

NBER WORKING PAPER SERIES

PREFERRED HABITATS AND TIMING IN THE WORLD'S SAFE ASSET

Alexandra M. Tabova
Francis E. Warnock

Working Paper 30722
<http://www.nber.org/papers/w30722>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2022, Revised May 2024

Alexandra Tabova (alexandra.m.tabova@frb.gov) is at the International Finance Division of the Board of Governors of the Federal Reserve System. Francis Warnock (few9t@virginia.edu) is at the University of Virginia's Darden Business School and National Bureau of Economic Research. The authors thank Ben Stagoff-Belfort for excellent research assistance and, for helpful comments, Leonardo Elias, Rohan Kekre, and Jesse Schreger (discussants) as well as Daniel Beltran, Carol Bertaut, John Burger, Ricardo Correa, Stephanie Curcuru, Wenxin Du, Rich Evans, Chris Gohrband, Robin Greenwood, Kinda Hachem, Karen Lewis, Gordon Liao, Juan Londono, Felipe Saffie, Rosen Valchev, Annette Vissing-Jorgensen, Veronica Cacad Warnock, Gwen Yu, Tony Zhang and participants in conferences or seminars at Bank of Italy, Bank of Finland, Bank of Canada, Federal Reserve Board, Darden Business School, European Economic Association, Midwestern Economic Association, Society for Economic Dynamics, CEBRA, NBER Asset Pricing and NBER International Finance and Macroeconomics meetings. Warnock thanks the International Finance Division of the Federal Reserve Board for ongoing hospitality. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, nor of any other person associated with the Federal Reserve System, nor of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Alexandra M. Tabova and Francis E. Warnock. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Preferred Habitats and Timing in the World's Safe Asset
Alexandra M. Tabova and Francis E. Warnock
NBER Working Paper No. 30722
December 2022, Revised May 2024
JEL No. F30,G11,G12

ABSTRACT

We build the first-ever comprehensive security-level dataset on the size, flows, coupon payments and returns of foreign and U.S. investors' Treasury portfolios. Private U.S. and private foreign investors hold longer-duration higher-return Treasuries, whereas foreign governments hold shorter-duration lower-return Treasuries, so all else equal private investors should earn higher returns than foreign governments. But all else is not equal. Foreign governments, even with their low-returns low-volatility portfolios, actually earn higher returns because U.S. and foreign private investors' returns are substantially reduced by poor timing. Moreover, we find that foreigners beat the (non- Fed) market, while U.S. investors earn a small and insignificant 19 basis points less than market returns. From our analysis, no investor-type earns significantly less than the (non-Fed) Treasury market, indicating that differential investors are not the source of the convenience yield. In terms of investor behavior more broadly, our novel dataset allows a direct comparison of the different investors in the Treasury market. Foreign private investors are similar to U.S. private investors and both behave differently and have different preferred habitats than governments.

Alexandra M. Tabova
Board of Governors of the Federal Reserve System
20th Street and Constitution Avenue N.W.
Washington, DC 20551
alexandra.m.tabova@frb.gov

Francis E. Warnock
Darden Business School
University of Virginia
Charlottesville, VA 22906-6550
and NBER
warnockf@darden.virginia.edu

1 Introduction

As the world transitions from quantitative easing to quantitative tightening and central banks reduce their sizeable bond portfolios, it would be useful to understand investors' behavior in the world's safe asset, U.S. Treasury bonds. Investors' actions help determine Treasury yields, which impact loan prices, not only in the United States (mortgages, corporate bonds, etc.) but all over the world (e.g., dollar bonds issued by foreign entities) and, relatedly, how U.S. monetary policy is transmitted domestically and to the rest of the world. Moreover, investors' actions determine the U.S. government's fiscal space and the country's international budget constraint. If investors have unlimited demand for Treasuries, the government could run larger budget deficits for longer, as price insensitive investors might willingly fund budget shortfalls. And if *foreign* investors have unlimited demand for Treasuries, the U.S. would not have to worry about its negative \$18 trillion net foreign asset position knowing that foreigners would continue to fund it at any price.

Investors' behavior in U.S. Treasuries is also crucial for our understanding of possible sources of *convenience yields*. There are many views, definitions and ways of calculating convenience yields. We group them into two. The traditional *liquidity premium* (plus collateral and safety) view - see, among others, Nagel (2016), Krishnamurthy and Vissing-Jorgensen (2012), van Binsbergen, Diamond and Grotteria (2022), Diamond and Van Tassel (2023), and Engel and Wu (2023) - is well represented by the opening sentence of Du, Im and Schreger (2018): "If investors value the liquidity and safety of government bonds, they may be willing to accept a lower yield to hold them over alternative investments that offer the same cash flows." Estimates vary, but papers in this liquidity premium view typically find that the U.S. Treasury convenience yield is 50 bps or less (and, relative to other advanced economies, the U.S. is not special).

Another portion of the convenience yield literature promotes a form of *investor differen-*

tiation in which some investors value the liquidity/safety much more than others and thus have a substantially higher convenience yield. Most notable contributions are Krishnamurthy and Lustig (2019) and Jiang, Krishnamurthy and Lustig (2021, 2023). In this line of the literature, "some investors" are foreigners, poor timing and poor returns are put forward as evidence of both inelastic demand and the convenience yield, and the estimates can be quite large (as large as 500 bps, often 200+ basis points).

That there is such a wide discrepancy between convenience yield estimates in the liquidity premium and investor differentiation views suggests deeper analysis is warranted. We do so by examining the investor segmentation view using techniques similar to its advocates (e.g., Krishnamurthy and Lustig (2019) and Jiang et al. (2023)). Specifically, for the Treasury portfolios of foreign private investors, foreign governments, U.S. private investors, the Federal Reserve and the market (i.e., all issuance), we compare buy-and-hold (or index) returns with actual returns, computed as the internal rate of returns (IRR) of cash flows. The former is based purely on the composition of Treasury portfolios; the latter takes into account both portfolio composition and the timing and magnitude of purchases (and, for the market, of issuance). The difference between actual and buy-and-hold returns is a measure of timing; in the investor differentiation view, actual returns much lower than buy-and-hold returns is taken as evidence of poor timing and a large convenience yield. Another useful comparison, an investor-type's returns versus market returns, gauges which investors beat or does substantially worse than the market.

While investors' behavior in the world's safe asset determines the transmission of Fed policy domestically and around the world and U.S. budget and international constraints, and would help understand sources of convenience yields, we actually know little about it for two main reasons. One, analysis is often of aggregate data, but *portfolio compositions* differ: within their Treasury bond portfolios, investors have preferred habitats (Figure 1, top

graph).¹ Over the period from 2003 to 2021 - the time span of most analysis in this paper - the duration of U.S. investors' Treasury bond portfolio fluctuated between 7 and 8 years, whereas foreigners' duration was shorter at 4 to 5 years. All else equal (although we will show that it is not), these preferred habitats translate directly into differential returns in the world's safe asset. For example, over the past two decades, average annual returns on 3-, 5- and 7-year Treasuries (computed using information from the Federal Reserve's H.15 release) were 2.0, 2.6 and 3.2 percent, respectively. That is, over the past two decades, a Treasury of 7-year maturity earned on average 120 basis points more per year than a 3-year Treasury. Since investor behavior is often inferred from an assessment of returns, not taking into account preferred habitats can lead to incorrect inference.

Two, while since 2003 annual security-level surveys of foreigners' Treasury portfolios can guide analysis, in general researchers have not had accurate information on the *timing and magnitude* of purchases of Treasuries. For example, the oft-used TIC S capital flow data (Figure 1, bottom graph) indicate that the bulk of foreign purchases of Treasuries have been by private foreign investors: in the TIC S data, every year from 2005 to 2019 private foreign investors purchased substantially more U.S. Treasuries than foreign governments, and the 15-year cumulative gap is substantial, amounting to \$3.5 trillion. But, in reality, that is not true: Each year in the decade starting 2005 foreign governments purchased more Treasuries than private foreign investors. And the gap is not \$3.5 trillion more purchases by private foreign investors, but \$1 trillion more by foreign governments. Prior to 2000, the TIC S data formed the basis for *any analysis* of U.S. portfolio flows. TIC S flows were the primary input into the Bureau of Economic Analysis' (BEA) Balance of Payments (BOP) flows that were, in turn, used to calculate the U.S. International Investment Position

¹We are pleased to have inserted in annual Treasury Department surveys a similar graph showing the remaining maturity of foreigners' (but not U.S. investors') Treasury portfolios, starting with the 2021 survey. See Exhibit 15 in <https://ticdata.treasury.gov/resource-center/data-chart-center/tic/Documents/shla2021r.pdf>.

(IIP), and the Federal Reserve's Financial Accounts of the United States (known as Flow of Funds, or FOF) were based on the BEA's BOP and IIP data. Unfortunately, Griever, Lee and Warnock (2001) and Warnock and Cleaver (2003) showed that TIC S reported foreign flows into Treasuries (and, hence, FOF and BEA flows too) were \$179 billion (or 44 percent) too high over the period January 1995 - March 2000. TIC S flow data were discontinued in February 2023.² Bertaut and Tryon (2007, henceforth BT), making use of infrequent positions surveys, created alternative *benchmark consistent* series in which flows are forced to be consistent with reported positions, but this is done in an ad hoc way: proportionally scaling flows up or down in an inter-survey period to eliminate any flows-positions discrepancy. Prior to 2000, surveys were every five years, meaning that BT had, for example, two data points - holdings at December 1994 and March 2000 - on which to base the (proportional) rescaling of TIC S reported flows to get flows plus valuation adjustments to equal March 2000 positions. For the purpose of assessing investors' timing, any flaws in the now-discontinued TIC S flow data are also in the BT data. And those flaws extend to U.S. investors' portfolios when they are calculated as a residual.

Our study addresses these two issues by building the first-ever comprehensive security-level dataset of the positions, flows, and income streams associated with the Treasury bond portfolios of the four investor types that span the entire market: U.S. investors, private foreign investors, foreign governments, and the Federal Reserve. Our results are both simple (there is nothing complicated in this paper) and striking. We find, perhaps not surprisingly once one thinks about it, that investors' preferred habitats determine the mean and volatility of annual buy-and-hold returns. U.S. investors had long-duration high-return high-volatility portfolios, foreign governments had the opposite (short duration, low returns, low volatility), and for most of the past two decades private foreign investors were somewhere in between.

²See <https://home.treasury.gov/data/treasury-international-capital-tic-system-home-page/tic-forms-instructions/tic-s-form-and-instructions>.

Perhaps more surprisingly, volatility-adjusted returns (i.e., Sharpe ratios) show that the portfolios of foreigners - both governments and private investors - performed better than the high-return high-volatility portfolios of U.S. investors; although these differences, while interesting, are not statistically significant.

Here is where all else is not equal. Yes, portfolio composition is important. But no one, not even the market, earns the buy-and-hold rate of return (RoR). To do so one of two conditions should hold: returns would have to be constant throughout the entire period under consideration or all purchases would have to have been made prior to the initial period. If either of those conditions hold, buy-and-hold (or index) returns would be ensured. Typically, however, neither condition holds. Returns vary through time and buying and selling occur. So to compute the returns actually earned by different investor types, we calculate IRRs that take into account both the composition of portfolios and the market, plus the timing and magnitude of purchases (and, for the market, of issuance). Our security-level dataset has the four components needed to compute IRRs: the market value of initial positions (the initial outlay), flows (interim outlays), face value of positions (to which we apply security-specific coupon rates to calculate interim payouts), and the market value of terminal positions (the final payout). We do this for all four investor types, which allows for a direct comparison of the returns earned by foreign and U.S. investors. Moreover, our IRR computations embed the composition effect, as they are built from the security-level portfolios. Finally, for each investor type, a comparison of its IRR and RoR is a measure of timing; if actual returns (IRR) are lower than the investor's own benchmark (its RoR), timing is poor because it reduces returns.³

Taking into account both investors' preferred habitats and the timing and magnitude of

³As always with an assessment of timing and performance, an appropriate benchmark is needed. We, following Barber and Odean (2000) and many others, use an 'own benchmark' that is determined by each investor type's portfolio composition. To do otherwise would be akin to gauging insurance companies' long-dated Treasury portfolio by comparing it with a 3-month bill, or grading a growth fund's performance by comparing with a portfolio of value stocks.

purchases, we find that the Treasury bond market does not earn market returns, as aggregate returns (the market IRR) are 77 basis points less per year than buy-and-hold returns (the market RoR).⁴ Moreover, this difference is much larger for private investors, whether U.S. (who earn 199 basis points less than buy-and-hold returns) or foreign (135 basis points less). In contrast, foreign officials' returns are only 36 basis points lower than their buy-and-hold returns.

In summary, investors' actual Treasuries returns differ materially from buy-and-hold returns, and this difference is quite large (and negative) for private investors. In fact, taking into account the composition of portfolios as well as the timing and magnitude of purchases, we find that foreign investors, whether private or official, earn no less on their Treasury portfolios than U.S. investors. And foreign governments earn significantly more than U.S. investors.

In a later portion of the paper we venture away from the pristine world of security-level data and bring monthly estimates of flows into the analysis. We view everything in the first portion of the paper as rock solid, whereas the analysis with monthly flows estimates is more speculative. We include it because it is useful for at least starting to ask whether different investors' demand can be described as inelastic. That question might seem narrow but is not: As the spectre of higher Treasury yields for longer looms, and one large entity (the Federal Reserve) reduces the size of its Treasury portfolio, which investors will step in and buy Treasuries? In this descriptive portion of the paper we find evidence that is suggestive of private foreigners - who are behind the bulk of foreign purchases over the past 5-10 years - and U.S. investors have elastic demand for Treasuries, but foreign governments do not. More specifically, using the higher frequency (monthly) dataset we show that U.S. and private foreign investors, but not foreign officials, increase their flows into U.S. Treasuries

⁴This is another way of saying that after the Treasury Department bond issuance increases, Treasury returns tend to be lower. A similar effect is evident in equities. For example, IPO issuance is elevated prior to equity market peaks (Ritter 1991, Pastor and Veronosi 2005)).

when CIP deviations (the synthetic sovereign dollar yield minus the Treasury yield) are low.⁵ The analysis of monthly flows suggests that private investors, whether U.S. or foreign, who together were behind 90 percent of net purchases in the 6-year pre-pandemic period (2014-2019), have demand for Treasuries that can be described as elastic. Additional analysis of the historical data suggests that as the Fed reduces the size of its balance sheet, private investors - especially U.S. investors - will be the counterparty.

The paper proceeds as follows. In the next section we discuss the related literature as well as caveats and constraints of our study. In Section 3 we create survey-consistent flows and assess the many different sources of data on foreigners' flows into and positions in U.S. Treasury bonds; the heavy lifting on this public service portion of our paper is relegated to an appendix. In Section 4 we calculate the returns on Treasuries earned by different investors, differentiating between the composition and timing effects, and compare with market returns. We also analyze how one could obtain very different results when assessing the 40-year period starting in 1980. In Section 5 we incorporate higher frequency flow data to examine two issues of current interest: how investors' flows adjust to changes in the spread between non-U.S. sovereign yields and Treasury yields and, when the Fed pares back its Treasury portfolio, who buys. Section 6 includes Treasury bills, which were excluded from our main analysis because annual surveys are ill-suited for examining securities with less than one year original maturity (results change in nuanced ways but not materially). Section 7 concludes. An online appendix includes much robustness analysis.

⁵For the higher frequency flow analysis we rely on survey-consistent monthly estimates of foreign flows based on Bertaut and Tryon (2007), Bertaut and Judson (2014) and our own calculations, coupled with data on monthly flows for the entire market and the Fed.

2 Related Literature and Caveats

We contribute to the literature on convenience yields, the non-pecuniary benefit of holding safe and liquid assets that can serve as substitute for money. In a closed-economy setting, Bansal and Coleman (1996) showed the importance of bond supply on convenience yields, while Krishnamurthy and Vissing-Jorgensen (2012) showed that the convenience yield on U.S. Treasuries - relative to other U.S. debt such as corporate debt bonds - is an important component of bond yields. Subsequent work - notably Engel (2016) and Du et al. (2018) - brought the notion to an open economy setting. There are important differences within this large and growing literature, as some, such as Engel and Wu (2023), do not posit that U.S. Treasuries are particularly special and instead focus more generally on all governments' bonds, whereas others, such as Krishnamurthy and Lustig (2019) and Jiang et al. (2021), emphasize the specialness of the dollar and Treasuries (and that foreign investors are different). And many if not all papers in this literature embed assumptions about investor behavior in safe assets. For example, demand for safe assets is at the heart of Kekre and Lenel (forthcoming), Bianchi et al. (2022), Valchev (2020), and Choi et al. (2024). Our study of investors' Treasuries portfolios can inform this literature.

Our paper is related to any work that differentiates between foreign and domestic investors. Some of this has focused on whether foreign demand for Treasuries is elastic.⁶ For example, in the model of Greenwood et al. (2023), foreign demand is elastic - a decrease in foreign bond yields prompts foreigners to purchase U.S. Treasuries - whereas central to the model of Jiang et al. (2021) is that foreign demand is inelastic. In the literature, at times the elasticity of demand is inferred from an analysis of performance and timing; if foreigners have poor timing in Treasuries (buying prior to price declines) or, relatedly, exhibit poor

⁶Our work is also tangentially related to literature on the effects of foreign investment on Treasury yields. In the theoretical model of Caballero, Farhi, and Gourinchas (2008), emerging market demand for Treasuries depresses U.S. long rates. Empirical estimates of foreigners' impact on Treasury yields are in Warnock and Warnock (2009), Beltran et al. (2013), and Wolcott (2020), among others.

performance over time (earning below market returns), it is inferred that their demand is inelastic (e.g., Krishnamurthy and Lustig 2019). Our work can directly assess whether there is a meaningful difference in the performance of domestic and foreign investors, something not yet done in the literature, and thus can speak to timing and the elasticity of demand, issues that remain unsettled in the literature. More generally, our assessment of data quality informs any study using capital flows data. For example, our assessment indicates that the flow data that underlie IRR calculations in Krishnamurthy and Lustig (2019) and subsequent papers through Jiang et al. (2021) are not up to the task (and have indeed been discontinued).

Our paper is also related to the broad set of research that differentiate bonds by maturity, such as the long literature on preferred habitats and the shorter but growing literature on the effects of quantitative easing (QE). On preferred habitats, the seminal papers are Culbertson (1957) and Modigliani and Sutch (1966). More recently, Vayanos and Vila (2021) deviate from a frictionless asset pricing model by assuming investors have preferred habitats;⁷ Krishnamurthy and Vissing-Jorgensen (2011) note that, when using that model to think about the effects of QE, it is an open question as to whether there are preferred habitats within an asset class (such as Treasuries). In a similar vein, our paper is related to QE analysis on how one large market participant - the central bank - affects the slope of the yield curve; see Swanson (2011), Greenwood and Vayanos (2014), and many others. Of course, that Philippon (2011, page 275) had to state that "[t]here was a time when macroeconomic textbooks used only one interest rate" indicates, as we all know, that the assumption of a single interest rate is common in papers ranging from single-country closed economy models to multi-country international finance models that assume interest rates are equalized across countries (e.g., Caballero et al. 2021).

Methodologically, our focus on IRRs follows the Dichev (2007) seminal work on the equity

⁷Kekre et al. (2022) incorporate preferred habitats to study the effects of monetary policy.

premium and Krishnamurthy and Lustig (2019), which applied IRR analysis to Treasuries. Dichev (2007), focusing on large U.S. equities, found that investors earned 130 basis points less per year than buy-and-hold returns. Our findings for the Treasury portfolios of private investors, whether U.S. or foreign, are similar in magnitude. Krishnamurthy and Lustig (2019), which helped propel an important part of the convenience yield literature (including but not limited to Jiang et al. (2021)), has eye-opening results and conclusions that support the notion that there is a large convenience yield that is due to foreigners. For example, in Krishnamurthy and Lustig (2019) foreigners accept returns that are 500 basis points per year lower than Treasury returns, and this finding becomes the heart of the theoretical model of Jiang et al. (2021). Such results are worth revisiting as they rely on since-discontinued inaccurate TIC S data, and those inaccuracies are mirrored in U.S. investors' flows that are calculated as a residual. We find that revisiting this analysis with better, more comprehensive data produces quite different results. Computing the convenience yield as the difference between investors' actual returns and the non-Fed market's actual returns - that is, an apples-to-apples comparison of IRRs - the convenience yield is not positive for foreigners and a paltry and insignificant 19 basis points for U.S. investors. There are other reasonable ways of computing convenience yields, for example (among many others) the van Binsbergen et al. (2022) focus on risky assets. We show that attributing convenience yield to foreigners' non-pecuniary benefits is not supported by the evidence. By measuring whose demand is (not) driving the Treasury convenience yield, our work can help inform theories of its deeper source and future.

Finally, our work informs literatures on the U.S. government and international budget constraints. For both, if investors - especially foreign investors - have inelastic demand for Treasuries, the budget constraint might be relaxed.

Our analysis comes with caveats and constraints. Our main dataset is constrained by the availability of confidential and mandatory security-level surveys on *foreigners'* U.S. Treasury

portfolios. The so-called SHL surveys have been conducted as of each June 30 since 2003 through the U.S. Treasury International Capital (TIC) System, so our main sample is annual (as of June 30) from 2003 to 2021.⁸ While it would be desirable to have a longer time series, we will show that data that pre-date our sample period are problematic.

That our observations are just once per year presents a few constraints. One, it precludes the inclusion of Treasury bills - Treasury securities with less than one year original maturity, which prior to the pandemic were only 15 percent of outstanding Treasuries - from much of our analysis. Thus, other than in Section 6 where we explicitly include Treasury bills, when we refer to Treasuries, we mean Treasuries of greater than one year original maturity.⁹ Two, annual data preclude the analysis of trading during the year. Three, among foreign investors the identification of the ultimate investor's country is not accurate if the investor uses a third-country custodian. More specifically, data are reported on a resident basis; that is, we observe the direct owner of these investments as reported by the custodians, but not necessarily the ultimate owner. This has two practical implications. One, while the survey data allow us to distinguish between holdings of foreign official institutions (e.g., central bank reserve managers) and holdings of private foreign investors, the distinction is not perfect because in some cases a central bank might use a custodian in a third country and, if so, some official holdings may be reported as private. In our opinion, informed by over twenty years working with these surveys, this blurring of official and private is not severe, but just in case we run a number of robustness checks. Two, a bigger issue is that because country attributions are subject to the nationality vs. residency issue studied in Warnock and

⁸See <https://www.treasury.gov/resource-center/data-chart-center/tic/Pages/shlreports.aspx> for the SHL reports. Over our sample period, TIC survey data are the primary input for BEA's IIP, so when summed our security-level holdings data add up to the official U.S. IIP data as reported by the BEA. The Fed's FOF relies on BEA's presentation, so our holdings data also sum to FOF holdings.

⁹Our dataset is of holdings. We do not have information on derivatives positions or repo activity. We exclude, unless otherwise noted, FRNs and TIPs, in part because duration cannot be computed for these securities, which are a small share (15 percent as of mid-2019) of outstanding Treasuries. In an online appendix we show that including these securities wherever possible does not alter our results.

Cleaver (2003), Bertaut, Bressler, and Curcuru (2019), and Coppola et al. (2021), country attribution can be problematic (as with any liabilities-based survey). To avoid the country-attribution problem we aggregate across foreign countries rather than assess country-level data.

The reader should be aware that while we use the term *investors*, we fully recognize that two Treasury market participants - the Fed and foreign governments - have motivations that differ from private investors'. The Fed does not attempt to maximize returns on Treasuries. As noted by the Federal Reserve Bank of New York System Open Market Account (SOMA) desk (<https://www.newyorkfed.org/markets/treasury-reinvestments-purchases-faq>): “In general, the Desk seeks to operate in a manner that is relatively neutral to the securities available for purchase and in a way that limits the potential for operations to affect normal market functioning, unless otherwise appropriate for efficient and effective implementation under the directive. As such, purchases of Treasury securities are conducted across a range of maturities and security types in rough proportion to the universe of Treasury securities outstanding.” In one period - Operation Twist from September 2011 until June 2012 - the Fed deviated from this by purchasing the long end and selling the short end. Similarly, foreign governments have broad objective functions that typically do not include maximizing returns on their Treasury portfolios. Foreign governments' demand for Treasuries are due to the world's main reserve asset's role in countries' reserve management and FX interventions, with motivations often divided into precautionary (building up defenses to reduce the probability of a future crisis), mercantilist (managing the exchange rate), and those that are natural by-products of other goals (e.g., managing inflation, smoothing business cycle fluctuations); see Arslan and Cantu (2019) for a useful discussion.

3 U.S. Treasury Portfolios: Stylized Facts

Security-level information on the overall Treasury bond market and the Fed’s portfolio is readily available.¹⁰ More difficult, even at the aggregate level, are data on foreigners’ Treasuries portfolios, in part because there are multiple sources of data on foreigners’ purchases of and positions in U.S. Treasury securities that, as we discussed with respect to Figure 1, can provide conflicting information.

Our study is based on data from the best source, annual security-level TIC SHL surveys on foreigners’ positions. From the surveys we can observe security-level positions and construct flows and coupon payments. Briefly - and please see the appendix for complete details - the security-level holdings data, collected annually (as of each June 30) since 2003, are mandatory, represent the universe to the extent it is known, and include a general security description and identifier, issue and maturity dates, coupon rate, and amount held reported at both face (which excludes price change effects) and market value. To confirm that we have included all bonds, for each annual survey from June 2003 through June 2021 we sum the holdings and compare to the surveys’ published aggregate amounts.¹¹ They match exactly. Confident we have all Treasury bond holdings, we use the security-level survey data to accurately calculate security-level flows as the change in the face value of holdings. With comprehensive data on foreigners’ and the Fed’s positions and flows and knowing the universe of all marketable Treasury bonds outstanding, we construct U.S. investors’ portfolios (at the security level) as a residual: total outstanding less foreign and Fed holdings.

The novel dataset enables us to provide a public service by creating survey-consistent

¹⁰See the U.S. Treasury Monthly Statement of the Public Debt on the TreasuryDirect website (<https://www.treasurydirect.gov>) and the Federal Reserve Bank of New York SOMA site (<https://www.newyorkfed.org/markets/soma-holdings>).

¹¹Over our sample period, the data are a primary input for U.S. official statistics; while the security-level data are confidential, various aggregations of *positions* are made publicly available in annual survey reports and those aggregations inform the BEA’s IIP as well as the U.S. current account (via income streams computed using the aggregate positions estimates).

series that can be used to check the quality of publicly available data on foreigners' positions in and net purchases of U.S. Treasuries. For *positions*, this assessment is straightforward. We find that since 2003 positions reported by other sources line up with the survey data. But to check data on *flows* requires another step: creating survey-consistent flows. From the survey data we directly, and therefore accurately, compute the security-level flows (i.e., net purchases). Summing the security-level flows produces a survey-consistent aggregate flow measure that can help assess reported flow data, and based on that we show which publicly available data on foreigners' flows should (and should not) be used. While we relegate much to an appendix, we briefly summarize here: Researchers should not use TIC S transactions data, which have been so problematic that they were discontinued in 2023. This warning applies to other sources that prior to 2000 were based on TIC S data, such as the BEA's BOP and the Fed's FOF data. Since 2000 and especially since 2003, data sources have improved in part because they have become based on the TIC SHL survey positions data rather than TIC S transactions data.

Some basic facts on positions and flows emerge from the dataset. On positions, for much of the past two decades, foreigners as a group have had the largest Treasury bond holdings (Figure 2, top graph). Their holdings, as a share of all marketable Treasury bonds outstanding, have mostly fluctuated between 40 and 50 percent, reaching a high of 58 percent in 2008. Over this period Fed holdings were often roughly 15 percent of the market, with noticeable sharp shifts during QE1 and the pandemic. U.S. investors' holdings, which we calculate as the residual, have fluctuated between 30 and 40 percent of the Treasury market over the past two decades. Within the set of foreign investors (Figure 2, bottom graph), foreign officials' Treasury holdings have been consistently larger than private foreigners'. That said, since 2012 foreign officials' holdings have been largely flat, mimicking global international reserves (which peaked in 2014Q2 and did not reach that level again until 2020Q3), while foreign private positions have been steadily increasing. As a result, as a

percent of the expanding Treasury market, the relative size of foreign official positions has steadily fallen and foreign private positions have increased a bit.

For flows, which for readability we depict in Figure 3 as 3-year moving sums, foreigners were the largest source of flows until 2016, when U.S. investors' flows began increasing and exceeding foreign flows. Fed purchases peaked 2011-2013 with the initial QE programs and by 2017 were zero or negative until large purchases resumed during the pandemic. Within foreign flows (Figure 3, bottom graph), every year since 2014 private foreign purchases of Treasuries have exceeded officials' purchases. Foreign governments, arguably the most important group early in our sample, have had zero net purchases for a full decade.

4 Portfolio Performance

In this section we present the actual returns earned by private U.S. and private foreign investors, as well as the Fed and foreign governments. We then assess two factors behind the observed returns: portfolio composition and timing.

4.1 Actual Returns

We assess investors' actual returns by calculating the internal rate of return (IRR); see Fabozzi (2014) for a textbook exposition, Dichev (2007) for an application to equities, and Krishnamurthy and Lustig (2019) for an application to Treasuries. Security-level data on each bond's market value, face value, and coupon yield, as well as flows that we compute from the security-level data, allow us to compute IRRs for a sample period of mid-2003 through mid-2021 using the following components:

- the initial investment is the sum of the market value of an investor type's holdings of each Treasury bond at the end of June 2003 (assumed to all be purchased at that time),

- intermediate contributions are our computed bond-by-bond net purchases from mid-2003 through mid-2021,
- intermediate distributions are coupon payments each period from mid-2004 through mid-2021, computed at the security-level from the bond's coupon yield and face value, and
- the final payout is equal to the sum of the market value of holdings of each Treasury bond as of end-2021.

Column 1 of Table 1 presents IRRs for foreigners (separating out private and official foreign investors), U.S. investors, and the Fed over the period from mid-2003 to mid-2021. We also show market and float (i.e., non-Fed) IRRs computed in the same manner as the other IRRs.

Result 1: Foreign governments have better returns than U.S. investors.

The first result is that foreign governments earn higher returns (2.66%) than private U.S. investors (2.27%). Based on bootstrapped standard errors in which the order of flows is shuffled 1000 times, the difference is statistically significant at the 1% level. In the rest of column 1, only the Fed's returns is statistically different from the return on the float (i.e. non-Fed) portfolio.¹²

We next delve into the portfolios and flows to understand the returns each investor-type earned. We do this in two steps. First, we compute average annual returns (RoR, or rate of return), which are based solely on portfolio composition; these are presented in Column 2. Then we examine the difference, by investor, of IRR and RoR (column 4). An IRR that is less than the corresponding RoR is indicative of poor timing (i.e., relatively more purchases prior to low returns).

¹²Note that the Fed remits any excess profits on its security holdings to the U.S. Treasury and does not attempt to maximize returns on Treasuries.

4.1.1 The Effect of Portfolio Composition

Figure 1 (top graph) foreshadowed this section: Investors have reasonably well defined preferred habitats in Treasury bonds. The differences in duration naturally produce differences in average annual returns.

Figure 4 shows the duration of the portfolios of private and official foreign investors, U.S. investors, the Fed and the overall market. Over the period 2003-2021, market duration (thick solid line) was between 5 and 6.5 years (with an average of 5.78). Among investor types, U.S. investors (top dashed-dotted line) had the longest duration for most of the period and in general stayed between 7 and 8.5 years (average 7.33). The duration of foreign private investors (dashed line) fluctuated between 6 and 7 with an average of 6.60), but since 2014 has trended up. Foreign governments (lowest line) had by far the shortest duration (mostly between 3.5 and 4.5 years, average 4.09). The Fed's duration (hashes) extended in the 2012-2014 period (following Operation Twist) and averages 6.05 years.

The different durations preferred by the different investor groups mean that, all else equal, average annual returns will differ. For the period from mid-2003 until mid-2021, average annual returns for 3-, 5- and 7-year constant maturity Treasury bonds were 2.01, 2.63 and 3.17 percent. Thus, based on preferred habitats we could guess that U.S. investors have substantially higher average annual returns - perhaps 100 basis points more per year - than foreign governments.

But we can be more precise by using the security-level data to calculate returns using standard returns index construction methodology and assuming investors reallocate on June 30th of each year (the day we observe foreign portfolios). In index construction, the returns of each bond are typically weighted by the relative size of the bond. Analogously, we weight the return on each Treasury bond by its size in each investor type's portfolio. That is, the weighted average rate of return on investor type i 's portfolio of U.S. Treasury securities from year t to year $t+1$, $\text{RoR}_{i,t+1}$, is calculated as follows:

$$\text{RoR}_{i,t+1} = \sum_{b=1}^n \omega_{b,i,t} \text{RoR}_{b,t+1} \quad (1)$$

where i denotes foreign official, foreign private, U.S. private, the Fed or the entire market. $\text{RoR}_{b,t+1}$ is bond b 's annual total rate of return (price changes plus interest) from year t to year $t+1$. The weight $\omega_{b,i,t}$ is investor i time t holdings of particular bond $H_{b,i,t}$ relative to her total holdings:¹³

$$\omega_{b,i,t} = \frac{H_{b,i,t}}{\sum_{b=1}^n H_{b,i,t}} \quad (2)$$

Results are in column 2 of Table 1. It is immediately apparent that the ordering of returns is very similar to that of the durations. That is, the composition of investors' Treasury portfolios are such that the Treasuries held by U.S. investors, who have the longest duration portfolio, had high average annual returns (4.25 percent). Foreign officials' Treasuries portfolio, with the shortest duration, had the lowest returns, 3.02 percent, while private foreign investors' Treasuries produced returns of 3.76 percent per annum.¹⁴ The returns orderings are in line with what we would expect from the preferred habitats, and the differences (for example, the portfolios of foreign officials earning 123 basis points less than U.S. investors) are in line with returns from the various constant maturity Treasury series.

The preferred habitats also determine volatility: The higher returns earned by U.S. and foreign private investors came with higher volatility. In fact, abstracting from the Fed, the ordering of Sharpe ratios (column 3) is reversed. Foreign officials' low-duration portfolios delivered the highest Sharpe ratio, while private investors' (both U.S. and foreign) higher-

¹³Our weights being at time t means that the bond must exist at time t to be included in our returns calculations. This is similar to standard bond index inclusion rules that require the bond to have been issued prior to index rebalancings.

¹⁴We perform two robustness checks for the return calculations. First, we shift Belgium's private holdings of Treasuries to foreign governments' portfolio; and second, we shift Cayman Islands' holdings from the foreign to the U.S. portfolio. The results are very similar.

duration portfolios had lower Sharpe ratios.

Result 2: The composition of Treasury portfolios is such that U.S. investors' high duration portfolios produce high returns and high volatility, while foreign officials' low duration portfolios have lower returns and lower volatility.

To get the highest return, private U.S. and foreign investors hold long-duration portfolios that had the most volatility and hence low Sharpe ratios, while the short-duration foreign official portfolio had a high Sharpe ratio. That said, preferred habitats notwithstanding, returns are volatile enough that no RoR or Sharpe ratio is statistically different from the market's.¹⁵

4.1.2 The Effect of Timing

That U.S. investors have the highest RoR but lowest IRR suggests that they have poor timing. Their portfolio composition (long duration bonds) would provide high (and volatile) returns, but their timing of purchases reduced those returns. More specifically, a comparison of actual and average annual returns provides a measure of timing. An IRR that is lower than the corresponding RoR indicates poor timing due to, for example, purchases made prior to low returns.¹⁶ Column 4 shows the IRR minus RoR. A number of striking results emerge.

Result 3: The market does not earn market returns.

The market IRR is 77 basis points lower than the market RoR. This is comparable to the Dichev (2007) finding that equity index returns are higher than IRRs - 130 basis points for large cap stocks – and hence the equity premium is lower than previously thought. Our

¹⁵These results should not be interpreted that one should hardwire that all investors' RoRs are equivalent to the market's, as is done in Jiang et al. (2023).

¹⁶To see this, consider a simple two-period example in which returns are positive 10% in the first period and negative 10% in the second period. With no trading, actual and buy-and-hold returns are identical (roughly zero). However, if the investor made additional purchases just prior to the low returns in period 2 - which would be poor timing - actual returns would be lower than buy-and-hold returns.

result indicates that actual Treasury bond returns are lower than index returns. Another way to put this: increased Treasury issuance tends to be followed by low returns.¹⁷

Result 4: Investors' timing is such that foreign governments' returns are reduced only slightly by timing, while U.S. investors lose almost 200 bps per year to poor timing.

Investors' returns are reduced by poor timing: In all non-Fed rows in Table 1 the difference between IRR and RoR is negative and significant. Private investors have particularly poor timing, with the IRRs of U.S. and foreign private investors being 199 and 135 basis points, respectively, lower than their RoRs. The timing of foreign officials, in contrast, lowers their returns only 36 basis points. And, abstracting from the Fed, as with Sharpe ratios, the ordering of IRRs is foreign official, foreign private, and U.S. private, the exact opposite ordering of RoRs. U.S. and private foreign investors have volatile high duration portfolios that should produce high returns, but poor timing substantially (and significantly) reduces their actual returns.

Graphically, the poor timing of private investors is depicted in the top row of Figure 5, which shows positions, flows and subsequent returns. U.S. private investors' strong purchases in 2005, 2010, 2012 and 2016 were all followed by low returns. For foreign private investors, strong purchases 2010-2014 were followed by relatively low returns. In contrast, foreign officials' sizeable purchases (middle left chart) were early in the sample, when Treasury returns were quite strong; since 2014 foreign governments' purchases have been roughly zero, meaning that in the second half of the sample (when Treasury returns were lower) foreign governments essentially received the buy-and-hold return. The Fed's good timing (middle right graph) is due largely to pre-pandemic QE purchases, with purchases in 2010 and 2014 being followed by strong returns. The figure also shows in the bottom chart that substantial Treasury issuance, such as in 2010/11 and 2020/21, is followed by relatively low

¹⁷An equity analogy is that IPOs tend to be high just prior to low returns (Ritter 1991, Pastor and Veronesi 2005).

returns. Another way of putting this is that the market does not earn the market return, suggesting that one should not compare investors' returns with market RoRs, as is done in Krishnamurthy and Lustig (2019) and Jiang et al. (2023).

4.2 Reconciliation with Existing Findings

Our findings are quite different from prominent findings in the convenience yield literature. In a series of papers, Krishnamurthy and Lustig (2019) and Jiang et al. (2021, 2023) have argued that the convenience yield on U.S. Treasuries is quite large, from 500 basis points in early estimates to a still sizeable 200 basis points in later work. Jiang et al. (2023) uses similar techniques to ours and finds that since 1980 foreigners' annual returns are 326 bps lower than the market's index returns (that is, foreigners have very poor timing) and 162 bps lower than U.S. investors' returns (that is, other investors perform better).

Results that suggest poor timing by foreigners are due to a combination of two things. One, any sample to pre-dates 2003 uses data derived from TIC S flow data that were flawed enough that the U.S. government discontinued them in 2023. There are adjustments researchers have made to TIC S flows to make them "benchmark consistent" – that is, restating the flows so that flows plus returns equals the change in positions. That flows plus returns did not equal the change in positions led to inaccurate calculations of U.S. exorbitant privilege (see Curcuru et al. (2008, 2013)), so having a benchmark-consistent dataset of flows, positions and returns is a positive development. However, adjustments to flows prior to 2000 were based on one observation of *positions* every five years. Here it is instructive to quote Bertaut and Tryon (2007, page 12): "the unobserved measurement errors...about which we have essentially no prior information. We simply do not know whether these errors are serially correlated or not, whether they have zero mean or not, or what their comparative magnitudes might be. In this situation any solution is likely to be almost entirely ad hoc." Because the goal of benchmark-consistent series is to ensure that flows plus returns equal the

change in position, any ad hoc adjustment that produces that result would suffice. Bertaut and Tryon chose to assume that the measurement errors (in flows) are uniformly distributed throughout each inter-survey period. That is, any problem with the now-discontinued TIC S flow series is also in benchmark-consistent series, with the sole exception being that in the latter flows plus returns equals the change in positions. Benchmark-consistent flows uniformly adjust reported TIC S flows up or down based on the inter-survey period. Appendix Figure A6 shows one such inter-survey period, January 1995 - March 2000. The upshot: If we do not think TIC S flows are appropriate for assessing investors' timing, we should not believe benchmark-consistent flows (TIC S times a uniform error) are.

The second reason analysis that starts in 1980 suggests poor foreign timing is that foreign positions were miniscule in the 1980s, when Treasury returns were extremely strong. Figure 6 depicts this. In the early 1980s, when Treasury annual returns were 15 percent or more, foreigners' Treasury positions were virtually non-existent (upper left chart). The first meaningful purchases were in the mid- to late 1990s. Thus, analysis of investors' actual returns from 1980 to 2021 would include a decade-long period when foreigners had near zero positions and the Treasury market produced high returns. In fact, if one examines performance in each half of that four-decade period (Table 2), foreigners performed fine in the first half (1980-2000) and in the second half (2001-2021).

Thus, poor performance by foreigners is in part due to asking too much from existing data series but also because foreign positions were very small in the 1980s. Any analysis that assesses the early period (1980s and 1990s) separately from the later period would come to a different conclusion, as would analysis that starts once foreigners meaningfully entered the Treasury market.

5 Addressing Two Timely Questions with Higher Frequency Data

In this section we step away from our annual security-level data and use the higher frequency benchmark-consistent data to address two timely questions. One, while our main analysis is entirely U.S.-based as it examines investors' performance in the U.S. Treasury market, because we do not have the data necessary to do such careful analysis of foreign markets, here we ask a basic question: What is the relationship between investor flows and past spreads between foreign and U.S. yields? Two, in periods when the Fed reduced the size of its Treasury portfolio, as it has recently started to do, who steps in and buys?

5.1 Flows and Yields

Does the volume of investors' flows into Treasuries react to changes in the spread between non-U.S. sovereign yields and Treasury yields? Answering this requires data at a higher frequency than annual, so we turn to monthly benchmark-consistent flows; see the appendix for details. In the monthly flows dataset, all long-term Treasuries are included (i.e., including FRNs and TIPs, as aggregate data do not permit excluding them). We follow the Krishnamurthy and Lustig (2019) analysis of foreign flows but use benchmark-consistent flows and expand the analysis to include U.S. investors. Specifically, we run separate regressions for each investor type (foreign private, foreign official, U.S. private, and the Fed) flows and for each lag k of the CIP deviation - synthetic dollar sovereign yield minus U.S. Treasury yield - as follows:¹⁸

$$\text{Flow}_t/\text{Position}_{t-1} = \alpha + \beta \text{CIP}_{t-k} + \epsilon_t \quad (3)$$

¹⁸Data on CIP deviations, also called the synthetic yield difference, are from Du and Schreger (2016) and Du, Im, and Schreger (2018), where country j 's synthetic dollar sovereign yield is constructed using data on the forward premium.

In Figure 7 we plot the β coefficients from each time series regression we run for each lag k . Each plot in Figure 7 shows, from 13 time series regressions of flows regressed on lagged 10-year G10 CIP deviation (CIP) and a constant, the β coefficients of CIP deviation lags from 0 to 12 months. The flows are scaled by lagged positions for each investor type, and the results per investor type are shown as follows: foreign private (top left), foreign official (top right), U.S. private (bottom left), and the Fed (bottom right). The sample period is January 1999 through March 2021.¹⁹ CIP deviation is computed as a weighted average of G10 CIP deviations, with the weights coming from each country's holdings of Treasuries. Plotted are coefficients (the dot in each graph), two standard deviation error bands (the vertical lines), and the statistical significance at the 1, 5 or 10 percent level (the stars).

The evidence indicates that private investors, whether foreign or U.S., purchase Treasuries after the CIP deviation is low, that is when the synthetic dollar yield is low relative to the Treasury yield. For private foreigners (top left), this relationship holds at 0-5-month lags, while for U.S. investors (bottom left) the relationship is evident for all lag lengths. In contrast to private U.S. and foreign investors, foreign officials (top right) do not seem to adjust their net purchases based on CIP deviations.

Overall, the evidence in Figure 7 indicates that private foreign and U.S. investors, but not foreign governments, react to changes in CIP deviations. Specifically, when Treasury yields are high relative to the synthetic sovereign yield, private investors, both U.S. and foreign, purchase more Treasury bonds. To put this another way, when U.S. Treasuries are inexpensive relative to synthetic sovereign yields, private investors increase their Treasury purchases.²⁰

¹⁹The CIP deviations data, available at <https://sites.google.com/view/jschreger/CIP>, begin in 1997, but data for the euro area do not start until the euro's inception, so our analysis starts in January 1999. The data are currently available through March 2021. Note that we show results that exclude the initial Covid shock (March 2020 - June 2020); including that period modifies the results only very slightly, making the Fed's reaction insignificant.

²⁰In untabulated results, we find similar results for duration. That is, private investors alter the duration of their portfolios when sovereign spreads change, foreign governments do not.

5.2 Who Buys from Whom?

A timely question as central banks unwind QE policies and hence reduce the size of their bond portfolios is who will step in and buy. We cannot answer this question definitively, but our comprehensive data and analysis of investor behavior help shed light on this question.

First, we can learn from a graphical representation of positions and flows by investor type, shown in Figure 8 (best viewed in color) using quarterly flow data for foreigners and our computed U.S. investor series. The top panel shows positions, which remind us of relative sizes in the Treasury market. Foreign officials are quite large, but their positions have plateaued since 2012; foreign governments have not been actively adding to their Treasury portfolio for about a decade because they have not been accumulating international reserves. On the other hand, private holdings, both U.S. and foreign, have grown substantially. Foreign private holdings, quite small until about 2010, are now sizeable but still much smaller than the largest holder, private U.S. investors. This evolution of Treasury holdings suggests that when thinking about who might step in as the Fed reduces its bond portfolio, the focus should be on private investors, whether U.S. or foreign.

Flows data and analysis add to our understanding. First, Figure 8 (bottom panel) shows that foreign governments (red bars) provided the majority of Treasury purchases until the GFC and still had sizeable purchases until 2012. More recently, however, foreign governments' purchases and sales have offset through time and private (both U.S. and foreign) investors' purchases have become more important: between 2014 and 2019 private investors were behind 90 percent of net purchases. The graph also shows the Fed's flows and highlights that periods of negative Fed flows are limited, with the QT periods from October 2017 until August 2019 and again starting June 2022 being most prominent. Note that in these QT periods the Fed did not actively sell Treasuries but allowed maturing securities to roll off their balance sheet (which is a negative flow). Treasury supply was increasing, so someone must have been buying; Figure 9 shows that during QT periods the primary purchasers were U.S.

investors. Second, Table 3 provides correlations that are embedded in the flows data. From these correlations, the most obvious counterparts to the Fed are private investors, whose flows are significantly negatively correlated with Fed flows whether the Fed is expanding or reducing its portfolio.²¹

Our earlier findings on preferred habitats in the Treasury market can further support the point that private investors, both U.S. or foreign, are most likely to step in as the Fed reduces its Treasury holdings. The Fed's Treasury portfolio has a long duration, a portion of the yield curve favored by private investors (Figure 4). One large group - foreign governments - has not purchased Treasuries in meaningful amounts for about a decade and focus on the shorter end, so they are not likely to step in and buy in a meaningful way when the Fed reduces its balance sheet.

A related timely question is: Given heightened geopolitical tensions and recent U.S. financial sanctions impacting access to foreign reserves, how important might a potential diversification of reserves away from Treasuries be going forward? Our data and analysis shed some light on this question. We show that the landscape has changed quite a lot: Foreign governments' positions have been flat for more than a decade while private positions (U.S. and foreign) have grown substantially (Figure 8, top panel). As a result, as a percent of the expanding Treasury market, the relative size of foreign official positions has steadily fallen from almost 50 percent in 2008 to less than 20 percent in 2022 (Figure 10). Relatedly, since roughly 2014 the majority of Treasury purchases have been by private investors, while foreign governments have not purchased in meaningful amounts because they have stopped accumulating additional foreign exchange reserves. Given the steadily diminishing role of foreign officials in the Treasury market, their diversification might matter less for Treasury

²¹The flows in Table 3 are monthly. Quarterly flows provide similar conclusions but with the correlation between U.S. private and foreign private flows becoming positive and weakly significant and that between U.S. private and foreign official increasing in significance (and still negative). For either monthly or quarterly flows, results excluding the 2017-2019 QT period are similar but moderately weaker in significance.

demand going forward than is commonly thought.

6 Including Treasury Bills

We have omitted Treasury bills for the simple reason that annual surveys are ill-suited to shed light on holdings of securities that have less than one year original maturity. No Treasury bill spans two surveys, and any 30-day bill bought between survey dates will usually fall out of the portfolio by the next survey. So, while Treasury bills are included in the survey data, because surveys are annual we do not include them in our main analysis.

In this section we include bills. We first note that they are small (Figure 11), both in terms of outstanding Treasuries (14 percent in 2019; 19 percent in 2021) and in investors' holdings. That said, during crises they have increased importance, with their share peaking around the GFC - when they were 30 percent of all Treasuries outstanding, 43 percent of U.S. investors' Treasuries portfolio, and 36 percent of foreign private investors' - and increasing again during the pandemic.

Including bills alters the duration pictures a bit (Figure 12). Not surprisingly, including short-dated securities reduces duration. And, while foreign investors' portfolios still have lower duration than U.S. investors', for much of the period private foreign and U.S. investors had very similar durations. This was true, even for bond and notes, at the end of our sample - see the last observation of Figure 4 - but is true throughout the sample once bills are included. The takeaway is that foreign private investors (15 percent of holdings) are similar to U.S. investors, a point that comes through in our analysis of bonds and notes and even more strongly when bills are included.

Table 3 replicates the Table 1 RoR and IRR calculations but includes bills. We cannot include bills at the security level in these annual calculations, so they enter in aggregate

weighted by their by-investor-type shares.²² RoRs are lower than in Table 1, but the RoR and Sharpe ratio orders are unchanged and IRRs are pretty similar. As in Table 1, U.S. and foreign private timing erodes their returns by substantial amounts of 136 and 97 basis points, respectively.

Overall, including bills alters the duration picture a bit - private investors, whether U.S. or foreign, have similar durations - but does not materially affect the returns and timing analysis.

7 Conclusion

Investors' behavior in the U.S. Treasury bond market - the nature and evolution of Treasury portfolios - is important but has been inferred from aggregate statistics that can be less than pristine. We directly observe Treasury portfolios at the security level and find that preferred habitats determine buy-and-hold returns - average annual returns are a function of duration and hence risk - while the timing and magnitude of purchases substantially reduces the actual returns earned by *private* investors. Specifically, U.S. investors have long duration Treasury portfolios that produce high returns with high volatility, but the timing of their purchases reduces their actual returns by about 200 basis points a year. Foreign governments are at the other end of the spectrum, with short duration portfolios that produce lower returns with less volatility (so that their Sharpe ratios are no lower than U.S. investors), and their timing is substantially better. Private foreign investors are in between when analysis is based on Treasury bonds and notes, and much more like U.S. investors when bills are included. Moreover, private foreign investors, like U.S. investors, appear to be sensitive to sovereign spreads, whereas foreign governments are not.

²²Specifically, RoR is security-level from Table 1 column 2 times the long-term share plus a bills return times the bills share. IRR cash flows are as in Table 1 column 1, with bills added (initial and ending positions, interim purchases, and "streams" computed as lagged face value times a bills return).

Our results have implications for a pressing question as we transition to a post-QE world: Who will buy Treasuries as the Fed reduces the size of its portfolio? Perhaps foreign governments, but they have not materially added to their Treasury portfolios in almost a decade. Moreover, they tend to hold shorter duration bonds, while the Fed's portfolio is tilted more towards longer durations. More likely it will be private investors, whether U.S. or foreign, whose purchases react to yields and whose portfolios are tilted towards longer duration bonds.

Our results have implications for many literatures. The most obvious connection is the blossoming convenience yield literature which, while well entrenched since Krishnamurthy and Vissing-Jorgensen (2007, 2012), is still evolving. A straightforward takeaway is that the convenience yield literature should not rely on foreigners being different. Private investors, whether U.S. or private, behave differently and have different preferred habitats than governments. A stronger statement is that differential investors within the Treasury market is not the source of the convenience yield: From our analysis, no investor-type earns significantly less than the (non-Fed) Treasury market. Rather than relying on investor differentiation, our results point to notions (and calculations) of convenience yields that apply to all investors, such as van Binsbergen, Diamond, and Grotteria (2022), Engel and Wu (2023), Piazzesi, Rogers, and Schneider (2021), and Krishnamurthy and Vissing-Jorgensen (2012). Similarly, the Du, Im, and Schreger (2018) CIP deviations apply to all investors. There may well be a convenience yield, but our evidence questions whether it applies differentially.²³

Our paper also has important implications for anyone using data on international flows and positions. There are multiple data sources, some that we show should not be used (and, in fact, have recently been discontinued), others that are more accurate but can still be

²³Another notion of convenience yield is a comparison of returns on U.S. and other governments' bonds. We do not have the data necessary to compute IRRs in non-U.S. markets. But a simple examination of index returns suggest that since 1986 global exUS bonds returned 6.49 percent per year, compared with a Treasury index RoR of 6.09 percent, a small gap of 40 basis points that is in line with the van Binsbergen et al. (2022) estimates.

confusing (see, for example, data discussions in Meng and van Wincoop (2020) and Chari, Dilts Stedman, and Lundblad (2021)). Our advice for anyone using data on international flows and positions is straightforward: Check data quality, if possible, and know what the series represent.

References

- Arslan, Yavuz and Carlos Cantú (2019). The size of foreign exchange reserves. In Reserve management and FX intervention, BIS Papers No. 104: 1-24.
- Bansal, Ravi, and Wilbur John Coleman (1996). A monetary explanation of the equity premium, term premium, and risk-free rate puzzles. *Journal of Political Economy* 104(6): 1135-1171.
- Barber, Brad M., and Terrance Odean (2000). Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *Journal of Finance* 55: 773–806.
- Beltran, Daniel O., Maxwell Kretchmer, Jaime Marquez, and Charles P. Thomas (2013). Foreign Holdings of U.S. Treasuries and U.S. Treasury Yields. *Journal of International Money and Finance* 32(1): 1120-1143.
- Bertaut, Carol C., Beau Bressler, and Stephanie Curcuru (2019). Globalization and the Geography of Capital Flows. FEDS Note, Washington: Board of Governors of the Federal Reserve System, September 6, 2019, <https://doi.org/10.17016/2380-7172.2446>.
- Bertaut, Carol C., and Ruth A. Judson (2014). Estimating U.S. Cross-Border Securities Positions: New Data and New Methods. International Finance Discussion Papers 1113. Board of Governors of the Federal Reserve System (U.S.).
- Bertaut, Carol C., and Ralph W. Tryon (2007). Monthly Estimates of U.S. Cross-Border Securities Positions. International Finance Discussion Papers 910. Board of Governors of the Federal Reserve System (U.S.).
- Bianchi J., S. Bigio, and C. Engel (2022). Scrambling for Dollars: International Liquidity, Banks and Exchange Rates. National Bureau of Economic Research Working Paper

29457.

Bureau of Economic Analysis (2019). How BEA Aligns and Augments Source Data From the U.S. Treasury Department for Inclusion in the International Transactions Accounts. Survey of Current Business 99(7), <https://apps.bea.gov/scb/2019/07-july/0719-treasury-international-transactions.htm>

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2008). An Equilibrium Model of “Global Imbalances” and Low Interest Rates. American Economic Review 98(1): 358 - 393.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2021). Global Imbalances and Policy Wars at the Zero Lower Bound. Review of Economic Studies 88: 2570 - 2621.

Chari, Anusha, Karlye Dilts Stedman, and Christian Lundblad (2021). Taper Tantrums: Quantitative Easing, Its Aftermath, and Emerging Market Capital Flows. Review of Financial Studies 34(3): 1445–1508.

Choi, Jason, Rishabh Kirpalani, and Diego Perez (2024). US Public Debt and Safe Asset Market Power. working paper.

Coppola, Antonio, Matteo Maggiori and Brent Neiman, and Jesse Schreger (2021). Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens. Quarterly Journal of Economics 136(3): 1499-1556.

Culbertson, J. (1957). The term structure of interest rates. Quarterly Journal of Economics 71:485–517.

Curcuru, Stephanie, Tomas Dvorak, and Francis E. Warnock (2008). Cross-border Returns Differentials. Quarterly Journal of Economics 123(4): 1495 - 1530.

- Curcuru, S., C. Thomas, and F. Warnock (2013). On Returns Differentials. *Journal of International Money and Finance* 36: 1-25.
- Diamond, W., and P. Van Tassel, 2023. Risk-Free Rates and Convenience Yields around the World. working paper.
- Dichev, Ilia (2007). What Are Stock Investors' Actual Historical Returns? Evidence from Dollar-Weighted Returns. *American Economic Review* 97(1): 386-401.
- Du, Wenxin and Jesse Schreger (2016). Local Currency Sovereign Risk. *Journal of Finance*, 71, 1027-1070.
- Du, Wenxin, Joanne Im, and Jesse Schreger (2018). The U.S. Treasury Premium. *Journal of International Economics*, 112, 167-181.
- Engel, Charles (2016). Interest Rates, Exchange Rates and the Risk Premium. *American Economic Review* 106(2): 436-474.
- Engel, Charles, and Steve Pak Yeung Wu, 2023. Liquidity and Exchange Rates: An Empirical Investigation. *Review of Economic Studies* 90(5), 2395-2438.
- Fabozzi, Frank (2014). *Bond Markets, Analysis, and Strategies*. 9th edition. Pearson.
- Greenwood, Robin, Samuel G. Hanson, Jeremy C. Stein, Adi Sunderam (2023). A Quantity-Driven Theory of Term Premia and Exchange Rates. *Quarterly Journal of Economics* 138(4): 2327 - 2389.
- Greenwood, Robin, and Dimitri Vayanos (2014). Bond Supply and Excess Bond Returns. *Review of Financial Studies* 27(3): 663-713.
- Griever, William, Gary Lee, and Francis E. Warnock (2001). The U.S. System for Measuring Cross-Border Investment in Securities: A Primer with a Discussion of Recent Developments, *Federal Reserve Bulletin* 87: 633-650.

- Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig (2021). Foreign Safe Asset Demand and the Dollar Exchange Rate. *Journal of Finance* 76(3): 1049-1089.
- Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig (2023). The Rest of the World's Dollar-Weighted Return on U.S. Treasurys. *IMF Economic Review*.
- Kekre, Rohan, and Moritz Lenel (forthcoming). The flight to safety and international risk sharing. *American Economic Review*.
- Kekre, Rohan, Moritz Lenel, and Federico Mainardi (2022). Monetary Policy, Segmentation, and the Term Structure. Working paper.
- Krishnamurthy, Arvind and Hanno Lustig (2019). Mind the Gap in Sovereign Debt Markets: The U.S. Treasury basis and the Dollar Risk Factor. In Jackson Hole Economic Symposium (2019): 445-507.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen (2007). The Demand for Treasury Debt. NBER Working Paper 12881.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen (2011). The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity*.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen (2012). The Aggregate Demand for Treasury Debt. *Journal of Political Economy* 120(2): 233 - 267.
- Meng, Gao, and Eric van Wincoop (2020). A Decomposition of International Capital Flows. *IMF Economic Review* 68: 362-389.
- Modigliani, F., and R. Sutch (1966). Innovations in interest-rate policy. *American Economic Review* 56:178-197.

- Nagel, S., 2016. The Liquidity Premium of Near-Money Assets. *Quarterly Journal of Economics* 1927 - 1971.
- Pastor, Lubos, and Pietro Veronesi (2005). Rational IPO Waves. *Journal of Finance*: 60(4): 1713-1757.
- Philippon, Thomas (2011). Comment on The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity*.
- Piazzesi, Monika, Ciaran Rogers, and Martin Schneider (2021). Money and Banking in a New Keynesian Model. working paper.
- Ritter, Jay R. (1991). The Long-Run Performance of initial Public Offerings. *Journal of Finance* 46(1): 3-27.
- Swanson, Eric (2011). Let's twist again: A high-frequency event-study analysis of Operation Twist and its implications for QE2. *Brookings Papers on Economic Activity* 1:151–88.
- Thomas, Charles P., Francis E. Warnock, and Jon Wongswan (2004). The Performance of International Portfolios. *International Finance Discussion Papers* 817. Board of Governors of the Federal Reserve System (U.S.).
- Valchev, Rosen (2020). Bond Convenience Yields and Exchange-Rate Dynamics. *American Economic Journal: Macroeconomics* 12: 124-166.
- van Binsbergen, Jules H., William Diamond and Marco Grotteria (2022). Risk Free Interest Rates. *Journal of Financial Economics* 143(1): 1-29.
- Vayanos, Dimitri, and Jean-Luc Vila (2021). A Preferred-Habitat Model of the Term Structure of Interest Rates. *Econometrica* 89(1): 77-112.

Warnock, Francis E. and Chad Cleaver (2003). Financial Centers and the Geography of Capital Flows, *International Finance* 6(1): 27 - 59.

Warnock, Francis E., and Veronica Cacadac Warnock (2009). International Capital Flows and U.S. Interest Rates, *Journal of International Money and Finance* 28: 903-919.

Wolcott, Erin L. (2020). Impact of Foreign Official Purchases of US Treasuries on the Yield Curve. *AEA Papers and Proceedings* 110: 535 - 40.

Appendix: Data on U.S. Treasury Portfolios

We divide this discussion into pre- and post-2003, because data on foreigners' portfolios fundamentally changed in 2003.

We note at the outset that two components of Treasury portfolios are easily obtained at the security level. The "market" - that is, security-level data on the every Treasury bond - is available at <https://www.treasurydirect.gov>. And security-level data on the Fed's Treasury portfolio - both the portfolio at different points in time and the new purchases or sales - are publicly available at <https://www.newyorkfed.org/markets/soma-holdings>. With those two components, data on foreigners' portfolios would enable the calculation of U.S. investors' security-level portfolios as the residual.

However, data on foreigners' portfolios has been more difficult to obtain and understand. We start with the gold standard, security-level data from annual surveys that started in 2003 and are the basis for much of the analysis in this paper. Then, because there are multiple additional sources of data on foreigners' purchases of and positions in U.S. Treasury securities and no clear way for researchers to ascertain data quality, we present and assess the various sources.

Data Since 2003

Security-Level Portfolio Data on Foreigners' Treasury Portfolios

Since 2003, security-by-security data on foreigners' Treasury portfolios - the amount foreigners hold of each and every Treasury security - are collected annually by the U.S. Department of the Treasury as part of the TIC reporting system.²⁴ The security-level data underlie the annual U.S. TIC reports (the so-called SHL reports) of foreign holdings of U.S. securities and feed into official BEA data on the U.S. international investment position (and thus the Fed's FOF data). The main reporters are U.S.-resident custodians (including brokers and dealers), which must report all U.S. securities they hold on behalf of foreign residents (including in their own foreign subsidiaries and affiliates). Given the mandatory reporting, the holdings data are comprehensive; they capture the entire foreign portfolio of U.S. Treasuries at the individual security level.

The security-level data are annual since 2003, reported as of June 30 of each year for each foreign country's holdings of each security. The survey data distinguish between holdings of foreign official institutions (e.g., central bank reserve managers) and holdings of private investors. Survey data are reported on a resident basis; that is, on the direct owner of these investments as reported by the custodians, but not the ultimate owner.

We conduct a number of checks to ensure that we have a comprehensive and accurate security-level dataset. For amounts issued and outstanding, we cross check against data from ICE BofA Merrill Lynch U.S. Treasury indices and TreasuryDirect. We also make

²⁴Griever, Lee, and Warnock (2001) discuss the origins of the TIC system. Briefly, in the early 1970s public concern about the rise in European and Japanese investors' U.S. investments, as well as about accumulation by oil-producing countries, prompted the first modern benchmark survey to measure foreign holdings of U.S. long-term securities as of year-end 1974. It was recognized that without benchmark surveys, the TIC system could not accurately identify the countries that were holding U.S. securities or provide much information on the actual securities being purchased. To address these shortcomings, Congress passed the Foreign Investment Study Act of 1974 (Public Law 93-479), which evolved into the International Investment and Trade in Services Survey Act (22 U.S.C. 3101 et seq.). The latter act stipulates, among other things, that a comprehensive benchmark survey of foreign portfolio investment in the United States be conducted at least once every five years and that information collected under the authority of the act be published for use by the general public and by U.S. government agencies. Such surveys were conducted every 5 years from 1974 through 1999 (with the 1999 survey being conducted in March 2000 to avoid possible Y2K complications). Since 2003 the surveys have been conducted annually.

sure we correctly capture security re-openings (when the U.S. Treasury issues additional amounts of a previously issued security; the reopened security has the same maturity date and coupon interest rate). We do this for each annual survey from June 2003 through June 2021. To confirm that we have included all bonds, we first sum the holdings and compare to the surveys' published aggregate amounts. They match exactly. We also, from the surveys' prices and payment terms, calculate each bonds' total return, yield-to-maturity, and duration at each end-June date. The prices and payment terms are the same or very close to the ones reported per security in the constituent files of the BofA Merrill Lynch index. The security-level data we use represent the universe, to the extent it is known.

Non-Security Level Data on Positions

Since 2003 the end-June TIC SHL data have been the primary input for BEA's International Investment Position, or IIP, data.²⁵ As Figure A1 (top panel) shows, since the SHL surveys became annual (in 2003), for *positions* there are no discrepancies between these series. These sources also provide a split between foreign official and private foreign investors, and are in agreement not only on the overall amount of foreign holdings, but also on the split between official and private investors (Figure A1, bottom panel). There are minor differences - BEA puts a little more in private holdings, because it moves holdings of international and regional organizations to private - but data on foreigners' holdings of U.S. Treasuries are pretty similar regardless of the source. This makes sense because since 2003 aggregations from the annual security-level TIC SHL have been the primary inputs into other sources, so minor differences arise only when a source creates estimates at a frequency higher than annual or when a source publishes data prior to the release of the annual survey data. But, because each source recognizes that the TIC SHL is the most comprehensive and accurate measure of foreign holdings of U.S. securities, differences tend to be small and short-lived.

²⁵BEA's IIP data is at <https://www.bea.gov/data/intl-trade-investment/international-investment-position>. FOF, which follows from BEA's data, is at <https://www.federalreserve.gov/releases/z1/>.

Flow Data Since 2003

While the TIC SHL *holdings* surveys obtain positions data from global custodians, there is no equivalent for data on transactions (i.e., flows). That is, there is no comprehensive security-level *transactions* survey, so each source must decide how to compute flows. One way - directly collect flow data through the TIC S monthly transactions reports - has been so inaccurate that it was discontinued in 2023. Since 2013 BEA, recognizing issues with the TIC S transactions data, imputes transactions from the reported TIC positions. *However, prior to 2013, TIC S flows were the main source for BEA and, hence, FOF transactions data.*²⁶

Figure A2 (top panel) shows that since 2012 TIC S flows differ substantially from BEA's BOP flows. For some years prior to the GFC, TIC S flows exceed BOP-reported flows, but since 2012 TIC S flows have been far below. Moreover, the split between official and private foreign flows is dramatically different. BEA BOP data (Figure A2, bottom left) show that through 2012 the bulk of foreign flows into U.S. Treasury bonds were from official investors, whereas TIC S (Figure A2, bottom right) suggests that private flows have exceeded official flows every year. The difference between TIC S and the other sources is vast. Over the period from 2005 through 2014, TIC S reported that private foreign flows into Treasury bonds were \$3.5 trillion greater than official flows, while for the same period BEA's BOP reported that official flows were \$1 trillion greater. Overall, since 2003 private foreign flows into Treasury bonds are \$1.6 trillion greater in TIC S than in BEA's BOP data.

An Assessment of Flow Data Since 2003

Researchers are confronted with a number of series on what is ostensibly the same thing - foreigners' purchases of Treasury bonds - and have no obvious way to discern which is most

²⁶For more details on the relationship between BEA and TIC S flows, see Bureau of Economic Analysis (2019). BEA's presentation is, in turn, the official source for FOF. Note that in FOF tables all flow series are seasonally adjusted, but unadjusted series are also available. BEA does not seasonally adjust portfolio flows, nor do we.

accurate. While there are official descriptions of each measure, an outsider might find it difficult to judge how to use the data.²⁷ But there is a direct way to ascertain which series should be used, albeit one that is only available since 2003. The comprehensive security-level annual surveys of foreigners' holdings is the single most accurate source. The problem for the researcher is that annual holdings data do not easily translate into flow series.

We can assess flows data sources by creating implied flows from the TIC SHL surveys. The confidential security-level data from the comprehensive annual surveys include various security characteristics, such as a general security description and identifier, issue and maturity dates, coupon rate, and amount held. The data also include both the face (which excludes price change effects) and market values of holdings. With face values, we can compute a direct measure of flow as the change in face value of holdings.²⁸

We do this for each annual survey from June 2003 through June 2021. To confirm that we have included all bonds, we first sum the holdings and compare to the surveys' published aggregate amounts. They match exactly. Confident that we have every bond, we then sum the flows; we will call the flows calculated from the security-level holdings and calculated valuation changes *implied flows*. As Figure A3 (top panel) shows, our implied flows differ greatly from TIC S flows and are closest to the BEA flow series, although in certain years, such as 2005 and 2018, the gap is sizeable. The evidence suggests that *since 2003* researchers should use BEA data (or FOF, which are identical).

Turning to the split of foreign holdings into those by governments (i.e., foreign official holdings) and those by private foreign investors (Figure A3, bottom panel), since 2003 the BEA series and our implied flows are very similar for the foreign official and foreign private

²⁷See, for example, the TIC FAQ page (<https://www.treasury.gov/resource-center/data-chart-center/tic/Pages/ticfaq1.aspx>), as well as Bertaut and Judson (2014).

²⁸We can also calculate flows an alternative way. Because we can accurately calculate the valuation change on each and every U.S. Treasury bond, security-level flow can be computed as the change in the position (which is observed) less the valuation adjustment (which is directly computed). Because survey information is complete and consistent, the two methods produce identical flow series.

flows. The only material deviation was in 2016, when the BEA series has more official and less private flows (but the same total). TIC S flows differ substantially from our implied flows and BEA-reported flows.

Data Prior to 2003

Prior to 2003, benchmark surveys were less frequent (every 5 years through 2000) and BEA's IIP data, based on the since-discontinued TIC S flow data, could deviate substantially from the surveys. By 2000 BEA tended to revise IIP estimates to be in line with the surveys.

A difficulty for flow analysis using data from the 1980s and 1990s is that TIC S was the primary source for BEA and FOF flows. Figure A4 shows that FOF and TIC S flows were almost identical in those decades; BEA flows (not shown) were very similar. *In the 1980s and 1990s, any problem with TIC S flows fed directly into BEA BOP flows and the Fed's FOF series.* And problems were substantial. Grier, Lee, and Warnock (2001) and Warnock and Cleaver (2003) showed for the period January 1995 - March 2000 that TIC S reported foreign flows into U.S. bonds were \$179 billion (or 44 percent) too high. TIC S was discontinued in February 2023 (<https://home.treasury.gov/data/treasury-international-capital-tic-system-home-page/tic-forms-instructions/tic-s-form-and-instructions>).

Figure A5 brings in another source. The Federal Reserve's International Finance Division has long-used the high-quality but infrequent positions surveys to create internally consistent monthly datasets on positions, flows, and valuation adjustments (Thomas, Warnock, and Wongswan (2004), Bertaut and Tryon (2007)).²⁹ More recently, since 2012 a new monthly holdings series, the TIC SLT, became available and forms the basis for another internally consistent monthly dataset on positions, flows, and valuation adjustments; see Bertaut and Judson (2014).³⁰ We call the spliced series - Bertaut and Tryon (2007) data through 2011,

²⁹Internally consistent in this context means that flows plus valuation changes equal the change in positions or, equivalently, that so-called *other adjustments* are zero. For more on this point, see Curcuru et al. (2008).

³⁰The new SLT data, collected at an aggregate level, are similar (but not identical) to the aggregated security-level SHL annual data and are an additional input to the BEA IIP and FOF published data.

Bertaut and Judson (2014) starting in 2012, and since 2021 series that we construct following the technique of Bertaut and Tryon (2007) - *benchmark consistent* series. For Treasuries, the benchmark consistent series for private and official foreign investors start December 1979. As Figure A5 shows, since the advent of annual holdings survey benchmark consistent positions are identical to SHL positions (top panel) and benchmark consistent flows are nearly identical to survey implied flows (bottom panel). For historical data series available at the monthly frequency, the benchmark-consistent data provide an internally consistent series on positions, flows, and valuation adjustments.

Users should note that the quality of benchmark consistent data may well vary through time, as quality is a function of the frequency of the underlying security-level surveys that provide the methodology's important fixed points for the positions. That is, when security-level surveys were conducted approximately every five years (1974 - 2000) the benchmark consistent data, while the best available for that time period, are likely lower quality than when the surveys have been conducted more frequently (annually since 2003). Specifically, the serious flaws in the TIC S flows data that have been known since Griever et al. (2000) - which showed that in the late 1990s TIC S overestimated foreign purchases by 44% - are, to a large extent, in the pre-2000 Bertaut and Tryon (2007) dataset. The reason is that Bertaut and Tryon (2007) force flows to be consistent with reported positions but do so in an ad hoc way: proportionally scaling flows up or down in an inter-survey period to eliminate any flows-position discrepancy. Pre-2000, surveys were every five years, meaning that BT had, for example, two data points - holdings at end-1994 and at March 2000 - and rescaled TIC S reported flows (proportionally) to get flows plus valuation adjustments to equal March 2000 positions. The resulting benchmark-consistent flows look very similar to the now-discontinued TIC S flows (Figure A6), even though cumulated flows differ by \$171 billion (or 44%).

To summarize, holdings data are very similar across all data sources, but flow data can

vary substantially. For current flows estimates, researchers should not use TIC S data and instead should turn to the quarterly BEA data (or unadjusted FOF data). For historical (monthly) time series, researchers should use benchmark-consistent series that have been formed by Thomas, Warnock, and Wongswan (2004), Bertaut and Tryon (2007) and Bertaut and Judson (2014), while recognizing that data quality substantially improved in 2003 with the advent of annual holdings surveys and that prior to 2000 the flow series are TIC S with a proportional scaling up or down.