Online Appendix for ’Shorting the Dollar When Global Stock Markets Roar: The Equity Hedging Channel of Exchange Rate Determination’

Nadav Ben Zeev*
Ben-Gurion University of the Negev

Daniel Nathan†
University of Pennsylvania and Bank of Israel

April 30, 2024

Abstract

This online appendix consists of the following five appendices: an appendix depicting a simple model that serves as theoretical motivation for the paper; an appendix detailing the estimation of the GIV and Bartik forward supply shocks; a robustness appendix; an appendix providing institutional details that support the distinction between forward flows and rebalancing-induced spot flows; and an appendix that discusses the external validity of our baseline results.

*Department of Economics, Ben-Gurion University of the Negev, Beer-Sheva, Israel. *E-mail:* nadavbz@bgu.ac.il.

†Finance Department, The Wharton School, University of Pennsylvania, Philadelphia, United States, and the Research Department, Bank of Israel, Jerusalem, Israel. *E-mail:* nathad@wharton.upenn.edu.

Appendix A Theoretical Motivation

In what follows we lay out a simple structural framework which is meant to fix ideas and form a suitable conceptual base for our paper's empirical analysis. The framework is a partial equilibrium of the FX forward market consisting of two time periods (t and $t+1$) and three agents. The first is a local institutional investor (II) who sells foreign currency forwards so as to hedge its position in foreign equity markets. The second is a local importer (IM) who demands foreign currency forwards for its import activity. And the third is a global arbitrageur (GA) whose activity produces violations from CIP that are unaffected by foreign equity prices, in line with our empirical evidence.

We start our depiction of the model with a presentation of the supply side of the forward market by presenting the local II's supply of foreign currency forwards. We then show demand for foreign currency forwards by the local IM followed by an exposition of GA's activity. We end the section by defining equilibrium and presenting the model's main prediction.

A.1 Supply of Foreign Currency Forwards

Local II's Hedging. We assume that the local II hedges a share h of the FX risk of its period t foreign equity position, which we denote by A_t. (This position can be thought of as the product of some fixed quantity of foreign stocks and the price of these shocks.) In particular, this hedging is done by the local II through the selling of $FCF_{t,II} = hA_t$ foreign currency forwards on the forward market to the IM at FX forward rate $F_{t,t+1}$.

Local II's Supply of Foreign Currency Forwards. $FCF_{t,II} = hA_t$ represents local II's supply of foreign currency forwards. Note that this supply is perfectly inelastic given that it has no dependence on $F_{t,t+1}$. Importantly, a positive shock to global stock prices induces a rightward shift in the supply of foreign currency forwards because it produces a rise in A_t.

A.2 Demand for Foreign Currency Forwards

General Setting. The demand side of the forward market is governed by a local importer (IM) who buys in period t $FCF_{t,IM} = P_{t,W}Q_{t,IM}$ foreign currency forwards at forward rate $F_{t,t+1}$ to fund the purchase of its imports of intermediate input quantity $Q_{t,IM}$ at foreign price $P_{t,W}$ (in foreign currency units).[1] It is effectively assumed here that the actual payment of this purchase will be made in period $t+1$ (i.e., the import deal is made with trade credit). The local IM's imported intermediate inputs are in turn used to produce and sell output quantity $M(Q_{t,IM})$ at local price $P_{t,L}$ (in local currency units) in the local economy, where $M(Q_{t,IM})$ is an increasing and concave function.

Local IM's Expected Profit. Given the setting described above, we can write local IM's profit as

$$\Pi_{t,IM} = P_{t,L}M(Q_{t,IM}) - P_{t,W}Q_{t,IM}F_{t,t+1}. \tag{A.1}$$

Optimal Demand for Foreign Currency Forwards. To derive the optimal demand for foreign currency forwards, we let the local IM maximize its expected profit from Equation (A.1) with respect to $Q_{t,IM}$. The solution to this maximization problem obtains local IM's optimal demand for imported intermediate inputs from which it is straightforward to compute the demand for foreign currency forwards $FCF_{t,IM} = P_{t,W}Q_{t,IM}$. The FOC of this problem is

$$P_{t,L}M'(Q_{t,IM}) = P_{t,W}F_{t,t+1}. \tag{A.2}$$

To see that the demand for foreign currency forwards is downward-sloping, we implicitly differentiate Equation (A.2) with respect to $F_{t,t+1}$ so as to obtain the first derivative of $Q_{t,IM}$ with respect to $F_{t,t+1}$ and then insert this derivative in the derivative of $FCF_{t,IM}$ with respect to $F_{t,t+1}$ to obtain the effect of the latter on the former:

$$\frac{\partial Q_{t,IM}}{\partial F_{t,t+1}} = \frac{P_{t,W}}{P_{t,L}M''(Q_{t,IM})} < 0, \forall Q_{t,IM}, \tag{A.3}$$

$$\frac{\partial FCF_{t,IM}}{\partial F_{t,t+1}} = \frac{\partial \left\{ P_{t,W}Q_{t,IM} \right\}}{\partial F_{t,t+1}} = \frac{P_{t,W}^2}{P_{t,L}M''(Q_{t,IM})} < 0, \forall Q_{t,IM}, \tag{A.4}$$

[1] Our assumption that the IM is the local II's counterparty is backed by both unconditional and conditional evidence shown later in the paper. For simplicity, we assume that the local IM funds its import purchases entirely through the forward market. While it is possible to extend this framework to allow for some of the purchases to be made at the realized future spot rate, the latter simplifying perfect hedging assumption is consistent with the fact that the real sector in Israel has bought on a net basis over our sample period six times more foreign currency on the forward market than on the spot market, indicating that most of importers' FX flow activity takes place on the forward (rather than spot) market.

where the assumed concavity of M was used to establish the negative relation between $F_{t,t+1}$ and $Q_{t,IM}$, which in turn ensures the downward-sloping nature of the demand for foreign currency forwards. This constitutes an important result because it allows us to interpret the effect of a rise in A_t on the supply of foreign currency forwards (discussed in Section A.1) through the lens of a demand-supply framework in which a perfectly inelastic supply curve intersects a downward-sloping demand curve in the forward market. In particular, the prediction that a shock to global stock prices will produce a rightward shift in the (perfectly inelastic) supply of foreign currency forwards can now be interpreted as happening along a downward-sloping demand curve and thus will lead in equilibrium in the forward market to a rise in foreign currency forward flows along with a decline in the FX forward rate.

A.3 Global Arbitrageur

We now introduce into the model a global arbitrageur (GA) that facilitates the determination of the FX spot rate, which we denote by S_t. This facilitation is an outcome of the following cross-currency swap. (While left unmodeled, the counterparty to this swap trade can be thought of as a broker-dealer institution.) The GA buys spot $Q_{t,GA}$ foreign currency units and sells spot $Q_{t,GA}S_t$ local currency units while simultaneously buying forward $Q_{t,GA}S_t(1 + i_{t+1,L})$ local currency units and selling $Q_{t,GA}(1 + i_{t+1,W})$ foreign currency units at forward rate $F_{t,t+1}$ (with $i_{t+1,L}$ and $i_{t+1,W}$ representing the local and foreign risk-free interest rates, respectively).

Haircut. We follow Ivashina et al. (2015) and Liao and Zhang (2020) and assume that a haircut is applied to GA's swap trade in the amount of $\kappa Q_{t,GA}$, with $0 < \kappa < 1$. That is, the GA is required to deposit a share κ of its swap position to its (unmodeled) broker-dealer counterparty. This initial margin requirement constitutes a cost for the GA that is equal the foregone interest earnings that it would be able to earn absent this requirement (i.e., $\kappa Q_{t,GA}i_{t+1,W}$). This haircut-induced cost has merit in producing a violation of CIP that accords with that we see in our data[2] in that it exists unconditionally but does not play a role in the equity hedging channel.

[2]This violation is of course not specific to our data given the robust finding from the post-GFC sample for various currencies on negative cross-currency basis with respect to the dollar (see, e.g., Du et al. (2018)).

GA's Profit Maximization. We are now in position to write GA's profit from its arbitrage activity as

$$Q_{t,GA}\frac{S_t}{F_{t,t+1}}(1+i_{t+1,L}) - Q_{t,GA}(1+i_{t+1,W}) - \kappa Q_{t,GA}i_{t+1,W}. \tag{A.5}$$

The FOC that results from maximizing the profit from Equation (A.5) with respect to $Q_{t,GA}$ is

$$\frac{S_t}{F_{t,t+1}}(1+i_{t+1,L}) = 1 + i_{t+1,W} + \kappa i_{t+1,W}, \tag{A.6}$$

where $\frac{S_t}{F_{t,t+1}}(1+i_{t+1,L})$ represents the synthetic, CIP-implied foreign (gross) risk-free interest rate which is clearly higher than the actual one owing to the haircut-induced cost. In other words, Equation (A.6) implies a negative cross-currency basis that is caused by the swap trade's haircut-induced friction with this basis unaffected by A_t. Also noteworthy is the fact that this equation implies a positive relation between the FX spot rate and the FX forward rate; this is important for our purposes as it implies that in our model the sign (as well as magnitude in percentage terms) of the FX spot rate's response to changes in foreign equity prices is the same as that of the forward rate.

A.4 Model Equilibrium

We define equilibrium in the FX forward market as the equality $FCF_{t,II} = FCF_{t,IM} = FCF_t$, with FCF_t denoting the equilibrium level of FX forward flows and where $FCF_{t,II} = hA_t$ and $FCF_{t,IM} = P_{t+1,W}F_{t,t+1}Q_{t,IM}$. The latter two equations, integrated with the equilibrium condition $FCF_{t,II} = FCF_{t,IM} = FCF_t$, join the FOCs of the local IM's, and GA's problems (i.e., Equations (A.2), and (A.6)) in forming a system of four equations in four unknowns (FCF_t, $Q_{t,IM}$, $F_{t,t+1}$, and S_t) which represents our model's equilibrium.[3]

Relation Between A_t and FCF_t, $F_{t,t+1}$, and S_t. A rise in A_t (as a result of a shock to global stock prices) implies a rightward shift in the perfectly inelastic supply of foreign currency forwards that takes place along a downward-sloping corresponding demand curve, where the latter is not affected by either A_t or h. This implies in turn that in equilibrium there must be a rise (fall) in quantity (price) of foreign currency forwards (i.e., a rise (fall) in FCF_t ($F_{t,t+1}$)).

[3] It is noteworthy that a proof that relies on a fixed-point argument for the existence and uniqueness of a solution to this four-equation system is available upon request from the authors.

Moreover, Since FOC (A.6) implies a positive and proportional relation between S_t and $F_{t,t+1}$ which is not dependent on A_t, the equilibrium prediction just noted for $F_{t,t+1}$ must also carry over to S_t (and in a one-to-one relation in percentage terms). Hence, in sum, we can deduce that a shock to global stock prices is predicted to reduce the spot rate in the same magnitude (in percentage terms) as it does the forward rate.

A.5 Basis for Abstraction from FX Swaps

This appendix's theoretical framework centers the equity hedging channel around forward contracts between IIs and importers, where the latter posses downward-sloping demand curves for foreign currency forwards. An alternative framework for the equity hedging channel would center around FX swap contracts between IIs and global arbitrageurs, where the latter have a concave net return from their alternative non-swap-related investment activities which in turn produces for them a downward-sloping demand for foreign currency forwards. Such frictional FX swap setting is used in Ivashina et al. (2015) to study the effect of non-U.S. banks' credit quality shocks on these banks' capacity to lend in dollars relative to euro; and in Liao and Zhang (2020) to study a *debt* hedging channel. At the core of this frictional FX swaps setting is a deviation from CIP, i.e., non-zero cross-currency basis; hence, a decline in the forward rate in this setting occurring in response to a rightward shift in IIs' supply of foreign currency swap-linked forwards is equivalent to a decline in cross-currency basis. (Equation (2) from Ivashina et al. (2015) is a formal demonstration of this equivalence.)

With regards to an effect on the spot rate, it is crucial to understand that an FX swap has two opposing legs. A long (short) spot transaction and a short (long) forward transaction. These transactions cancel each other in terms of the effect of excess demand on the exchange rate. Moreover, in our econometric analysis we have found that IIs' USD/ILS cross-currency basis does not meaningfully move in response to GIV value shock (see Table 3 from the text). Hence, the data rejects a meaningful role for FX swaps in the equity hedging channel while favoring a meaningful such role for standard, one-leg forward contracts between IIs and importers.[4]

[4] The real sector in Israel is a minor (and net seller, not buyer, of swap-linked forwards) player in the swap market, indicating that importers do very little hedging via FX swaps. IIs' *net* swap average daily volume activity amounts to 133.7 million dollars (11.4% of the entire swap market), relative to a 62.2 million dollar non-swap-linked forward volume activity (25.7% of the entire forward market). (The *net* swap number is the absolute value of the difference between swap trades that are long on the dollar (i.e., those whose first leg is a selling of spot dollars and second leg is a buying of forward dollars) and those that are short on it

Notwithstanding the irrelevance of FX swaps to the mechanism of our equity hedging channel, Appendix D exploits the transaction-level FX swap and forward data available to us to establish that IIs tend to unconditionally roll over their foreign asset positions by having the bulk of their forward dollar selling funded by their buying of FX swap-linked spot dollars. Such FX swap buying is a substitute for selling foreign asset positions as a way to fund maturing forward contracts' payments and thus its dominant role in funding the latter payments serves as meaningful support for our argument that IIs' forward selling of dollars is not done to rebalance their portfolios but rather to hedge their foreign equity profits without selling foreign stocks.

Appendix B Estimation of Forward Supply Shocks

First Step Estimation. In the first step, we identify idiosyncratic forward supply shocks from 13 micro-level regressions of our IIs' forward flows on a fixed effect, day-dummies, time trend, lags of IIs' aggregate (volume-weighted average) forward rate, own lags, and current and lagged values of the exogenous controls used in the constituent-level regressions from the text.[5,6] (The only exception is that we also include the 3- and 12-month Telbor rates to control for local interest rate changes.). This rich specification ensures that the micro-level innovations to the IIs' forward flow series from the micro-level regressions represent well idiosyncratic forward supply shocks. Based on numerous conversations we have had with FX market traders from two local banks, the microstructure of the USD/ILS forward market is such that IIs' forward flows are IIs-initiated transactions. I.e., these flows represent seller-initiated forward orders placed by IIs with market making local banks. Hence, after appropriately controlling for potential exogenous drivers of these flows, the remaining variation in these seller-initiated forward flows constitutes a valid measure of idiosyncratic II-level forward supply shocks.

(i.e., those whose first leg is a buying of spot dollars and second leg is a selling of forward dollars).)

[5]The number of lags for flows' lags, forward rate, and exogenous controls is common and determined as the average of the chosen lag specifications from the AIC, corrected AIC, BIC, and HQIC lag length criteria tests for each II-level regression. The mean and standard deviation of lags across the 13 regressions are 5.7 and 3.9, respectively.

[6]While there are 175 funds in our sample, they are managed by only 14 investment banks and insurance companies. Two funds managed by the same firm are governed by the same investment strategy and policy, including those pertaining to hedging practices. Hence, to obtain a meaningfully distinct enumeration of IIs, it is vital to resort to looking at the list of distinct 14 IIs. As explained in the text in Footnote 9, we omit one very small II from this list for the estimation of the equity hedging channel because, unlike the other IIs who systematically sell forward dollars, it systematically buys forward dollars.

Second Step Estimation. In the second step of our estimation, we run the same regressions from the GIV and Bartik *value* shock cases - Equations (2)-(4) in the text - only that now our IV is the GIV/Bartik forward supply shock. The GIV shock is constructed as the difference between the size-weighted- and inverse-variance-weighted-average of the 13 idiosyncratic micro-level forward supply shocks from the first estimation step, where the size weights are calculated from the share of forward flows average volume of each II in total IIs' average volume. As shown in Gabaix and Koijen (2023), this GIV construction is optimal in the sense that the resulting estimation possesses the highest precision. Moreover, it provides an aggregate shock to IIs' forward supply in line with Gabaix and Koijen (2023). In particular, it alleviates the concern that each of the idiosyncratic forward supply shocks still contains some common component (e.g., variation coming from contemporaneous forward rate variation) because the gap between the size-weighted- and inverse-variance-weighted-average of these shocks removes any potential remaining such common variation and hence produces a shock (shifter) to IIs' aggregate forward supply.

As in Gabaix and Koijen (2023), we define the Bartik instrument shock as the cross-sectional mean (equally-weighted-average) of our 13 IIs' idiosyncratic forward supply shocks. Gabaix and Koijen (2023) define and compare the Bartik instrument in relation to their GIV instrument and emphasize that the two instruments should be viewed as complementary identification approaches. Specifically, while the GIV approach seems to be the natural and preferable method for identification when there are large idiosyncratic shocks driving the aggregate, the Bartik instrument may be more suitable when there are no such large shocks. It is therefore of value to confirm that both identification approaches yield similar results. To summarize, the second step of our estimation allows us to estimate the effects of the GIV- and Bartik-based forward supply shocks on our main outcome variables.

Rationale Behind Identification. By purging II-level forward flow variation of innovations to various common drivers of the USD/ILS forward market (including global equity price innovations as captured by the MSCI), the estimated II-level residuals from the first estimation step provide a valid measure of II-level forward supply shocks. One can structurally interpret these shocks as arising from exogenous, discretionary behavior on the part of IIs' FX traders or the investment committees that guide these traders. And any potentially remaining common variation in our 13 idiosyncratic shocks is likely removed owing to our GIV approach which subtracts the inverse-variance-weighted-average of the shocks from the size-weighted one. While the Bartik approach does not do this subtraction and merely takes the simple mean of the idiosyncratic shocks, the fact that the average absolute pairwise correlation among these shocks is low (standing at 3.8% - this number is further discussed below) indicates that the Bartik forward supply shock also provides a valid aggregate forward supply shock.

A crucial element of our estimation's first step, also shared by the corresponding first estimation step from our equity hedging channel's econometric model, is that we allow for all of the estimated coefficients to be II-specific. Technically, this implies that we separately estimate this first step's equation for each of our 13 IIs. Substantively, this heterogenous coefficient setting allows us to remove common variation in IIs' forward supply arising not only from the common exogenous control variables but also from the *way* by which IIs' respond to these variables. This is important because in addition to time-invariant differences across IIs' forward supply (captured by the fixed effect) there are also time-varying such differences stemming from heterogenous sensitivities of IIs' to both lagged II-specific forward flows and common forward market drivers. The latter heterogeneity is what our heterogenous coefficient setting precisely accounts for, resulting in a panel of 13 idiosyncratic demand shocks that exhibit a mere average absolute pairwise correlation of 3.7% with a standard deviation of 3.4%.

Appendix C Robustness Checks

This section examines the robustness of the baseline results from the text (presented in first panel of Table 2) along three dimensions. First, using different lag specifications for the constituent-level regressions. Second, truncating the baseline sample at 2/28/2020 so as to exclude the COVID-19 period. And third, replacing the constituent-level, granular approach with an aggregate one by defining the instrumented value shock as an MSCI residual purged of variation from the same exogenous controls used in the constituent-level regressions. The presentation of all of the results follows the same exposition and structure underlying the baseline results from the first panel of Table 2 from the text.

C.1 Alternative Lag Specifications

The lags for the constituent-level regressions from the text (Equation (1)) were chosen optimally as the average of the chosen lag specifications from the AIC, corrected AIC, BIC, and HQIC lag length criteria tests for each constituent-level regression. An alterative approach to this standard lag choice approach is to additionally check for each regression if the associated residuals pass the Ljung–Box auto-correlation test and then increase the lag for the regressions that do not pass this test until they do. Note that this extended lag choice approach is more stringent than our baseline one because optimal lag criteria tests already have embedded in them the residual white noise assumption. Nevertheless, using a threshold p-value of 10%, we have found that in 268 of our 774 regressions the residual series did not pass the Ljung–Box auto-correlation test; hence, in those regressions we increased the number of lags until the test was passed for the 10% p-value threshold.

Table C.1 shows the first stage effect of the GIV (Bartik) value shock on IIs' aggregate forward flows (second column) and reduced form effect on IIs' aggregate (volume-weighted average) forward rate (fourth column) as well as the 2SLS-estimated second stage estimate of the forward demand semi-elasticity (third column) conditional on the GIV value shock, where we now use the above-described approach to lag choice in the constituent-level regressions. (The mean lag across the 774 regressions is 5.6 instead of 4.4.) For completeness, we also report in the first column the OLS-estimated semi-elasticity. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units.

The results are very similar to the baseline ones. The significance of the first stage estimate (second column) is very strong with an F-statistic of 20.3, reflecting a -0.13 standard deviation units of forward selling, and implying that the GIV value shock is a sufficiently strong instrument. The GIV value shock produces significant forward rate appreciation of 0.07%, with the 2SLS estimation producing a significant estimate of 0.51.

C.2 Excluding the COVID-19-Period

While the COVID-19-related period clearly provides increased volatility to our baseline sample and thus has the potential of improving identification of the equity hedging channel, one may also argue that its uniqueness makes the case for showing that the baseline results are not driven by its inclusion. Toward this end, Table C.2 shows the first stage effect of the GIV (Bartik) value shock on IIs' aggregate forward flows (second column) and reduced form effect on IIs' aggregate (volume-weighted average) forward rate (fourth column) as well as the 2SLS-estimated second stage estimate of the forward demand semi-elasticity (third column) conditional on the GIV value shock, where we now use a sample that omits the COVID-period (resulting in sample size of 1,697 observations instead of 2,042 observations in the baseline case).

The results confirm the robustness of the baseline results to the omission of the COVID-19 period. The significance of the first stage estimate (second column) is very strong with an F-statistic of 18.9, reflecting a 0.14 standard deviation units of forward selling, and implying that the GIV value shock is a sufficiently strong instrument also in the non-COVID-period sample. The GIV value shock produces significant forward rate appreciation of 0.05%, with the 2SLS estimation producing a significant estimate of 0.33.

C.3 Aggregate Approach

Our granular econometric approach to studying the equity hedging channel was premised on the notion that the associated granular constituent-level shocks and resultant GIV construction would generate global equity market value variation that is not coming from macro forces but rather from idiosyncratic large companies' value shocks. Nevertheless, it is worthwhile to also confirm that the baseline results are robust to using an aggregate estimation approach which define the value shock instrument as the variation in MSCI purged of the variation from the exogenous controls we used in the constituent-level regressions. In other words, the aggregate estimation approach first regresses the MSCI return on an intercept, time trend, day-dummies, own lags, as well as current and lagged values of the exogenous controls vector used in the constituent-level regressions (3- and 12-month Libor rates, broad dollar index, EBP, and FTSE US Government Bond Index);[7] and then it estimates Regressions (1)-(3) from the text where the IV is the MSCI residual from the first estimation step.

Table C.3 shows the results from the above-described aggregate estimation approach, keeping the same exposition as in the previous two tables. The results continue to support the equity hedging channel with an F-statistic of 13.8 for the first stage, reflecting a 0.11 standard deviation

[7] The number of lags is 3 and is computed as the average of the chosen lag specifications from the AIC, corrected AIC, BIC, and HQIC lag length criteria tests for the MSCI regression.

units of forward selling, and implying that the GIV value shock is a sufficiently strong instrument also for the aggregate estimation approach (albeit lesser so than the baseline case). The aggregate MSCI value shock produces significant forward rate appreciation of 0.04%, with the 2SLS estimation producing a significant estimate of 0.35.

C.4 Examination of Currency Heterogeneity

Our baseline empirical analysis does not separately consider IIs' forward flows and rates by currency but rather uses the *total* dollar value of IIs' forward flows and corresponding volume-weighted forward rate by converting all non-dollar flows and rates to dollar values. This aggregate approach maximizes sample size as it exploits *all* variation in IIs' forward flows and rates data. And it is consistent with the equity hedging channel as the latter posits a mechanism that is not currency-specific.

Be that as it may, it seems worthwhile to consider as outcome variables both USD/ILS and non-USD/ILS IIs' forward flows and rates as the associated results can inform us about the potential presence of meaningful currency heterogeneity underlying our baseline results. 79.9% of IIs' forward flow daily average volume is conducted in USD, with the remaining 20.1% share conducted in various other currencies which include EUR, GBP, JPY, CAD, and SEK.[8] The two most commonly traded non-USD currencies are the EUR and GBP, with 14.4% and 2.4% average daily volume shares, respectively, with the remaining currencies making up the residual 3.3% volume share. We use IIs' transaction-level data to construct daily USD/ILS and non-USD/ILS forward flows and (volume-weighted) rates; all associated variables are in dollar terms owing to the con-

[8]The USD also appears to be the dominant in IIs' foreign equity investments, with the former being the dominant one in line with its dominance in IIs' forward volume. Specifically, while (unfortunately) we do not have foreign equity shares by currency for our IIs, we do have, however, starting in November 2015, monthly data on foreign asset shares by regions (not currencies). Making the somewhat crude (albeit arguably reasonable) assumption that foreign equity and non-equity portfolio weights are not meaningfully different by region, our data can speak to a reasonable extent to the issue of IIs' foreign equity allocation by region. In particular, an average 56.8% weight of IIs' foreign asset holdings is allocated to the U.S.; an average 12.9% weight is allocated to Europe; an average 2.8% weight is allocated to Australia; an average 2.7% weight is allocated to BRIC (Brazil, Russia, India, and China); and an average 1.5% weight is allocated to Asia. There is an average weight of 24.1% of IIs' foreign asset portfolios that is not allocated to any specific region because this share represents IIs' holding of global exchange traded funds whose specific regional holding weights are not identified by the BOI when constructing the regional weight shares. Making the somewhat crude (albeit arguably reasonable) assumption that this residual share is allocated in similar fashion as the complementary BOI-identified part, we can deduce that the U.S. equity market represents more than 70% of IIs' foreign equity investments.

version of each non-USD/ILS forward transaction to dollar value terms.

The first and second panels of Table C.4 show the results from using USD/ILS- and non-USD/ILS-only forward flows and rates as our outcome variables, respectively, keeping the same exposition as in the previous three tables. The dollar's dominant role in IIs' forward volume is also reflected in its commanding role in driving our baseline results, as manifested through its significant -0.14 standard deviation USD/ILS forward flows response, which is similar to the baseline corresponding response of -0.13 - although recall that the latter is in terms of standard deviation units of the *total* forward flows variable, as well as effectively identical OLS and forward rate appreciation estimates (-0.09% and -0.07%) relative to the baseline corresponding estimates. The associated USD-specific semi-elasticity of 0.52 is also very similar to the baseline corresponding 0.53 estimate.

The equity hedging channel is also present for non-USD/ILS currencies, which bear out a significant -0.06 standard deviation non-USD/ILS forward flows response as well a significant 0.09% currency appreciation. The high associated demand semi-elasticity estimate of 1.37 should be interpreted with caution. Specifically, the dollar's dominance in IIs' forward flow activity and its associated dominant role in driving the equity hedging channel results of this paper imply that, even in the absence of any change in IIs' supply of non-USD/ILS forward flows, triangular arbitrage is expected to appreciate non-USD/ILS rates. Hence, the estimated 1.37 non-USD/ILS demand semi-elasticity possibly reflects triangular arbitrage effects in addition to equity hedging ones.

C.5 Controlling for Latent Factors

The first stage of our estimation procedure (Equation (1) from the text) purges our individual returns of variation coming from endogenous macro forces, with this equation allowing for macro observables to affect each stock return in a differential manner in accordance with the notion that each company's sensitivity to these macro forces depends on various company-specific characteristics. While our GIV construction of our aggregate value shock should remove any additional variation coming from other unaccounted for macro forces, one may argue that further controlling for common later factors driving our returns can bolster additional confidence in the validity of our results.

Towards this end, we add the current and lagged values of the second and third factors from a principal component analysis of our returns as supplementary controls in our baseline Equation (1) from the text.[9] Table C.5 shows the first stage effect of the GIV (Bartik) value shock on IIs' aggregate forward flows (second column) and reduced form effect on IIs' aggregate (volume-weighted average) forward rate (fourth column) as well as the 2SLS-estimated second stage estimate of the forward demand semi-elasticity (third column) conditional on the GIV value shock, where we now add these two latent factors' current and lagged values to our individual return regressions from Equation (1) from the text.

The results confirm the robustness of the baseline results to controlling for latent factors, with a significant -0.09 forward flows standard deviation unit response and a corresponding 0.06% appreciation. These estimates reflect a 2SLS forward demand semi-elasticity of 0.73, which is comparable to our baseline 0.53 estimate and continues to support the existence of a meaningful equity hedging channel in the data.

[9]The first factor, which explains 35% of the common variation of our returns, is effectively the S&P 500 index (the latter's return series has a 95.1% correlation with this factor) and thus using it as a control in Equation (1) from the text mechanically voids our identification procedure of any value. (Recall the validity of our GIV shocks rests on its capacity of drive the aggregate stock market; hence, controlling for the latter in Equation (1) from the text mechanically eliminates this validity.) The second and third factors explain 3.9% and 3.1% of the returns' common variation with the next four factors' contributions dropping to mere contribution shares of 2.3%, 1.6%, 1.4%, and 1%.

Table C.1: **GIV Value Shock Estimation Results: Alternative Lag Specifications.**

Response	OLS	2SLS 1^{st} Stage	2SLS 2^{nd} Stage	Reduced Form
IIs' Forward Rate	0.09***		0.51***	-0.07***
	(0.02)		(0.09)	(0.01)
IIs' Forward Flows		-0.13***		
		(0.03)		
F-Stat		20.29		
R^2	2.91%	8.96%		1.73%
Obs	2,042	2,042	2,042	2,042

Notes: This table shows the OLS-estimated effects of the GIV value shock on IIs' aggregate forward flows (second column) from Equation (1) from the text and aggregate (volume-weighted average) forward rate (fourth column) from Equation (2) from the text as well as the 2SLS-estimated forward demand semi-elasticity (third column) from structural Equation (3) from the text conditional on the GIV value shock. The difference relative to the baseline case is that we increase the number of lags in constituent-level regressions whose regressions do not pass the Ljung–Box auto-correlation test with a 10% p-value threshold until they do. For completeness, we also report in the first column the OLS-estimated effect from structural Equation (3) from the text. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units. Numbers in parentheses represent standard errors computed from the heteroskedasticity- and autocorrelation-consistent procedure of Newey and West (1987) with the truncation lag selected from the data-driven procedure from Andrews (1991). *, **, and *** represent significance levels at the 10%, 5%, and 1% levels.

Table C.2: **GIV Value Shock Estimation Results: Excluding COVID-19 Period.**

Response	OLS	2SLS 1^{st} Stage	2SLS 2^{nd} Stage	Reduced Form
IIs' Forward Rate	0.03***		0.33***	-0.05***
	(0.01)		(0.09)	(0.01)
IIs' Forward Flows		-0.14***		
		(0.03)		
F-Stat		18.86		
R^2	0.53%	8.93%		0.90%
Obs	1,697	1,697	1,697	1,697

Notes: This table shows the OLS-estimated effects of the GIV value shock on IIs' aggregate forward flows (second column) from Equation (1) from the text and aggregate (volume-weighted average) forward rate (fourth column) from Equation (2) from the text as well as the 2SLS-estimated forward demand semi-elasticity (third column) from structural Equation (3) from the text conditional on the GIV value shock. The difference relative to the baseline case is that we exclude we truncate the sample at 2/28/2020. For completeness, we also report in the first column the OLS-estimated effect from structural Equation (3) from the text. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units. Numbers in parentheses represent standard errors computed from the heteroskedasticity- and autocorrelation-consistent procedure of Newey and West (1987) with the truncation lag selected from the data-driven procedure from Andrews (1991). *, **, and *** represent significance levels at the 10%, 5%, and 1% levels.

Table C.3: **GIV Value Shock Estimation Results: Aggregate Approach.**

Response	OLS	2SLS 1^{st} Stage	2SLS 2^{nd} Stage	Reduced Form
IIs' Forward Rate	0.09***		0.35***	-0.04***
	(0.02)		(0.12)	(0.01)
IIs' Forward Flows		-0.11***		
		(0.03)		
F-Stat		13.84		
R^2	2.82%	8.54%		0.66%
Obs	2,117	2,117	2,117	2,117

Notes: This table shows the OLS-estimated effects of the GIV value shock on IIs' aggregate forward flows (second column) from Equation (1) from the text and aggregate (volume-weighted average) forward rate (fourth column) from Equation (2) from the text as well as the 2SLS-estimated forward demand semi-elasticity (third column) from structural Equation (3) from the text conditional on the aggregate MSCI value shock from Section C.3. For completeness, we also report in the first column the OLS-estimated effect from structural Equation (3) from the text. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units. Numbers in parentheses represent standard errors computed from the heteroskedasticity- and autocorrelation-consistent procedure of Newey and West (1987) with the truncation lag selected from the data-driven procedure from Andrews (1991). *, **, and *** represent significance levels at the 10%, 5%, and 1% levels.

Table C.4: **GIV Value Shock Estimation Results: Breakdown by USD/ILS- and non-USD/ILS-Only Transactions.**

		USD/ILS-Only Transactions		
Response	OLS	2SLS 1^{st} Stage	2SLS 2^{nd} Stage	Reduced Form
IIs' Forward Rate	0.09***		0.52***	-0.07***
	(0.02)		(0.10)	(0.01)
IIs' Forward Flows		-0.14***		
		(0.03)		
F-Stat		20.21		
R^2	3.07%	8.02%		1.90%
Obs	2,001	2,001	2,001	2,001
		non-USD/ILS-Only Transactions		
IIs' Forward Rate	-0.01		1.38***	-0.09***
	(0.03)		(0.40)	(0.03)
IIs' Forward Flows		-0.06**		
		(0.03)		
F-Stat		4.69		
R^2	0.07%	2.22%		1.74%
Obs	1,484	1,484	1,484	1,484

Notes: The first (second) panel of this table shows the OLS-estimated effects of the GIV value shock on IIs' aggregate USD/ILS-only (non-USD/ILS-only) forward flows (second column) from Equation (1) from the text and aggregate (volume-weighted average) USD/ILS-only (non-USD/ILS-only) forward rate (fourth column) from Equation (2) from the text as well as the associated 2SLS-estimated forward demand semi-elasticity (third column) from structural Equation (3) from the text conditional on the GIV value shock. For completeness, we also report in the first column the OLS-estimated effect from structural Equation (3) from the text. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units. Numbers in parentheses represent standard errors computed from the heteroskedasticity- and autocorrelation-consistent procedure of Newey and West (1987) with the truncation lag selected from the data-driven procedure from Andrews (1991). *, **, and *** represent significance levels at the 10%, 5%, and 1% levels.

Table C.5: **GIV Value Shock Estimation Results: Controlling for Latent Factors.**

Response	OLS	2SLS 1^{st} Stage	2SLS 2^{nd} Stage	Reduced Form
IIs' Forward Rate	0.09***		0.73***	-0.06***
	(0.02)		(0.14)	(0.01)
IIs' Forward Flows		-0.09***		
		(0.03)		
F-Stat		9.27		
R^2	2.91%	8.07%		1.57%
Obs	2,042	2,042	2,042	2,042

Notes: This table shows the OLS-estimated effects of the GIV value shock on IIs' aggregate forward flows (second column) from Equation (1) from the text and aggregate (volume-weighted average) forward rate (fourth column) from Equation (2) from the text as well as the 2SLS-estimated forward demand semi-elasticity (third column) from structural Equation (3) from the text conditional on the GIV value shock. The difference relative to the baseline case is that we add to Equation (1) from the text the second and third latent factors driving our individual returns from a principal component analysis. For completeness, we also report in the first column the OLS-estimated effect from structural Equation (3) from the text. The Forward flows variable is divided by its standard deviation prior to entering the regressions for comparability purposes and hence its response is in terms of standard deviation units. Numbers in parentheses represent standard errors computed from the heteroskedasticity- and autocorrelation-consistent procedure of Newey and West (1987) with the truncation lag selected from the data-driven procedure from Andrews (1991). *, **, and *** represent significance levels at the 10%, 5%, and 1% levels.

Appendix D Are Forward Flows Distinct from Rebalancing-Induced Spot Flows?

A relevant concern regarding this paper's story is that IIs' forward selling of dollars merely acts as rebalancing-induced spot dollar selling as opposed to acting as a hedging device against their foreign equity positions' capital gains. I.e., one may worry that what we are picking up in the data in this paper is mainly a portfolio rebalancing mechanism where forward dollar selling substitutes spot dollar selling in the presence of foreign equity capital gains. This section's goal is to alleviate this concern in a twofold manner. (Ideally, we would want to remove this concern by looking directly at foreign equity holdings of IIs. However, we lack this data and hence turn to what we consider is a second-best way of accomplishing this goal.)

First, we emphasize the relatively long maturities typically characterizing IIs' forward contracts. Second, this section turns to inform us about whether IIs - when faced with the imminent need to obtain dollars to fund their maturing forward contracts - tend to roll over their foreign equity positions or rather realize the capital gains which are normally accrued to these positions. Understanding which is done is important to us because the latter alternative is consistent with a portfolio rebalancing framework while the former alternative implies that IIs use a different funding option for settlement of forward contracts and is therefore inconsistent with a meaningful portfolio rebalancing framework.

IIs' Forward Contracts' Maturity Distribution. To emphasize that the channel we identify does not merely serve as another method for investors to rebalance their portfolio through forward markets, we present in Table D.1 IIs' forward contracts' maturity distribution. Encouragingly, this table shows that the volume-weighted average and median maturity of forward contracts in our sample are 52 and 25 days, respectively. If IIs were simply using the forward market as a convenient means of rebalancing when the foreign equity market appreciates instead of simply selling in the spot market, we would expect to see a majority of forward contracts with close-to-zero maturity, as there would be no real incentive to wait almost a month to sell the asset.[10]

FX Swaps as a Central Funding Tool for Forward Contracts' Settlement. Investing abroad presents risks for IIs. They face FX risk because of the currency mismatch due to their liabilities being in the local currency while holding a portion of their assets in foreign currency. International investors such as IIs can mitigate potential losses by hedging against the FX risk. For instance, let's consider a foreign investor who invested in the S&P 500 in 2020, which saw a 16% increase. However, the US broad dollar index weakened by 3% during the same period. Without hedging the FX risk, the investor would have experienced approximately a 20% loss when converting their investment into their local currency.

IIs employ various strategies to hedge their FX risk. One approach - the focus of our paper - involves selling their dollar profits and purchasing the local currency through the FX forward mar-

[10]Conversations with market practitioners confirm that selling in the forward market and rolling over the forward contract is a common operational approach.

ket. By selling the profits in the forward market, investors safeguard themselves against potential losses arising from future fluctuations in the value of the dollar. This approach offers an additional crucial advantage compared to simply selling the asset in the spot market which also differentiates our paper from the previous literature: it allows the investors to *maintain* their exposure to the foreign asset without the FX exposure. When the forward contract matures, its settlement can get funded with an FX swap, eliminating the immediate need to sell the underlying asset. This allows for continuous full hedging, ensuring ongoing protection against FX risk.

The ideal way to test the validity of our proposed mechanism is using daily foreign equity holdings of IIs which is unfortunately not available to us. However, our mechanism does offer a testable implication: we anticipate a negative relationship between the amount of dollars being bought or sold in the FX forward and the amount of dollars being bought or sold in the FX swap around the maturity of the forward contract.

Consider the following hypothetical example. On 01/01/2020, an II sells dollars in the forward market for one million dollars, with a maturity date of 31/01/2020. This transaction implies that the II will need to supply one million dollars around the end of January. Consequently, the II has three funding options available to him on or around the maturity date: tapping into the FX swap market by purchasing swap-linked dollars; tapping into the spot market by purchasing spot dollars; and rebalancing its foreign asset portfolio positions by selling foreign assets. Given our data availability, we can estimate the meaningfulness of the first two of these funding options while using their quantification to learn about the relevance of the third option.

Econometric Specification. To rigorously examine the relationship between FX swaps and forward contracts, we employ a panel regression analysis at the fund level, utilizing our granular (transaction-level) dataset comprising of 175 funds. The dependent variable in our regression is the daily amount of dollars bought or sold in the FX swap market by fund i at time t ($\text{FX\_SWAP}_{i,t}$), while the independent variable is the corresponding amount of dollars bought or sold by fund i in the forward market at maturity time t ($FF_{i,t}$) and the surrounding leads and lags (indexed by l below). The regression equation can be formally expressed as follows:

$$\text{FX\_SWAP}_{i,t} = \alpha_i + \sum_{l=-7}^{l=7} \beta_l \text{FF}_{i,t+l} + \epsilon_{i,t}. \tag{D.1}$$

To examine the FX spot funding option, we merely replace the outcome variable from Equation (D.1) with the daily amount of dollars bought or sold by IIs in the FX *spot* market at time t.

Results. Our findings for FX swaps, displayed in Figure D.1, reveal notable point estimates pertaining to the maturity of the forward contract. Specifically, as anticipated, the estimates for the day prior to maturity (-1 on the x-axis), the day of maturity (0 on the x-axis), and the day after maturity (1 on the x-axis) all exhibit statistical significance and negative values, aligning with our initial expectations. Remarkably, these estimates cumulatively amount to approximately -0.78, suggesting that nearly 80% of forward contracts' settlement is funded by IIs' tapping into the FX swap market. This empirical evidence substantiates our hypothesis that II's do not engage in actual foreign stock sales but rather maintain exposure to the assets when using FX forwards. Furthermore, it is important to highlight that the estimates for the days preceding and following the maturity date are in close proximity to zero, indicating their non-significance. This outcome serves as a placebo test, further reinforcing the validity of our proposition

What accounts for the remaining 20%? One possibility is that II's acquire the remaining exposure by purchasing dollars in the spot market. However, when we conduct the same regression analysis as in Equation (D.1), employing the amount of dollars bought or sold in the spot market as the dependent variable (Figure D.2), we do not observe any discernible pattern. This suggests that the 20% may originate from alternative sources, such as IIs' utilization of their dollar reserves - with foreign equity portfolio rebalancing being one possible such utilization. Be that as it may, this relatively small 20% share represents a modest upper bound for the importance of foreign equity rebalancing as a way for settling forward contracts on the part of IIs.

This section has revealed valuable information on the institutional details of IIs' hedging. Given that IIs play a significant role in global financial markets, such revelation and studying of the institutional details of how these investors operate is crucial because their investment strategies and decision-making processes can significantly impact market dynamics and financial stability (see, e.g., Greenwood and Vayanos (2010), Klingler and Sundaresan (2019), and Czech et al. (2021)).[11] Hence, considering that FX swaps produce off-balance sheet debt - or 'missing dollar

[11] A recent example includes the significant turmoil in September 2022 in the UK sovereign bonds market brought about by pension funds as one of their methods to hedge their duration mismatch was through the use of interest rate swaps. When yields rose sharply in the UK, they were met with margin calls which threatened to destabilize their bond market. Policymakers and regulators can use this institutional knowledge to design rules and regulations that promote market stability and protect investors.

debt' (Borio et al. (2022)), our results about IIs' tendency to fund the settlement of their forward contracts with FX swaps may potentially have implications for such stability issues.

Table D.1: **Summary Statistics of IIs' Forward Contracts' Maturity Distribution.**

Volume-Weighted Mean	Median	25th Percentile	75th Percentile	Min	Max
52	25	3	64	0	1767

Notes: This table shows the volume-weighted mean, median, and upper and lower percentiles for IIs' transaction-level forward contracts' maturities in our sample, covering the baseline period from 4/26/2011 to 8/18/2021.

Figure D.1: **The Relation between FX Forwards and FX Swaps.**

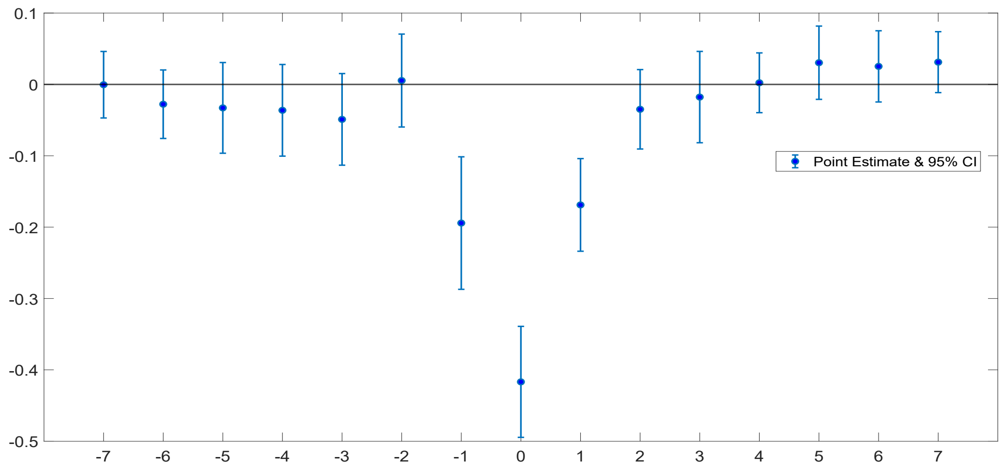

Notes: This table presents the point estimates obtained from a panel regression with fixed effects, following the specification of Equation (D.1). The x-axis represents days relative to the forward contract. To calculate standard errors, we employ Driscoll and Kraay (1998) standard errors with 7 lags (this standard error procedure accounts for arbitrary serial and spatial correlation of panel regression's residual). The regression analysis includes data from 175 funds, covering the period from 4/26/2011 to 8/18/2021.

Figure D.2: **The Relation between FX Forwards and FX Spots.**

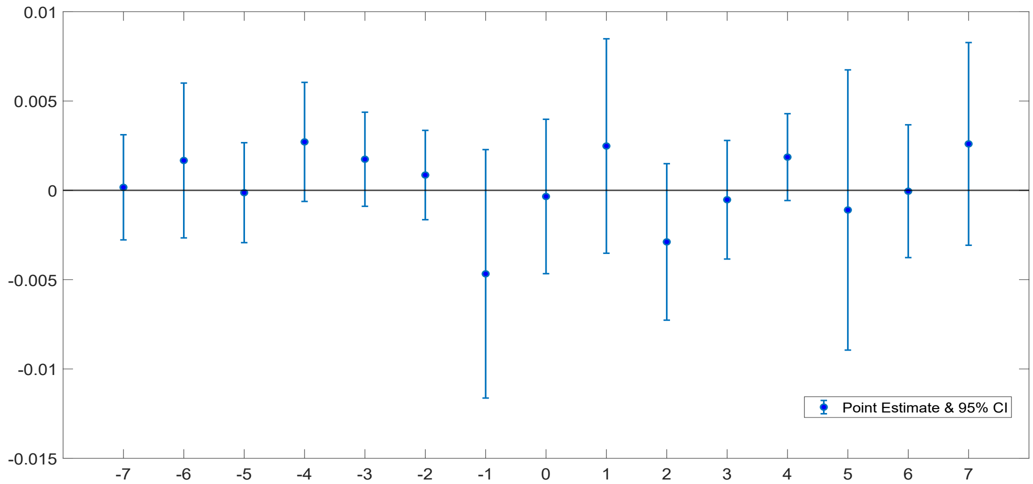

Notes: This figure presents the point estimates obtained from a panel regression with fixed effects, following the specification of Equation (D.1) but replacing the FX swap outcome variable with the FX spot one. The x-axis represents days relative to the maturity of the forward contract. To calculate standard errors, we employ Driscoll and Kraay (1998) standard errors with 7 lags (this standard error procedure accounts for arbitrary serial and spatial correlation of the panel regression's residual). The regression analysis includes data from 175 funds, covering the period from 4/26/2011 to 8/18/2021.

Appendix E External Validity: Conditions for a Meaningful Equity Hedging Channel

This section discusses the issue of our analysis's external validity, i.e., whether we can infer a broader conclusion regarding the equity hedging channel we uncover in Israel for other economies as well. We lay out three necessary conditions for a meaningful equity hedging channel along with some survey evidence supporting the likely relevance of these conditions for a broad sample of economies.

An important question arising from this paper's analysis is whether its obtained results can be considered as externally valid for broader sample of economies. While the answer to this question can not be unconditionally affirmative, in what follows we discuss three conditions which are met by a large sample of economies and whose possession by an economy is vital for there to be a meaningful equity hedging channel of exchange rate determination in this respective economy.[12]

Smallness. For a meaningful equity hedging channel of exchange rate determination, the economy at hand needs to be sufficiently small such that foreign IIs' FX exposure to this economy's currency is negligible and hence does not motivate foreign IIs to pursue the same hedging activity that is done by these economies' IIs.[13] This is an important condition because, under the fairly reasonable assumption of comovement across foreign and local stock markets, not meeting the smallness condition would facilitate a counteracting equity hedging channel that is coming from anchor currencies' large economies.

U.S. pension funds seem to have a limited position in world equities, holding only 16.7% of their total investment funds (i.e., indirect investment) in foreign (non-U.S.) equity funds (Yazdani (2020)). Making the reasonable assumption that U.S. pension funds are less inclined to make direct investment in foreign equities than they are with respect to domestic ones, the latter 16.7% is likely to go down when computing it in terms of U.S. pension funds' total investment (i.e., direct and indirect (through investment funds) investment). But even if this number were much higher, so long that the economy at hand is small, U.S. pension funds' position (or any other large economy's pension funds' position for that matter) in that economy's equities would represent a negligible share of their total assets and thus would be unlikely to warrant hedging of this position's FX exposure on the part of U.S. pension funds. Israel is a small economy that does not belong to a large monetary union and therefore meets the smallness condition.And clearly this condition is met by a large sample of economies.

Meaningful Foreign Equity Position. IIs in the economy at hand also need to hold a meaningful share of their assets in foreign equities so that their FX exposure would be sufficient to warrant hedging and so that this resulting hedging would also produce meaningful FX forward flows. Israeli IIs hold on average 10.9% of their assets in foreign equities over our sample period. Given the global nature of IIs' investments across the world, comparable values are expected to hold for the typical small economy. A recent report from Yazdani (2020) corroborates this reason-

[12]These three conditions do not include the obvious condition of having a flexible exchange rate regime in place.

[13]Note that this smallness condition is not necessarily implied by the smallness of an economy in real terms (e.g., in GDP terms) as a small economy that belongs to a large monetary union such as the Euro area would not meet this condition.

able expectation, documenting a 18.5% average share of foreign equities in total pension funds' assets across several small economies (Australia, Canada, Chile, Colombia, Denmark, Mexico, New Zealand, Norway, Peru, South Korea, Sweden, and Switzerland) along with a moderate standard deviation of 8%.

Hence, the foreign equity condition seems to be relevant for a broad sample of economies which includes as its subset the sample of economies adhering to the smallness condition. And the foreign equity condition is likely to become all the more applicable to the latter sample over time as IIs in small economies around the world are becoming more global in their investment strategies.

Meaningful Hedging. Clearly, IIs need to hedge a meaningful part of their foreign equity position for there to be a meaningful equity hedging channel of exchange rate determination. This third condition is also formalized in the motivating model from Section A. While direct data on hedging-related FX flows of IIs is quite scarce (with Israel and Chile being notable exceptions), we view this hedging condition as intertwined with the second one and we therefore expect economies possessing the foreign equity condition to also posses the hedging one. (It is not uncommon for some minimal hedging of pension funds' FX exposure to be required by government regulation in the form of a minimal currency match ratio between FX liabilities and assets. E.g., according to the OECD 2019 Survey of Investment Regulation of Pension Funds, such minimal ratios are required for pension funds in Chile (50%), Colombia (50%-85%), Denmark (80%), Mexico (70%-100%), Norway (70%), Sweden (80%-100%), and Switzerland (70%).)

In accordance with this expectation, Mercer (2020) provides survey evidence for 2020 from 927 IIs across 12 countries (with a total asset value of over 1.1 trillion dollars) indicating that 42% of the surveyed IIs hedge over 60% of their FX exposure in listed equity portfolios.[14,15] And Alfaro et al. (2021) report that Chilean pension funds are the largest holders of gross positions of FX derivatives, having the largest net short FX derivatives position and, at times, being the only net suppliers of U.S. dollars in the forward market. By the end of 2018, they held 41.3 billions of U.S

[14] 2 countries out of the 12 that were surveyed meet the smallness condition and have a floating exchange rates (Norway, and Switzerland), with the remaining economies consisting of the UK (which violates the smallness condition owing to it economy's relatively large size), 8 Euro area member economies, and Denmark whose exchange rate is fixed to the Euro.

[15] Also see Melvin and Prins (2015) for a good summary of additional survey evidence on IIs' foreign equity portfolio hedging practices.

dollars in FX derivatives, which is equivalent to 30% of the commercial banking credit and 15% of GDP.

References

Alfaro, L., Calani, M. and Varela, L.: 2021, Currency hedging: Managing cash flow exposure, *Working Paper 28910*, National Bureau of Economic Research.

Andrews, D. W. K.: 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* **59**(3), 817–858.

Borio, C., McCauley, R. N. and McGuire, P.: 2022, Dollar debt in FX swaps and forwards: huge, missing and growing, *BIS quarterly review*, Bank for International Settlements.

Czech, R., Huang, S., Lou, D. and Wang, T.: 2021, An unintended consequence of holding dollar assets, *Bank of England working papers 953*, Bank of England.

Driscoll, J. C. and Kraay, A. C.: 1998, Consistent covariance matrix estimation with spatially dependent panel data, *The Review of Economics and Statistics* **80**(4), 549–560.

Du, W., Tepper, A. and Verdelhan, A.: 2018, Deviations from covered interest rate parity, *The Journal of Finance* **73**(3), 915–957.

Gabaix, X. and Koijen, R. S. J.: 2023, Granular instrumental variables, *Journal of Political Economy (forthcoming)* .

Greenwood, R. and Vayanos, D.: 2010, Price pressure in the government bond market, *American Economic Review* **100**(2), 585–90.

Ivashina, V., Scharfstein, D. S. and Stein, J. C.: 2015, Dollar Funding and the Lending Behavior of Global Banks, *The Quarterly Journal of Economics* **130**(3), 1241–1281.

Klingler, S. and Sundaresan, S.: 2019, An explanation of negative swap spreads: Demand for duration from underfunded pension plans, *The Journal of Finance* **74**(2), 675–710.

Liao, G. Y. and Zhang, T.: 2020, The Hedging Channel of Exchange Rate Determination, *International Finance Discussion Papers 1283*, Board of Governors of the Federal Reserve System (U.S.).

Melvin, M. and Prins, J.: 2015, Equity hedging and exchange rates at the London 4p.m. fix, *Journal of Financial Markets* **22**(C), 50–72.

Mercer: 2020, European asset allocation insights 2020, *Mercer Investing in the Future Paper Series*.

Newey, W. K. and West, K. D.: 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**(3), 703–708.

Yazdani, D.: 2020, Beyond their borders - evolution of foreign investments by pension funds, *Market research centre partner, PwC*.