The Review of Financial Studies

# Are Intermediary Constraints Priced?

**Wenxin Du**
University of Chicago, Federal Reserve Bank of New York, and NBER, USA

**Benjamin Hébert**
Stanford University and NBER, USA

**Amy Wang Huber**
Stanford University, USA

Violations of no-arbitrage conditions measure the shadow cost of intermediary constraints. Intermediary asset pricing and intertemporal hedging together imply that the risk of these constraints tightening is priced. We describe a "forward CIP trading strategy" that bets on CIP violations shrinking and show that its returns help identify the price of this risk. This strategy yields the highest returns for currency pairs associated with the carry trade. The strategy's risk substantially contributes to the volatility of the stochastic discount factor, is correlated with both other near-arbitrages and intermediary wealth measures, and appears to be consistently priced across various asset classes. (*JEL* F31, G12, G15)

Received March 12, 2021; editorial decision March 20, 2022 by Editor Stefano Giglio. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Covered interest rate parity (CIP) violations following the Global Financial Crisis (GFC) have been interpreted as a sign that intermediaries are constrained

---

We thank Nina Boyarchenko, John Campbell, David Chapman, Thummim Cho, Patrick Dennis, Darrell Duffie, Federico Gavazzoni, Pierre-Olivier Gourinchas, Valentin Haddad, Tarek Hassan, Francis Longstaff, Hanno Lustig, Matteo Maggiori, Arvind Krishnamurthy, Peter Kondor, Gordon Liao, Stavros Panageas, Hélène Rey, Ken Singleton, Valeri Sokolovski, Andreas Stathopoulos, Christian Wagner, Jonathan Wallen, Dimitri Vayanos, and Vish Viswanathan for helpful comments. We would like to particularly thank Svetlana Bryzgalova and Isaiah Andrews for help understanding issues related to weak identification. We also would like to thank seminar and conference participants at the American Finance Association Annual Meeting, Arrowstreet Capital, Bank of Canada-SF Fed Conference on Fixed Income Markets, Canadian Derivatives Institute Annual Conference, JHU Carey Finance Conference, LSE Paul Woolley Centre Annual Conference, Macro Finance Society, NBER IFM Fall Meeting, NBER AP Fall Meeting, NBER Summer Institute (IAP), Northern Finance Association Annual Meeting, Stanford Institute for Theoretical Economics, UVA McIntire, the Vienna Symposium on FX Markets, and Virtual Derivatives Seminar. We would like to thank Christian Gonzales Rojas, Sylvia Klosin, Haviland Sheldahl-Thomason, and Lulu Wang for outstanding research assistance. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of New York or any other person associated with the Federal Reserve System. All remaining errors are our own. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Wenxin Du, wenxin.du@chicagobooth.edu.

*The Review of Financial Studies* 36 (2023) 1464–1507

https://doi.org/10.1093/rfs/hhac050 Advance Access publication August 9, 2022

(e.g., Du, Tepper, and Verdelhan 2018; Avdjiev et al. 2019; Fleckenstein and Longstaff 2018; Hébert 2018). The intermediary asset pricing literature argues that constraints on intermediaries have important implications for asset prices (for a survey, see Kondor and Vayanos 2019 or He and Krishnamurthy 2017). In this paper, we use CIP violations to measure the extent to which constraints bind, and provide direct evidence that the risk of constraints becoming tighter is priced. Our results offer novel evidence in support of intermediary-based asset pricing.

Following the GFC, new regulations (e.g., the Basel III leverage ratio rule and the U.S. supplementary leverage ratio) were introduced that require banks to maintain a minimum capital ratio against *all* assets, *regardless* of their risk characteristics. These leverage ratio constraints have been one of the most binding constraints facing large global banks post-GFC (Duffie 2017). As a result, low-margin, balance-sheet-intensive, risk-free arbitrage conditions, such as CIP, can fail to hold. We use the term "balance sheet cost" to refer to the shadow cost associated with these constraints.[1]

The existence of liquid foreign exchange (FX) and interest rate derivatives across granular maturities allows us to directly measure innovations to the shadow cost of the relevant constraint from the term structure of CIP deviations. In particular, the difference between the ex post realized short-term CIP deviation and the ex ante forward-implied short-term CIP deviation is a measure of the "shock" to the shadow cost of the constraint.

We begin by considering a standard intermediary asset pricing model, augmented with a regulatory constraint. Because this model features arbitrage (CIP violations), not a single stochastic discount factor (SDF) prices all assets. Nevertheless, every asset can be priced using an SDF (not the same for each asset) that is a function of the return on intermediary wealth and the magnitude of a cross-currency basis (i.e., a CIP violation). These SDFs differ only in terms of their mean, and that mean is a function of the assets' weight in the relevant regulatory constraint. Consequently, shocks to the cross-currency basis are innovations to the SDFs that price each asset. These shocks can be measured by a proposed "forward CIP trading strategy." We argue that the most straightforward test of this model is whether these forward CIP trading strategies, which bet on arbitrages becoming smaller, can earn excess returns.

We then proceed to the data and estimate the excess returns of these forward CIP trading strategies. We define the forward CIP trading strategy as using FX forwards and forward-starting interest rate swaps to conduct a forward-starting CIP trade, and then unwinding the trade at its forward starting date. Consider

[1] Because short-term CIP arbitrage trades have little mark-to-market risk, we interpret these deviations as primarily reflecting the shadow cost of non-risk-weighted capital constraints. However, nothing in our analysis proves that these constraints, not other constraints, bind with respect to cross-currency basis trades, and we do not rely on this interpretation in our empirical analysis.

a trader who, at time $t$, first enters into a forward-starting CIP trade to go long Japanese yen and short Australian dollars for 3 months between $t+1$ and $t+4$, with the currency risk fully hedged. We refer to this trade as a 1-month forward 3-month CIP trade. Then in a month, at $t+1$, the trader unwinds the forward CIP trade by going long Australian dollars and short Japanese yen for 3 months, cancelling all the promised cash flows of the forward CIP trade. The profits of this two-step forward CIP trading strategy are proportional to the difference between the market-implied 1-month forward 3-month CIP deviation observed at $t$ and the actual 3-month CIP deviation realized 1 month later at $t+1$. The forward CIP trading strategy has a positive (negative) return if the future CIP deviation is smaller (bigger) than the market-implied forward CIP deviation today.

The expected return on the forward CIP trading strategy offers a direct test of intermediary asset pricing theories in which large CIP deviations indicate that intermediaries are very constrained. Because the forward CIP trading strategy pays off poorly in these constrained states, if the constraints of financial intermediaries are indeed a priced factor, we should expect the forward CIP trading strategy to earn positive excess returns on average, as a risk premium to compensate investors for bearing the systematic risk exposure to variations in the shadow cost of intermediary constraints.

We find a significant risk premium for certain forward CIP trading strategies during the post-GFC period. Specifically, we study our forward CIP trading strategy for seven of the most liquid currencies: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), Japanese yen (JPY), and U.S. dollar (USD). We consider both the cross-currency basis vis-à-vis the USD and the basis between two non-USD cross-pairs.

Our model emphasizes the importance of the currency pairs with the largest spot cross-currency bases. We show that the average returns of the forward CIP trading strategies for these pairs are generally sizable and statistically significant post-GFC. In particular, the forward CIP trading strategy for the "classic carry" AUD-JPY pair has annualized average profit equal to 16 basis points (bps) and an annualized Sharpe ratio of roughly 1.2. In contrast, the mean return of the AUD-JPY forward CIP trading strategy pre-GFC is negligible.

We also examine the performance of the forward CIP trading strategies for portfolios of currency pairs. The returns of the forward CIP trading strategy (henceforth "forward CIP returns") are significant and positive post-GFC for portfolios of currency pairs with large interest rate differentials ("carry") or large spot cross-currency bases ("basis"). In contrast, we do not find evidence of risk premiums when using a dollar strategy that equally weights all currencies vis-à-vis the USD. The strong performance of the carry and the basis portfolios and the lack of significance of the dollar portfolio are consistent with our model's prediction that the largest CIP violations are most informative about intermediary constraints. We also find that a portfolio based on the first principal component of the largest bases exhibits return characteristics that are similar to portfolios based on carry.

Intermediary constraints, if present, should affect many asset markets beyond the FX market. We show that CIP deviations are correlated with the first principal component of various other near-arbitrages. We also show that the returns of the forward CIP trading strategy are correlated with the proxies for intermediary wealth returns of He, Kelly, and Manela 2017.

However, the correlation between the forward CIP return and the intermediary equity return measure of He, Kelly, and Manela 2017 cannot explain the risk premium we uncover. We demonstrate this in regressions and more formally using the Bayesian factor model comparison method of Chib, Zeng, and Zhao 2020. These results justify the inclusion of the forward CIP return as an additional factor (along with the intermediary wealth return) in the SDF, consistent with our model. In particular, our results suggest that intermediaries are risk-tolerant and perceive strategies that perform poorly when investment opportunities are best to be especially risky.

We then test whether the excess returns of the tradable factors in this SDF (intermediary equity and the forward CIP return) are consistent with the prices of risk implied by the cross-section of assets, in an exercise building on He, Kelly, and Manela 2017 and related to Hu, Pan, and Wang 2013 and Pasquariello 2014.[2] We cannot reject the hypothesis that this risk is priced consistently across the various asset classes we consider, even when pooling across asset classes.

Our paper sits at the intersection of literature on arbitrage and on intermediary asset pricing. Recent empirical work on CIP violations (e.g., Du, Tepper, and Verdelhan 2018) has documented the existence and time-series properties of spot CIP arbitrages, as well as the quarter-end dynamics of these arbitrages (the leverage ratio calculation only relies on quarter-end bank balance sheets in many non-U.S. jurisdictions).[3] Spot CIP arbitrage opportunities exist at very short horizons (e.g., overnight), making it difficult for any risk-based story to explain the existence of these arbitrages. This differentiates our work from the large literature on "limits to arbitrage" that focuses the convergence risk.[4] Instead,

[2] We focus on arbitrage opportunities post-GFC, which we attribute to constraints on financial intermediaries resulting from post-GFC regulations, whereas Hu, Pan, and Wang 2013 and Pasquariello 2014 study mostly pre-GFC price dislocations attributable to transaction costs, stale prices, and related issues. Their results can be seen as demonstrating that marginal utility is high when transaction costs are high.

[3] Besides Du, Tepper, and Verdelhan 2018, a large recent literature explores CIP deviations post-GFC. For example, Borio et al. 2016 argue that hedging demand of different national banking systems can help explain cross-sectional variations in CIP deviations. Rime, Schrimpf, and Syrstad 2022 discuss the role of market segmentation in explaining CIP violations. Anderson, Du, and Schlusche 2019 measure the amount of potential arbitrage capital available to global banks for CIP arbitrage. Liao 2020 finds that CIP deviations post-GFC affects the corporate sector's funding currency decision. Avdjiev et al. 2019 examine the relationship between CIP deviations, the dollar exchange rate, and the cross-border bank flows in dollars. Du, Im, and Schreger 2018 and Jiang, Krishnamurthy, and Lustig 2021, and Krishnamurthy and Lustig 2019 use CIP deviations in government bond yields to measure convenience yield differentials between safe-haven government bonds and study implications for exchange rate dynamics. Augustin et al. 2020 model the term structure of CIP deviations.

[4] See, for example, Shleifer and Vishny 1997, Liu and Longstaff 2003, Duarte, Longstaff, and Yu 2007, and Duffie 2010.

short-dated CIP deviations can exist because of *constraints* on intermediaries, and in particular, non-risk-weighted total leverage constraints in the post-GFC regulatory environment. Other authors, including Boyarchenko et al. 2018, also attribute the existence a broad class of arbitrages post-GFC to the leverage ratio constraints. Fleckenstein and Longstaff 2018 link the cash-derivative basis in the interest rate future market to the cost of renting financial intermediary balance sheet space. Hébert 2018 interprets these arbitrages through an optimal policy framework.

We broaden this burgeoning literature on constraints-induced arbitrage by studying the term structure of arbitrage violations, as opposed to spot arbitrage violations, and emphasizing the general asset pricing implications of these deviations. In particular, we show that a significant fraction of the time-series variation in spot CIP violations is anticipated by the forward curve of CIP violations. This is true both generally and with respect to the quarter-end spikes documented in Du, Tepper, and Verdelhan 2018.

Within the intermediary asset pricing framework (surveyed by He and Krishnamurthy 2017), the models of Gabaix and Maggiori 2015 and Fang 2018 feature intermediary constraints as explanations of exchange rates dynamics. Much of this literature considers constraints that limit intermediaries' ability to access investments with favorable risk/return trade-offs, whereas we emphasize constraints (such as non-risk-weighted leverage constraints) that inhibit true arbitrages. In this respect, our model builds on Garleanu and Pedersen 2011. We also contribute to this literature by emphasizing the importance of intertemporal hedging considerations, following Campbell 1993 and Kondor and Vayanos 2019, whereas much of the literature (e.g., He and Krishnamurthy 2011; Garleanu and Pedersen 2011; He, Kelly, and Manela 2017) relies on log utility for intermediaries and neglects these considerations.

In taking the model to the data, we are building on He, Kelly, and Manela 2017, Adrian, Etula, and Muir 2014, Hu, Pan, and Wang 2013, and Haddad and Muir 2021. Our model can be thought of as nesting the SDFs discussed by Adrian, Etula, and Muir 2014 and He, Kelly, and Manela 2017. When risk aversion is equal to one, the intertemporal terms in our SDF vanish, and the intermediary wealth return is the SDF (as in He, Kelly, and Manela 2017). When risk aversion is equal to zero (the risk-neutral case), the SDF consists only of intertemporal hedging terms, which are proxied for by the shadow cost of intermediary constraints (as in Adrian, Etula, and Muir 2014). We measure these shadow costs using CIP violations, which we argue in the context of our model is a clean and valid measure. In contrast, Adrian, Etula, and Muir 2014 measure these shadow costs using leverage. It is not clear, however, whether the price of leverage risk comes from its correlation with intermediary wealth returns, its correlation with intermediary shadow costs, or some combination thereof. Closer in spirit to our exercise is Hu, Pan, and Wang 2013, who measure intermediary constraints using Treasury yield curve dislocations. A comparison to this approach reveals the second key advantage of using CIP violations to measure shadow costs: we can directly estimate the price of risk from our forward arbitrage trading strategy, instead of relying on the usual cross-sectional asset pricing analysis. We focus our analysis on this direct estimate of the price of risk, and verify in our cross-sectional analysis that the price of risk we infer from the cross-section is consistent with the price of risk we estimate directly.

## 1. Hypothesis and Model

In the empirical analysis that follows, we will test the hypothesis that changes in the magnitude of cross-currency bases (i.e., CIP violations) are priced. This hypothesis is motivated by a specific intermediary asset pricing model, which we outline below and detail in the Internet Appendix B. Our paper is primarily an empirical study; the purpose of the model is to motivate our hypothesis and to provide a framework to interpret our results. However, at the outset, we should acknowledge other possible interpretations of our empirical results, some of which we will discuss after presenting those results.

Our model is designed to capture three key ideas:

1. Arbitrage violations measure investment opportunities. This is true in our model because arbitrage violations can exist only if constraints on intermediaries bind, and constraints on intermediaries bind only if they prevent those intermediaries from taking advantage of investment opportunities (this is the definition of “bind”).
2. Investment opportunities are likely to be best when intermediary wealth is low. For this reason alone, if our empirical proxies for intermediary wealth are imperfect, we should expect changes in arbitrage violations to be a priced risk controlling for imperfect wealth proxies.
3. Even holding wealth fixed, changes in investment opportunities can be a priced factor. In the presence of good investment opportunities, the marginal value of wealth might be high, because it allows intermediaries to take advantage of those opportunities, or it might be low, because good investment opportunities enable larger payouts. The sign of this effect is determined by the intermediary’s intertemporal hedging concern.

The model is a discrete-time version of He and Krishnamurthy 2011 that incorporates a regulatory constraint (building on He and Krishnamurthy 2017) and intertemporal hedging considerations (following Campbell 1993). Under this regulatory constraint, each asset that the intermediary can hold (indexed by $i \in I$) is subject to an asset-specific weight $k^i$. As a result, each asset the intermediary owns is priced by the log SDF $m_{t+1}$ of the following form:

$$m_{t+1} = \mu_t(k^i) - \gamma r^w_{t+1} + \xi |x_{t+1,1}|, \tag{1}$$

where $r^w_{t+1}$ is the return on the manager of an intermediary’s wealth portfolio and $|x_{t+1,1}|$ is the absolute value of a one-period cross-currency basis. The dependence of the mean of the SDF on $k^i$ reflects the effect of the regulatory constraint, and $|x_{t+1,1}|$ serves as a proxy for future investment opportunities. Note that assets with different risk weights, $k^i$, will be priced by different SDFs, leading to arbitrage opportunities; however, all of these SDFs agree on the risk prices $\gamma$ and $\xi$.

The SDFs in (1) nest the SDFs discussed in Adrian, Etula, and Muir 2014 and He, Kelly, and Manela 2017. When $\gamma = 1$, $\xi = 0$, and the SDF is exactly the intermediary wealth return as in He, Kelly, and Manela 2017. When $\gamma = 0$, $\xi > 0$, meaning that the marginal value of wealth is high when future investment opportunities are best (as in Adrian, Etula, and Muir 2014).[5]

Our hypothesis is that $\xi$ is economically and statistically distinguishable from zero. The key idea behind this hypothesis is that the cross-currency basis $|x_{t,1}|$ is both a literal arbitrage and a measure of the investment opportunities available to intermediaries at time $t$. An arbitrage can exist only if intermediaries are constrained and cannot take advantage of an otherwise attractive investment opportunity. In the presence of such constraints, an intermediary concerned with hedging against changes in future investment opportunities should perceive assets whose returns are correlated with $|x_{t+1,1}|$ as particularly risky or safe, depending on the sign of the intermediary's intertemporal hedging concerns.

Campbell 1993 shows that SDFs with the form of (1), interpreting $|x_{t+1,1}|$ as an arbitrary random variable as opposed to a cross-currency basis and without a mean that depends on $k^i$, can be derived using CRRA or Epstein-Zin preferences (and assuming lognormality and homoscedasticity). In this case, $|x_{t+1,1}|$ must proxy for the revision in expectations about future investment opportunities. That is,

$$|x_{t+1,1}| - E_t[|x_{t+1,1}|] \propto \sum_{j=1}^{\infty} \rho^j (E_{t+1} - E_t)[r^w_{t+1+j}].$$

The sign of the coefficient $\xi$ depends on whether the relative risk aversion[6] coefficient $\gamma$ is greater or smaller than one ($\gamma < 1 \Leftrightarrow \xi > 0$). Intertemporal hedging concerns on their own can be used to justify any SDF, including the ones we consider. Our point is that for specific reasons we expect arbitrage violations to predict future investment opportunities.

Suppose the manager of an intermediary has CRRA or Epstein-Zin preferences and holds an equity claim on the intermediary.[7] The intermediary

[5] Adrian, Etula, and Muir 2014 build on Brunnermeier and Pedersen 2009 (effectively a three-period model), and therefore summarize investment opportunities with single future return. Adrian, Etula, and Muir 2014 also assume investment opportunities are negatively correlated with intermediary leverage; it is not a priori clear this should be the case, but it holds in the Brunnermeier and Pedersen 2009 model.

[6] As discussed in Campbell 1993, this result holds for both CRRA and Epstein-Zin preferences. That is, it is $\gamma$ and not the elasticity of intertemporal substitution coefficient that determines the sign of $\xi$.

[7] We follow He and Krishnamurthy 2011 in assuming that the manager must hold an equity claim of a certain size to avoid moral hazard. For the remainder of this section, we will assume that this constraint does not bind. We

is subject to a regulatory constraint,

$$\sum_{i\in I} k^i |\alpha_t^i| \leq 1. \tag{2}$$

Here, $\alpha_t^i$ is the intermediary's holding of asset $i$ at time $t$ as a share of the intermediary's equity, and $k^i$ is the asset-specific weight mentioned above. This constraint captures some of the key features of leverage ratios and risk-weighted capital requirements. First, to the extent that the $k^i$ differ across assets, the constraint can capture risk-weights. Second, the constraint is relaxed by increasing the level of equity financing relative to debt financing, holding fixed the dollar holdings of each asset. Third, the constraint can entirely omit certain assets, such as derivatives, consistent with how some leverage constraints and risk-weighted capital constraints operate. To simplify our exposition, we will assume in what follows that derivatives are not included in the regulatory constraint.[8]

The manager's first-order condition for the portfolio share $\alpha_t^i$ is

$$E_t[\exp(m_{t+1})(R_{t+1}^i - R_t^b)] = \lambda_t^{RC} k^i sgn(\alpha_t^i), \tag{3}$$

where $m_{t+1}$ is the manager's log SDF, $R_{t+1}^i$ is the gross return on asset $i$, $R_t^b = \exp(r_t^b)$ is the gross rate on the intermediary's debt between dates $t$ and $t+1$, $sgn(\cdot)$ is the sign function, and $\lambda_t^{RC}$ is the (scaled) multiplier on the regulatory constraint.[9] Let us apply this equation to two portfolios of assets: the cross-currency basis arbitrage and the wealth portfolio.

Let $S_t$ denote the exchange rate at time $t$ (in units of foreign currency per U.S. dollar), and let $F_{t,1}$ denote the one-period ahead forward exchange rate. We define the spot one-period cross-currency basis as

$$X_{t,1} = \frac{R_t^b}{R_t^c} \frac{F_{t,1}}{S_t} - 1$$

where $R_t^c$ is the foreign currency risk-free rate, and let $x_{t,1} = \ln(1 + X_{t,1})$ be the log version. The first order condition is, taking absolute values,

$$E_t[\exp(m_{t+1} + r_t^b)]|1 - \exp(-x_{t,1})| = \lambda_t^{RC} k^c, \tag{4}$$

where $k^c$ is the risk-weights of the foreign currency risk-free bond.

make this assumption both for simplicity and to emphasize that the regulatory constraint can bind even if the equity constraint does not. For formulas that extend to the case with a binding constraint, and a more detailed discussion of this issue, see Internet Appendix B.

[8] This particular functional form follows He, Kelly, and Manela 2017. Its details are not essential for our result; in particular, we could easily accommodate a constraint that treats long ($\alpha_t^i > 0$) and short ($\alpha_t^i < 0$) positions asymmetrically. Considering regulatory constraints that include derivatives complicates the analysis but does not alter the main predictions of the model that we will take to the data.

[9] In the particular case in which $\alpha_t^i = 0$, we have the usual inaction inequalities, $-\lambda_t^{RC} k^i \leq E_t[\exp(m_{t+1})(R_{t+1}^i - R_t^b)] \leq \lambda_t^{RC} k^i$ (see Internet Appendix B).

The key takeaway from this equation is that the absolute value of the cross-currency basis can be used to measure the shadow cost of the regulatory constraint. Intuitively, if an arbitrage opportunity is available to the intermediary, the intermediary would take advantage of it if the intermediary could; therefore, the intermediary must be constrained. The size of the arbitrage opportunity can be used to measure the degree to which the constraint binds (a point emphasized by Hébert 2018).

Let us now consider the first-order condition applied to the entire wealth portfolio (i.e., taking the $\alpha_t^i$-weighted sum of Equation (3) across the various assets). In this case, by the definition of the constraint,

$$E_t[\exp(m_{t+1})(\exp(r_{t+1}^w)-\exp(r_t^b))]=\lambda_t^{RC}. \tag{5}$$

This equation captures the intuition that the shadow cost of the constraint is equal to the marginal value of the forgone investment opportunities. The constraint binds only if the intermediary has valuable investment opportunities it cannot exploit due to the constraint.

Combining these two equations to eliminate the shadow cost,

$$\frac{E_t[\exp(m_{t+1})(\exp(r_{t+1}^w-r_t^b)-1)]}{E_t[\exp(m_{t+1})]}=\frac{|1-\exp(-x_{t,1})|}{k_c}.$$

That is, the arbitrage available at time $t$ can measure the investment opportunities available at time $t$. Log-linearizing and assuming homoscedasticity (see (A1) in Internet Appendix B for details),

$$(E_{t+1}-E_t)[r_{t+1+j}^w]=(E_{t+1}-E_t)[r_{t+j}^b+k_c^{-1}|x_{t+j,1}|].$$

Thus, revisions in expected future cross-currency bases measure revisions in future investment opportunities more generally. Moreover, these effects are amplified by leverage $k_c^{-1}$. Because innovations to the cross-currency basis are persistent, we can proxy for revisions in expectations about $|x_{t+j,1}|$ with the innovation to $|x_{t+1,1}|$. This result, combined with intertemporal hedging, justifies the SDFs of Equation (1).

This argument (described in more detail in the Internet Appendix) motivates our exercise, which attempts to measure price of cross-currency basis risk ($\xi$). The most direct way to estimate this price of risk is to study a derivative contract whose payoff is linear in $|x_{t+1,1}|$. If such a contract has an excess return that cannot be explained by the covariance between $|x_{t+1,1}|$ and the other parts of the hypothesized SDF (i.e., $r_{t+1}^w$), we should conclude that innovations in the cross-currency basis are indeed a priced risk factor (or at least correlated with an omitted factor). The forward CIP trading strategy that we construct in our empirical analysis is exactly this derivative contract. The following remarks discuss some basic insights from the model that guide our empirical analysis.

**Correlation between factors.** The two factors in our SDFs (the intermediary wealth return and the basis) likely move together. Because our model treats asset prices as exogenous, it makes no predictions about this comovement. Most general equilibrium intermediary asset pricing models (e.g., He and Krishnamurthy 2011) predict that investment opportunities are best for intermediaries precisely when intermediaries have lost wealth, and hence we should expect a negative correlation between the two factors.

**Omitted factors.** Equation (1) likely omits important elements of the SDF. Any factor that predicts revisions in expectations about future investment opportunities (about the future cross-currency basis, future risk-free rates, or, in a heteroscedastic model like Campbell et al. 2018, future volatility) should also enter the SDFs.

**Noisy intermediary wealth return measures.** Our proxies for intermediary wealth returns are measured with noise. For this reason, we will never be able to definitively prove that the excess returns we document are caused by intertemporal hedging concerns as opposed to by correlation with intermediary wealth returns that is not captured by our proxies. Moreover, in light of the point on omitted factors above, any atheoretical factor that has explanatory power above and beyond our two SDF factors can be rationalized as either providing a better measure of intermediary wealth returns or a proxy for future investment opportunities. We refrain from running "horse race" regressions with additional factors because of this lack of a clear interpretation.

**Sources of variation.** Our model shows that the shadow cost of regulatory constraints can be measured with CIP violations, but is silent on why CIP violations vary over time. We expect that supply shocks (low intermediary net worth), demand shocks (e.g., changing customer preferences), and changes in the structure of the regulatory constraint will all affect the shadow cost of the constraints on intermediaries. Our results demonstrate that, regardless of what is driving changes in these shadow costs, the SDFs of Equation (1) should price the assets available to the intermediary.

**CIP versus other arbitrages.** Our model places no special emphasis on CIP violations. Any arbitrage that intermediaries engage in could be used to measure $\lambda_t^{RC}$. In subsection 4.1, we argue that among various arbitrages and near-arbitrages documented in the literature, CIP violations are unique in terms of our ability to accurately measure the spot arbitrage $x_t$ and to construct a trading strategy that directly bets on $x_t$ becoming larger or smaller in the future. We also document that spot CIP violations are highly correlated with other arbitrages, consistent with our model.

**Magnitudes.** Shocks to the cross-currency basis are small (basis points). However, intermediaries are quite levered, meaning that $k^c$ might be small, consumption-wealth ratios for managers are likely small (meaning $\rho$ is close to one), and innovations to the basis are persistent. These forces increase the price of cross-currency basis risk, and might cause a significant fraction of the volatility of the SDFs to be attributable to innovations in the basis.

**Leverage constraints post-GFC.** In our model, CIP violations can arise only if the regulatory constraint binds for this riskless arbitrage. In the absence of a binding constraint, CIP violations cannot exist in equilibrium, even if the inside equity constraint binds. The lack of CIP violations pre-GFC in the data is therefore consistent with the absence of non-risk-weighted leverage constraints for many banks prior to the GFC. The persistence of CIP violations and other short-term arbitrages (such as the interest rate on reserve arbitrage) post-GFC is consistent with binding leverage constraints under Basel III.[10]

**Heterogeneity across currencies.** Our description of the model has emphasized a single cross-currency basis, whereas our empirical analysis will consider a variety of currency pairs. In the context of the model, all of the cross-currency bases that the manager invests in will have the maximal arbitrage per unit risk weight available. Any basis that offers an inferior level of arbitrage per unit risk weight will receive a zero portfolio weight. In particular, if the risk-weights are identical across currencies, the manager would invest only in the basis with the largest arbitrage violation. In reality, a smaller measured basis might be nevertheless actively traded by intermediaries for several reasons.

- Intermediaries may have some degree of market power, and face different demand curves across currencies (Wallen 2020). Intermediaries are also heterogeneous, and in particular have different deposit bases and access to wholesale funding markets across currencies.[11]
- Some regulatory metrics, such as the liquidity coverage ratio (LCR), are monitored on the currency-by-currency basis.[12] In addition, the allocation of CIP arbitrage activities across currency pairs also affects the distribution of liquidity across different entities and jurisdictions, which could make liquidity stress tests and resolution planning rules more binding (Correa, Du, and Liao 2020). These considerations would lead to the intermediary having different shadow costs for different currencies.

---

[10] Non-U.S. banks did not face a non-risk-weighted leverage ratio requirement prior to the 3% leverage ratio requirement under Basel III. U.S. banks had a 3% leverage ratio requirement prior to Basel III, and a 5%-6% leverage ratio requirement under Basel III. During the pre-GFC period, other kinds of constraints might have been binding, and CIP violations are not the right way to measure the shadow cost of such constraints.

[11] See, for example, Rime, Schrimpf, and Syrstad 2022 on the impact of money market segmentation on CIP deviations.

[12] Even though the Basel III LCR requirement is calculated at the aggregate level across all currencies, currency-specific LCRs are nevertheless actively monitored by bank examiners and bank internal managers.

- The meaning of the benchmark OIS rate varies across currencies, and it may be a better proxy for the banks' borrowing/lending costs in some currencies than in others. Heterogeneity in the basis might be caused in part by a lack of perfect comparability of interest rates across currencies. We will discuss this point in more detail in Section 3.

In our empirical analysis, we strike a balance between the literal interpretation of our model and these real-world considerations by focusing on the currency pairs with the largest and most robust bases.

## 2. Forward CIP Arbitrage

We describe the forward CIP trading strategy that bets on the size of the future cross-currency basis in three steps. First, we revisit "spot" cross-currency bases (as in Du, Tepper, and Verdelhan 2018), and describe the cross-currency bases based on overnight index swap (OIS) rates that we use in our empirical analysis.[13] Second, we discuss "forward" cross-currency bases, constructed from forward-starting OIS swaps and FX forwards. Third, we introduce our forward CIP trading strategy, which initiates a forward-starting cross-currency basis trade but then unwinds the trade once it becomes a spot trade. This trading strategy is not itself an arbitrage, but rather a risky bet on whether available arbitrages will become bigger or smaller.

We study cross-currency bases in seven major currencies, AUD, CAD, CHF, EUR, GBP, JPY, and USD.[14] We examine the bases of both individual currency pairs and portfolios of currency pairs, although our portfolios exclude CHF because of data limitations.[15] All data on spot and forward FX rates, interest rate swaps, and FRAs are daily data obtained from Bloomberg using London closing rates. Our data set begins in January 2003 and ends in December 2020. We divide our data into three periods based on potentially different regulatory environments facing the intermediaries: pre-GFC, January 1, 2003, to June 30, 2007, GFC, July 1, 2007, to June 30, 2010, and post-GFC, July 1, 2010, to December 31, 2020.[16]

---

[13] For robustness, in Internet Appendix Tables A5 and A6, we also consider a forward CIP trading strategy based on interbank offer rates (IBOR) and forward rate agreements (FRAs) indexed to these IBOR rates. The OIS and FRA data for the pre-GFC period appear less reliable (more missing or erroneous values) than the data for the GFC and post-GFC periods.

[14] We begin with the G10 currencies and exclude the Norwegian krona (NOK) and Swedish krona (SEK) because of limited data availability on OIS rates and IBOR FRAs. We also exclude the New Zealand dollar (NZD) because the OIS floating leg for the NZD is not a market rate but rather an administered central bank policy rate, the Official Cash Rate (OFR). The OFR is not equal to the actual overnight rate in the financial market, which generally fluctuates 0.25% around the OFR.

[15] The CHF OIS reference was changed at the end of 2017 because of a lack of liquidity in the underlying market. OIS swaps on the new index are not liquid enough for the purposes of our analysis. For this reason, we present single currency-pair results with CHF through the end of 2017 but do not include CHF in our results based on portfolios.

[16] We focus on monthly returns in our main analysis, and hence the last trading date in our sample is November 30, 2020.

Our main analysis focuses on the post-GFC period, which features a non-risk-weighted leverage ratio constraint under the Basel III regulatory environment. This stands in sharp contrast to the pre-GFC and GFC samples, during which bank capital constraints were largely based on risk and riskless short-term CIP arbitrages faced no capital charge. An important lesson from the GFC turmoil was that the ex ante risk weights could inaccurately reflect risk, and in 2010 the non-risk-weighted leverage ratio requirement was drafted as an important pillar of Basel III. Since then, the Basel III regulations have been finalized and gradually implemented. Even before the final rules took effect, early compliance of Basel III was common among large banking organizations, as it takes time to reorganize complex business activities. Regulators and bank shareholders may also have taken Basel III regulatory metrics into account even before the regulations were formally implemented. In addition to the Basel III implementation, our post-GFC sample also features two major financial crises, the European debt crisis and the COVID-19-induced financial turmoil in March 2020. CIP deviations widened significantly during both crises.[17]

## 2.1 OIS-based spot cross-currency bases

We first define the $\tau$-month tenor OIS-based spot cross-currency basis vis-à-vis the USD. Let $R^c_{t,0,\tau}$ denote the annualized spot gross $\tau$-month interest rate in foreign currency $c$ available at time $t$, and let $R^\$_{t,0,\tau}$ denote the corresponding spot rate in U.S. dollars. The middle subscript "0" denotes a spot rate (as opposed to a forward rate). We express exchange rates in units of foreign currency per USD. That is, an increase in the spot exchange rate at time $t$, $S_t$, is a depreciation of the foreign currency and an appreciation of the USD. The $\tau$-month forward exchange rate at time $t$ is $F_{t,\tau}$.

Following convention (e.g., Du, Tepper, and Verdelhan (2018)), we define the $\tau$-month tenor spot cross-currency basis of foreign currency $c$ vis-à-vis the USD as

$$X^{c,\$}_{t,0,\tau} = \frac{R^\$_{t,0,\tau}}{R^c_{t,0,\tau}} \left(\frac{F_{t,\tau}}{S_t}\right)^{\frac{12}{\tau}} - 1, \tag{6}$$

and the log version as $x^{c,\$}_{t,0,\tau} = \ln(1+X^{c,\$}_{t,0,\tau})$. This definition is identical to the one employed in our model, except that we now consider an arbitrary tenor $\tau$ and use annualized interest rates.

The classic CIP condition is that $x^{c,\$}_{t,0,\tau} = X^{c,\$}_{t,0,\tau} = 0$. If the cross-currency basis $x^{c,\$}_{t,0,\tau}$ is positive (negative), then the direct U.S. dollar interest rate, $R^\$_{t,0,\tau}$, is higher (lower) than the synthetic dollar interest rate constructed from the foreign currency bond and exchange rate transactions.

The CIP condition is a textbook no-arbitrage condition if the U.S. and foreign interest rates used in the analysis are risk-free interest rates. For our main

[17] See Section 3.5 for discussion on sample splits of the post-GFC period.

analysis, we choose OIS rates as our proxy for risk-free interest rate. The OIS rate is the fixed rate of a fixed-for-floating interest rate swap in which the floating rate is an overnight unsecured rate.[18]

The OIS is a good proxy for the risk-free rate across maturities for several reasons. First, the OIS allows investors to lock in fixed borrowing and lending rates for a fixed maturity, by borrowing and lending at the nearly risk-free floating overnight rate each day over the duration of the contract. Second, the interest rate swaps themselves have very little counterparty risk, because there are no exchanges of principal, only exchanges of interest. These derivative contracts are also highly collateralized and in recently years have been centrally cleared in most major jurisdictions. Third, OIS swaps are generally very liquid and traded at a large range of granular maturities (unlike, e.g., repo contracts).

Internet Appendix Figure A1 shows the 3-month OIS-based cross-currency basis for the six sample currencies vis-à-vis the USD between January 2003 and December 2020. The 3-month OIS basis was close to zero pre-GFC and deeply negative during the peak of the GFC. After the GFC, OIS-based CIP deviations persisted. Among our sample currencies, AUD has the most positive OIS basis, and JPY, CHF, and EUR have the most negative OIS bases. Internet Appendix Figure A2 shows 3-month IBOR cross-currency bases, which follow similar patterns.

We define the spot cross-currency basis between two non-USD currencies $c_1$ and $c_2$ as the difference in their respective log cross-currency basis vis-à-vis the USD,

$$x_{t,0,\tau}^{c_1,c_2} = x_{t,0,\tau}^{c_1,\$} - x_{t,0,\tau}^{c_2,\$}. \tag{7}$$

We use this definition, as opposed to directly constructing the cross-currency basis between $c_1$ and $c_2$, both because most trades in currency forwards involve a USD leg and to restrict our sample to U.S. FX trading days on which the U.S. federal funds market is open.[19]

### 2.2 Forward bases

Next, we define a forward-starting cross-currency basis. Trading a forward starting cross-currency basis allows an agent to lock-in the price of a cross-currency basis trade that will start in the future.

We define a forward-starting cross-currency basis using forward interest rates and FX forwards. Let $R_{t,h,\tau}^{c}$ be the $h$-month forward-starting annualized

[18] The list of overnight reference rates for the OIS and their day count conventions for the seven major currencies we study can be found in Internet Appendix Table A1. For two currencies, the OIS rate is nonstandard. For CAD, the overnight rate is a repo (secured) rate; for CHF, the unsecured overnight rate had volumes so low that the OIS rate was changed to reference a secured rate in 2017.

[19] According to recent BIS FX derivatives statistics, 90% of global FX swaps have the USD on one leg. Some cross-pairs, such as EURJPY and EURCHF, are actively traded. There are only negligible differences between the cross-currency basis calculated directly using the FX swap rates for the cross pairs and the basis calculated using Equation (7). The triangular arbitrage for the cross-currency basis holds quite well post-GFC because the arbitrage only involves trading FX derivatives with limited balance sheet implications.

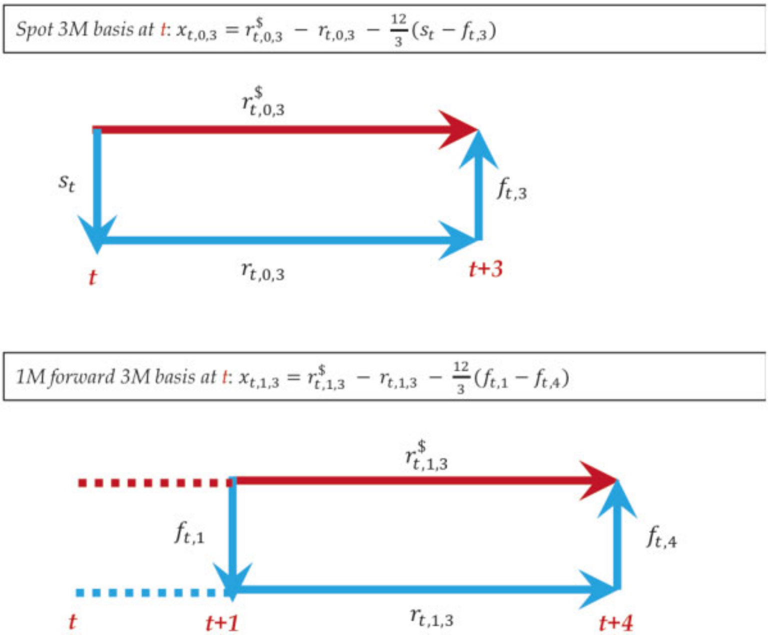


**Figure 1**
**Illustration of spot versus forward cross-currency basis**
This figure illustrates the spot 3-month cross-currency basis and the 1-month forward 3-month cross-currency basis. The spot basis $x_{t,0,3}$ is defined in the text, and the forward basis $x_{t,1,3}$ is defined in the text.

$\tau$-month gross interest rate in currency $c$ at time $t$, and let $R^{\$}_{t,h,\tau}$ be the equivalent rate in the USD. The forward-starting cross-currency basis of foreign currency $c$ vis-à-vis the USD is

$$X^{c,\$}_{t,h,\tau} = \frac{R^{\$}_{t,h,\tau}}{R^{c}_{t,h,\tau}} \left( \frac{F_{t,h+\tau}}{F_{t,h}} \right)^{\frac{12}{\tau}} - 1, \tag{8}$$

and the log version is $x^{c,\$}_{t,h,\tau} = \ln(1 + X^{c,\$}_{t,h,\tau})$. Figure 1 illustrates the definitions of the spot and forward cross-currency basis.

Equivalently, we can define the logarithm of $h$-month forward $\tau$-month cross currency basis at time $t$ in terms of two spot cross-currency bases under the assumption of no-arbitrage between forward interest rate swaps and the term structure of spot interest rate swaps:

$$x^{c,\$}_{t,h,\tau} = \frac{h+\tau}{\tau} x^{c,\$}_{t,0,h+\tau} - \frac{h}{\tau} x^{c,\$}_{t,0,h}. \tag{9}$$

The equivalence between Equations (8) and (9) is shown in Internet Appendix C. Equation (9) also expresses a close analogy between forward cross-currency bases and forward interest rates. As in Equation (7), we define the forward

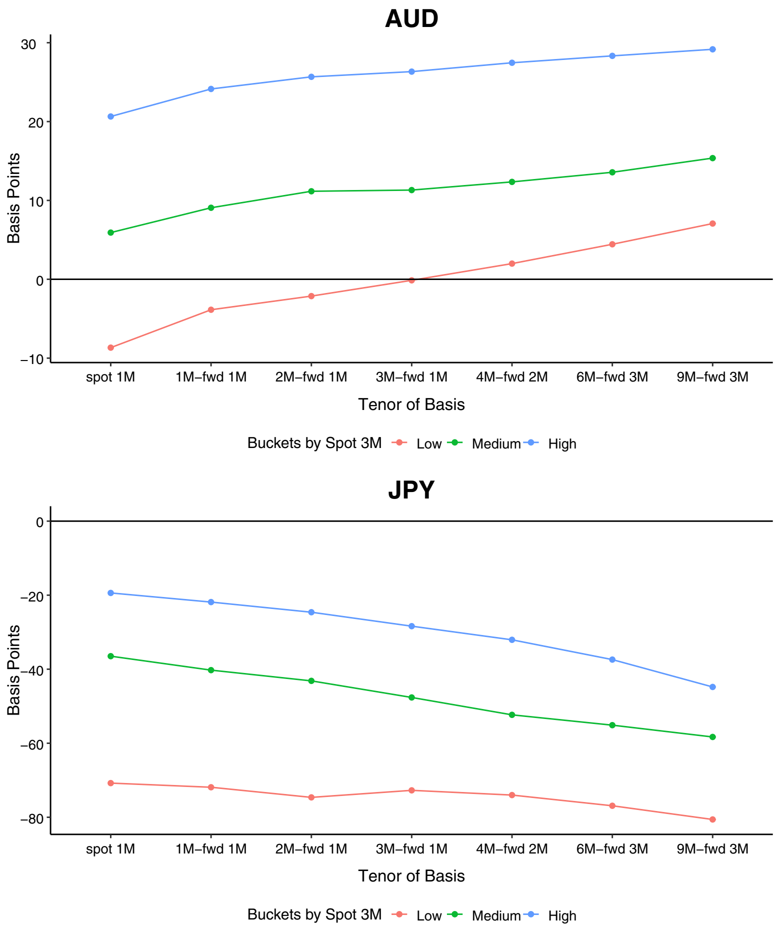


**Figure 2**
**Term structure of the forward cross-currency basis**
This figure illustrates the time-series average spot and forward-starting cross-currency bases in AUD and JPY, vis-à-vis the USD, respectively, as defined in Equation (8). For each currency, the sample from July 2010 to December 2020 is split into three subsamples based on the tercile of the level of the spot 3-month OIS cross-currency basis. Within each subsample, the time-series average of the relevant spot/forward OIS cross-currency basis is shown.

cross-currency basis between non-USD currencies $c_1$ and $c_2$ as

$$x_{t,h,\tau}^{c_1,c_2} = x_{t,h,\tau}^{c_1,\$} - x_{t,h,\tau}^{c_2,\$}. \tag{10}$$

Next, we consider the typical shape of the term structure of CIP violations, that is, the shape of the cross-currency basis forward curve. It is possible to construct forward CIP trades of many different horizons $h$ and tenors $\tau$. However, the most liquid and reliable OIS tenors are 1, 2, 3, 4, 6, 9, and 12 months. In Figure 2, we present the forward curves of AUD and JPY vis-à-vis the USD for all reliable horizons: spot and 1, 2, 3, 4, 6, and 9 months. The tenor $\tau$ of these forward CIP trades differs, beginning at one month and increasing to 3 months. Internet Appendix Figure A3 presents an alternative version of the forward curve that uses only 3-month tenors.

We present these forward basis curves as time-series averages for two currencies, AUD and JPY. These two currencies stand out in the data as having very positive/negative spot cross-currency bases vis-à-vis the USD during our post-GFC sample period, respectively. For each currency, we divide our sample into three subsamples based on the tercile of the level of the spot 3-month tenor basis. We then compute the time-series average of the spot and forward-starting cross-currency basis within each subsample.

From these forward curves, it is immediately apparent that the forward cross-currency bases tend to be larger (more positive) than the spot cross-currency basis for AUD, and smaller (more negative) for JPY. This fact is somewhat analogous to the tendency of the term-structure of interest rates to be upward sloping. If we think of forward cross-currency bases as being equal to expectations under a risk-neutral measure (an approach that is valid in our model despite the presence of arbitrage), then this suggests that the absolute value of spot cross-currency basis is generally expected to increase under the risk-neutral measure.

This raises the question of whether the spot cross-currency basis is also expected to increase in absolute value under the physical measure. That is, do the slopes of these forward curves reflect expectations, risk premiums, or some combination thereof?

## 2.3 Forward CIP trading strategy

The forward CIP trading strategy consists of a forward cross-currency basis trade and a spot cross-currency basis trade at a later date. At time $t$, an agent enters into the $h$-month forward $\tau$-month cross-currency basis trade. After $h$ months, at time $t+h$, the agent unwinds the trade by shorting the then-spot $\tau$-month cross-currency basis.

Although the forward CIP trading strategy involves two potential arbitrage opportunities, it is itself risky in that the spot $\tau$-month cross currency basis at time $t+h$ is not guaranteed to be equal to the $h$-month forward $\tau$-month cross-currency basis at time $t$. Figure 3 illustrates the mechanics this trading strategy.

The profits from this trading strategy are primarily a function of the realized cross-currency basis at time $t+h$, compared to the $\tau$-month forward cross-currency basis at time $t$. To a first-order approximation, the annualized profit per dollar notional (which can be thought of as an excess return) is

$$\pi^{c_1,c_2}_{t+h,h,\tau} \approx \frac{\tau}{h}\left(x^{c_1,c_2}_{t,h,\tau} - x^{c_1,c_2}_{t+h,0,\tau}\right). \tag{11}$$

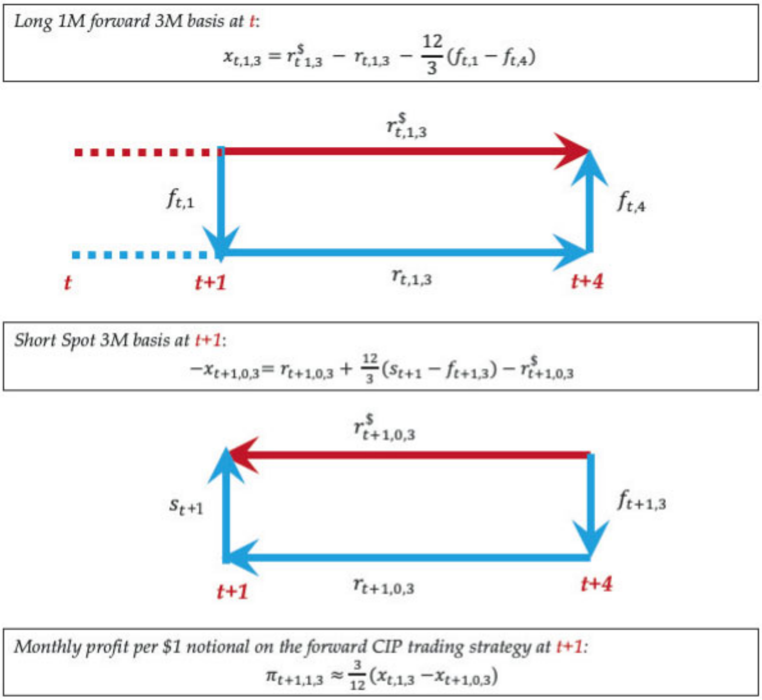


**Figure 3**
**Illustration of the forward CIP trading strategy**
This figure illustrates the return on a 1-month forward 3-month forward CIP trading strategy. At time $t$, the trader enters the forward basis, $x_{t,1,3}$, which is the forward direct interest less the forward synthetic interest. At time $t+1$, the trader unwinds the spot basis, $-x_{t+1,0,3}$, which is the spot synthetic interest less the spot direct interest. The realized monthly profit per dollar notional on this forward CIP trading strategy is approximately the sum of the two bases: $x_{t,1,3}+(-x_{t+1,0,3})$, normalized by the duration 3/12.

The term $\frac{\tau}{h}$ plays the role of a duration, converting the difference between the forward and realized basis, $x^{c_1,c_2}_{t,h,\tau} - x^{c_1,c_2}_{t+h,0,\tau}$, into an annualized dollar profit per unit notional.[20]

The key property of the forward CIP trading strategy for our purposes is that it allows an intermediary to bet on whether the cross-currency basis will be higher or lower than implied by the forward cross-currency basis. Our model equates the magnitude of the basis with the degree to which regulatory constraints binds. Consequently, this strategy allows intermediaries to bet on whether constraints will be tighter or looser in the future.

The forward CIP trading strategy is a valid trading strategy even if the underlying cross-currency basis is not actually tradable or not a pure arbitrage. For example, individual arbitrageurs may not have direct access to the OIS

[20] We derive this expression, which is a first-order approximation, from a more exact calculation in Internet Appendix D.

floating leg.[21] Nevertheless, the forward CIP trading strategy is a valid trading strategy that bets on whether the basis as measured by OIS swaps referencing this rate becomes larger or smaller.

Moreover, the forward CIP trading strategy *per se* does not materially contribute to the balance sheet constraints of financial intermediaries, especially in comparison with the spot CIP arbitrage. This is because interest rate forwards and FX derivatives have zero value at inception. The required initial and variation margins for the derivative positions are generally a few percent of the total notional of the trade. In contrast, the spot CIP arbitrage requires actual cash market borrowing and lending, and is therefore balance sheet intensive.[22]

### 2.4 Forward CIP returns during financial turmoil

In the next section, we study the average returns of forward CIP trading strategies in various currencies and portfolios of currencies. To understand the patterns in the data behind our results, it is useful to first understand when the trading strategy does particularly well and particularly poorly. To highlight these patterns, we focus our discussion on the AUD-JPY currency pair (which is representative of other portfolios to be discussed in the next section), and on three distress periods in our data sample: the 2008–2009 GFC, the European debt crisis during 2011–2012, and the COVID-19-induced turmoil in March 2020.

We begin with the events of March 2020. By March 19th, 2020 (roughly the peak of the financial turmoil), spot CIP bases widened to levels seen in prior crises in almost all currencies. The OIS-based AUD-JPY 3M spot basis reached 222 bps, and this spike was not anticipated by the forwards. As a result, the profit per dollar notional of the 1M-forward 3M AUD-JPY forward CIP trading strategy (defined in Equation (11)) was on average -103 bps for trades initiated in the 30 days prior to March 19, 2020.

However, in April 2020 the spot bases converged back to roughly the levels observed in February 2020, and this normalization was also not fully anticipated by the forwards.[23] As a result of these movements, the 1-month forward 3-month AUD-JPY forward CIP trading strategy experienced large positive

---

[21] In the United States, the floating leg of the OIS is the federal funds rate. Only banks with reserve accounts at the Federal Reserve can trade in the federal funds market.

[22] For example, a $100 million spot CIP trade requires borrowing $100 million in the cash market and lending $100 million in the FX swap market, which expends the size of the total exposure of the bank by $100 million for the leverage ratio calculation. However, the impact of $100 million interest rate swap on the leverage ratio is significantly smaller. The total exposure includes initial and variation margins (typically a couple percent of total notional), and an additional 0%–1.5% of the swap notional calculated for off-balance-sheet interest rate derivative exposure using the Current Exposure Method, depending on the maturity of the interest rate swaps (Haynes, McPhail, and Zhu 2018).

[23] On March 19, 2020, the 1-month forward 3-month basis was 107 bps. Thus, in contrast to the usual upward-sloping pattern, at the peak of the crisis the forwards anticipated a substantial decline in the spot basis. The realized decline, however, was even larger than what was priced in, by April 21st, 2020, the spot basis was back to 56 bps and the 1-month forward basis was 61 bps (returning to the usual upward-sloping pattern).

profits, 189 bps on average, on trades initiated in the 30 days following March 19, 2020.

We observe similar patterns during other periods of financial turmoil. Trades initiated in the 30 days prior to September 15, 2008 (the bankruptcy of Lehman Brothers), lost on average 290 bps, and trades initiated in the subsequent 30 days gained on average 189 bps. Similarly, trades initiated in the month before November 21, 2011 (roughly the peak of the European debt crisis, although this is difficult to date precisely), lost on average 29 bps, and traded initiated in the subsequent 30 days gained on average 63 bps.

In summary, our strategy experiences its largest losses on trades initiated right before the peak of financial turmoil, but experiences its largest gains on trades initiated in the immediate aftermath of that peak. We view these patterns as consistent with our interpretation that the forward CIP trading strategy returns capture unexpected tightening or loosening of financial constraints.

## 3. Forward CIP Trading Strategy's Excess Returns

In this section, we present evidence that the forward CIP trading strategy is profitable on average. The excess returns are observed in certain individual currencies against the USD, in trades between cross-currency pairs, and in portfolios. We find that the currency pairs with the highest excess returns are the currency pairs associated with the "FX carry trade." These currency pairs have high interest rate differentials, large CIP violations, and unhedged currency returns that are positively correlated with returns on the S&P 500 index. We also study a 1-day forward CIP arbitrage to examine balance sheet constraints on quarter-end regulatory reporting dates.

### 3.1 USD-based currency pairs

We begin by discussing results for individual currencies. Panel A of Table 1 reports the profits per dollar notional on the 1-month forward 3-month tenor forward CIP trading strategy in each of the six sample currencies vis-à-vis the USD. For each forward CIP trading strategy, we present the annualized mean profit per dollar notional and the Sharpe ratio, by period. Standard errors of the statistics are reported in parentheses.[24]

Beginning with the pre-GFC period, we observe that for all sample currencies vis-à-vis the USD, the pre-GFC profits are very close zero. Post-GFC, the profits in most currencies are larger in absolute value. Some currencies, such as JPY, have marginally statistically significant Sharpe ratios in the pre-GFC

[24] Means and Sharpe ratios are calculated using overlapping monthly profits per dollar notional from daily data and then scaled up by 12 and $\sqrt{12}$, respectively. We use Newey-West standard errors and the Newey and West 1994 bandwidth selection procedure, and use the "delta" method to compute standard errors for the Sharpe ratios Lo 2002. Internet Appendix Table A7 presents for robustness virtually identical results for portfolios (Table 3) using nonoverlapping monthly data.

**Table 1**
**Forward CIP trading profits and additional properties for USD-based currency pairs**

*A. Summary statistics of returns on OIS 1-month fwd 3-month forward CIP trading strategy*

| | Mean | | | Sharpe ratio | | |
|---|---|---|---|---|---|---|
| | Pre-GFC | GFC | Post-GFC | Pre-GFC | GFC | Post-GFC |
| AUD_USD | 0.98 | 4.74 | 7.45*** | 0.26 | 0.12 | 1.24*** |
| | (0.98) | (19.76) | (1.81) | (0.27) | (0.51) | (0.31) |
| CAD_USD | −0.76 | 1.79 | 5.55*** | −0.21 | 0.06 | 0.90*** |
| | (1.59) | (15.54) | (1.62) | (0.44) | (0.48) | (0.27) |
| GBP_USD | −1.50** | 9.36 | 4.62*** | −0.50* | 0.28 | 0.68*** |
| | (0.73) | (15.67) | (1.72) | (0.26) | (0.41) | (0.22) |
| EUR_USD | −1.10* | 13.68 | −1.73 | −0.60* | 0.33 | −0.17 |
| | (0.59) | (19.47) | (2.56) | (0.31) | (0.40) | (0.27) |
| CHF_USD | −1.37 | 5.33 | −3.59 | −0.36 | 0.13 | −0.25 |
| | (0.87) | (17.43) | (4.78) | (0.27) | (0.40) | (0.36) |
| JPY_USD | −1.88* | 8.92 | −8.79** | −0.79* | 0.18 | −0.63* |
| | (0.97) | (22.98) | (3.57) | (0.41) | (0.44) | (0.35) |

*B. Post-GFC properties*

| | Mean fwd CIP Trad. ret. | Average basis | Average slope | Avg. Interest diff. vs. USD | Corr SPX and FX |
|---|---|---|---|---|---|
| AUD_USD | 7.45 | 9.43 | 6.37 | 1.63 | 0.56 |
| CAD_USD | 5.55 | −12.34 | 5.46 | 0.33 | 0.53 |
| GBP_USD | 4.62 | −15.99 | 4.65 | −0.19 | 0.36 |
| EUR_USD | −1.73 | −34.08 | −1.24 | −0.67 | 0.16 |
| CHF_USD | −3.59 | −51.84 | −3.55 | −0.85 | −0.03 |
| JPY_USD | −8.79 | −44.89 | −7.86 | −0.62 | −0.27 |

Panel A of this table reports the annual profits and annualized Sharpe ratios from the OIS 1-month forward 3-month forward CIP trading strategy vis-à-vis the USD. Pre-GFC is January 1, 2003, to June 30, 2007; GFC is July 1, 2007, to June 30, 2010; and Post-GFC is July 1, 2010, to December 31, 2020 (CHF ends December 31, 2017). Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West 1994 selection procedure. Panel B of this table reports additional characteristics the USD-based currency pairs post-GFC. "Mean fwd CIP trad. ret." is the annualized profit from the forward CIP trading strategy; "Average basis" is the average spot 3-month OIS cross-currency basis; "Average slope" is the average spread between the 1-month forward 3-month and spot 3-month OIS cross-currency basis; "Avg. int. rate diff." is the average spread between the 3-month foreign OIS rate and the U.S. OIS rate; and "Corr SPX and FX" is the correlation between weekly currency returns of going long the foreign currency and shorting the USD and the weekly returns on the S&P 500 index. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

period, a fact that reflects small mean profits and even smaller standard deviations. In contrast, post-GFC, four currencies vis-à-vis USD have nontrivial mean profits and both statistically and economically significant Sharpe ratios.

Panel B of Table 1 illustrates that the sign of a currency's forward CIP trading profits (vis-à-vis the USD) is related to a number of other economically important properties of the currency. AUD, CAD, and GBP have positive forward CIP arbitrage profits, while EUR, CHF, and JPY have negative forward CIP arbitrage profits. We define the "investment" and "funding" currency based on the average interest rate of the foreign currency vis-à-vis the U.S. dollar in our sample, consistent with the standard FX carry trade literature, such as in Lustig and Verdelhan 2007, Lustig, Roussanov, and Verdelhan 2011. The former group (AUD, CAD, and GBP) are high-interest-rate "investing currencies" and the latter group (EUR, CHF, and JPY) are low-interest-rate "funding currencies" for the unhedged FX carry trade.[25] In bad times (proxied by low S&P 500 returns), the "funding currencies" tend to appreciate against the USD, while the "investing currencies" depreciate. CIP deviations make "funding currencies" more appealing in terms of their synthetic dollar interest rates. These currencies have substantial negative cross-currency bases (higher synthetic dollar interest rates), whereas "investing currencies" have less negative or even positive cross-currency bases vis-à-vis USD.

Moreover, the AUD, CAD, and GBP all have an upward-sloping CIP term structure on average.[26] In contrast, EUR, CHF, and JPY have a downward-sloping CIP term structure on average. Put another way, the increases in the absolute value of the basis implied by the forward curves do not actually occur, on average. This result is analogous to the existence of the term premium in the term structure literature.[27]

As mentioned in Section 1, OIS rates are not directly comparable across currencies. For example, the CAD OIS rate is collateralized, whereas most other rates are not. As another example, the CHF OIS rate changed by around 20 bps as part of an overnight benchmark rate reform in 2017. If we used the new index instead of the old index to compute the spot basis, the CHF-USD basis would be less negative by about 20 bps.[28] In USD, institutional factors cause the OIS rate (fed funds) to fall below the rate on excess reserves (Bech and Klee 2011), whereas the EONIA rate generally exceeds the ECB deposit rate. For these reasons, we do not view (for example) the sign of the CAD-USD basis or the exact ranking of the bases in B of Table 1 as particularly meaningful. Instead, we note that the investing currencies and funding currencies are notably different across all five of the dimensions considered in panel B of Table 1.

### 3.2 Currency pairs with the largest bases

Our model suggests that intermediaries will actively trade the bases with the maximal arbitrage per unit risk weight. If risk weights are roughly equal across currencies, this suggests focusing on the currency pairs with the largest bases. As discussed earlier ("Heterogeneity across currencies" in Section 1), for a variety of reasons we do not take a stand on which currency pair truly has the largest basis but instead present results for the 10 currency pairs with the largest

---

25 We group GBP with "investment" currencies, even though the average GBP interest rate is slightly below the U.S. rate in the post-GFC sample, because GBP interest rates are generally higher than U.S. interest rates over longer samples (e.g., post-2000). The low average GBP interest rate in the recent sample is driven primarily by low rates in the post-Brexit period.

26 We define slope as the difference between the 1-month forward 3-month basis and the spot 3-month basis.

27 Taking this analogy further, Internet Appendix E provides suggestive evidence that the slope of the forward curve predicts forward CIP trading profits, just as the slope of the term structure predicts bond returns Campbell and Shiller 1991. However, the statistical significance is sensitive to the inclusion of the COVID-19 crisis period in the sample.

28 As mentioned previously, swaps on the new CHF OIS index are not liquid enough for our purposes, which is why we present results with the old index.

**Table 2**
**Summary statistics of currency-pair returns on OIS 1-month forward 3-month forward CIP trading strategy**

| | Mean | | | Sharpe ratio | | | Post-GFC metrics | |
|---|---|---|---|---|---|---|---|---|
| | Pre-GFC | GFC | Post-GFC | Pre-GFC | GFC | Post-GFC | Avg basis | Int. rate diff. |
| AUD_CHF | 2.44** | −0.73 | 9.85* | 0.49* | −0.03 | 0.65 | 60.22 | 2.47 |
| | (1.07) | (10.33) | (5.25) | (0.26) | (0.37) | (0.41) | | |
| AUD_JPY | 2.36* | −3.20 | 16.25*** | 0.58* | −0.12 | 1.18*** | 54.32 | 2.25 |
| | (1.38) | (9.78) | (3.57) | (0.35) | (0.37) | (0.39) | | |
| USD_CHF | 1.37 | −5.33 | 3.59 | 0.36 | −0.13 | 0.25 | 51.84 | 0.85 |
| | (0.87) | (17.43) | (4.78) | (0.27) | (0.40) | (0.36) | | |
| USD_JPY | 1.88* | −8.92 | 8.79** | 0.79* | −0.18 | 0.63* | 44.89 | 0.62 |
| | (0.97) | (22.98) | (3.57) | (0.41) | (0.44) | (0.35) | | |
| CAD_CHF | 1.09 | −4.80 | 9.02** | 0.21 | −0.17 | 0.66 | 44.22 | 1.18 |
| | (1.85) | (8.17) | (4.48) | (0.40) | (0.28) | (0.41) | | |
| AUD_EUR | 2.12** | −8.91 | 9.18*** | 0.58** | −0.39 | 0.89*** | 43.51 | 2.30 |
| | (0.88) | (8.30) | (2.76) | (0.25) | (0.34) | (0.30) | | |
| GBP_CHF | 0.20 | 3.77 | 8.02* | 0.04 | 0.17 | 0.60 | 38.48 | 0.66 |
| | (0.94) | (6.92) | (4.35) | (0.21) | (0.32) | (0.40) | | |
| USD_EUR | 1.10* | −13.68 | 1.73 | 0.60* | −0.33 | 0.17 | 34.08 | 0.67 |
| | (0.59) | (19.47) | (2.56) | (0.31) | (0.40) | (0.27) | | |
| CAD_JPY | 0.04 | −7.19 | 14.26*** | 0.01 | −0.24 | 1.20*** | 32.55 | 0.96 |
| | (1.40) | (11.13) | (2.86) | (0.36) | (0.36) | (0.40) | | |
| GBP_JPY | 0.02 | 1.41 | 13.40*** | 0.00 | 0.06 | 1.32*** | 28.91 | 0.43 |
| | (0.88) | (9.78) | (2.66) | (0.27) | (0.41) | (0.39) | | |

This table reports the annual profits and annualized Sharpe ratios from the OIS 1-month forward 3-month forward CIP trading strategy for the 10 currency pairs with the largest average Post-GFC spot 3-month bases. Pre-GFC is January 1, 2003, to June 30, 2007; GFC is July 1, 2007, to June 30, 2010; and Post-GFC is July 1, 2010, to December 31, 2020 (CHF ends December 31, 2017). "Avg. basis" is the average spot 3-month OIS cross-currency basis post-GFC, and "Int. rate diff." is the average spread between the 3-month foreign OIS rate and the U.S. OIS rate post-GFC. Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West 1994 selection procedure. $*p<.1$; $**p<.05$; $***p<.01$.

bases. In the context of our model, if intermediaries actively trade all of these bases, the differences in the magnitude of the basis across pairs must be due to either measurement issues or differences in risk weights. In both of these cases, we would expect positive expected returns and Sharpe ratios from our forward arbitrage strategy for all 10 bases.[29]

Table 2 presents the forward CIP returns associated with the ten currency pairs with the largest average spot 3-month bases post-GFC.[30] The mean returns of these currency pairs are linear combinations of the mean returns for each currency leg vis-à-vis USD presented earlier; the Sharpe ratios are not. The mean average profits are positive for all 10 pairs post-GFC, and the annualized Sharpe ratio is above 0.5 for 8 of 10 currency pairs. The average interest rate differential for each of these 10 pairs is positive, once again suggesting a relationship between carry and the profits of the forward CIP trading strategy.

[29] If each of the 10 forward arbitrage strategies were an exact (noiseless) bet on the size of the shadow cost $\lambda_{t+1}^{RC}$, scaled by the risk weight $k^i$, they would all have positive expected returns and the same Sharpe ratios. In the presence of currency-specific noise, the Sharpe ratio will be attenuated for each currency pair based on the magnitude of the noise.

[30] These 10 pairs also all have a positive spot 3M basis on virtually every day in our sample.

Consider as an example the “classic carry” currency pair of long AUD, short JPY. This pair has one of the largest spot bases, and is particularly associated with the carry trade.[31] The AUD-JPY forward CIP trading strategy earns an a post-GFC average profit equal to 16 bps and its annualized Sharpe ratio is 1.18. Both results are highly statistically significant, and the magnitude of the Sharpe ratio is high compared to many documented trading strategy returns in the literature. For comparison, the traditional un-hedged FX carry trade has an annualized Sharpe ratio of 0.48 for developed market currencies from 1987 to 2009, and the annualized Sharpe ratio of a value-weighted portfolio of all U.S. stocks from 1976 to 2010 is 0.42 (Burnside et al. 2010, Burnside, Eichenbaum, and Rebelo 2011). Note, however, that our analysis is limited to the post-GFC period, which is a short sample. During this period, the developed market carry trade and U.S. stock portfolio had annualized Sharpe ratios (0.18 and 1.05, respectively) that are not representative of the longer sample.[32]

The connection between interest rate differentials and the spot cross-currency basis was documented in Du, Tepper, and Verdelhan 2018. As discussed in the survey of Du and Schreger 2021, CIP deviations are induced by the interaction between steady demand for funding and hedging services and intermediary constraints. The funding and hedging demand for high-interest-rate currencies is particularly strong from low-interest-rate countries, as the search-for-yield investors demand assets denominated in high-interest-rate currencies. One implication of that story, through the lens of our model, is that the risk that the classic carry basis becomes larger is priced because it correlates with intermediary constraints more broadly. Consistent with this, our results reveal a relationship between the spot basis, interest rate differentials, and forward CIP trading profits.

### 3.3 Portfolios of forward CIP trading strategies

In addition to the “classic carry” AUD-JPY currency pair, we examine five portfolios of forward CIP trading profits: “dollar-neutral carry”, “dynamic top-five basis”, “static top-six basis”, “top-six first PC”, and “dollar”. The first four of these are based on interest rate differentials and the size of the spot cross-currency basis; they generate positive mean returns and high Sharpe ratios post-GFC. In contrast, the return for the dollar portfolio is insignificant. As mentioned previously, all of these portfolios exclude CHF because of data limitations.

The portfolios are defined as follows. The “dollar-neutral carry” portfolio is a dollar-neutral carry strategy. The portfolio goes long in the forward CIP trading

31 AUD-CHF has a slightly larger average spot basis than AUD-JPY. However, as mentioned previously, because of the significant problems with the CHF OIS index, our CHF sample ends in 2017.

32 That said, the correlation of the AUD-JPY forward CIP return to the U.S. market is not very high (about 0.3 monthly post-GFC), so we have no particular reason to think that whatever forces caused the market to have high returns post-GFC also influenced the forward CIP return.

strategy for the AUD, CAD, and GBP vis-à-vis the USD, each with one-third weight, and short in the forward CIP trading strategy for the EUR and JPY vis-à-vis the USD, each with one-half weight. The "dynamic top-five basis" portfolio equally weights the forward CIP trading strategies for the largest five spot cross-currency basis pairs and is rebalanced monthly. The "static top-six basis" portfolio equally weights the forward CIP trading strategies in the six non-CHF currency pairs listed in Table 2, which were selected based on the average spot 3-month cross-currency basis in the post-GFC sample. The "top-six first PC" portfolio is the first principal component of those same six pairs. The sign and scale of the first PC (which are arbitrary) are chosen to match the volatility of the AUD-JPY currency pair and to ensure a positive correlation between the first PC and AUD-JPY.[33] The "dollar" portfolio places equal weights on each of the individual non-CHF sample currencies vis-à-vis the USD. All of these portfolios, with the exception of the dollar portfolio, are have highly correlated returns (the lowest pairwise correlation is 0.93 in our nonoverlapping monthly sample).

We report the annualized mean profit and the Sharpe ratio for these five portfolios, together with the performance for the "classic carry" strategy in Table 3. The pre-GFC mean profits of these portfolios are all close to zero. The post-GFC mean profits are significantly positive at about 11 to 16 bps for all of the portfolios except the "dollar" portfolio, with significant annualized Sharpe ratios between 1 and 1.3. In contrast, the post-GFC profits and the Sharpe ratio for the "dollar" portfolio remain close to zero. In robustness checks, we show similar patterns hold for 1-month tenors, 3-month horizons, and for strategies based on IBOR bases (Internet Appendix Tables A3, A4, A5, and A6).[34]

### 3.4 Quarter-ends

Quarter-ends offer an interesting window to examine forward CIP deviations and the profits of our forward trading strategy. As documented in Du, Tepper, and Verdelhan 2018, there are large spikes in short-term CIP deviations for contracts that cross the quarter-ends. Those authors attribute the quarter-end spikes to the fact that the Basel III leverage ratio is calculated using quarter-end snapshots of bank balance sheets in many non-U.S. jurisdictions, tightening leverage constraints for intermediaries on quarter-ends. We show an additional

[33] The weights are roughly AUD-JPY 0.26, AUD-EUR 0.16, USD-JPY 0.27, CAD-JPY 0.22, USD-EUR 0.17, and GBP-JPY 0.18. The sum of the weights exceeds one because AUD-JPY is more volatile than most other currency pairs; matching the volatility of AUD-JPY facilitates a comparison between results with the PC1 and results with AUD-JPY. We have constructed versions that use weights based on the first principal component of the spot bases, as opposed to of the forward CIP returns, as well as versions that use all pairwise combinations of forward CIP returns, and not just the six listed in Table 2. The resultant portfolios have virtually identical returns. The "dynamic top-five" portfolio is tradable as it is formed based on ex ante available interest rate information, whereas the "static top-six basis" and "top-six first PC" are constructed using the full sample. The choice of five versus six pairs is arbitrary and has minimal impact on the results.

[34] Our IBOR data has missing data in CAD, but not CHF; as a result, our IBOR portfolios differ from our OIS portfolios because they include CHF, but not CAD.

**Table 3**
**Summary statistics of portfolio returns on OIS 1-month forward 3-month forward CIP trading strategy**

| | Mean | | | Sharpe ratio | | |
|---|---|---|---|---|---|---|
| | Pre-GFC | GFC | Post-GFC | Pre-GFC | GFC | Post-GFC |
| Classic carry (AUD-JPY) | 2.36* | −3.20 | 16.25*** | 0.58* | −0.12 | 1.18*** |
| | (1.38) | (9.78) | (3.57) | (0.35) | (0.37) | (0.39) |
| Dollar-neutral carry | 0.16 | −5.84 | 11.16*** | 0.07 | −0.30 | 1.27*** |
| | (0.83) | (7.56) | (2.29) | (0.37) | (0.36) | (0.40) |
| Dynamic top-five basis | 0.73 | −5.95 | 11.27*** | 0.32 | −0.20 | 1.03*** |
| | (0.84) | (12.45) | (2.80) | (0.37) | (0.38) | (0.39) |
| Static top-six basis | 0.77 | −6.63 | 10.65*** | 0.36 | −0.24 | 1.00** |
| | (0.78) | (11.46) | (2.74) | (0.36) | (0.38) | (0.39) |
| Simple dollar | −0.82 | 7.47 | 1.41 | −0.44 | 0.20 | 0.20 |
| | (1.15) | (18.15) | (1.88) | (0.60) | (0.46) | (0.25) |
| Top-six first PC | 1.06 | −8.00 | 13.87*** | 0.38 | −0.23 | 1.00** |
| | (1.04) | (14.83) | (3.55) | (0.37) | (0.39) | (0.39) |

This table reports the annual profits and annualized Sharpe ratios from the OIS 1-month forward 3-month forward CIP trading strategy. All statistics are reported by period: Pre-GFC is January 1, 2003, to June 30, 2007; GFC is July 1, 2007, to June 30, 2010; and Post-GFC is July 1, 2010, to December 31, 2020. The "Classic carry" portfolio is the forward CIP trading strategy for the AUD-JPY pair. The "Dollar-neutral carry" portfolio longs the forward CIP trading strategy in AUD, CAD, and GBP and shorts the forward CIP trading strategy in JPY and EUR. The "Dynamic top-five basis" portfolio has equal weight in the 5 currency pairs that exhibit the highest spot 3-month basis, rebalanced monthly. The "Static top-six basis" portfolio has equal weights in the six non-CHF currency pairs with the largest average spot 3-month basis post-GFC shown in Table 2. The "Simple dollar" portfolio puts equal weights on the forward CIP trading strategy for all non-CHF sample currencies vis-à-vis the USD. The "Top-six first PC" portfolio is the first principal component of the six non-CHF currency pair returns with the largest average spot 3-month basis post-GFC shown in Table 2. Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West 1994 selection procedure. *$p < .1$; **$p < .05$; ***$p < .01$.

premium in our forward trading strategy associated with the quarter-end turn. To the extent that the quarter-end effect in the spot CIP deviation offers clean evidence on the effects of leverage constraints, our finding of a positive risk premium associated with the quarter-end turn offers additional support for the idea that leverage constraints are priced.

We have thus far side-stepped the issue of quarter-ends in our analysis by studying forward CIP trading strategies with a 3-month tenor, ensuring that the contracts in question always cross quarter-end. In this subsection, we instead study a forward CIP trading strategy that uses tenors of a single business day and focus on quarter-end effects. One advantage of examining 1-day forward trading strategy is that we have many more nonoverlapping daily observations to calculate the mean returns over the post-GFC sample.

We follow Correa, Du, and Liao 2020 and construct overnight (ON) and tomorrow/next (TN) CIP deviations. These are constructed from central bank deposit rates as opposed to OIS rates. The ON CIP deviation is a 1-day spot CIP violation; the TN CIP deviation is a 1-day forward starting 1-day CIP violation. From these, we can construct a one-day forward CIP trading strategy by betting on whether the TN CIP deviation traded at time $t$ is larger than the subsequent realized spot ON CIP deviation traded at $t+1$. We provide details on ON and TN basis calculations in Internet Appendix F.

In Table 4, we regress the annualized ON CIP deviation, 1-day-lagged TN CIP deviation, and the profit on the 1-day forward CIP trade on a constant

**Table 4**
**Returns on ON-TN forward CIP trade and quarter-ends**

| | (1) ON basis | (2) Lagged TN basis | (3) ON-TN forward profit |
|---|---|---|---|
| QE dummy | 154.1*** | 207.4*** | 53.25** |
| | (32.77) | (34.83) | (25.17) |
| Constant | 10.98*** | 15.17*** | 4.194*** |
| | (0.580) | (0.545) | (0.537) |
| Observations | 4,953 | 4,953 | 4,953 |
| $R$-squared | .112 | .186 | .020 |

This table reports regression results for the overnight CIP deviations (column 1), 1-day-lagged tomorrow/next CIP deviations (column 2) and the return on the ON-TN forward CIP trade (column 3), or the difference between columns 2 and 3. The independent variable is a quarter-end (QE) dummy, which is equal to one if the date is equal to the last business date of the quarter. The sample currencies include CHF, EUR, and JPY. The CIP deviations are calculated as the difference between swapped foreign central bank deposit rate into U.S. dollars and the U.S. interest rate on excess reserves. The sample period is post-GFC from July 1, 2010, to December 31, 2020. Robust standard errors are reported in the parentheses. See Internet Appendix F for details about the ON and TN CIP deviations.

and quarter-end dummy, pooled across the funding currencies (CHF, EUR and JPY) vis-à-vis the USD.[35] Column 1 shows that the ON basis averages to 11 bps and jumps by 154 bps on average when crossing quarter-ends. Column 2 shows that the TN basis averages to 15 bps and jumps by 207 bps on average 1 (business) day before the quarter-ends. The large jump in the TN basis right before quarter-ends suggests that the quarter-end effects are anticipated and in part priced into forward CIP deviations. Since the average TN basis is higher than the average ON basis and the quarter-end jump in the TN premium is on average larger than ON premium, there is a positive average profit for our 1-day forward CIP trading strategy and an additional positive risk premium crossing the quarter-ends. In column 3, we can see that the average profit on the ON-TN forward CIP trading strategy outside is about 4 bps, and about 53 bps higher on quarter-ends.

In Internet Appendix Figure A4, we show the shape of the forward curve of the 1-month tenor AUD-JPY basis with three subsamples based on whether the next quarter-end is within the next month, between 1 and 2 months in the future, or more than 2 months in the future. We observe that all three lines exhibit a spike, precisely when the interest tenor in the basis crosses the quarter end. This further illustrates that intermediary's quarter-end constraints are anticipated and priced. However, we cannot detect additional risk premium for the 1-month tenor forward associated with quarter-end crossings in Internet Appendix Table A12, likely because of limited power when using the longer-tenor contracts in detecting the risk premium associated with quarter-ends.

35 Because of our limited sample of quarter-ends, pooling across currencies is necessary to precisely estimate quarter-end effects. We focus on the funding currencies vis-à-vis the USD because of data availability and the relationship between forward CIP returns and the carry trade documented thus far. Using central bank deposit rates allows us to include CHF.

**Table 5**
**Pre-/Post-COVID-19 statistics of portfolio returns on OIS 1-month forward 3-month forward CIP trading strategy**

| | Mean | | | Sharpe Ratio | | |
|---|---|---|---|---|---|---|
| | Pre-2020 | Post-2020 | Overall | Pre-2020 | Post-2020 | Overall |
| Classic carry (AUD-JPY) | 15.82*** | 20.72 | 16.25*** | 1.58*** | 0.62 | 1.18*** |
| | (3.01) | (25.80) | (3.57) | (0.33) | (0.89) | (0.39) |
| Dollar-neutral carry | 11.15*** | 11.29 | 11.16*** | 1.72*** | 0.54 | 1.27*** |
| | (2.04) | (15.14) | (2.29) | (0.37) | (0.81) | (0.40) |
| Dynamic top-five basis | 11.33*** | 10.68 | 11.27*** | 1.51*** | 0.38 | 1.03*** |
| | (2.32) | (20.81) | (2.80) | (0.34) | (0.82) | (0.39) |
| Static top-six basis | 10.76*** | 9.54 | 10.65*** | 1.48*** | 0.35 | 1.00** |
| | (2.27) | (20.43) | (2.74) | (0.33) | (0.82) | (0.39) |
| Simple dollar | 0.87 | 6.99 | 1.41 | 0.19 | 0.38 | 0.20 |
| | (1.49) | (14.57) | (1.88) | (0.32) | (0.71) | (0.25) |
| Top-six first PC | 14.02*** | 12.31 | 13.87*** | 1.50*** | 0.35 | 1.00** |
| | (2.92) | (26.77) | (3.55) | (0.33) | (0.83) | (0.39) |

This table reports the annual profits and annualized Sharpe ratios from the OIS 1-month forward 3-month forward CIP trading strategy by subsamples. All statistics are reported by period: Pre-2020 is January 1, 2003, to December 31, 2019; Post-2020 is January 1, 2020, to December 31, 2020; and the overall is the full sample period. The "Classic carry" portfolio is the forward CIP trading strategy for the AUD-JPY pair. The "Dollar-neutral carry" portfolio longs the forward CIP trading strategy in AUD, CAD, and GBP and shorts the forward CIP trading strategy in JPY and EUR. The "Dynamic top-five basis" portfolio has equal weight in the five currency pairs that exhibit the highest spot 3-month basis, rebalanced monthly. The "Static top-six basis" portfolio has equal weights in the six non-CHF currency pairs with the largest average spot 3-month basis post-GFC shown in Table 2. The "Simple dollar" portfolio puts equal weights on the forward CIP trading strategy for all non-CHF sample currencies vis-à-vis the USD. The "Top-six first PC" portfolio is the first principal component of the six non-CHF currency pair returns with the largest average spot 3-month basis post-GFC shown in Table 2. Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West 1994 selection procedure. *$p$ <.1; **$p$ <.05; ***$p$ <.01.

### 3.5 Additional subsample analysis and transaction costs

**COVID-19** Our subsample analysis in Table 5 presents results with and without the market turmoil induced by the COVID-19 pandemic. The average returns of our portfolios are similar in the two subsamples.

As discussed in Section 2.4, the AUD-JPY forward CIP trading strategy experienced large losses at the onset of the pandemic, and subsequently experienced large gains as the turmoil subsided. In total, the trading strategy earned roughly average profits during 2020 relative to previous years, despite the COVID-19 crisis. A similar pattern appears for other portfolios: the large losses incurred on trades initiated in February 2020 were offset by large gains over the next few months, and the overall returns for 2020 are similar to the average of prior years. The Sharpe ratios for the sample including 2020 are slightly lower than the sample ending in 2019 due to the higher volatility of returns in 2020.

Consistent with the idea that CIP violations measure the tightness of constraints, our strategy experiences negative profits at the beginning of financial stress. However, during periods of financial stress, the risk premiums earned by our strategy are particularly large, which is to say that the risk that constraints tighten further carries a large price during these episodes. As a result, sample periods that include both the onset and resolution of financial stress have an ambiguous effect on the average returns of our strategy.

**Basel III** Our post-GFC sample begins in 2010, which marks the beginning of the lengthy process for the introduction and implementation of Basel III banking regulations. These regulations substantially affected banks ability to engage in balance-sheet intensive activities.[36] Notably, the public disclosure of the Basel III leverage ratio starts on January 1, 2015. Given that the leverage ratio requirement under Basel III is the most relevant regulatory constraint for CIP arbitrage, we split the post-GFC period into pre- and post-2015. Internet Appendix Table A8 shows that our main results are robust in both subsamples post-GFC, with mean returns that are perhaps slightly larger in the post-2015 subsample. The Sharpe ratios in the post-2015 sample are smaller because of the effects of the COVID-19 crisis discussed above.

**Transactions costs** We have limited data on the transactions costs associated with implementing the forward CIP trading strategy. Large intermediaries are likely to implement the strategy at low cost (either collecting the bid-offer when trading with clients or trading at close to the mid-price in interdealer transactions). Anecdotal evidence suggests that some large hedge funds use interest rate and FX derivatives to arbitrage the term structure of CIP violations, suggesting that the transaction costs are not prohibitively large.

We study these forward CIP trading strategies because they reveal interesting information about currencies and intermediaries, and not because we advocate them as an investment strategy. It may well be the case that a typical trader in a small hedge fund paying the bid-offer on the various instruments used to implement the trading strategy would not find it profitable. We provide a conservative estimate of transaction costs by assuming the full quoted bid-offer spreads from Bloomberg are paid on every single instrument involved in the trading strategy. Note this approach likely substantially overstates total transaction costs as our trade can be easily structured as an asset package. Taking the USD-JPY OIS trade as an example, the annualized transaction costs based on full Bloomberg bid-offer spreads are almost three times the size of the annualized profit for the 1-month horizon forward trade and roughly equal to the profit for the 3-month horizon forward trade (Internet Appendix Table A13).

## 4. Implications for the Price of Risk

In the preceding section, we found a substantial risk premium associated with the risk that AUD-JPY and other bases become larger. We interpret this, through the lens of our model, to mean that this basis is a measure of intermediary constraints and that the risk that intermediary constraints tighten is a priced risk factor.

---

36 Besides CIP deviations, higher balance sheet costs are also manifested in the repo market. In Internet Appendix Figure A5, we show that the average gross repo position of primary dealers declined by more than 40% in the post-GFC sample compared to their pre-GFC peak, which is consistent with more binding non-risk-weighted leverage constraints. The spread between large banks' lending rate and borrowing rate in repo markets also widened significantly post-GFC.

This interpretation has several implications that we explore in this section. First, it suggests that the level of the spot CIP basis should be correlated with other arbitrages affected by constraints on intermediaries. Second, it suggests that the basis should be correlated with measure of intermediary wealth and other measures of intermediary constraints. Third, it implies (assuming an intertemporal hedging motive) that the forward CIP risk premium should exist even after controlling for intermediary wealth, and that the forward CIP return should be included along with intermediary wealth in the SDF. Fourth, the forward CIP risk premium should be consistent with the prices of risk extracted from other assets that intermediaries trade. We explore each of these implications in turn.

### 4.1 CIP versus other arbitrages

We interpret the AUD-JPY basis and other CIP deviations as measures of intermediary constraints. However, our model implies that intermediary constraints, if present, should affect many no-arbitrage relationships and not just CIP. To verify that the bases we study are indeed measures of intermediary constraints, we begin by confirming that they comove with other documented near-arbitrages. Specifically, we show that the AUD-JPY cross-currency basis comoves with the first principal component of other near-arbitrages from outside the FX market.

We consider seven types of near-arbitrages: the bond-CDS basis, the CDS-CDX basis, the USD Libor tenor basis, 30-year swap spreads, the Refco-Treasury spread, the KfW-Bund spread, and the asset-swapped TIPS/Treasury spread. Bai and Collin-Dufresne 2019, Boyarchenko et al. 2018, Fleckenstein, Longstaff, and Lustig 2014, Jermann 2019, Longstaff 2004, Schwarz 2018, among others, have examined these near-arbitrages in the recent literature. We describe these near-arbitrages in more detail in Internet Appendix G.

Each of these near-arbitrages is subject to measurement errors and idiosyncratic supply and demand shocks. We use a principal component analysis to extract the common component. Our model implies that variation in the balance sheet capacity of financial intermediaries should affect all these near-arbitrages, and we therefore view this common component as an alternative measure of intermediary constraints.

Figure 4 shows that our benchmark AUD-JPY cross-currency basis and the first principal component (PC) of the other near-arbitrages follow broadly similar trends. We find that the first PC explains 51% of total variation in the level of the seven near-arbitrages between January 2005 and December 2020, and has a 41% correlation with the level of the AUD-JPY cross-currency basis post-GFC.[37]

[37] All seven near-arbitrages are long term (5 years or longer), while the cross-currency basis we use has a 3-month tenor, so the correlation between the two series should not be perfect.

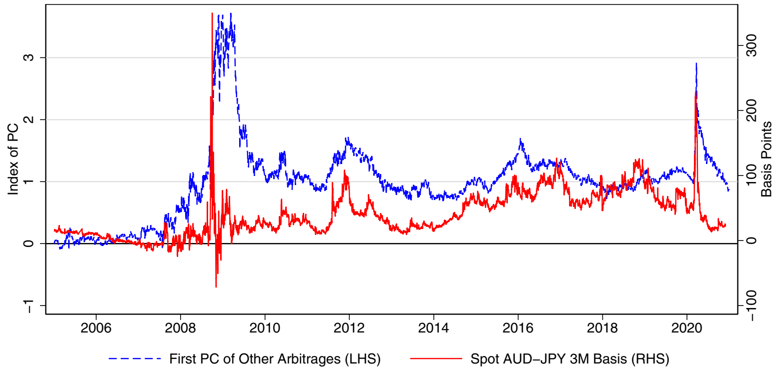


**Figure 4**
**Cross-currency basis and other near-arbitrages**
This figure plots the daily spot 3-month AUD-JPY cross-currency basis and the scaled first principal component of seven other near-arbitrages: the bond-CDS basis, the CDS-CDX basis, the U.S. Libor tenor basis, the 30-year Treasury-swap spread, the Refco-Treasury spread, the KfW-Bund spread, and the TIPS-Treasury spread. See Internet Appendix G for details about these other near-arbitrages.

However, CIP deviations have several advantages over these other near-arbitrages. First, they are close to true arbitrages, unlike some of these other measures. For example, the Libor tenor basis and Treasury swap spread may reflect credit risk rather than intermediary constraints, while the bond-CDS basis has a cheapest-to-deliver option and other complications. Second, they are precisely measured, exhibiting less high frequency volatility than most of these other measures. Third, and most importantly for our empirical strategy, they have a rich term structure that allows us to construct our forward CIP trading strategy. For these reasons, we use CIP violations as our as our preferred measure of intermediary constraints.

## 4.2 CIP versus existing measures of intermediary constraints

If the cross-currency basis measures intermediary constraints, then the general equilibrium models of He and Krishnamurthy 2011 and Kondor and Vayanos 2019 imply that it should comove with intermediary net worth. However, demand and regulatory shocks should also affect the tightness of intermediary constraints, so we do not expect a perfect correlation. Similarly, Adrian, Etula, and Muir 2014 argue for broker-dealer leverage as an alternative measure of intermediary constraints. Of course, if broker dealers are subject to leverage constraints, the tightness of those constraints might change without leverage changing (e.g., in response to changing investment opportunities). Again, for this reason, we do not expect a perfect correlation between our measure of the tightness of intermediary constraints and broker-dealer leverage.

We explore the relationship between existing intermediary constraint measures and the AUD-JPY cross-currency basis in Figure 5. We use as our primary measure of intermediary wealth the intermediary equity value of

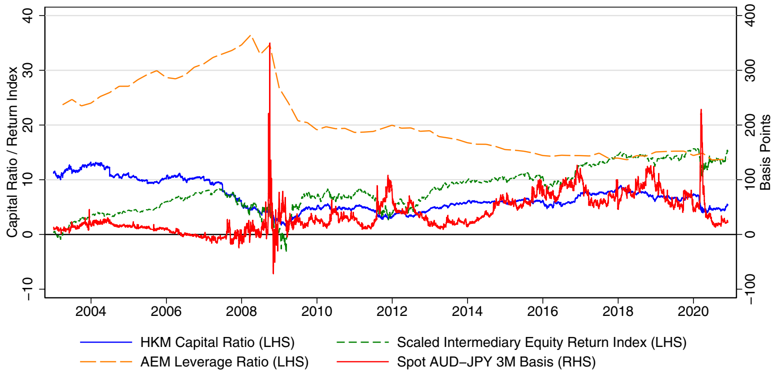


**Figure 5**
**Cross-currency basis and intermediary wealth**
This figure plots the monthly spot 3-month AUD-JPY cross-currency basis and measures of intermediary wealth and constraints from 2003 to 2020. The HKM Capital Ratio (in percent) is the equity capitalization ratio of the primary dealer. The cumulative intermediary equity return is based on the value-weighted return of the equity of primary dealers calculated from January 2003 and scaled by 10. The AEM leverage ratio (in percent) is calculated as the ratio of book assets to book equity for the broker-dealer sector from the Flow of Funds.

He, Kelly, and Manela 2017 (henceforth HKM), which is the cumulative return of value-weighted equity of primary dealers. We also consider, following He, Kelly, and Manela 2017, the equity capitalization ratio of the dealers (HKM Capital Ratio) and the broker-dealer leverage ratio used in Adrian, Etula, and Muir 2014 (defined as broker dealer book asset over book equity).

Figure 5 presents a time series of the spot 3M AUD-JPY basis and these intermediary wealth and constraint measures. The cross-currency basis and the proxies for intermediary wealth measures appear to be (negatively) correlated. This suggests that variations in the spot basis are in part driven by shocks to intermediary wealth. However, in recent years, we observe an upward trend in the basis, likely attributable to changes in regulation (e.g., the implementation of Basel III). During this period, intermediary wealth also increases. As discussed earlier, we expect the basis to capture regulatory and demand shocks in addition to changes to intermediary wealth. The widening of the CIP deviations post-GFC lines up with a decline in the broker-dealer leverage, consistent with the idea that both variables capture intermediary constraints.

### 4.3 The basis and the SDF

Can the correlation between CIP violations and intermediary equity returns explain the risk premiums we documented in the previous section? If the intermediary's SDF consisted only of the intermediary wealth return, then regressions of excess asset returns on proxies for this factor should generate intercepts of zero (Cochrane 2009). In the analysis that follows, instead of focusing on one particular forward CIP trading strategy, we focus on the first

**Table 6**
**Pricing Fwd CIP returns with intermediary wealth**

| | One-month fwd 3-month classic pc forward CIP returns | | | | | | |
|---|---|---|---|---|---|---|---|
| | Monthly returns | | | | Quarterly returns | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Market | | 0.011* | | 0.009 | 0.006 | | 0.020*** |
| | | (0.006) | | (0.009) | (0.008) | | (0.006) |
| Int. equity | | | 0.007** | 0.002 | | | |
| | | | (0.003) | (0.004) | | | |
| HKM factor | | | | | 0.004 | | |
| | | | | | (0.003) | | |
| AEM factor | | | | | | | 0.0001 |
| | | | | | | | (0.0001) |
| Constant | 0.050*** | 0.036*** | 0.045*** | 0.037*** | 0.041*** | 0.150*** | 0.079** |
| | (0.010) | (0.010) | (0.009) | (0.011) | (0.011) | (0.032) | (0.034) |
| Observations | 126 | 126 | 126 | 126 | 126 | 42 | 42 |

In this table, we regress the returns of the "Top-six first PC" forward CIP trading portfolio on a constant and the intermediary wealth and constraint proxies described in the text: Market, Intermediary Equity, the HKM Factor, and the AEM factor. Regressions (1) through (4) use monthly returns. Regressions (5) and (6) use quarterly returns. Standard errors are computed using the Newey-West kernel with a 12-month (monthly) or four-quarter (quarterly) bandwidth. *$p$ <.1; **$p$ <.05; ***$p$ <.01.

principal component portfolio ("Top-Six First PC") described in the previous section.[38]

Table 6 reports these regressions for four configurations (columns 2 through 5) of intermediary wealth return proxies: Market only, Intermediary Equity only, Market and Intermediary Equity, and Market and HKM Factor. Here, Market and Intermediary Equity refer to the return on the stock market and the value-weighted equity of primary dealers, respectively, and HKM Factor refers to innovations to the AR(1) process of primary dealers' equity capital ratio, as defined in He, Kelly, and Manela 2017. The outcome variable is the profits of the 1M-forward 3M-tenor "Top-Six First PC" forward CIP trading strategy, scaled by one-third to convert the units from annualized profits per dollar notional to bps per month (see Equation (11)).

All four configurations generate statistically significant intercepts ("$\alpha$"), rejecting the null that the HKM wealth-portfolio factors are sufficient to explain the risk premiums on the forward CIP trading strategy. Moreover, the point estimates for the intercepts range from 3.6 to 4.5 bps per month when using nonoverlapping monthly data, which is close to the average excess return of 5.0 bps/month (see column 1). That is, proxies for intermediary wealth returns explain only a small part of the excess returns associated with the forward CIP trading strategy. Column 7 presents a quarterly regression that uses the Market and AEM broker-dealer leverage factor as proxies for intermediary wealth and constraints. We again find that the forward CIP return's exposure to the Market

38 Prior versions of this paper instead focused on AUD-JPY and obtained qualitatively similar results: the two portfolios are 97% correlated in the post-GFC period.

and AEM factor does not explain the bulk of its excess returns; however, with only 42 quarterly data points, this result is substantially noisier.

These results suggests that SDFs which include both intermediary wealth returns and forward CIP returns (as in Equation (1)) should fit the data better than SDFs that use only intermediary wealth. The Bayesian approach of Chib, Zeng, and Zhao 2020 (which corrects the approach of Barillas and Shanken 2018) provides a method of comparing SDFs to formalize this intuition. We consider three sets of possible SDFs: (1) combinations of the market, intermediary equity return, and forward CIP return, (2) combinations of the market, HKM factor, and forward CIP return, and (3) combinations of the AEM leverage factor, market, and forward CIP return. The second and third sets of SDFs always include the nontradable HKM and AEM factors in the SDF. In Table 7, for each of these groups of SDFs, we calculate the posterior probabilities (given our data and the prior over models described by Chib, Zeng, and Zhao 2020) for each SDF. The total posterior probabilities of models that do and do not include the forward CIP return are presented in the rightmost column. In the monthly data specifications, the models that include the forward CIP return have a total posterior probability that is substantially higher than models that do not include the forward CIP return. That is, the Chib, Zeng, and Zhao 2020 procedure recommends including the forward CIP return as a factor in the SDF. The quarterly data sample is insufficient to distinguish between models that do and do not include the forward CIP return.

We would caution readers, however, that the exact level of the posterior probabilities can be sensitive to changes in the sample period or variables in question. The probabilities associated with models featuring a forward CIP return are substantially higher when that forward CIP return is the AUD-JPY return as opposed to the "Top-Six First PC" portfolio, despite the high

**Table 7**
**Bayesian posterior prob of SDF models with Fwd. Ret. PC**

| Factor space | Model | Probability | Subtotal |
|---|---|---|---|
| {Market, Int. Equity, Fwd. ret. PC} | Market | 0.074 | |
| | Int. Equity | 0.002 | |
| | Market + Int. Equity | 0.046 | 0.123 |
| | Fwd. ret. PC | 0.380 | |
| | Market + Fwd. ret. PC | 0.262 | |
| | Int. equity + Fwd. Ret. PC | 0.048 | |
| | Market + Int. Equity + Fwd. ret. PC | 0.187 | 0.877 |
| {HKM factor, Market, Fwd. ret. PC} | HKM factor | 0.001 | |
| | HKM factor + Market | 0.111 | 0.112 |
| | HKM factor + Fwd. ret. PC | 0.037 | |
| | HKM factor + Market + Fwd. ret. PC | 0.851 | 0.888 |
| {AEM Factor, Market, Fwd. Ret. PC} | AEM factor | 0.228 | |
| | AEM factor + Market | 0.212 | 0.440 |
| | AEM factor + Fwd. ret. PC | 0.449 | |
| | AEM factor + Market + Fwd. ret. PC | 0.112 | 0.560 |

In this table, we report posterior probabilities for factor models that do and do not include the forward CIP return on the "Top-six first PC" portfolio, using the method of Chib, Zeng, and Zhao 2020 (see Internet Appendix J for details).

correlation between those two returns (see Internet Appendix Table A9). This difference is driven by the small differences in the mean return and in the correlation between those two returns and the Market portfolio. Likewise, the procedure's preference for models with the market return as opposed to the intermediary equity return is driven the high return on the market, but not intermediary equity, in the post-GFC period; the two have similar average returns over longer samples.

If the forward CIP return is part of the SDF, its mean return can help identify the coefficients of the SDF. By construction, the returns of our forward CIP trading strategy are also the (negative of) innovations to the magnitude of the cross-currency basis. Our "Top-Six First PC" forward CIP trading strategy earns 4.6 bps per month on average in the post-GFC period.[39] Take the Intermediary Equity return as the tradable proxy for intermediary wealth returns; the mean excess return of Intermediary Equity from 1970 to 2018 is 59 bps per month. Defining $\lambda$ as the vector containing these mean excess returns, we can extract estimates of $\gamma$ and $\xi$ (the coefficients in the SDF of Equation (1)) by multiplying these means by the inverse of the variance-covariance matrix ($\Sigma$) of the two factors (see Cochrane 2009). We estimate the standard deviation of the "Top-Six First PC" forward CIP return at 14 bps per month, and the standard deviation of the Intermediary Equity return at 6.3% per month. The correlation between these two factors is 0.31 in our post-GFC sample, meaning that the bases with the largest magnitudes tend to shrink when intermediary equity returns are positive. Using these estimates, with GMM standard errors in parentheses,

$$\begin{bmatrix} \gamma \\ \xi \end{bmatrix} = \Sigma^{-1}\lambda = \begin{bmatrix} -0.2\,(1.5) \\ 247\,(63.0) \end{bmatrix}.$$

Our results are consistent with $\gamma < 1$ in two ways. First, our direct estimate of the $\gamma$ parameter is less than one, and in fact essentially zero. Second, the sign of our estimate of $\xi$ is greater than zero, which should be expected if $\gamma < 1$. The basic fact driving this result is that that exposure to the forward CIP return appears to explain essentially all of the excess returns on intermediary equity. As a result, our point estimates are consistent with the model that motivates Adrian, Etula, and Muir 2014, with risk-neutral but constrained intermediaries. However, these are point estimates and subject to estimation error; our estimate for $\xi$ is statistically significant at the 1% level, but we cannot reject $\gamma$ coefficients that exceed one. See Internet Appendix H for a discussion of the estimation and standard errors.

Our estimate of $\xi > 0$ (implying $\gamma < 1$) is driven by the fact that the forward CIP strategy achieves a risk premium that is larger than would be expected given its beta to the intermediary equity factor. Recall that a large basis indicates better future investment opportunities. Intermediaries will view exposure to

[39] This is the average of the daily sample of overlapping monthly returns. The nonoverlapping monthly return sample has an average of 5.0 bps per month, as shown in column 1 of Table 6.

basis shocks as risky if they prefer to hoard wealth to take advantage of those better investment opportunities, which occurs when $\gamma < 1$. We emphasize that this is not a quirk of our model, but rather a general fact about investment opportunities and intertemporal hedging.[40]

However, we cannot rule out the alternative possibility that the forward CIP return is a better proxy for the true intermediary wealth return. If the risk premiums we document is entirely caused by this effect and not intertemporal hedging, the forward CIP return must be a much better proxy for the true intermediary wealth return than the HKM intermediary equity measure. Suppose that $\gamma = 1$, so there is no intertemporal hedging concern. Our Sharpe ratio estimate for the "Top-Six First PC" forward CIP trading strategy implies (by the Hansen-Jagannathan bound) that the true intermediary wealth return must have an annual volatility of at least 100%, far higher than the annualized volatility of the Intermediary Equity return. Suppose instead that $\gamma > 1$. In this case, the intermediary manager should view the forward CIP trading strategy as not being very risky at all, because it offers low returns only when future arbitrage opportunities are available. To overcome this intertemporal hedging effect, we would have to suppose that our forward CIP trading returns are strongly correlated with the true intermediary wealth returns.[41]

We believe our results are interesting regardless of which of these interpretations is preferred. Either intertemporal hedging considerations are large and can be proxied for by the forward CIP return or the forward CIP return is a better way of measuring intermediary wealth returns (the main component of the SDF). Under either of these interpretations, we would be justified in using the forward CIP return as an asset pricing factor.

### 4.4 Cross-sectional asset pricing

Next, we present a cross-sectional analysis, which provides an additional test of our theory. If the SDF in Equation (1) is correctly specified and the traded factors are good proxies of the true factors, the prices of risk estimated from the cross-section of asset returns should be the same as the unconditional risk premiums of the traded factors.

We view this analysis as a complement to our earlier direct estimates of the forward arbitrage return. The key advantage our approach enjoys over other empirical intermediary asset pricing exercises is that we are able to directly estimate the price of the risk that constraints tighten. The cross-sectional exercise allows us to verify that the price of risk we directly estimate is consistent with the price of risk we infer from the cross-section. Because we are focused on the post-GFC sample, the power of our cross-sectional exercise

---

[40] See, for example, Kondor and Vayanos 2019, p. 1157.

[41] Given the persistence of the cross-currency basis, which causes the intertemporal term to be large, it is not obvious that even perfect correlation with the intermediary wealth return solves the problem. However, without an exact quantification of the intertemporal terms, we cannot rule out this possibility.

is limited. Despite this limitation, in some cases (when pooling across asset classes) we are able to reject the hypothesis of a zero risk price, but not the hypothesis that the risk price in the cross-section matches the risk price we directly estimate.

Our exercise directly builds on HKM. We study equities (FF6, the Fama-French six size by value portfolios, Fama and French 1993), currencies (FX, developed and EM currencies sorted on forward premiums, Lustig, Roussanov, and Verdelhan 2011), U.S. bonds (United States, six maturity-sorted CRSP "Fama Bond Portfolios" of Treasury bonds and five Bloomberg corporate bond indexes), sovereign bonds (Sov, sorted on credit rating and beta to the market, Borri and Verdelhan 2015), equity options (Opt, 12 portfolios of S&P 500 calls and puts, Constantinides, Jackwerth, and Savov 2013), credit default swap indexes (CDS, five traded CDS indexes), and commodities (Comm, six Bloomberg commodity futures return indexes). We also study single-currency forward CIP returns with OIS and IBOR rates (FwdCIP).

Note that many of these test portfolios have a factor structure to their returns. For example, the currency portfolios of Lustig, Roussanov, and Verdelhan 2011 can be summarized by the "carry" and "dollar" factors. Following HKM, we do not include these factors as additional factors in our model. That is, we are asking whether the risk premium associated with the carry trade is explained by its exposure to intermediary wealth and the forward CIP return, not whether the proposed factors of the SDF predict the part of currency returns that is not explained by the carry and dollar factors.

The conjecture we are testing, which follows from our hypothesized form of the stochastic discount factor,[42] is that

$$E[R^i_{t+1} - R^f_t] = \alpha + \beta^i_w \lambda_w + \beta^i_x \lambda_x, \tag{12}$$

where $\beta^i_w$ is the beta of asset $i$ to the intermediary wealth return and $\beta^i_x$ is the beta to the negative of the forward CIP return. These betas can be estimated in the standard way using a time-series regression,

$$R^i_{t+1} - R^f_t = \mu_i + \beta^i_w (R^w_{t+1} - R^f_t) + \beta^i_x r^{|x|}_{t+1} + \epsilon^i_{t+1}, \tag{13}$$

where $r^{|x|}_{t+1}$ is the negative of the return of the "Top-Six First PC" of forward CIP returns.[43]

Our preferred specification uses the Intermediary Equity return as our proxy for $R^w_{t+1}$. As discussed in Cochrane 2009, with tradable factors, if we included the factors as test assets and used GLS or two-step GMM to estimate the risk

[42] Our hypothesis is expressed as a linear form for the log SDF, but we test a linear SDF to stay closer to the procedure of He, Kelly, and Manela 2017.

[43] Our model implies that the risk-free rate in this time-series regression is misspecified, and should include an adjustment proportional to $x_t$ (see Equation (A2) in the Internet Appendix). However, because $x_t$ has only a little ability to predict $R^w_{t+1} - R^f_t$ or $r^{|x|}_{t+1}$, omitting it has almost no effect our results. See Internet Appendix Table A22 for a version of Table 8 with a risk-free rate adjustment.

**Table 8**
**Cross-sectional asset pricing tests, two-factor**

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| | US | Sov | FX | Opt | FwdArb | CDS | FF6 | Comm | 1-8 |
| Int. equity | −0.309 | 0.757 | 2.692 | 0.612 | −1.476 | 1.217 | 0.984 | 0.465 | 0.478 |
| | (0.409) | (1.060) | (0.795) | (0.767) | (8.123) | (1.068) | (0.613) | (1.195) | (0.550) |
| Fwd. CIP ret. PC1 | −0.0595 | −0.0322 | −0.0586 | −0.0751 | −0.0405 | −0.101 | 0.0299 | −0.146 | −0.0485 |
| | (0.0353) | (0.0721) | (0.0416) | (0.0396) | (0.0310) | (0.0730) | (0.152) | (0.110) | (0.0199) |
| Intercepts | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MAPE (%/mo.) | 0.047 | 0.047 | 0.052 | 0.058 | 0.014 | 0.006 | 0.097 | 0.043 | |
| H1 $p$-value | .051 | .970 | .016 | .825 | .943 | .302 | .306 | .639 | .966 |
| KZ $p$-value | .000 | .545 | .217 | .016 | .001 | .069 | .077 | .864 | .000 |
| N (assets) | 11 | 6 | 11 | 12 | 13 | 5 | 6 | 6 | 70 |
| N (beta, mo.) | 126 | 126 | 126 | 126 | 126 | 126 | 126 | 126 | |
| N (mean, mo.) | 388 | 312 | 446 | 300 | 126 | 198 | 1,133 | 359 | |

This table reports estimates for the price of risk ($\lambda$ in Equation (12)) from the various asset classes described in Internet Appendix K. "Fwd. CIP Ret. PC1" is the negative of the return on the "Top-six first PC" forward CIP trading portfolio as described in the text and "Int. equity" is the intermediary equity return of He, Kelly, and Manela 2017. Intercepts indicates an intercept is used for each asset class (one intercept per column, with six in (9), which pools the eight asset classes). MAPE is mean absolute pricing error of monthly returns. H1 $p$-value tests whether the price of the Fwd. CIP Ret. PC1 and Int. Equity risks are equal to the mean excess returns of those trading strategies. KZ $p$-value is the $p$-value of the Kleibergen and Zhan 2020 test rejecting reduced rank for the betas of the test assets. Standard errors are computed using GMM with the Newey-West kernel and a 12-month bandwidth. Units are in percentage points.

prices $\lambda_x$ and $\lambda_w$, we would recover the mean excess returns of those factors. We ask instead whether the price of risk implied by the cross-section of other asset returns is consistent with the mean excess returns on our tradable factors. For this reason, we estimate Equation (12) as an OLS regression, with GMM standard errors to account for the estimation of the betas in Equation (13), following chap. 12 of Cochrane 2009.[44] Both regressions use monthly data. For each asset class, and for a combination of all eight asset classes, we report the "H1 $p$-value" from testing whether $\lambda_w$ and $\lambda_x$ are equal to the mean excess return of the corresponding factor (59 bps per month and -5.0 bps per month in our nonoverlapping monthly sample, respectively). This $p$-value, as opposed to the usual test of whether the coefficients are zero, is the focus of our analysis.

One difference between our main specification and the textbook procedure is that the samples we use to estimate the betas and the mean excess returns are different. Our model argues that the cross-currency basis enters the SDF because it measures the degree to which regulatory constraints bind, a viewpoint relevant for the post-GFC period. Ten years of data, however, is generally too short of a time period to reliably determine whether one test portfolio has a higher expected return than another test portfolio. To overcome this difficulty, we estimate the cross-sectional regression using the longest available sample for each test portfolio, while estimating the betas using only the post-GFC

[44] More efficient (in an asymptotic sense) procedures estimate Equations (12) and (13) jointly as moment conditions. These procedures have advantages and disadvantages relative to the cross-sectional approach (see Cochrane 2009).

sample. This approach increases the likelihood of rejecting our "H1" hypothesis (biasing against our main finding), and is valid if the long-sample expected excess returns are also the expected excess returns in the post-crises period. We present broadly similar results using only post-GFC data in the Internet Appendix Table A28.[45]

Our setting has the potential for weak identification. Weak identification arises when there is not significant cross-sectional variation in the betas of the test assets to the factors. Kleibergen and Zhan 2020 develop a "pretest" that tests whether the estimated betas of the test assets are different from each other. We report the $p$-value associated with this test. Low $p$-values suggest rejection of the hypothesis that the betas are equal.[46] Note that the primary focus of our exercise is whether we can reject our "H1" hypothesis. Spurious rejection induced by weak identification makes it more difficult for us to find cross-sectional results that are consistent with our direct estimates of the risk premiums for our tradable factors.

Note that our estimates depart in a variety of ways from HKM. Our results are estimated on monthly data, and our betas are estimated only in the post-GFC period. Our test portfolios in each asset class are also different, in some cases only slightly and in some cases more substantially. We describe these details in Internet Appendix K.

Our main cross-sectional results are shown in Table 8. These results use the two-factor specification of Equation (12), with the Intermediary Equity return as the empirical proxy for $R^w_{t+1}$. The first eight columns show results for individual asset classes. For several asset classes, we cannot reject the hypothesis of weak identification, and the problem is particularly severe for commodities. Pooling our eight asset classes improves identification, and we report pooled results in column 9.

Our main outcome of interest is the H1 hypothesis that the prices of risk are equal to the mean excess returns of Intermediary Equity and negative of the forward CIP return. We are unable to reject this hypothesis when we pool across asset classes to achieve more precise identification. Our point estimates in the pooled specification ($\lambda_w = 0.472, \lambda_x = -0.0485$) are in fact quite close to the mean excess returns. Note also that our point estimates for the price of

45 One difference between our main results and our results using post-GFC means is the risk price of the basis shock with FX test assets. We find that carry trade returns are correlated with the basis shock, but in the post-GFC period, carry trade returns are smaller than in the pre-GFC period. This example illustrates the costs and benefits of using the full sample for mean returns. If we believe carry still earns a large risk premium, but happens to have not done as well during the post-GFC period, using the long sample provides a better estimate of the price of risk for the basis shock. If instead we believe that the risk premium of carry has declined, then using only the post-GFC sample is preferable.

46 Specifically, we use the multifactor version of the test described in the appendix of Kleibergen and Zhan 2020. We report the $F$ test associated with their statistic to account for the "large $N$, small $T$" nature of some of our regressions. We modify their test slightly to account for the fact that when we pool asset classes, we have one intercept for each asset class in the asset pricing equation as opposed to a single intercept. Unfortunately, the other robust inference methods described by Bryzgalova 2020 and Kleibergen and Zhan 2020 could not be directly applied to our setting.

the basis risk, $\lambda_x$, are strikingly consistent across the asset classes (except for equities, which has large standard error).

The H1 hypothesis is rejected at conventional thresholds in the U.S. bond and FX asset classes due to (respectively) very low and very high estimates for the risk price $\lambda_w$. With regard to the FX asset class, three points are worth mentioning. First, the KZ $p$-value is high, indicating that there may be an insufficient spread in the betas of the interest-rate-sorted portfolios to the intermediary asset pricing factors to identify the relevant prices of risk, in which case the standard errors on our estimates might be misleading. Second, some authors (e.g., Burnside, Eichenbaum, and Rebelo 2011) have argued that the factors that price the carry trade are disconnected from the factors that price other assets; consistent with this, our point estimates can be interpreted as saying that the carry factor is correlated with intermediary equity returns but carries a higher risk price than can be justified by that correlation. Third, as noted in footnote 45, the carry trade has performed less well in recent years; our results that use only the post-GFC sample to estimate mean returns (Internet Appendix Table A28) estimate prices of risk for the FX asset class that are close to our directly estimated risk prices.

With regard to U.S. bonds, the negative price of risk estimated for $\lambda_w$ appears only when including the March 2020 COVID-19 crisis. As noted by, for example, He, Nagel, and Song 2021, during this episode, longer-dated Treasury bonds initially fell in price, in contrast to the price increases observed during prior crisis episodes. Moreover, these movements were large relative to the usual degree of volatility in Treasury yields during the post-GFC period. As a result, the inclusion of March 2020 in the sample has a large impact on our estimates of the beta between Treasury bonds and our risk factors; of course, it also affects the betas estimated for other, more volatile assets, but to a smaller degree. Internet Appendix Table A14 presents results for a data sample that ends in December 2019. The pooled results are similar to those obtained using the full data sample, but for some asset classes, and in particular for U.S. bonds, the point estimates differ substantially.

We consider a variety of alternative specifications in the Internet Appendix. The finding of a low (negative) intermediary risk price for U.S. bonds and an excessively large intermediary risk price for currencies appear in almost all of these specifications, including specifications that replicate the HKM analysis and do not include our basis factor.

Our model emphasizes that the SDF should include proxies for both intermediary wealth returns and future investment opportunities. For this reason, we view our forward CIP return factor as complement to the HKM intermediary equity return measure. However, as discussed above, we cannot rule out the possibility that our measure is instead a better proxy for intermediary wealth returns. Internet Appendix Tables A15, A16, A17, and A18 all run two-factor models involving the Market and an intermediary-related factor (the HKM factor, the HKM intermediary equity return, the AEM broker-dealer leverage factor, and the "Top-Six First PC" portfolio return, respectively). The price of the HKM factors in the pooled regressions are not significantly different from zero. The AEM leverage factor helps price equity portfolios, but is not priced consistently across asset classes. In contrast, our CIP risk factor is consistently priced across asset classes and yields a significant price of risk in the pooled regressions. More discussion on how CIP risk factor compares with the HKM and AEM factors can be found in Internet Appendix I.

The Internet Appendix also presents a number of other variants on Table 8. There is a variant (Table A21) in which we use the HKM capital ratio innovation (along with the Market and forward CIP return), which is nontradable and is the primary specification in HKM. In this case, we can only test whether the risk prices of the traded factors are consistent with their excess returns, and our results are noisier both in terms of standard errors and with respect to the weak identification test. The Internet Appendix also contains variants that use alternative measures in the place of the "Top-Six First PC" portfolio return. Tables A23 and A24 use, respectively, AUD-JPY and USD-JPY, and Tables A25, A26, and A27 use the other portfolios of forward CIP returns described in Table 3. Table A29 uses the AR(1) innovation of the 3m OIS AUD-JPY spot basis instead of a forward CIP return.[47] Table A30 replaces the forward CIP return in Table 8 with the AR(1) innovation (following HKM) of the first principal component of the near-arbitrages described in Section 4.1, scaled to match the volatility of the AUD-JPY forward CIP return. These variants generate results that are similar to those of Table 8.

## 5. Conclusion

We provide direct evidence that innovations to the cross-currency bases are correlated with the SDF. These results are consistent with our motivating hypothesis, derived from an intermediary-based asset pricing framework and intertemporal hedging considerations. They are also consistent with the correlation between the basis and other near-arbitrages, the correlation between the basis and measures of intermediary wealth, and with our cross-sectional asset pricing tests. Taken together, we view our results as strongly supportive of intermediary asset pricing theory.

More broadly, we view this paper as beginning an investigation in the dynamics and pricing of arbitrages induced by regulatory constraints. If intermediaries play a central role in both asset pricing and the broader economy, then the question of how to measure the constraints they face and the properties of those constraints is of first-order importance.

---

47 An AR(1) model provides a reasonable description of the spot basis post-GFC: the correlation between the AR(1) innovation and the forward CIP return for AUD-JPY is 0.79 in the post-GFC data. However, the forwards contain information not captured by the spot basis. For example, there was a large, correctly anticipated spike in the 3-month spot basis across year-end 2019; such spikes were small or nonexistent for the 3-month tenor in prior years. As a result, the AR(1) innovation and forward CIP returns differ sharply in fall 2019.

**References**

Adrian, T., E. Etula, and T. Muir. 2014. Financial intermediaries and the cross-section of asset returns. *Journal of Finance* 69:2557–96.

Anderson, A., W. Du, and B. Schlusche. 2019. Arbitrage capital of global banks. Working paper, Federal Reserve Bank.

Augustin, P., M. Chernov, L. Schmid, and D. Song. 2020. The term structure of cip violations. Working paper, McGill University.

Avdjiev, S., W. Du, C. Koch, and H. S. Shin. 2019. The dollar, bank leverage, and deviations from covered interest parity. *American Economic Review: Insights* 1:193–208.

Bai, J., and P. Collin-Dufresne. 2019. The cds-bond basis. *Financial Management* 48:417–39.

Barillas, F., and J. Shanken. 2018. Comparing asset pricing models. *Journal of Finance* 73:715–54.

Bech, M. L., and E. Klee. 2011. The mechanics of a graceful exit: Interest on reserves and segmentation in the federal funds market. *Journal of Monetary Economics* 58:415–31.

Borio, C. E., R. N. McCauley, P. McGuire, and V. Sushko. 2016. Covered interest parity lost: Understanding the cross-currency basis. *BIS Quarterly Review September* .

Borri, N., and A. Verdelhan. 2015. Sovereign risk premia. Working paper, LUISS.

Boyarchenko, N., T. M. Eisenbach, P. Gupta, O. Shachar, and P. Van Tassel. 2018. Bank-intermediated arbitrage. New york fed staff report.

Brunnermeier, M. K., and L. H. Pedersen. 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22:2201–38.

Bryzgalova, S. 2020. Spurious factors in linear asset pricing models. Working paper, Stanford.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo. 2010. Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24:853–91.

Burnside, C., M. Eichenbaum, and S. Rebelo. 2011. Carry trade and momentum in currency markets. *Annual Review of Financial Economics* 3:511–35.

Campbell, J. 1993. Intertemporal asset pricing without consumption data. *American Economic Review* 83:487–512.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley. 2018. An intertemporal capm with stochastic volatility. *Journal of Financial Economics* 128:207–33.

Campbell, J. Y., and R. J. Shiller. 1991. Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies* 58:495–514.

Chib, S., X. Zeng, and L. Zhao. 2020. On comparing asset pricing models. *Journal of Finance* 75:551–77.

Cochrane, J. H. 2009. *Asset pricing: revised edition*. Princeton, NJ: Princeton University Press.

Constantinides, G. M., J. C. Jackwerth, and A. Savov. 2013. The puzzle of index option returns. *Review of Asset Pricing Studies* 3:229–57.

Correa, R., W. Du, and G. Y. Liao. 2020. Us banks and global liquidity. Working paper, Federal Reserve Board.

Du, W., J. Im, and J. Schreger. 2018. The us treasury premium. *Journal of International Economics* 112:167–81.

Du, W., and J. Schreger. 2021. CIP deviations, the dollar, and frictions in international capital markets. In G. Gopinath, E. Helpman, and K. Rogoff, eds., *Handbook of International Economics*, vol. 5. Amsterdam, the Netherlands: Elsevier.

Du, W., A. Tepper, and A. Verdelhan. 2018. Deviations from covered interest rate parity. *Journal of Finance* 73:915–57.

Duarte, J., F. A. Longstaff, and F. Yu. 2007. Risk and return in fixed-income arbitrage: Nickels in front of a steamroller? *Review of Financial Studies* 20:769–811.

Duffie, D. 2010. Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance* 65:1237–67.

———. 2017. Post-crisis bank regulations and financial market liquidity. Working paper, Banca d'Italia.

Fama, E. F., and K. R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3–56.

Fang, X. 2018. Intermediary leverage and currency risk premium. Working paper, the University of Hong Kong.

Fleckenstein, M., and F. A. Longstaff. 2018. Shadow funding costs: Measuring the cost of balance sheet constraints. Working paper, University of Delaware.

Fleckenstein, M., F. A. Longstaff, and H. Lustig. 2014. The tips-treasury bond puzzle. *Journal of Finance* 69:2151–97.

Gabaix, X., and M. Maggiori. 2015. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics* 130:1369–420.

Garleanu, N., and L. H. Pedersen. 2011. Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies* 24:1980–2022.

Haddad, V., and T. Muir. 2021. Do intermediaries matter for aggregate asset prices. *Journal of Finance* 76:7219–61.

Haynes, R., L. McPhail, and H. Zhu. 2018. When leverage ratio meets derivatives: Running out of options? Working paper, Commodity Futures Trading Commission.

He, Z., B. Kelly, and A. Manela. 2017. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics* 126:1–35.

He, Z., and A. Krishnamurthy. 2011. A model of capital and crises. *Review of Economic Studies* 79:735–77.

———. 2017. Intermediary asset pricing and the financial crisis. *Annual Review of Financial Economics* 173–97.

He, Z., S. Nagel, and Z. Song. 2021. Treasury inconvenience yields during the covid-19 crisis. *Journal of Financial Economics* 143:57–79.

Hébert, B. 2018. Externalities as arbitrage. Working paper, Stanford.

Hu, G. X., J. Pan, and J. Wang. 2013. Noise as information for illiquidity. *Journal of Finance* 68:2341–82.

Jermann, U. 2019. Negative swap spreads and limited arbitrage. *Review of Financial Studies* 33:212–38.

Jiang, Z., A. Krishnamurthy, and H. Lustig. 2021. Foreign safe asset demand and the dollar exchange rate. *Journal of Finance* 76:1049–89.

Kleibergen, F., and Z. Zhan. 2020. Robust inference for consumption-based asset pricing. *Journal of Finance* 75:507–50.

Kondor, P., and D. Vayanos. 2019. Liquidity risk and the dynamics of arbitrage capital. *Journal of Finance* 74:1139–73.

Krishnamurthy, A., and H. N. Lustig. 2019. Mind the gap in sovereign debt markets: The us treasury basis and the dollar risk factor. In *2019 Jackson Hole Economic Symposium*.

Liao, G. Y. 2020. Credit migration and covered interest rate parity. *Journal of Financial Economics* 138:504–25.

Liu, J., and F. A. Longstaff. 2003. Losing money on arbitrage: Optimal dynamic portfolio choice in markets with arbitrage opportunities. *Review of Financial Studies* 17:611–41.

Lo, A. W. 2002. The statistics of sharpe ratios. *Financial Analysts Journal* 58:36–52.

Longstaff, F. A. 2004. The flight-to-liquidity premium in us treasury bond prices. *Journal of Business* 77:511–26.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *Review of Financial Studies* 24:3731–77.

Lustig, H., and A. Verdelhan. 2007. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97:89–117.

Newey, W. K., and K. D. West. 1994. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies* 61:631–53.

Pasquariello, P. 2014. Financial market dislocations. *Review of Financial Studies* 27:1868–914.

Rime, D., A. Schrimpf, and O. Syrstad. 2022. Covered interest parity arbitrage. *Review of Financial Studies* Advance Access published May 18, 2022, 10.1093/rfs/hhac026.

Schwarz, K. 2018. Mind the gap: Disentangling credit and liquidity in risk spreads. *Review of Finance* 23:557–97.

Shleifer, A., and R. W. Vishny. 1997. The limits of arbitrage. *Journal of Finance* 52:35–55.

Wallen, J. 2020. Markups to financial intermediation in foreign exchange markets. Working paper, Harvard Business School.