfolio

In this section, I provide a definition of the externality-mimicking portfolio as a portfolio of arbitrage-able assets (as opposed to a portfolio of replicating portfolios). Defining

$$\tilde{\theta}^{A^*,r} = \tilde{\theta}(\theta^{A^*,r}),$$

where $\theta^{A^*,r}$ is the externality-mimicking portfolio of definition 1, and

$$\tilde{\chi}_a = \frac{-Q_a + \sum_{a' \in A'} w_{a'}(a) Q_{a'}}{Q_a},$$

we can see that the following defines a portfolio of arbitrage-able assets with payoffs identical to those of $\theta^{A^*,r}$.

Definition 5. The alternative externality-mimicking portfolio is a portfolio arbitrage-able assets in A^* , with weights on the risky assets equal to

$$\tilde{\theta}^{A^*,r} = (\Sigma^{A^*,r})^{-1} \left(\tilde{\chi}^{A^*} - \tilde{\chi}_f \frac{\mathbf{v}^{A^*,r}}{R_f} \right), \quad (21)$$

and a weight on the risk-free asset equal to

$$\tilde{\theta}_f^{A^*,r} = -(\tilde{\theta}^{A^*,r})^T \frac{\mathbf{v}^{A^*,r}}{R_f} + \frac{1}{(R_f)^2} \tilde{\chi}_f. \quad (22)$$

This definition is identical to definition 1, except that the arbitrage is now normalized by

the asset price (as opposed to the replicating portfolio price) and the expected returns, variance-covariance matrix, and risk-free rates are for the arbitrage-able assets as opposed to the replicating portfolios.

F Proofs

F.1 Proof of Proposition 1

For a formal definition of the planner's problem discussed in this proof, see appendix section C.

Consider a perturbation to the solution of the planner's problem in which the planner re-allocates an asset $a \in A$ from agent i to agent h , or vice versa. If such perturbations are feasible, we must have (by differentiability)

$$-\lambda^{h_0} V_{I,s_0}^{h_0} \sum_{s \in S} \sum_{j \in J_s} \mu_{j,s} [X_{I,j,s}^h - X_{I,j,s}^i] Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}) = \sum_{s \in S} (\lambda^i V_{I,s}^i - \lambda^h V_{I,s}^h) Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}),$$

where λ^i is the multiplier on intermediary i 's participation constraint and $\mu_{j,s} \lambda^{h_0} V_{I,s_0}^{h_0}$ is the multiplier on the goods-market clearing constraint. Note that the derivatives $X_{I,j,s}^h = \frac{\partial}{\partial I} X_{j,s}^h(I, \{P_{j,s}^*\}_{j \in J_s})|_{I=I^{h^*}}$ and $V_{I,s}^h = \frac{\partial}{\partial I} V_s^h(I, \{P_{j,s}^*\}_{j \in J_s})|_{I=I^{h^*}}$ are evaluated at the solution to the planner's problem. Note also that I have normalized the multiplier to units of dollars, rather than social utility, using the marginal utility of an arbitrary agent h_0 .

By the definition of the wedges, this is

$$\lambda^{h_0} V_{I,s_0}^{h_0} \sum_{s \in S} \sum_{j \in J_s} (\pi_s^r \tau_{j,s}^r - \frac{1}{|J_s|} \sum_{j' \in J_s} \frac{\mu_{j',s}}{P_{j',s}^*}) P_{j,s}^* [X_{I,j,s}^h - X_{I,j,s}^i] Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}) = \sum_{s \in S} (\lambda^i V_{I,s}^i - \lambda^h V_{I,s}^h) Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}).$$

By non-satiation, the identity $\sum_{j \in J_s} X_{I,j,s}^h P_{j,s}^* = 1$ holds, and this simplifies using the definition of the externalities to

$$\lambda^{h_0} V_{I,s_0}^{h_0} \sum_{s \in S} \pi_s^r \Delta_s^{h,i,r} Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}) = \sum_{s \in S} (\lambda^i V_{I,s}^i - \lambda^h V_{I,s}^h) Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}).$$

By the first-order condition for the transfer between h_0 and h ,

$$\lambda^{h_0} V_{I,s_0}^{h_0} = \lambda^h V_{I,s_0}^h$$

and likewise $\lambda^{h_0} V_{I,s_0}^{h_0} = \lambda^i V_{I,s_0}^i$, and therefore

$$\begin{aligned} \sum_{s \in S} \pi_s^r \Delta_s^{h,i,r} Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}) &= \sum_{s \in S} \left(\frac{V_{I,s}^i}{V_{s_0}^i} - \frac{V_{I,s}^h}{V_{I,s_0}^h} \right) Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}) \\ &= \sum_{s \in S} \pi_s^r (M_s^{i,r} - M_s^{h,r}) Z_{a,s}(\{P_{j,s}^*\}_{j \in J_s}), \end{aligned}$$

where

$$M_s^{h,r} = \frac{V_{I,s}^h}{\pi_s^r V_{I,s_0}^h}.$$

F.2 Proof of Proposition 2

Observe first that, via quantity constraints on portfolio choices, the planner can dictate the asset allocations of all agents. So any feasible allocation (including any optimal allocations) can be implemented via portfolio constraints.

Moreover, for each asset, the planner can implement the allocation without regulating one agent's trade in that asset. If the planner regulates the quantities for all other agents who can trade that asset, market clearing will ensure the unregulated agent holds the desired asset allocation. In this case, the asset price will be determined by the valuation of the unregulated agent.

Applying this logic to all assets $a \in A^h$ for some $h \in \mathcal{H}$, the planner can implement any desired allocation without regulating the household h . All other households cannot trade that asset, and hence do not need to be regulated either. Generically, the planner will need to regulate the trades of all intermediaries in that asset.

Applying the same logic to the assets in A^I , which no household can trade, it is without loss of generality to designate one intermediary, i^* , as unregulated for all the assets in A^I . Note, as discussed in the text, that this in fact requires only that the regulations do not bind for i^* (which is to say i^* is not necessarily "unregulated.")

In such an implementation, equations (1) and (2) hold. The result immediately follows by proposition 1 and those equations.

F.3 Proof of Proposition 3

The first two of these claims are simply the definition of the least-squares projection. That is

$$\begin{aligned}
\left(\sum_{s \in S} \pi_s^r R_s^I (R_s^I)^T\right)^{-1} \left(\sum_{s \in S} \pi_s^r R_s^I \Delta_s^{h,i,r}\right) &= \left(\sum_{s \in S} \pi_s^r R_s^I (R_s^I)^T\right)^{-1} \begin{bmatrix} \chi_a \\ \chi_f \end{bmatrix} \\
&= \begin{bmatrix} (\Sigma^{A^*,I,r} + \mathbf{v}^{A^*,I,r} (\mathbf{v}^{A^*,I,r})^T) & R_f^I \mu^{A^*,I,r} \\ R_f^I (\mathbf{v}^{A^*,I,r})^T & (R_f^I)^2 \end{bmatrix}^{-1} \begin{bmatrix} \chi_a \\ \chi_f \end{bmatrix} \\
&= \begin{bmatrix} (\Sigma^{A^*,I,r})^{-1} & -(R_f^I)^{-1} (\Sigma^{A^*,I,r})^{-1} \mathbf{v}^{A^*,I,r} \\ -(R_f^I)^{-1} (\mathbf{v}^{A^*,I,r})^T (\Sigma^{A^*,I,r})^{-1} & (R_f^I)^{-2} (1 + (\mathbf{v}^{A^*,I,r})^T (\Sigma^{A^*,I,r})^{-1} \mathbf{v}^{A^*,I,r}) \end{bmatrix} \begin{bmatrix} \chi_a \\ \chi_f \end{bmatrix} \\
&= \begin{bmatrix} \theta^{A^*,r} \\ \theta_f^{A^*,r} \end{bmatrix}.
\end{aligned}$$

where R_s^I is the vector form of $\{R_{a,s}^I\}_{a \in A^*}$. The mean externalities, by construction, are

$$\sum_{s \in S} \pi_s^r \Delta_s^{h,i,r} = \frac{\chi_f}{R_f^I},$$

and the claim about the variance follows.

I prove the third claim below. Define the Lagrangian, using $R_{a,s}(1 - \chi_a) = R_{a,s}^I$, as

$$\begin{aligned}
\min_{m \in \mathbb{R}^{|S_1|}} \max_{\theta \in \mathbb{R}^{|A^*|}} \frac{1}{2} \sum_{s \in S_1} \pi_s^r (m_s - m_s^{I,r})^2 - \\
\sum_{a \in A^*} \theta_a [(1 - \chi_a) - \sum_{s \in S_1} \pi_s^r m_s R_{a,s}^I].
\end{aligned}$$

The FOC for m is

$$\pi_s^r (m_s - m_s^{I,r}) + \pi_s^r \sum_{a \in A^*} \theta_a R_{a,s}^I = 0.$$

Plugging this back into the problem,

$$\begin{aligned}
\max_{\theta \in \mathbb{R}^{|A^*|}} \frac{1}{2} \sum_{s \in S_1} \pi_s^r \left(\sum_{a \in A^*} \theta_a R_{a,s}^I \right)^2 - \\
\sum_{a \in A^*} \theta_a \left[\sum_{s \in S_1} \pi_s^r ((1 - \chi_a) - m_s^{I,r} R_{a,s}^I + R_{a,s}^I \sum_{a' \in A^*} \theta_{a'} R_{a',s}^I) \right],
\end{aligned}$$

which simplifies, using the assumption that $m_s^{I,r}$ prices the replicating portfolios ($\sum_{s \in S_1} \pi_s^r m_s^{I,r} R_{a',s}^I =$

1), to

$$\max_{\theta \in \mathbb{R}^{|A^*|}} \sum_{a \in A^*} \theta_a \tilde{\chi}_a - \frac{1}{2} \sum_{s \in \mathcal{S}_1} \pi_s^r \left(\sum_{a \in A^*} \theta_a R_{a,s}^I \right)^2. \quad (23)$$

It follows immediately that θ^* , the solution to this problem, is the projection of χ_a onto the space of returns, and hence is the externality-mimicking portfolio. Observe, by construction, that the mean return of the externality-mimicking portfolio is $\frac{\chi_f}{R_f^I} = \left(\frac{1}{R_f^I} - \frac{1}{R_f} \right)$. Therefore,

$$m_s^r = m_s^{I,r} - \sum_{a \in A^*} \theta_a^{A^*,r} R_{a,s}^I$$

is the household SDF that minimizes the variance of the difference between SDFs subject to the constraint that the SDFs are consistent with the observed arbitrages.

Lastly, consider the fourth claim: the externality-mimicking portfolio maximizes the Sharpe ratio due to arbitrage,

$$\begin{aligned} \hat{S}^{A^*,I,r}(\theta) &= S^{A^*,r}(\tilde{\theta}(\theta)) - S^{A^*,I,r}(\theta) \\ &= \frac{\tilde{\theta}(\theta)^T v^{A^*,r}}{R_f} - \sum_{a \in A^*} \tilde{\theta}_a(\theta) - \frac{\theta^T v^{A^*,I,r}}{R_f^I} - \sum_{a \in A^*} \theta_a \\ &= \frac{(\tilde{\theta}(\theta)^T \Sigma^{A^*,I,r} \tilde{\theta}(\theta))^{\frac{1}{2}}}{(\theta^T \Sigma^{A^*,I,r} \theta)^{\frac{1}{2}}} - \frac{\sum_{a \in A^*} \theta_a (\chi_a - \chi_f \frac{v^{A^*,I,r}}{R_f^I})}{(\theta^T \Sigma^{A^*,I,r} \theta)^{\frac{1}{2}}}. \end{aligned}$$

Suppose not; let $\hat{\theta}$ be some portfolio with a higher ratio. Note that the Sharpe ratio due to arbitrage is homogenous of degree zero. Moreover, mixing in some amount of the risk-free rate does not change this ratio,

$$\frac{\sum_{a \in A^*} \hat{\theta}_a (\chi_a - \frac{v^{A^*,I,r}}{R_f^I} \chi_f)}{(\theta^T \Sigma^{A^*,I,r} \theta)^{\frac{1}{2}}} = \frac{\sum_{a \in A^*} (\hat{\theta}_a + \mathbf{1}(a=f)) (\chi_a - \frac{v^{A^*,I,r}}{R_f^I} \chi_f)}{((\hat{\theta}_a + \mathbf{1}(a=f))^T \Sigma^{A^*,I,r} (\hat{\theta}_a + \mathbf{1}(a=f)))^{\frac{1}{2}}}.$$

It is therefore without loss of generality to suppose that

$$\sum_{a \in A^*} \hat{\theta}_a v^{A^*,I,r} = \sum_{a \in A^*} \theta_a v^{A^*,I,r}$$

and

$$\sum_{a \in A^*} \hat{\theta}_a \chi_a = \sum_{a \in A^*} \theta_a \chi_a.$$

But in this case, $\hat{\theta}$ must achieve a higher payoff than the externality-mimicking portfolio in (23), a contradiction.

Internet Appendix References

- Carol Alexander. Moving average models for volatility and correlation, and covariance matrices. *Handbook of finance*, 2008.
- Y Peter Chung. A transactions data test of stock index futures market efficiency and index arbitrage profitability. *The Journal of Finance*, 46(5):1791–1809, 1991.
- Wenxin Du, Benjamin M Hébert, and Amy Wang Huber. Are intermediary constraints priced? *The Review of Financial Studies*, Forthcoming.
- Mark B Garman and Steven W Kohlhagen. Foreign currency option values. *Journal of international Money and Finance*, 2(3):231–237, 1983.
- John Geanakoplos and Heraklis M Polemarchakis. Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. *Uncertainty, information and communication: essays in honor of KJ Arrow*, 3:65–96, 1986.
- Eric Ghysels, Pedro Santa-Clara, and Rossen Valkanov. Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131(1-2):59–95, 2006.
- Bruce C Greenwald and Joseph E Stiglitz. Externalities in economies with imperfect information and incomplete markets. *The quarterly journal of economics*, 101(2):229–264, 1986.
- Lukas Kremens and Ian Martin. The quanto theory of exchange rates. *American Economic Review*, 109(3):810–43, 2019.
- Hanno Lustig, Nikolai Roussanov, and Adrien Verdelhan. Common risk factors in currency markets. *The Review of Financial Studies*, 24(11):3731–3777, 2011.
- A Craig MacKinlay and Krishna Ramaswamy. Index-futures arbitrage and the behavior of stock index futures prices. *The Review of Financial Studies*, 1(2):137–158, 1988.
- William Margrabe. The value of an option to exchange one asset for another. *The journal of finance*, 33(1):177–186, 1978.
- Ian Martin. What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1):367–433, 2017.
- Merton H Miller, Jayaram Muthuswamy, and Robert E Whaley. Mean reversion of standard & poor’s 500 index basis changes: Arbitrage-induced or statistical illusion? *The Journal of Finance*, 49(2):479–513, 1994.
- Roger B Myerson. Virtual utility and the core for games with incomplete information. *Journal of Economic Theory*, 136(1):260–285, 2007.
- Jules Van Binsbergen, Michael Brandt, and Ralph Koijen. On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618, 2012.
- Jules H Van Binsbergen, William F Diamond, and Marco Grotteria. Risk-free interest rates. *Journal of Financial Economics*, 143(1):1–29, 2022.