Exchange Rate Disconnect Revisited*

Ryan Chahrour
Cornell University

Vito Cormun
Santa Clara University

Pierre De Leo
University of Maryland

Pablo Guerrón-Quintana
Boston College

Rosen Valchev
Boston College

June 16, 2025

Abstract

We find that variation in expected U.S. productivity explains over half of U.S. dollar/G7 exchange rate fluctuations. Both correctly-anticipated changes in productivity and expectational noise, which influences expectations of productivity but not its eventual realization, have large effects. This "noisy news" is primarily related to medium-to-long-run TFP growth, and causes significant deviations from uncovered interest parity. Together, these disturbances generate many well-known exchange rate puzzles, including predictable excess returns, low Backus-Smith correlations, and excess volatility. Our findings suggest these puzzles have a common origin linked to productivity expectations. We draw a number of implications for leading classes of open-economy models.

JEL Codes: D8, F3, G1
Keywords: Exchange Rate Disconnect, TFP News, Excess Returns, Excess Volatility

*We thank Gurdip Bakshi, Max Croce, Charles Engel, Tarek Hassan, Rohan Kekre, Oleg Itskhoki, Yang Liu, Dmitry Mukhin, Francesco Pappadà, Mikkel Plagborg-Møller, Stephanie Schmitt-Grohé, Andreas Stathopoulos and seminar and conference participants at Boston College, UC Berkeley, UT Austin, University of Florida, University of Wisconsin, Banco de Mexico, IMF, CEPR ESSIM 2021, NBER IFM Spring 2022, Norges Bank, 2021 International Finance and Macro Conference at BFI, T2M 2022, VSFX 2021, 2021 Workshop on Exchange Rates at Swiss National Bank, 2021 Workshop on Macroeconomic Dynamics at Università Cattolica, and the 2024 Oxford Macroeconomics workshop for helpful comments. This paper previously circulated under "Exchange Rate Disconnect Redux".
Emails: ryan.chahrour@cornell.edu; vcormun@scu.edu; deleop@umd.edu; pguerron@gmail.com; valchev@bc.edu

The real exchange rate—the relative price of consumption across countries—plays a crucial role in clearing markets for both goods and financial assets in international macroeconomic models. As a result, models typically imply that exchange rates are tightly linked to cross-country differentials in macroeconomic quantities, real interest rates, and other asset prices in the economy. However, in the data, real exchange rates appear largely disconnected from such macro fundamentals: they exhibit virtually no correlation with current and past macro quantities (*e.g.*, Meese and Rogoff, 1983, Engel and West, 2005) or interest rates (*e.g.*, Fama, 1984), and at the same time are surprisingly volatile (*e.g.*, Rogoff, 1996). As a consequence, a tremendous amount of theoretical work seeks to understand what mechanisms could generate these puzzling patterns in models. Yet, there is relatively little direct *empirical* evidence on the origins of these puzzles in the data.

To close this gap, in this paper we aim to uncover the main empirical drivers of exchange rate fluctuations in a model-agnostic way. Our key finding is that noisy news about future TFP accounts for more than half of the overall variation in *both* real exchange rates and macroeconomic variables. Thus, our results indicate that the exchange rate is indeed intrinsically connected to macroeconomic fundamentals, and that the apparent disconnect the previous literature has documented arises because the impact of news manifests at *different horizons* for exchange rates and macro quantities. Specifically, we find that the response of the exchange rate leads the response of many macroeconomic variables by two years or more, resulting in low contemporaneous correlation between the exchange rate and macro aggregates, even though ultimately both sets of variables react to the same underlying source of fluctuations. Moreover, the results indicate that the conditional responses to the noisy news we identify generate a number of famous exchange rate puzzles, suggesting that many well-known anomalies have a common, fundamental origin in noisy news about future TFP.

Our analysis proceeds in two steps. First, we estimate the basic comovement patterns associated with surprise changes in exchange rates in the data in reduced form. To do this, we follow the VAR procedure of Uhlig (2003) to recover a set of orthogonal reduced form shocks ordered by their importance in explaining exchange rate variation. When applied to the dollar/G6 exchange rate we find that the first shock—the one most important to exchange rate fluctuations—explains two-thirds of exchange rate variation *and* around 40% of the variation in macro aggregates. However, while the shock affects the exchange rate immediately, its effects on macroeconomic quantities such as consumption, output and TFP are delayed. Thus, this shock only generates a correlation between exchange rates and *future* macro aggregates, leaving exchange rates "disconnected" from contemporaneous macro aggregates.

In particular, the shock has no immediate impact on Fernald's (2012) utilization-adjusted U.S. TFP, but leads to a significant increase in TFP three to five years in the future.

Intuitively, these results suggest that the exchange rate, a forward-looking asset price, reacts to the arrival of news of future macro fundamentals. These results motivate the second step of our analysis, in which we directly identify and isolate disturbances to expectations about future U.S. TFP, and then estimate their effects on the real exchange rate and other cross-country macro variables. To do so, we follow the approach of Chahrour and Jurado (2021) to decompose and separately identify two orthogonal disturbances: (i) partially anticipated changes to actual future productivity and (ii) expectational "noise" disturbances that move expectations of productivity, *but never materialize in realized productivity*. Intuitively, this identification approach is designed to separately identify the signal from the noise component in noisy news. It is useful to stress that in this second step of our analysis, we first focus solely on TFP and extract noisy news about that (arguably) exogenous macro fundamental, and then turn around and study the impact of these noisy news on the exchange rate.

We find that both anticipated future TFP changes and "noise" in TFP expectations play an important role in driving exchange rates in the data. The noise disturbances give rise to volatile but relatively short-lived fluctuations in the exchange rate, while correctly-anticipated TFP changes impart persistent effects on the exchange rate. As a result, expectational noise appears relatively more relevant at higher frequencies, while lower-frequency movements in exchange rates predominantly reflect true technological disturbances. Taken together, the two disturbances account for more than 60% of the variation in the level of the real exchange rate, while explaining a comparable fraction of the variation in real macroeconomic quantities. However, as foreshadowed by our initial analysis, the impact on macro aggregates is delayed relative to the response of exchange rates. We also study the impulse responses of a number of other variables, such as the trade balance and equity prices, and find that the U.S. current account deteriorates and U.S. and foreign equity prices rise in anticipation of an improvement of future U.S. productivity. Thus, noisy news about future TFP prove to be the rare structural disturbances that drive a large portion of both real exchange rates, international business cycles and also other asset prices.[1]

Moreover, we find that the responses to these identified disturbances also generate a number of well-known "exchange rate puzzles," suggesting that a constellation of famous anomalies share a common, fundamental origin in noisy news about future TFP. In partic-

[1] To the contrary, we find that pure surprise changes to TFP contribute little to exchange rates.

ular, we find that the noisy news disturbances cause significant *predictable* fluctuations in excess currency returns, which generate violations of uncovered interest parity (UIP) that are consistent with both the classic forward premium puzzle (Fama, 1984) and the more recently documented reversal in this predictability pattern at longer horizons (Engel, 2016; Valchev, 2020). Second, the conditional responses of exchange rates and cross-country consumption differentials exhibit a weak negative correlation, in line with the famous Backus and Smith's (1993) puzzle. Third, the conditional responses also imply that real exchange rate dynamics are highly persistent and that the impact on exchange rates is large relative to macro quantities, two well-known features of exchange rate dynamics that the literature calls the "PPP Puzzle" and the excess volatility of the exchange rate (*e.g.*, Rogoff, 1996). Combined, these results suggest that the constellation of classic exchange rate puzzles may be explained by endogenous responses to noisy news about future TFP.

Two features of our analysis help explain why other studies that have examined the forward-looking nature of exchange rates have generally struggled to establish a robust correlation between exchange rates and future TFP (as well as with macro variables more broadly). First, our estimates suggest that the news that matters for the exchange rate are mainly *medium-to-long horizon* news—specifically news about TFP changes roughly 3 to 5 years in advance. The bulk of existing studies have concentrated on much shorter horizons, seeking lead-lag relationships at horizons of one to two years (Engel and West, 2005). Indeed, we find that extending the forecast horizon of Engel and West (2005) regressions reveals stronger evidence of Granger causality. Second, unlike previous literature, we separately identify and account for expectational noise. This sharpens our results because expectational noise weakens the correlation between expectations and *realized* fundamentals, reducing the statistical power of empirical approaches that focus on the lead-lag relationship between exchange rates and *realized* future fundamentals.

We conclude the paper with a discussion the implications of our results for several popular types of open-economy models. One key finding is that models with *complete asset markets* fail to replicate our finding that higher expected home TFP is associated with relatively high current home consumption and an appreciated exchange rate. For example, in Backus et al. (1994), consumption and the exchange rate do not move with the news of future TFP, but only after the actual realization of TFP. Modern complete markets models with volatile risk premia (*e.g.*, Colacito and Croce, 2013) do generate expectational effects. But, in these models, high expected home productivity cause a fall in current home consumption and a depreciation of the real exchange rate, qualitatively *the opposite* of what we find in the data.

By contrast, models with *incomplete markets* (*e.g.*, Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021) succeed in generating a current appreciation and an increase in the consumption differential when home productivity is expected to improve. Intuitively, with incomplete markets, productivity anticipation functions as a demand shock until the actual realization of TFP (Kekre and Lenel, 2024). Nevertheless, all of the models that we examine have the same counterfactual implication: they imply that the exchange rate depreciates strongly upon the actual realization of TFP, while our empirical findings show that the exchange rate only gradually depreciates after a home productivity improvement, remaining appreciated overall for up to three years after the actual productivity increase.

We conjecture that this failure in the models is due to the fact that none of them capture the significant increase in expected foreign currency returns that we estimate in the data following the realization of the TFP shock, which helps keep the USD appreciated. In the benchmark incomplete-markets setups the literature currently operates with, however, currency premia fluctuate primarily due to exogenous "UIP" shocks and do not react significantly in response to "macro" shocks (Itskhoki and Mukhin, 2021; Kekre and Lenel, 2024).[2] A fruitful avenue for future research would be to develop models that feature both incomplete markets and endogenously volatile currency premia in response to TFP shocks.

Related literature This paper is related to several different strands of the international finance and macro literatures. First, we speak to the exchange rate determination puzzle, that is the lack of correlation between exchange rates and macroeconomic fundamentals, both contemporaneously and in terms of forecasting future exchange rates using current macroeconomic fundamentals (Meese and Rogoff, 1983; Cheung et al., 2005; Rogoff and Stavrakeva, 2008; Miyamoto et al., 2023). A related observation is that the exchange rate is "excessively" volatile and persistent, as compared to macroeconomic fundamentals (Obstfeld and Rogoff, 2000; Chari et al., 2002; Sarno, 2005; Corsetti et al., 2008a; Steinsson, 2008).

Contrary to this literature, we find that there is in fact a strong connection between exchange rates and macroeconomic fundamentals, but one that relates current exchange rates to future fundamentals. Our evidence supports the basic point of Engel and West (2005) that exchange rates are forward-looking and should therefore predict, rather than lag behind, macroeconomic variables. Our results contribute to this discussion in a number of

[2]Exogenous UIP shocks are common in the literature; *e.g.*, Devereux and Engel (2002), Jeanne and Rose (2002) Kollmann (2005), Bacchetta and van Wincoop (2006), Farhi and Werning (2012), and Eichenbaum et al. (2020). Relatedly, Huo et al. (Forthcoming) find that international comovement between macro aggregates is likely explained by non-fundamental shocks, though they do not speak to correlation with exchange rates.

ways. First, we show that the link between current exchange rates and future fundamentals runs specifically through *imperfect* and *noisy* anticipation of future productivity. Second, we show that the noisy TFP news priced into exchange rates primarily concern TFP at medium-to-long horizons (3 to 5 years in the future). Third, our study is unique in explicitly accounting for expectational noise—fluctuations in expectations that are not associated with actual subsequent changes in fundamentals—and we emphasize that ignoring this component of expectations leads one to understate the forward-looking nature of exchange rates.[3]

Another contribution of this paper is to show that a large number of famous exchange rate puzzles have a common fundamental origin in noisy news about future TFP. In particular, we find that the noisy news we identify do not only play an important role in the exchange rate disconnect and excess volatility puzzles referenced above, but also in generating the UIP puzzle (Fama, 1984; Engel, 2014) and the Backus-Smith puzzle (Backus and Smith, 1993). These puzzles have also received extensive theoretical attention, and interestingly the great majority of proposed models consider classic "pure surprise" shocks to TFP, not news about future TFP.[4] Our results show, however, that in the data the expectations of future TFP are a more important driving force behind exchange rates, as compared to the surprise realization of higher than expected current TFP.

Relatedly, there is a small but growing literature that documents the effects of "news shocks" on international business cycles and develops international RBC models driven in part by news shocks. That literature, however, has typically focused on the question of comovement between macro aggregates across countries, and not on exchange rate dynamics and related puzzles. In that vein, Siena (2015) argues that news shocks only lead to a small amount of comovement between macro aggregates across countries, contrary to previous

[3]Another literature uses survey of expectations to measure the surprises in macroeconomic announcements and studies their effect on exchange rates (Andersen et al., 2003; Faust et al., 2007; Engel et al., 2008). In a recent paper, Stavrakeva and Tang (2020) find that the new information about *past* macroeconomic fundamentals that the market obtains upon a new statistical release is an important driver of exchange rate fluctuations. Our definition of "news" is different, however, as we specifically identify disturbances to the forecast of *future U.S. TFP* changes, as opposed to revision of beliefs about *past* variables such as output.

[4]For example, time-varying risk (Alvarez et al., 2009; Verdelhan, 2010; Bansal and Shaliastovich, 2012; Farhi and Gabaix, 2015; Gabaix and Maggiori, 2015), non-rational expectations (Gourinchas and Tornell, 2004; Burnside et al., 2011; Ilut, 2012; Candian and De Leo, 2023) and liquidity premia (Engel, 2016; Valchev, 2020) have been proposed as explanations of the UIP Puzzle. On the other hand, Corsetti et al. (2008b), and Karabarbounis (2014) develop models that explain the Backus-Smith puzzle. The Colacito and Croce (2013) model of "long-run risk" occupies an interesting middle ground, since their shocks to TFP have persistent, non-monotonic impulse responses and hence the main impact is in the future, which is akin to news.

evidence by Beaudry and Portier (2014).[5,6] Perhaps most closely related to us is the work of Nam and Wang (2015), who use Barsky and Sims' (2011) approach to identifying news-to-TFP shocks. In contrast to us, however, they do not separately identify the effects of correctly anticipated TFP improvements and expectational noise disturbances and they do not document that role the anticipation plays in explaining the plethora of FX puzzles. This additional analysis helps us to provide guidance for the development of exchange rate models.

Last, we note that a few recent papers document the presence of *contemporaneous* relationships between exchange rates and specific macro aggregates in the data. Lilley et al. (2020) find a contemporaneous correlation between U.S. purchases of foreign bonds and the U.S. dollar, but only in the post-2009 period. Over a longer time span, Engel and Wu (2024) show that monetary variables and measures of risk and liquidity account for a sizable portion of variation in U.S. dollar/G10 exchange rates. Kekre and Lenel (2024) document that higher U.S. yields relative to G10 yields (ranging from 3 months to 10 years tenors) are associated with an appreciated U.S. dollar/G10 exchange rate. Kekre and Lenel argue that this evidence, along with the Backus-Smith correlation, favors models in which exchange rates are driven by persistent demand shocks, and note that news about long-run productivity (like the ones we document) could in principle generate persistent movements in demand.[7]

1 Data and basic empirical framework

Our empirical analysis centers on a VAR

$$Y_t = C(L)Y_{t-1} + u_t, \tag{1}$$

where the vector Y_t contains data on the U.S. and a trade-weighted aggregate for the other G7 economies. Hereinafter, we will refer to the U.S. and the other G7 economies as the "home" and "foreign" economies, respectively, and use the * notation to denote non-U.S. variables. For our baseline analyses, the vector Y_t contains eight variables: (i) the nominal exchange

[5] Corsetti et al. (2014) use sign restrictions to identify the cross-country effects of unanticipated productivity and demand disturbances in the United States, including on the real exchange rate.

[6] Gornemann et al. (2020) develop an international model of endogenous TFP growth, and show that at their calibration it can account well for the low frequency movements in real exchange rates.

[7] Other papers document statistically significant relationships between exchange rates and commodity prices (Chen et al., 2010; Ayres et al., 2020), exchange rate and external imbalances (Gourinchas and Rey, 2007), as well as exchange rates and trade flows (Alessandria and Choi, 2021; Gornemann et al., 2020). In addition, Hassan et al. (2016) document a relationship between stochastic properties of exchange rates and differences in capital-output ratios across countries.

rate S_t expressed in units of U.S. dollar per foreign currency, (ii) Fernald's (2012) series on utilization-adjusted U.S. TFP, (iii) and (iv) are the U.S. real consumption and investment, (v) and (vi) are foreign real consumption and investment, (vii) the nominal interest rate differential, (viii) and the CPI price level differential between the U.S. and a trade-weighted aggregate for the other G7 economies. The variables in our VAR can be denoted as:

$$Y_t' \equiv \left[\ln\left(S_t\right), \ln\left(TFP_t\right), \ln\left(C_t\right), \ln\left(C_t^{\star}\right), \ln\left(I_t\right), \ln\left(I_t^{\star}\right), \ln\left(\frac{1+i_t}{1+i_t^{\star}}\right), \ln\left(\frac{CPI_t}{CPI_t^{\star}}\right)\right].$$

Using a trade-weighted aggregate of the other G7 economies as the "foreign" country is standard in the literature (*e.g.*, Engel, 2016). In any case, in Appendix B, we conduct the analysis using six separate bilateral VARs between the U.S. and each other G7 country separately. The results and the emerging conclusions of the bilateral VARs are very similar, hence we have found that a VAR with a trade-weighted aggregate serves as a useful benchmark to summarize the results.[8]

For our benchmark results, we use quarterly data for the time period 1978:Q4-2008:Q2. The sample stops in 2008:Q2 out of abundance of caution, to guard against a possible structural break in the aftermath of the financial crisis, which is a potential pitfall as argued by Baillie and Cho (2014) and Du et al. (2018). However, in Appendix B we conduct our analysis on an extended sample through the end of 2018 and the results remain very similar. Thus, we think the potential structural break is not a concern for our analysis, but to respond to potential concerns we use the sample that stops in 2008:Q2 in the baseline analysis.

We describe the data and their sources in detail in Appendix A. As a brief overview, the exchange rate is the average of the daily exchange rates within a quarter, obtained from *Datastream*. The interest rate differential is the average of daily Eurodollar rates within a quarter, obtained from *Datastream*.[9] The CPI indices and the consumption and investment series are from the *OECD* database. Lastly, the U.S. TFP is from John Fernald's website.

We do not have a comparable, utilization-adjusted quarterly TFP series for countries other than the US. Recently, a few papers have constructed novel utilization-adjusted TFP for foreign economies (*e.g.*, Huo et al., 2023; Comin et al., 2023), but we find that these measures are not appropriate for our purposes. Most importantly, most such measures of adjusted foreign TFP can only be constructed at annual frequencies (Huo et al., 2023) or

[8]Results when we use a simple average across the foreign economies, instead of a trade-weighted average, are virtually identical.

[9]Note that these interest rate differentials are not forward discount-implied interest rate differentials, but actual eurodollar rates.

over a very a short sample period (Comin et al., 2023), which limits the scope of the analysis. Moreover, the Huo et al. (2023) measure of utilization-adjusted TFP for the United States displays a fairly low correlation (0.42) with the widely-accepted Fernald measure for U.S. TFP, suggesting it is measuring some different notion of technology (and beyond that is only available annually). As a result, our benchmark analysis does not include a measure of foreign TFP, but as a robustness exercise we report in the Appendix B, we consider alternative specifications where we use the foreign Solow residual at quarterly frequency (to maximize data coverage). All of our main results and conclusions remain unchanged.

We estimate the VAR in (1) using four lags and Bayesian methods with a Minnesota prior. This commonly used prior assumes all series are separate random walks and that thus there is no relationship between different variables in Y_t. This choice of prior is particularly conservative for our purposes, given that we want to test for a potential connection between the exchange rate and the other macro variables.

Following the established convention (*e.g.*, Sims et al., 1990; Eichenbaum and Evans, 1995), we estimate the VAR in levels and do not impose ex-ante that there are any specific cointegration relationships. Nevertheless, in Appendix B we show that results remain unchanged if one instead estimates a Vector Error Correction Model (VECM) that imposes the same cointegration relationships as Engel (2016), where the real exchange rate and interest rate differential are assumed stationary. More generally, we have found the results to be robust to imposing a variety of other potential cointegration relationships.

Given the VAR estimates, our goal is to isolate structural shocks that are significant drivers of the real exchange rate. To that end, we note that the reduced-form VAR residuals, u_t can be expressed as a linear combination of leads and lags of the underlying structural disturbances ε_t

$$u_t = A(L)\varepsilon_t \tag{2}$$

where we stress that $A(L) \equiv \sum_{-\infty}^{\infty} A_k L^k$ is a potentially two-sided lead-lag polynomial. This expression generalizes more traditional treatments in allowing for the innovations in the VAR to depend on past, present, and future structural disturbances.

Structural assumptions are needed to identify $A(L)$, and different assumptions lead to identifying different structural disturbances series ε_t. For example, standard VAR treatments often assume the true data generating process is invertible, which implies that $A(L) = A_0$ is just a matrix. From there, one might assume the disturbances have a natural "Cholesky ordering" in which some disturbances affect the economy before others, so that A_0 is triangular. Or, alternatively, one might impose sign restrictions on either short-term or long-term impacts, which would in turn imply a different restrictions on A_0.

We conduct our analysis in two steps, going from fewer to more structural assumptions. First, we use a "max-share" approach in the vein of Uhlig (2003) which makes the usual assumption that $A(L) = A_0$, but otherwise makes minimal structural assumptions about A_0. The results of this analysis then motivate our second step, in which we allow for a general, potentially doubled-sided form for $A(L)$, but then impose stronger economic restrictions to specifically identify anticipated TFP disturbances and noise disturbances to expectations of future TFP.

2 The main driver of real exchange rate fluctuations

We begin with an agnostic empirical approach that aims to isolate the main driver of exchange rate fluctuations in the data while imposing minimal ex-ante assumptions on the nature of the underlying structural disturbances. To do so, we follow the max-share approach of Faust (1998) and Uhlig (2003) to extract the shock that explains the biggest share of the variation in the real exchange rate. This approach was recently applied to real macro quantities by Angeletos et al. (2020) to extract a so-called "main business cycle" shock. In turn, we apply it to exchange rates, and in parallel to the Angeletos et al.'s (2020) terminology, we refer to the shock we extract here as the "main exchange rate" (MFX) shock.[10]

The real exchange rate q_t is the difference between the log nominal exchange rate and the differential in log CPIs, $q_t = s_t + p_t^\star - p_t$, and hence given the included variables in Y_t,

$$q_t = e_q' Y_t,$$

where $e_q = [1, 0, 0, 0, 0, 0, 0, -1]'$.

To apply the max-share procedure, note that given the assumption of $u_t = A_0 \varepsilon_t$ it follows that the moving average representation of the VAR in (1) can be written as

$$Y_t = \sum_{k=0}^{\infty} B_k A_0 \varepsilon_{t-k}, \tag{3}$$

where the $\{B_k\}$ correspond to the coefficients in the $MA(\infty)$ representation of $B(L) \equiv$

[10] Kurmann and Otrok (2013), Basu et al. (2021), Cormun and De Leo (2024), and Chahrour et al. (2023) also contain applications of the max-share approach.

$(I - C(L))^{-1}$. Then, the h-step ahead forecast error over the real exchange rate is given by

$$q_{t+h} - \mathbb{E}_{t-1}q_{t+h} = e_q' \left[\sum_{\tau=0}^{h} B_\tau A_0 \varepsilon_{t+h-\tau} \right]$$

where $h \geq 0$. The forecast error variance (FEV) of the real exchange rate, $\mathrm{var}(q_{t+h} - \mathbb{E}_{t-1}(q_{t+h}))$, is a linear combination of the variances of the (orthogonal) elements of the vector ε_t, and in particular, the contribution of the first element ε_{1t} can be expressed as

$$\mathrm{var}\left(q_{t+h} - \mathbb{E}_{t-1}(q_{t+h}) \,|\, \varepsilon_{i,t+\tau} = 0 \;\forall\, i \neq 1,\; \tau = 0, \ldots, h\right) = e_q' \left[\sum_{\tau=0}^{h} B_\tau A_0 e_1 e_1' A_0' B_\tau' \right] e_q \quad (4)$$

where e_1 is the selection vector so that $A_0 e_1$ selects the first column of A_0, and ε_{kt} is the k^{th} shock in the vector ε_t.

We choose the rotation matrix A_0 by maximizing (4) subject to $A_0 A_0' = \mathbb{E}(u_t u_t')$. This requires us to specify a horizon h at which the forecast error variance in (4) is computed, and for that we choose $h = 100$ quarters, which effectively gives us the unconditional variance of q_t.[11] This procedure yields a partially identified system, in the sense that the above maximization problem will uniquely determine the first column of A_0 and thus the first element of the shock vector of ε_t (which is what we are interested in), but not the rest.

Intuitively, the resulting shock series ε_{1t} is the reduced form innovation that makes the largest contribution to the fluctuations in the log real exchange rate q_t. This is not a structural shock with a clear economic interpretation – ε_{1t} is potentially a linear combination of several underlying structural shocks. Rather, we view ε_{1t} as a reduced form way of capturing the dominant factor driving surprise changes in the exchange rate in the data, whatever its true deep origins might be (this is the topic of the second part of the paper).

Estimation results Table 1 reports the share of the forecast error variance of each macro variable that is explained by the "main exchange rate shock" shock, across horizons from 1 to 100 quarters. We find that this shock is indeed a dominant driver of exchange rate fluctuations – it explains roughly 70% of variance of the real exchange rate, in a variance decomposition sense. Given its empirical importance for the exchange rate, understanding the characteristics and the footprint this reduced form innovation leaves in the rest of the

[11]Our results are robust to choosing a variety of horizons h. Moreover, the same procedure can also be applied in the frequency domain, and the results remain very much the same if we target variation of q_t over specific frequencies instead.

Table 1: Share of forecast error variance explained by the Main FX shock (ε_1)

	Forecast Horizon (Quarter)					
	$Q1$	$Q4$	$Q12$	$Q24$	$Q40$	$Q100$
Home TFP	0.03	0.06	0.20	0.37	0.45	0.43
Home Consumption	0.02	0.04	0.21	0.47	0.51	0.40
Foreign Consumption	0.01	0.04	0.06	0.21	0.36	0.30
Home Investment	0.29	0.34	0.32	0.40	0.42	0.41
Foreign Investment	0.06	0.08	0.15	0.22	0.34	0.33
Interest Rate Differential	0.40	0.39	0.30	0.34	0.35	0.39
Real Exchange Rate	0.50	0.69	0.82	0.73	0.70	0.68
Expected Excess Returns	0.47	0.33	0.34	0.44	0.45	0.47

Notes: The table reports the estimated variance shares accounted for by the main exchange rate shock at different horizons.

economy could be very informative about the deep origins of exchange rate variation.

As a first step in this direction, we report the impulse response functions of the exchange rate, and also other macro variables, in Figure 1. The median impulse response to a one standard deviation shock is plotted with a solid blue line, and the shaded areas denote the 16-84th percentile bands.

The real exchange rate shows a significant response on impact, appreciating by about 2.5% after a one standard deviation increase in the MFX shock. The exchange rate response also displays persistent hump-shaped dynamics, where it continues to appreciate for another 5 quarters after the initial impact, peaking at a maximum appreciation of about 3.5%, and thereafter it steadily depreciates back to its long-run mean. These persistent non-monotonic dynamics – with a half-life of roughly three years – are similar to the estimates in Steinsson (2008) in the context of an *univariate* reduced form innovation to the exchange rate.[12]

The hump-shaped exchange rate dynamics also underlie a related cyclical pattern in the exhibited deviations from UIP – specifically the MFX shock generates non-monotonic movements in *expected* excess currency returns (bottom left panel of Figure 1). Expected excess currency returns are defined as $\mathbb{E}_t\lambda_{t+1} \equiv \mathbb{E}_t\Delta q_{t+1} + r_t^\star - r_t$, and computed using VAR-implied expectation. The impulse response of expected excess currency returns reveals that these are negative on impact and remains so up to five quarters after the shock, and

[12]Hump-shaped dynamics also emerge following an identified monetary policy innovation. This "delayed overshooting" result was initially shown by Eichenbaum and Evans (1995).

Figure 1: Impulse Response Functions to the Main FX shock (ε_1)

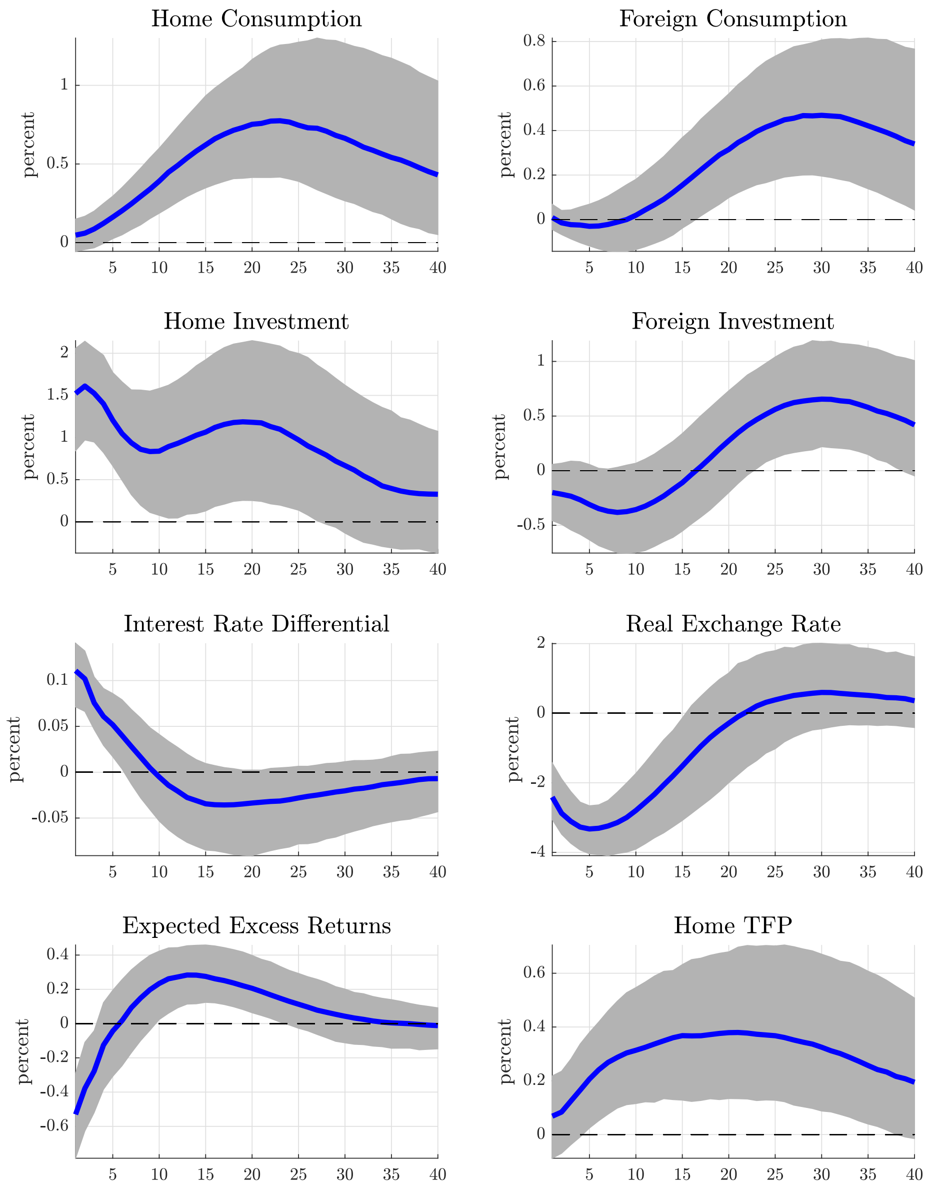

Notes: The figure reports the impulse responses to the main FX shock, along with the 16-84th (dark gray) percentile bands. Each period is a quarter.

then turns significantly positive and remains so for several years afterwards.

We also observe that in response to the MFX shock, the real interest rate differential increases on impact (*i.e.* high U.S. interest rate) and gradually returns to its long-run mean. As a result, in the immediate aftermath of the shock, the exchange rate response displays the classic version of the UIP puzzle where the high interest rate currency (the USD) is earning high returns (Fama, 1984). On the other hand, in the medium run, the direction of the UIP violation reverses, with the USD earning low returns for an extended period of time – and thus the MFX shock generates exchange rate dynamics that are consistent with the reversal of UIP violations at longer horizons documented by previous studies such as Engel (2016) and Valchev (2020).

Overall, the results suggest the MFX shock we are isolating is indeed associated with the key features of exchange rate behavior that the empirical literature has emphasized for decades – high volatility, and persistent, hump-shaped dynamics which are associated with UIP violations that switch direction from short to long horizons. Thus, whatever the ultimate structural source(s) of this reduced form innovation, it is indeed responsible for the defining characteristics of exchange rate fluctuations in the data.

Broader footprint of the MFX shock The procedure for extracting ε_{1t} imposes very few ex-ante assumptions on the data. The trade-off is that we cannot uniquely label the structural origins of the MFX shock. Still, given the prevailing view that exchange rates are disconnected from the broader macroeconomy, it is interesting to consider whether the MFX shock, which is identified purely off its consequences for exchange rates, is potentially associated with significant dynamic effects on any *other* macro aggregate.

And surprisingly, given the vast prior literature on the disconnect between exchange rates and macro fundamentals (*e.g.*, Meese and Rogoff, 1983 and Engel and West, 2005), we find that the MFX shock does explain a significant portion of the overall variation of several important macro aggregates. Most importantly, it accounts for around 40% of the forecast error variance of consumption (both home and foreign) and TFP at long horizons (see Table 1, rows 1 through 5). The reason this surprisingly strong connection with fundamentals has been largely overlooked so far is that the effects of our main exchange rate shock materialize at different horizons in exchange rates and macro aggregates – the exchange rate reacts strongly immediately, while most macro aggregates react with a significant lag. This difference in *timing* implies there is actually only a mild correlation between exchange rates and current and past macro aggregates (conditional on this shock), in line with previous findings.

We can see this most clearly from the impulse responses reported in Figure 1. Home consumption only responds in statistically significant terms to the MFX shock after a year, and foreign consumption does not exhibit a significant response until four years after the shock. The effect on home consumption peaks at around 22 quarters after the shock, while foreign consumption's response peaks at around 30 quarters after the shock (by this time the effect on exchange rates has completely died out). Moreover, the relatively stronger and quicker impact on home consumption implies that the MFX shock generates the celebrated empirical violation of the Backus and Smith (1993) condition, where high domestic consumption is associated with an appreciated exchange rate, rather than a weak exchange rate.

Similarly, the MFX shock causes no significant impact on U.S. TFP up to four quarters in the future, but productivity displays a significant and prolonged increase at longer horizons. The effect peaks at an increase of 0.4% around 20 quarters after the initial impulse. Thus, overall both consumption and TFP display a significant response in the medium-to-long run, but no response in the immediate aftermath of the shock.

The differences in the timing of the effects of the MFX shock among different variables can also be observed in the forecast error variance decompositions reported in Table 1. As can be expected given the shapes of the IRFs in Figure 1, while this shock is equally important for both short-run and long-run exchange rate fluctuations, it only explains 2% and 3% of the one-quarter-ahead forecast error variance of U.S. consumption and TFP, respectively. At the same time, the MFX shock explains roughly 20% of the forecast error of consumption and TFP at the 3 year horizon, and 40% of their forecast error variance at longer horizons.

The main exchange rate shock at different frequencies We highlight that the strength of the connection between exchange rates and macro aggregates depends on the frequency range over which the maximization is performed. When maximizing over all frequencies (1-1000 quarters periodicities), the MFX shock explains 37% of U.S. TFP variance at forecast horizon Q24 – and even more at longer horizons (see Table 1). By contrast, the explained fraction is about half as large when the maximization is restricted to business cycle frequencies (6-32 quarter periodicities, see Appendix Table B.1). This pattern is consistent with Miyamoto et al.'s (2023) finding that the relationship between exchange rates and fundamentals appears weak when the max-share approach targets the business cycle range.

In the following sections, we show that this empirical observation aligns with the notion that expectational noise — fluctuations in expected productivity but not its eventual realization — obscures the connection between exchange rates and realized fundamentals, especially at higher frequencies.

Takeaways Taken together, this evidence sheds important light on the "exchange rate disconnect puzzle," as broadly construed.

Indeed, our results confirm that the bulk of the variation in the real exchange rate is essentially unrelated *contemporaneously* to aggregate consumption or TFP, two key macro aggregates the prior literature has often studied in the context of exchange rates. Rather, we find that the exchange rate *leads* these macro aggregates. Thus, our results suggest that the canonical finding of a "disconnect" does not emerge because of an actual separation between exchange rates and fundamentals, but rather because of a difference in the timing of the responses in these variables.

The empirical approach so far offers an enlightening statistical summary of the data that relies on minimal assumptions, but it cannot sharply label the deep structural origins of the reduced form exchange rate innovation ε_{1t} we have uncovered. Still, the key finding that macroeconomic quantities such as consumption and TFP only respond with a significant delay, while forward looking variables such as asset prices (the exchange rate itself, but also interest rates) jump on impact, suggests the further hypothesis that the MFX is capturing (or at least heavily loading on) the classic notion of a news shock about TFP.

This is an interesting hypothesis, and challenges the emerging consensus in the literature that exchange rate fluctuations are primarily due to currency market-specific financial or risk shocks that are unrelated to macroeconomic fundamentals. In order to investigate this news about TFP hypothesis in detail, in the next section we introduce some simple structural assumptions that are sufficient to identify disturbances to expectations about TFP, and evaluate whether they are indeed an important driver of exchange rates.

3 Expectations of TFP and exchange rates

In this section, we use the strategy of Chahrour and Jurado (2021) to identify two distinct types of disturbances to expected TFP. The first type of disturbance captures all realized changes to TFP, and may be imperfectly anticipated at an arbitrary horizon. The second disturbance captures the noise in expectations that leads agents' forecasts about future TFP to fluctuate even when realized TFP does not change.

Our identification strategy can be understood in the context of a general information structure. Suppose that the process for home productivity is given by an MA(∞)

$$a_t = \sum_{k=0}^{\infty} \delta_k \varepsilon^a_{t-k}. \tag{5}$$

The ε^a_t are the innovations in TFP's univariate Wold representation; we impose no additional assumptions on the time-series dynamics of a_t.

One key goal of our analysis is to identify the degree to which the future increments to TFP, ε^a_{t+k}, are forecastable, i.e. whether $\mathbb{E}_t(\varepsilon^a_{t+k}) \neq 0$. There are different ways to model the time-t information set underlying this conditional expectation, and our empirical approach will be general and flexible.

As a motivational example that would be helpful to explain the intuition of our identification scheme, assume that each period agents perfectly observe current and past productivity, along with an additional signal about future productivity,

$$\eta_t = \sum_{k=1}^{\infty} \zeta_k a_{t+k} + v_t. \tag{6}$$

The process $v_t \perp a_{t+k}, \forall k$ represents informational noise. Again, we put no ex ante restrictions on the weighs ζ_k in the signal or on the dynamic process for v_t, which also follows an arbitrary MA(∞) process,

$$v_t = \sum_{k=0}^{\infty} \nu_k \varepsilon^v_{t-k}. \tag{7}$$

Evidently, the information structure in (5)-(7) allows for a wide range of signal structures with noisy information about future TFP. But this structure also encompasses a range of processes in which *there is no explicit expectational noise shock*, including the well-known "long-run risk" process that has played an important role in asset pricing research and international finance models with volatile currency premia such as Colacito and Croce (2011) and Colacito and Croce (2013). Beveridge-Nelson, long-run risk, and other processes that decompose a scalar fundamental into separate components observed by agents all provide those agents with information above-and-beyond the observation of the fundamental history, and they can all be represented as including noisy observations of future realizations (Chahrour and Jurado, 2018). In Appendix D.2 and D.3, we provide an explicit mapping to (5)-(7) for the information structures of Blanchard et al. (2013) and Colacito and Croce (2013).

We separately identify the true technological disturbances ε_t^a and the expectational noise disturbances ε_t^v using the method of Chahrour and Jurado (2021). The assumptions of their procedure are that (i) the productivity disturbances ε_t^a are orthogonal to other structural shocks (as is standard) and (ii) the expectational noise innovations ε_t^v are orthogonal to ε_t^a (intuitively, the procedure captures rational expectations errors, which are unrelated to actual TFP by construction).

To get some intuition for how the procedure works in practice, consider the illustrative case where productivity a_t and the signal η_t are both directly observed by the econometrician. As in standard VAR analysis, the joint process for $[a_t, \eta_t]$ has many potential representations that deliver the same autocovariances; identification requires some theoretical restrictions. We can impose the (very minimal) restrictions implied by (5)-(7) by placing zeros in the two-sided MA representation of $[a_t, \eta_t]$ in the following way:

$$\begin{bmatrix} a_t \\ \eta_t \end{bmatrix} = \cdots + \begin{bmatrix} 0 & 0 \\ * & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_{t+1}^a \\ \varepsilon_{t+1}^v \end{bmatrix} + \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \varepsilon_t^a \\ \varepsilon_t^v \end{bmatrix} + \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \varepsilon_{t-1}^a \\ \varepsilon_{t-1}^v \end{bmatrix} + \cdots \tag{8}$$

The identification assumption that the noise disturbances ε_t^v are not related to TFP at any lead or lag implies that the upper right corner of *all* matrices, leads and lags, is zero. In addition, since ε_t^a is the Wold innovation to TFP, the upper left corner of all *lead* matrices is zero – i.e. the technological shocks naturally move TFP only after they realize. Chahrour and Jurado (2021) show this gives enough zero restrictions to identify the system—intuitively, the number of remaining unrestricted coefficients is equal to the number of moments we can estimate from data. With estimates for the unrestricted * coefficients in (8), we can subsequently recover timeseries estimates of the disturbances ε_t^a and ε_t^v.

Of course, the above illustrative example implausibly assumes that the econometrician directly observes the relevant signal η_t. However, the same identification can be achieved by replacing η_t in the observation vector $[a_t, \eta_t]$ with forecasts of future TFP, $\mathbb{E}_t(a_{t+k})$, since under our working hypothesis the expectation is a linear combination of the history of signals and past TFP levels (which the econometrician observes). Thus, when we take this to the data, we use the forecast of future TFP implied by our estimated VAR (equation (1)). Formally, the key assumption needed for replacing η_t with the VAR-implied $\mathbb{E}_t(a_{t+k})$ is that the forward-looking variables in the VAR, *e.g.* exchange rates and interest rates, reflect the forward information about TFP that agents receive through the unobserved signals η_t. Moreover, if the data in our VAR fail to capture some of agents' forward information, this will only bias us *against* finding any anticipation effects in our estimation. So if anything, it will be harder to reject our null hypothesis that future TFP innovations are at least partly forecastable.

For the actual implementation of this procedure we need to specify a target "horizon" k for the VAR-expectation $\mathbb{E}_t(a_{t+k})$ that we will use to proxy for η_t. We choose $k = 20$ to match the peak in the impulse response of TFP in Figure 1. Based on that preliminary estimation, it seems that TFP is most forecastable at medium-to-long horizons and we pick k accordingly. However, in principle, if agents truly only observe one independent signal η_t about future TFP, then the choice of horizon k is irrelevant, as any choice of k will yield identical estimation results (see Chahrour and Jurado (2021)). And in practice, we find that our estimation yields very similar results for a wide range of k choices.

Lastly, it is interesting to note that when there are more than two variables in the VAR, as in our baseline application, the procedure imposes only a subset of the restrictions implied by the signal structure (5)-(7). In particular we do not impose the restriction that other variables besides the target expectation ($\mathbb{E}_t(a_{t+k})$) have a zero response prior to the realization of the noise disturbance ε_t^v. Because we do not impose these additional over-identifying restrictions *ex ante*, one can use these additional restrictions as an *ex post* test of the information assumptions that motivate our identification approach. In our main application, we find very small responses of other variables in this anticipation period, suggesting that equations (5)-(7) provide a good description of the expectations process driving the economy.

3.1 The dynamic effects of TFP and noise disturbances

After we separately identify the TFP and informational noise disturbances, ε_t^a and ε_t^v, we compute the resulting impulse responses of a number of variables of interest.

In order to understand first what our estimates suggest about how much forward information about future TFP there really is in the economy, we start by examining the impulse responses of TFP itself, a_t and also the IRF of the 20-quarter ahead expectation of TFP, $\mathbb{E}_t(a_{t+20})$ in Figures 2 and 3 respectively. Since our approach allows for productivity disturbances to potentially be anticipated, we plot each impulse response starting from 20 quarters *before* the actual change in productivity, which we normalize to be $t = 0$, and then also plot the evolution of the IRFs until 20 periods after the realization of the disturbance. Hence, x-axis ranges from -20 to +20. The extent to which TFP anticipation plays a role in the data can be evaluated by examining whether the estimated TFP forecast $\mathbb{E}_t(a_{t+20})$ responds

Figure 2: Impulse responses to Technology (ε^a) disturbances

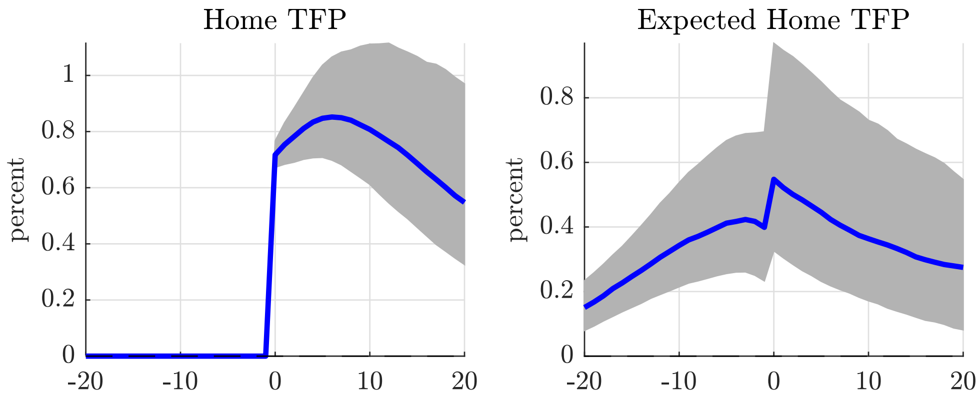

Notes: The figure displays IRFs to a one standard deviation impulse in the technological disturbance at time $t = 0$. The shaded areas are 16-84th percentile bands. Each period is a quarter.

significantly to ε_t^a prior to the shocks actual realization in period 0.

Consider first the response of TFP to an increase in ε_t^a, depicted in the left panel of Figure 2. This IRF is exactly zero in periods -20 to -1, reflecting the identification assumptions described above – *i.e.* the estimated ε_t^a is the true innovation to TFP, hence there is no impact on TFP before the actual realization. Consistent with typical empirical results on U.S. TFP dynamics, we estimate that technological disturbances have very persistent effects. A one standard deviation disturbance increases the level of TFP by around 0.75% on impact (at $t = 0$), and the level of TFP is still a full 0.50% above trend 20 quarters after that.

In the right panel of Figure 2 we plot the impulse response of the expectation of TFP 20-quarters ahead, $\mathbb{E}_t(a_{t+20})$. For this and other endogenous variables, our estimation imposes no *ex-ante restrictions* on whether anticipation effects are present (*i.e.* an effect before $t = 0$). The key result is that indeed expected TFP is significantly higher than its long-run mean even 20 quarters before the innovation actually realizes, manifesting a significant amount of anticipation. Specifically, we estimate that 20-quarters before the actual increase of 0.75% in TFP at time 0, agents expect that this future quarter's TFP will be 0.2% higher than average. Thus, roughly about one quarter of the actual improvement is anticipated a full five years ahead. Moreover, as we get closer to the actual realization of ε_t^a at $t = 0$, the TFP forecast steadily rises, as could be expected if forecasting near-term TFP is easier. Still, even one quarter ahead anticipation is not perfect, as can be deduced from the jump in the IRF at time 0, which indicates that the actual realization of ε_t^a still surprises agents and leads to an adjustment of expectations upwards.

Figure 3: Impulse responses to Noise (ε^v) disturbances

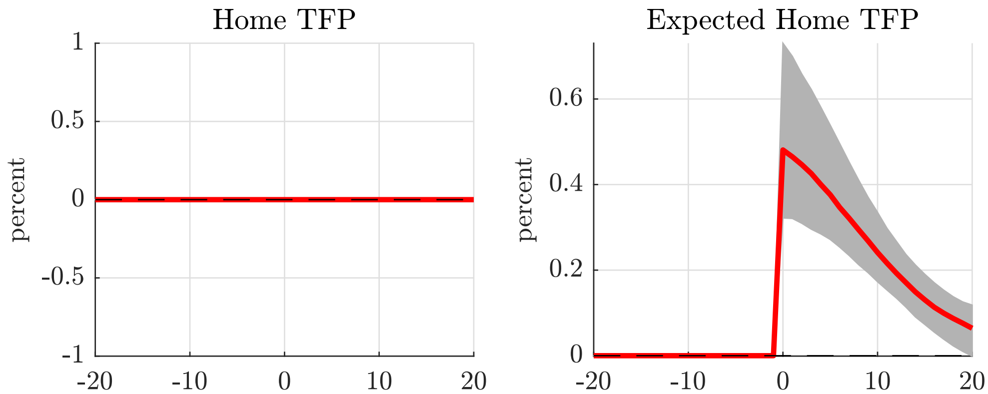

Notes: The figure displays IRFs to a one standard deviation impulse in the expectational disturbance at time $t = 0$. The shaded areas are 16-84th percentile bands. Each period is a quarter.

Another way to see that expectations are imperfect is by considering the impulse response to the pure expectational noise disturbance, ε_t^v, which we plot in Figure 3. The left panel reflects our identification assumption that the expectational noise disturbances have no effect on TFP at any lead or lag. Nevertheless, the right panel shows that these disturbances move *expected TFP* significantly, as we estimate that a one standard deviation increase in expectational noise (so an "optimistic" revision of future TFP), leads to a 0.5% increase in $\mathbb{E}_t(a_{t+20})$. The expectational effect is mean-reverting, which is consistent with the idea that over time agents learn that their initial optimism was misplaced and expectations gradually return back to their long-run mean in the absence of actual TFP changes.

Overall, the results in Figure 2 and Figure 3 support the hypothesis of a noisy-information paradigm, where future movements in TFP are partially anticipated, but expectations are noisy and sometimes move even though there is no actual change in productivity.

Broader effects We now turn to the effects of these two disturbances on the rest of the endogenous variables in the VAR, with a special attention paid to the response of the exchange rate. In Figure 4 we plot the responses of the interest rate differential, home consumption, the real exchange rate, foreign consumption and the expected currency returns, $\mathbb{E}_t(\lambda_{t+1})$, to a TFP improvement (an increase in ε_t^a). In the top left panel we report again the response of the level of TFP for reference.

We focus on the real exchange rate first. The response shows a pronounced V-shape, reflecting significant anticipation effects where the real exchange rate steadily appreciates *before* the actual TFP improvement. Moreover, there is only a small jump in the real exchange rate right at time 0, which suggests that the *surprise* component in the TFP change at time 0 does not matter much for the exchange rate. After the actual TFP improvement the real exchange rate steadily depreciates back to its long-run mean. Thus, we estimate that news about future TFP are indeed strongly reflected in the real exchange rate.

Moving onto the response of real interest rates and consumption, we see that the 3-month real interest rate differential also increases (meaning higher U.S. interest rates) before the TFP innovation, peaking at around 7.5 basis points higher than its long-run mean (which is 0.3% at an annualized basis), just before time 0. The interest rate differential then steadily declines after the TFP increase materializes, and is in fact significantly lower than its long-run mean for a prolonged period of time between 10 and 20 quarters after the TFP improvement. Similarly, there is also a U.S. consumption boom before the actual TFP improvement, and while foreign consumption also increases, that effect is noticeably weaker, and the consumption differential is large and positive (not pictured). There is also a sense in which the exchange rate leads consumption as the peak effect occurs sooner than the peak effect in consumption.

This difference in the timing of effects is in fact more straightforward to see in the impulse responses to the expectational shock ε_t^v, which we plot in Figure 5 and discuss next. First, upon the improvement in expectations (recall that is period 0 on the x-axis), the real exchange rate immediately appreciated strongly. This is another, and perhaps more direct, estimate of the anticipation effects on the exchange rate, as the expectational noise disturbance has no impact on TFP at any horizon. The exchange rate response to the fluctuations in expectations is also fairly persistent, with the exchange rate returning to its long-run mean only after around 12 quarters. The interest rate differential (another asset price), also jumps immediately upon impact of the upward revision in expectations.

On the other hand, the response in consumption is much more gradual and delayed, with no significant jump upon the shift in expectations at time 0, but rather a delayed increase in consumption peaking around 3 years after the shift in expectations. This is a direct evidence of the delayed effects on consumption which contrast with the immediate impact of news on the exchange rate. These differential responses suggest that the underlying information structure is one where the news agents receive is primarily about TFP at medium-to-long horizons. And while the exchange rate, as a forward-looking asset price, jumps on impact, consumption does not respond strongly until the expected TFP improvement becomes closer in time. Lastly, foreign consumption also increases, but weakly so, and thus the consumption

Figure 4: Impulse responses to Technology (ε^a) disturbances

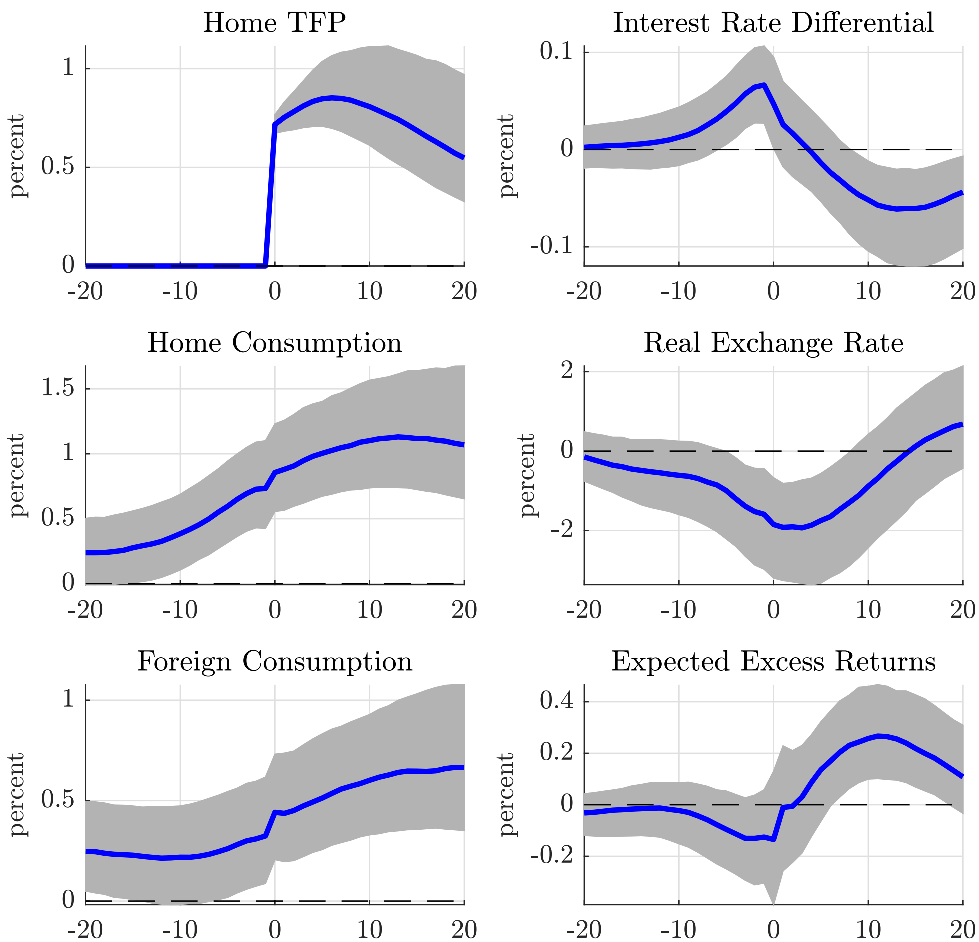

Notes: The figure displays IRFs to a one standard deviation impulse in the technological disturbance at time $t = 0$. The shaded areas are 16-84th percentile bands. Each period is a quarter.

differential is positive and significant throughout.[13]

In Appendix Figure B.1, we report the dynamics of the U.S. trade balance (as a % of U.S. GDP) in response to technological and noise disturbances. We find that the real exchange rate and the trade balance appear highly positively correlated in their conditional

[13] As described above, the responses of real variables prior to a noise disturbance are zero under the information assumptions summarized by (6)-(7), but our procedure does not impose these restrictions *ex ante*. The small responses before $t = 0$ in Figure 5 suggest that our baseline description of information captures the data well.

Figure 5: Impulse responses to Noise (ε^v) disturbance

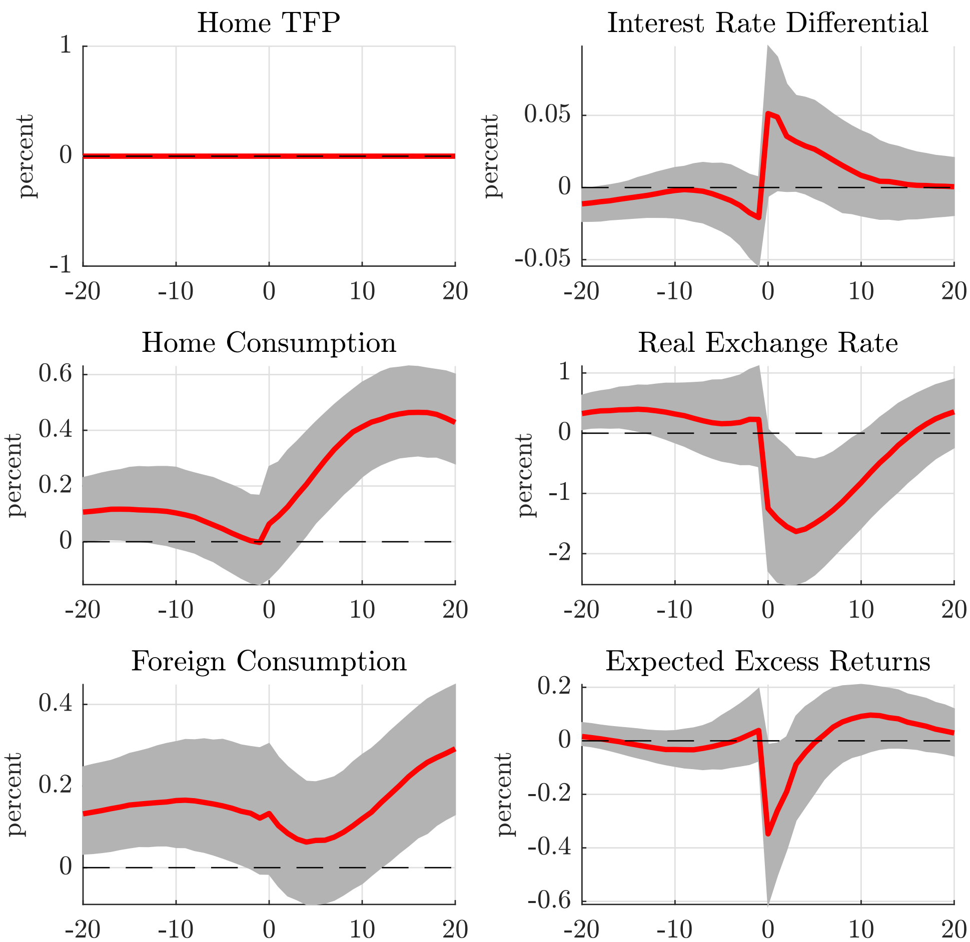

Notes: The figure displays IRFs to a one standard deviation impulse in the expectational disturbance at time $t = 0$. The shaded areas are 16-84th percentile bands. Each period is a quarter.

responses, suggesting that technological and noise disturbances are responsible, at least in part, for the positive comovement between the real exchange rate and the trade balance documented in the literature (Alessandria and Choi, 2021; Gornemann et al., 2020). We also find that the U.S. trade balance deteriorates in anticipation of future TFP improvements, albeit mildly, consistent with the so-called intertemporal approach to the current account, as well as with Hoffmann et al.'s (2019) observation that during the 1990s and 2000s, survey expectations of long-run output growth for the U.S. relative to the rest of the world were

Table 2: Variance Decomposition

	Periodicities of 2-100 Quarters			Periodicities of 6-32 Quarters		
	Both	Tech.	Noise	Both	Tech.	Noise
Home TFP	1.00	1.00	0.00	1.00	1.00	0.00
Home Consumption	0.70	0.54	0.16	0.30	0.10	0.20
Foreign Consumption	0.63	0.49	0.14	0.30	0.13	0.17
Home Investment	0.62	0.46	0.15	0.42	0.29	0.13
Foreign Investment	0.68	0.43	0.25	0.45	0.14	0.31
Interest Rate Differential	0.57	0.46	0.11	0.37	0.23	0.14
Real Exchange Rate	0.64	0.45	0.20	0.36	0.14	0.22
Expected Excess Returns	0.50	0.35	0.15	0.37	0.19	0.18

Notes: The table reports the estimated variance shares (at periodicities between 2 and 100 quarters and between 6 and 32 quarters) explained by technological disturbances ("Tech"), expectational disturbances ("Noise"), and the combination of both.

highly correlated with the U.S. current account (see also Nam and Wang, 2015).

Variance Decomposition To further quantify the effects of technology and noise disturbances, we compute the shares of the variance of the endogenous variables that these disturbances explain. Table 2 reports the decomposition of variation over both a wide band of frequencies (2-100 quarters) and the higher, business cycle frequency (6-32 quarters) many models target. As per our identification restrictions, technological disturbances account for 100% of the variation in TFP, while expectational noise disturbances are orthogonal to TFP.

Perhaps most importantly, our estimates indicate that the two disturbances together explain 60-70% of the lower-frequency variation in the levels of both the real exchange rate and also U.S. and foreign real consumption and physical investment. Thus, noisy news about future TFP appears to be a very important, common source of fluctuations for both exchange rates and real quantities in the data, reinforcing the notion that there is indeed a fundamental connection between exchange rates and the broader macroeconomy. Similar to these lower-frequency results, the two noisy-news disturbances we identify are also important at business cycle frequencies, as they explain between 30-45% of business cycles variation in the real macro aggregates and 36% of the business cycles variation in the exchange rate.

The relative importance of the true technological disturbance ε_t^a and the expectational disturbance ε_t^v differ across frequency bands in an interesting way. The true technological disturbances are more important for the lower frequencies, explaining about two-thirds of the exchange rate variation captured by the two disturbances combined (45% of 64% in total). On the other hand, at the business cycle frequencies the expectational noise is the dominant force and explains two-thirds of the variation that the two disturbances generate (22% vs 36% in total). This is a reflection of the characteristics of the IRFs we discussed above, where we saw that noise disturbances give rise to volatile and short-lived fluctuations in the exchange rate, while correctly-anticipated future TFP changes impart persistent effects on the exchange rate. A similar pattern in the relative importance of disturbances across frequencies can also be observed in the variance decompositions of all real quantities.

We stress that the variance contribution of technological disturbances are due to both movements in anticipation of the realized change in TFP, and also fluctuations that come *after* the actual change in TFP (that is, impacts both before and after time 0 in Figure 4). With this in mind, it is interesting to quantify the pure anticipation effect. To do so, we compute the variance contributed by just the leading terms in the impulse responses, *i.e.* the response of the variables in anticipation of the subsequent TFP change. We find that this anticipation effect is the main driver of the exchange rate fluctuations we estimate. Specifically, the exchange rate movements *before* the actual realization of TFP account for 34% of the lower-frequency variation of the exchange rate, while the impact after the actual realization of TFP is only responsible for 11% of exchange rate variation. Thus, the component after the realized change in TFP has little importance in explaining the real exchange rate, while the anticipation component and the expectational noise combined contribute up to 53% of the lower-frequency variation in the real exchange rate.

4 Further results and broader implications

Next, we look into a number of further results. First, Section 4.1 documents that news and noise disturbances give rise to exchange rate fluctuations that display many well known exchange rate puzzles, revealing that these puzzles share a common, fundamental origin. Second, Section 4.2 documents the effects of news and noise on other asset prices. Third, Section 4.3 discusses potential explanations for the challenges encountered by previous studies in establishing a robust correlation between exchange rates and macroeconomic fundamentals. And fourth, Section 4.4 documents that identified technological and noise disturbances are orthogonal to other economic disturbances such U.S. monetary policy shocks.

4.1 Common origin in exchange rate puzzles

The exchange rate literature is traditionally organized around the study of various "puzzles" in the empirical behavior of exchange rates. Given the large effect the two identified disturbances play in exchange rate dynamics, it is interesting to consider whether they are also generating any of the classic exchange rate puzzles. And indeed, the answer is yes they do, which will lead us to the conclusion that exchange rate puzzles seemingly have a common, and fundamental origin, in noisy news about TFP.

The Forward Premium Puzzle Our results indicate that our noisy news disturbances cause significant deviations from interest parity, and in particular generate the dynamics that characterize the "forward premium puzzle" (e.g. Fama (1984), Hassan and Mano (2018)). On the one hand, we can directly observe from Figure 4 that there is significant negative time-series comovement between interest rate differentials and the subsequent currency returns – the movements in "expected currency returns" in response to both ε_t^a and ε_t^v are closely mirrored (in the opposite direction) by the conditional dynamics of the interest rate differential.

But to formally showcase that our shocks generate the forward premium puzzle, consider the following version of the Fama regression, often used to document this pattern:

$$\lambda_{t+1} = \alpha + \beta_{UIP}(r_t - r_t^{\star}) + u_t. \tag{9}$$

Estimating this regression in our raw dataset, we find a significantly negative β_{UIP} of -2.43, in line with earlier evidence (*e.g.*, Engel, 2014). We then compute the β_{UIP} that emerges in a counter-factual dataset where only the two noisy-news disturbances we identify, ε_t^a and ε_t^v, are active. To obtain these series, we simulate from our estimated VAR by setting the variance of all other disturbances to zero.

In this counter-factual simulation we estimate a $\beta_{UIP} = -2.20$, revealing that the combination of disturbances to TFP and to expectations of future TFP qualitatively and quantitatively reproduces the classic UIP Puzzle relationship. Drilling down further, we construct similar counter-factual β_{UIP} based on either only-TFP disturbances (including anticipation effects) and only expectational noise disturbances. The results imply that the TFP disturbances by themselves generate a β_{UIP} of -2.07, while the β_{UIP} based on only expectational disturbances is -2.96, as we also report in Table 3.

The regression coefficients β_{UIP} are interesting because they are comparable to many previous findings, but it is also informative to consider how much of the raw covariance $\text{cov}(\lambda_{t+1}, r_t - r_t^\star)$—the numerator in the β_{UIP}—our disturbances can account for. Also in Table 3, we can see that our two disturbances account for roughly two-thirds of $\text{cov}(\lambda_{t+1}, r_t - r_t^\star)$—the raw covariance is -1.30, and in the counter-factual data based on our two disturbances we find a covariance of -0.86. Thus, the bulk of the negative correlation between interest rates and excess returns that underlies the UIP puzzle is due to the noisy-news disturbances we identify.

Table 3: Exchange Rate Related Puzzles and TFP Expectations

Panel A: UIP Puzzle Moments

	Technology	Noise	Both	Unconditional
Fama β_{UIP}	-2.07	-2.96	-2.20	-2.43
$\text{cov}(\lambda_{t+1}, r_t - r_t^*)$	-0.68	-0.14	-0.86	-1.30
Engel β_Λ	2.17	1.72	2.50	2.56
$\text{cov}(\sum_{k=0}^{\infty} \mathbb{E}_t(\lambda_{t+k+1}), r_t - r_t^*)$	0.52	0.06	0.60	1.08
$\sigma(r_t - r_t^\star)/\sigma(\Delta q_t)$	0.37	0.13	0.25	0.17
$\text{autocorr}(r_t - r_t^\star)$	0.99	0.93	0.98	0.95

Panel B: Backus-Smith Moments

	Technology	Noise	Both	Unconditional
$\text{corr}(\Delta q_t, \Delta(c_t - c_t^\star))$	-0.31	-0.38	-0.35	-0.27
$\text{cov}(\Delta q_t, \Delta(c_t - c_t^\star))$	-0.08	-0.12	-0.26	-0.70

Panel C: Excess Volatility and Persistence

	Technology	Noise	Both	Unconditional
$\text{autocorr}(\Delta q_t)$	0.90	0.33	0.58	0.29
$\sigma(\Delta q_t)/\sigma(\Delta c_t)$	3.99	8.14	5.65	6.04

Notes: The table reports the estimated moments conditional on technological disturbances (Technology), expectational disturbances (Noise), and the sum of both disturbances, along the moments estimated from the raw data (Unconditional). The moments in the table are defined in the text.

In addition to this "classic" Forward Premium Puzzle, the conditional responses of the exchange rate to our identified disturbances also exhibit the Engel (2016) observation that the puzzle essentially "reverses" direction at longer horizons. Namely, it has now been established that while the regression (9) finds a negative association between interest rate differentials and one quarter ahead currency excess returns, the correlation between today's interest rate differential and currency excess returns 2+ years into the future is actually *positive*.

We again see the same qualitiative pattern in the impulse responses of Figure 4: high excess currency returns in the periods following the actual TFP improvement are preceded, a few years beforehand, by high interest rates. Thus, at longer horizons, the correlation between interest rates and excess returns is positive, not negative, in our impulse responses (and this is especially pronounced in response to technological disturbances).

As a summary statistic of this phenomenon, we consider the same moment that Engel (2016) emphasizes, which is the coefficient β_Λ in the following regression:

$$\sum_{k=0}^{\infty} \mathbb{E}_t(\lambda_{t+k+1}) = \alpha_0 + \beta_\Lambda(r_t - r_t^\star) + \varepsilon_t.$$

In the raw data, we find $\beta_\Lambda = 2.56$, indicating that the whole sum $\sum_{k=0}^{\infty} \mathrm{cov}(\lambda_{t+k+1}, r_t - r_t^\star)$ is positive, even though the very first term of the sum (*i.e.* the coefficient in regression (9)) is negative. In our counter-factual simulation where only the two disturbances we identify are active, we find $\beta_\Lambda = 2.50$. Moreover, the two disturbances together generate around 60% of the overall $\mathrm{cov}\left(\sum_{k=0}^{\infty} \mathbb{E}_t\lambda_{t+k+1}, r_t - r_t^\star\right)$ in the data, though it is interesting to note that the expectational noise disturbances are responsible for only one tenth of this effect. Thus, the reversal in the UIP puzzle also largely emerges as a result of the dynamic responses to the noisy news disturbances we have identified.

In closing, we also note that our two disturbances not only generate empirically relevant regression β's, but the underlying dynamics of the interest rate differentials are also largely in line with their unconditional counterpart, as can be seen by the $\sigma(r_t - r_t^\star)/\sigma(\Delta q_t)$ and $\mathrm{autocorr}(r_t - r_t^\star)$ moments reported in Table 3.

Deviations from the Backus-Smith condition We now turn to the so called Backus-Smith puzzle. The classic international risk-sharing condition of Backus and Smith (1993) predicts that relative consumption across countries should be strongly positively correlated with the real exchange rates. However, this condition is largely violated in the data, where, in fact, this correlation is mildly negative (Kollmann, 1995; Corsetti et al., 2008b).

And indeed, we can see this negative relationships qualitatively in our IRFs. From Figure 4 we observe that in anticipation of the U.S. TFP improvement (*i.e.* $t < 0$) the U.S. dollar *appreciates* even though U.S. consumption is relatively high. The IRFs in response to the expectational noise disturbances in Figure 5 also showcase a similar relationship.

To quantify the effects, we consider the following moment:

$$\mathrm{corr}(\Delta q_t, \Delta c_t - \Delta c_t^{\star}),$$

and compute it both in the raw data and in a counter-factual simulation where only the noisy-news disturbances are active. The results are presented in Table 3 (Panel B). As in previous research the correlation in the raw data is mildly negative, equal to -0.27 in our sample. In the counter-factual sample generated by only the two disturbances we identify, this correlation is very similar and equals -0.35. Moreover, computing the covariance $\mathrm{cov}(\Delta q_t, \Delta c_t - \Delta c_t^{\star})$ which underlies the negative correlation of the Backus-Smith puzzle, we find that our two disturbances combined explain around 40% of this moment in the data. Thus, the noisy-news disturbances we find indeed play an important role in generating the celebrated Backus-Smith puzzle as well.

Excess volatility and persistence Two other well-known exchange rate "puzzles" are the excess persistence and volatility of the real exchange rate, and next we ask to what extent these phenomena are also accounted for by the noisy-news disturbances we identify.

In Table 3, we consider a few related moments. We find that the exchange rate dynamics conditional on the two disturbances we extract are indeed highly persistent. This is exemplified by the autocorrelation of the quarterly change in the real exchange rate: conditional on both identified disturbances the autocorrelation is 0.58 versus 0.29 in the unconditional data. Thus, our two disturbances generate an even higher degree of persistence than the exchange rate exhibits on average (with the unconditional persistence already being "puzzling"). It suggests that all other disturbances driving the exchange rate (*e.g.* monetary shocks) have relatively transitory effects (as is true in standard models). Thus, the puzzling persistence of the exchange rate is indeed mainly due to its responses to noisy-news about TFP.

Second, we find that the volatility of the exchange rate generated by the noisy-news disturbances is also very high, relative to macro aggregates. For example, the ratio of the standard deviation of the quarterly growth in the exchange rate and consumption conditional on both identified disturbances is 5.65, while the same ratio is 6.04 in the raw data. Notably, noise disturbances generate twice the amount of excess volatility in exchange rates compared to technological disturbances. This outcome is not surprising, as noise disturbances affect expectations, and thus asset prices, without being linked to subsequent changes in productivity. Consequently, they give rise to volatile fluctuations in exchange rates while causing only minor shifts in macroeconomic aggregates, such as consumption and output.

4.2 Technology, noise, and other asset prices

Do other asset prices react to the noisy news we identify? In Figure 6, we present the response of domestic and foreign equity prices to technological and noise disturbances. Indeed, we find a significant increase in equity prices, both domestically and internationally, in anticipation of future improvements in home productivity.[14] This finding is consistent with Beaudry and Portier's (2006) finding that equity prices incorporate, at least in part, news about future economic fundamentals. From the right column of the Figure (the IRF in response to ε_t^v) we can also directly see the intuitive result that the US equity prices react first to the arrival of news about future US TFP.

In turn, Figure 7 reports the response of *risk premia* across home and foreign assets, calculated as the VAR-implied expected quarterly excess return differentials for both equity and long-term bonds. Analogously to the excess currency returns we analyzed earlier, excess return differentials in equity and bonds are calculated by subtracting the *expected* foreign return on equity or long-term bonds (expressed in foreign currency) from the corresponding home return (expressed in home currency), in excess of the foreign interest rate differential. We estimate significant fluctuations in cross-country returns across both equity and long-term bonds in response to our shocks, indicating that the news we estimate indeed impact risk premia of other assets as well.

Nevertheless, the fact that risk premia across many different assets respond to the noisy news about TFP does not mean that the implied contemporaneous correlation between these variables is necessarily high. In fact, the variation generated by technology and noise disturbances implies a correlation between equity and currency expected excess returns of only around -0.10. This is consistent with the fact that unconditionally (in the raw data), some studies find that equity and currency returns are only mildly negatively correlated (*e.g.*, Verdelhan, 2010, and Hau and Rey, 2005) while other studies find no significant correlation (Chernov et al., 2023). The low correlation is due to the fact that equity and currency returns are negatively correlated due to pure fluctuations in TFP expectations (*e.g.* see IRF

[14]We estimate the impulse response of equities and long-term bonds (which are not directly included in the VAR), by projecting those returns on the VAR and its lags, and then using the VAR impulse responses. In any case, the results reported in this section are unchanged if instead we alternatively add equity or bond prices in our VAR and repeat the whole identification procedure.

Figure 6: Technology, noise, and equity prices

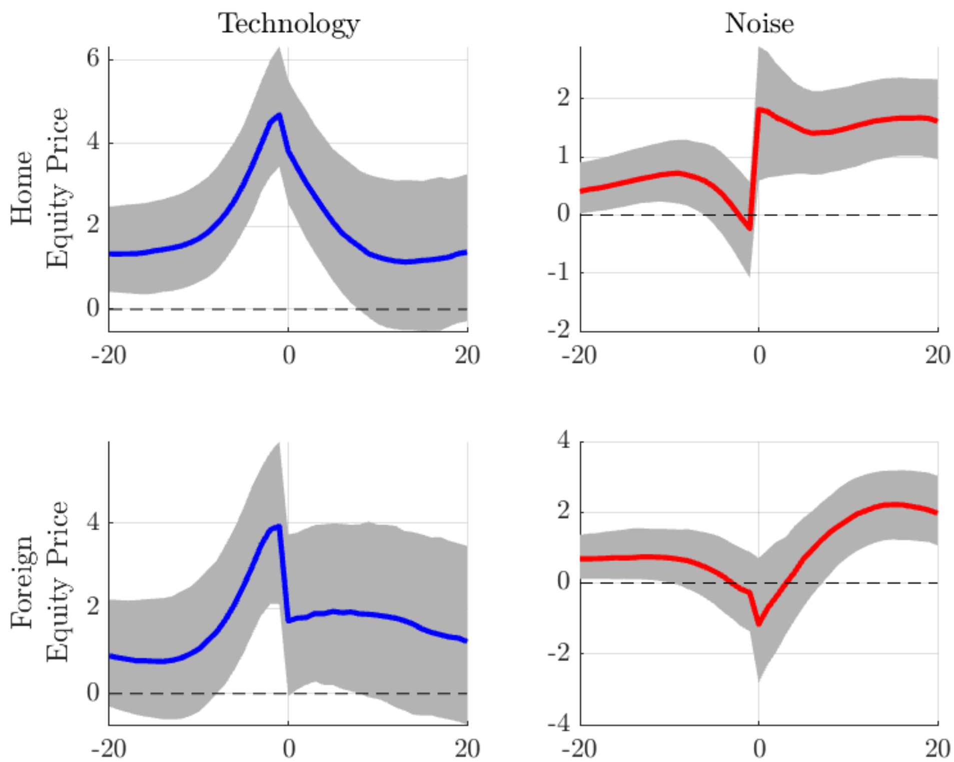

Notes: The figure displays the IRF a one standard deviation impulse in the technological disturbance (left column) and the expectational disturbance (right column) at time $t = 0$. The shaded area is the the 16-84th. Each period is a quarter.

to ε_t^v), but are *positively* correlated following an actual improvement in TFP (see IRF to ε_t^a for $t \geq 0$). This showcases that while the link between noisy news and assert prices is highly significant, it is also subtle. Our results provide sharp guidance and discipline for the development of future models, as they require equity and currency risk premia to comove differently in anticipation of future TFP improvements and following actual changes in TFP.

Contrary to the low-implied correlation between the generated fluctuations in equity and currency risk premia, the implied correlation between bond and currency returns, on the other hand, is a robust -0.50. This implies that noisy-news disturbances could also underlie the strong negative correlations observed in studies examining the joint relationship between long-term bonds and exchange rates (*e.g.*, Lustig et al., 2019, Greenwood et al., 2023, Gourinchas et al., 2022, and Lloyd and Marin, 2020).

Figure 7: Technology, noise, and other asset prices

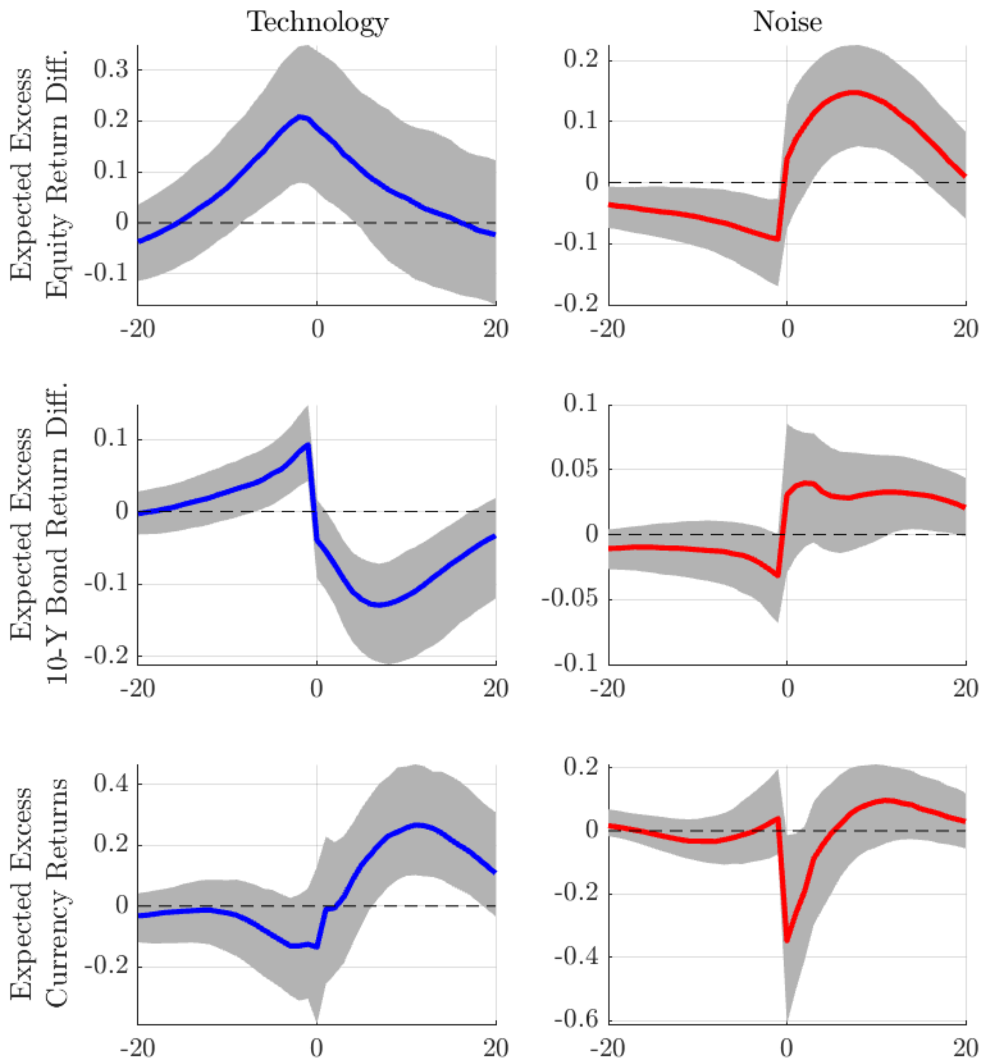

Notes: The figure displays the IRF a one standard deviation impulse in the technological disturbance (left column) and the expectational disturbance (right column) at time $t = 0$. The shaded area is the the 16-84th. Each period is a quarter.

Thus, overall, noisy news disturbances are reflected across a variety of asset classes, as is to be expected. And the significant effects on all these different asset prices is indeed consistent with prior results on both the low unconditional correlation between currency and equity risk premia, and the strong correlation between currency and bond risk premia. At the same time, it's important to note that these disturbances account for only a portion of the variation in risk premia: the two identified disturbances explain around 50% of the business cycle fluctuations in equity and bond returns. Their residual variation could be attributed to idiosyncratic factors, potentially explaining the overall weak relationship between exchange rate returns and returns in other asset classes (Burnside, 2011; Chernov and Creal, 2023; Chernov et al., 2023).

4.3 Exchange rates and future fundamentals

Overall, our headline result is that exchange rate fluctuations are significantly related to predictable changes in future TFP (*i.e.* "news"). As TFP is the quintessential macroeconomic "fundamental" in most models and empirical studies, it is perhaps surprising that the strong connection we uncover has gone unnoticed until now. We discuss possible explanations, which also sheds more intuition on the key empirical features that underlie our results.

First, many of the previous studies that have tried to find a link between exchange rates and macro fundamentals, and TFP in particular, take as a working hypothesis the standard model formulation where all TFP disturbances are pure surprises. From that point of view, one would only look for a relationship between the exchange rate and current and past TFP (and macro aggregates more generally). Studies looking for a relationship between the exchange rate and past fundamentals show a very low empirical connection between the two.

The lead-lag relationships between exchange rates and TFP is the key feature of the data that underlies our VAR identification. To show this clearly, consider the following simple exercise, where we regress the annual change in the real exchange rate at time t on leads and lags of the annual change in TFP:

$$\Delta q_t = \alpha + \beta_0 \Delta TFP_t + \sum_{k=1}^{h} \beta_{-k}^{lag}(\Delta TFP_{t-k}) + \sum_{k=1}^{h} \beta_{k}^{lead}(\Delta TFP_{t+k}) + \varepsilon_t \tag{10}$$

If we included just the constant and the first term on the right-hand side, the regression would estimate the standard relationship between contemporaneous changes in the exchange rate and TFP. If we include also the first summation term on the right-hand side, then we would also consider the additional explanatory power of lagged changes in TFP of up to h-years in the past. Once we include the second summation term, we also consider a potential correlation with *future* TFP changes, of up to h-years ahead.

Figure 8: Real exchange rate growth and leads and lags of TFP growth

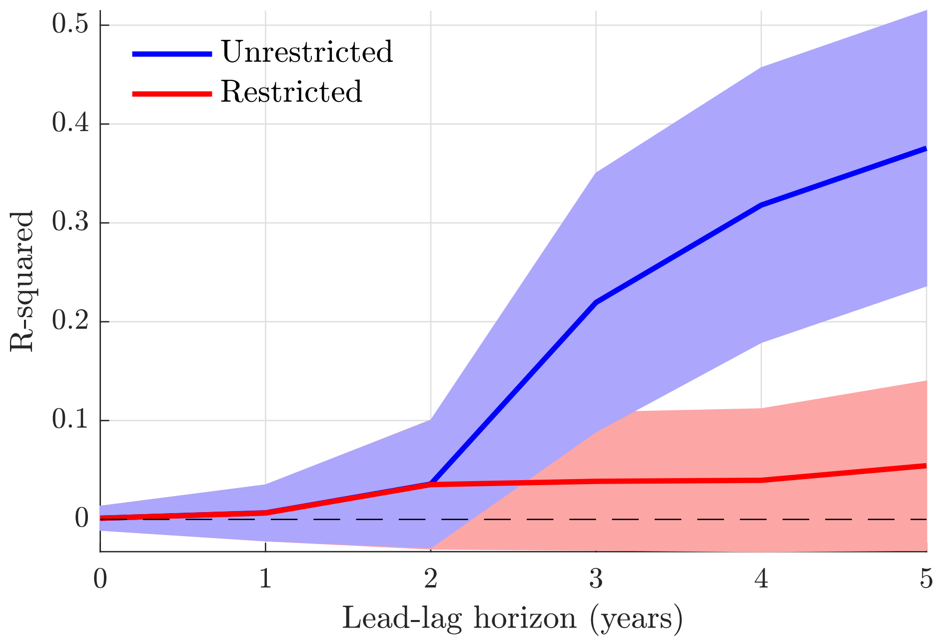

Notes: The figure reports the R^2 of a regression of exchange rate changes on present and past TFP (Restricted), and the R^2 of a regression of exchange rate changes on present, past and future TFP (Unrestricted). See regression equation (10).

Figure 8 reports the R^2 of two versions of the regression equation (10): a "Restricted" backward-looking version that includes only current and lagged TFP growth terms (red line), and an "Unrestricted" version that includes all terms on the right-hand side of (10) (blue line).

The R^2 of the purely backward looking regression is statistically insignificant no matter how many lags of TFP growth one includes, embodying the typical "disconnect" result. On the other hand, the message changes substantially once one also includes terms capturing *future* TFP growth. The relationship between real exchange rate changes and TFP growth is similarly insignificant if we only include TFP growth of up to 2 years in the future, but becomes increasingly significant as we include TFP growth 3 to 5 years out. Thus, just a simple regression can show that exchange rates do lead TFP, and this predictive nature of the exchange rate is what underlies our VAR results.[15]

[15] Appendix Table B.4 reports the p-values from a Wald test of Granger causality based on equation (10), for both the G6 aggregate and the individual countries, across different annual horizons. We reject the null hypothesis in favor of exchange rates Granger-causing TFP for most countries at horizons longer than two

More specifically, the above results also highlight again that the exchange rates contain a substantial amount of information about future TFP growth in the *medium-run to long-run*—a message that also emerged from our more detailed analysis above. This long-horizon nature of the noisy news we estimate is important in the context of the existing literature. A number of existing papers have also made the observation that the exchange rate, a forward looking asset price, ought to lead macroeconomic fundamentals. Yet, this literature has generally struggled to find robust correlation between exchange rates and future macroeconomic fundamentals.

Engel and West (2005) find a weak predictive relationship between exchange rates and future "macroeconomic fundamentals" using Granger causality regressions. Their regressions focus on quarterly exchange rate changes and forecast horizons of up to one year for variables such as GDP and interest rate differentials.[16] We highlight two main observations in regard to the evidence in Engel and West (2005). First, given the low frequency nature of the news we estimate, *short-horizon* regressions may tend to yield limited predictive power in finite samples. Second, we will note that Granger causality does not need to manifest for *all* macroeconomic fundamental considered by Engel and West (2005), even when TFP expectations are important drivers of exchange rates.

Figure 9 summarizes the significance of Granger causality tests for a range of variables and countries, using the 1974:Q1-2001:Q3 sample and variables originally considered by Engel and West. All regressions are estimated with Newey-West standard errors to account for serial correlation in errors. Consistent with our hypothesis that lower-frequency relationships are key, the Figure shows that the predictive relationships studied in Engel and West (2005) strengthen at longer horizons, which they did not consider. A similar pattern emerges when we extend the sample through 2018:Q4 (see Appendix Figure B.8). This pattern is intuitive and consistent with noisy-news view: long-horizon exchange rate changes reflect accumulated information about future macro movements, when the influence of transitory noise tends to average out.

That said, Granger causality results are not uniformly significant, and the strength of the exchange rate-fundamentals relationship differ by variable. In particular, the evidence of Granger causality is stronger for output, money, and price differentials, while the relationship is weaker for interest rate differentials.

years.

[16] Sarno and Schmeling (2014) adopt a more non-parametric approach, but still limit their null hypothesis to testing predictive relationships to a maximum of two years ahead.

Figure 9: Bivariate Granger Causality Tests, Different Macro Aggregates (Sample: 1974:Q1-2001:Q3)

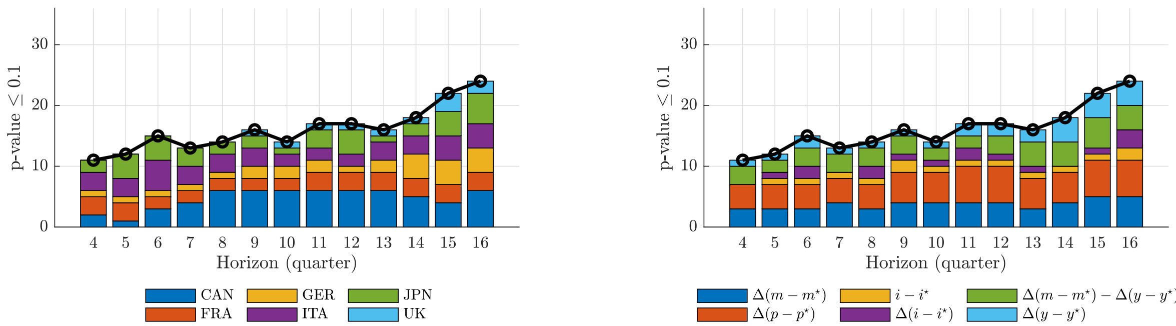

(a) P-values of Granger Causality tests, by country

(b) P-values of Granger Causality tests, by variable

Notes: Number of p-values below 0.1 from a Wald test of Granger causality using the data and procedure of Engel and West (2005). We compute Newey-West standard errors using a data-driven bandwidth, a Bartlett kernel, and prewhitening. Each test is based on a bivariate VAR estimated using the quarterly change in the nominal exchange rate, Δs, and a given fundamental, Δf. The horizon on the x-axis denotes the number of lags included in the VAR. The p-value corresponds to the Wald test of the joint significance of the lags of Δs in predicting Δf. Specifically, Δm is the percentage change in M1 (M2 for the United Kingdom); Δy is the percentage change in real GDP; Δp is the percentage change in consumer prices; and i is the short-term interest rate on government debt.

In this regard, we note that even if anticipated TFP movements drive exchange rates, this does not necessarily imply that exchange rate changes should Granger-cause *all* macro variable. To illustrate this point, Appendix C draws a simple-model example where agents receive advance information about future TFP. In that model, both exchange rates and interest rate differentials respond to news about future TFP. As a result, the exchange rate does *not* contain information about future interest differentials that is independent of current interest differentials. Hence, exchange rate changes do *not* Granger-cause interest differentials, despite the presence of advance information about future TFP.

The intuition is straightforward: interest differentials are themselves forward-looking, like exchange rates, and thus already incorporate the same future information. This is consistent with our empirical findings, where interest differentials comove contemporaneously with exchange rates (see, e.g., Figure 1). It also aligns with the evidence in Kekre and Lenel (2024), who show a tight contemporaneous comovement between bond yields and exchange rates. Therefore, the exchange rate-fundamentals lead-lag relationship should be more evident, and it is so, for slow-moving fundamentals, such as TFP or consumption differentials.

Lastly, it is important to stress that our analysis is also unique in that it directly accounts for the noise in expectations of future TFP and this increases the statistical power of our approach. The reason is that the expectational noise causes changes in TFP expectations that do not subsequently result in actual TFP. Since exchange rates react to expectations regardless of whether those expectations are correct ex post, the raw correlations between realized fundamentals and exchange rates is likely to systematically understate the importance of fluctuations in expected TFP (since a significant portion of the expectations is noise). By accounting for the expectational noise directly, rather than relying on a simple correlation between exchange rates and ex-post realizations of future TFP, we obtain much greater statistical power in estimating the potential impact of noisy news. This same reasoning could help explain why studies employing direct measurement of expectations through survey data tend to reveal a much higher degree of connection between exchange rates and fundamentals (*e.g.*, Engel et al., 2008; Stavrakeva and Tang, 2020).

4.4 Other economic shocks

In order to interpret our empirical results as picking up disturbances that are solely about TFP and its expectations, and not potentially also mixing up other economic shocks that could be endogenously related to TFP growth, in Sections 3 and 4 we make the (common) assumption that the Fernald utilization-adjusted TFP is orthogonal to any other economic shock. One may be concerned that other economic disturbances such as shocks to R&D productivity, monetary policy, or may be even FX noise trader/sentiment shocks in the vein of Itskhoki and Mukhin (2021) may lead to a future change in TFP through endogenous investment in research and development (R&D). In that case, our identified technological and noise disturbances might be contaminated by other economic shocks.

On the one hand, we view shocks to R&D productivity as a type of "news" disturbances about future TFP. R&D investment is a small proportion of the overall macroeconomy, so a productivity shock to such a small sector is unlikely to have any direct impacts on the exchange rate itself, and thus our estimates are still correctly picking up that exchange rate variation is to a large extend driven by predictable fluctuations in future aggregate TFP.

On the other hand, if current *contractionary* monetary shocks spur R&D investment and thus future TFP growth, then indeed our identified noisy news disturbances might be related to such monetary shocks. We stress, however, that for this to be a viable alternative explanation it must be the case that contractionary monetary shocks spur R&D investment and improve future TFP, since our empirical results in Figures 4 and 5 show that the exchange rate *appreciates* in anticipation of TFP improvements, and if that appreciations is to be driven by monetary shocks it must be due to monetary tightening.

We see the hypothesis of contractionary monetary policy spurring R&D as unlikely, but for robustness we also directly test whether our procedure picks up U.S. monetary policy shocks. The measure of U.S. monetary policy shocks we consider is the one identified through the "high frequency approach" by Gertler and Karadi (2015) and recently updated by Jarocinski and Karadi (2020). In Table 4 we report the correlation between our technology and expectational noise disturbances, and these U.S. monetary policy shocks. We find that both technology and noise disturbances are orthogonal to U.S. monetary policy disturbances.

These findings also complement the empirical results of Kim et al. (2017), who find that while the exchange rate responds significantly to monetary policy shocks, the response to such shocks *does not* display any violations of UIP. However, the impulse responses to our disturbances imply significant deviations from UIP, which is another reason to think that we are picking up something different than monetary policy shocks.

Lastly, we should not lose sight of the fact that while our estimated noisy news about TFP account for a significant portion of RER variation (up to 66% overall, and roughly a third of the variation of q_t at business cycle frequencies), our identified disturbances still leave a non-trivial fraction of the exchange rate variation unexplained. Certainly a number of other shocks also play an important role in the exchange rate fluctuations in the data, with monetary shocks a prime example of such disturbances.

Table 4: Correlation between Technology, Noise and U.S. Monetary Policy disturbances

	Technology	Noise
U.S. Monetary Policy disturbances	0.09 (p-value = 0.46)	0.06 (p-value = 0.62)

Notes: The table reports the correlation between technological disturbances (Technology) and expectational disturbances (Noise) with U.S. monetary policy disturbances from Gertler and Karadi (2015) and recently updated by Jarocinski and Karadi (2020).

Estimation of VAR without exchange rate We have also performed the baseline analysis after removing the exchange rate from the set of observables we use to estimate the VAR. We do so to guard from the possibility that some FX noise trader/sentiment shock is leaking into our estimated noisy news disturbances. The key result here is that with this restricted VAR, we still recover essentially the same set of technology and noise disturbances (Figure B.6), and, in turn, these disturbances have a similar effect on the exchange rate (Figure B.7).[17] This highlights again that the key empirical variation our analysis exploits is that a substantial portion of TFP growth is predictable, and the relevant (noisy) TFP expectations are priced into the exchange rate.

5 Implications for Models

A key conclusion of our analysis is that a collection of famous exchange rate puzzles have a *common* and *fundamental* origin in noisy news about future TFP. These results have important implications about the development of general equilibrium open-economy models. To this end, this section confronts three leading classes of models with our findings.

To evaluate these models, we parameterize them with a commonly used noisy information specification for the TFP process and information available to agents. As a baseline, we use the noise representation of the structure in Blanchard et al. (2013). We use this process because it fits the data on TFP dynamics and expectations well (as we will explicitly show below). But we also consider the long-run risk type of formulation in the Appendix, and find the same qualitative conclusions.[18] To close the models, we assume that foreign TFP is cointegrated with home TFP, that is $\Delta a_t^\star = \tau(a_{t-1} - a_{t-1}^\star)$, and we calibrate the parameter $\tau \in (0, 1)$ to a small number to capture slow cross-country convergence.

We study three two-country endowment economies, described in detail in Appendix D.1. These models differ in their international asset-market structure and household preferences. The first model features complete markets and CRRA preferences, in the spirit of Backus et al. (1994). The second model also assumes complete markets but adopts Epstein and Zin (1989) preferences, as in Colacito and Croce (2011; 2013; 2018). The third model features incomplete international asset markets with CRRA preferences, following Corsetti et al. (2008b), Itskhoki and Mukhin (2021), and Kekre and Lenel (2024).

To evaluate these models, we study the impulse responses of key international macro

[17] We estimate the impulse response of the exchange rate in the restricted VAR by first projecting the exchange rate on the VAR variables and their lags, and then constructing the impulse response from this approximation of the exchange rate. The procedure is akin to using Jorda-projections to estimate the response of q_t to the shocks we identify from the restricted VAR that does not include the exchange rate.

[18] We provide an explicit mapping from the original Blanchard et al. (2013) and Colacito and Croce (2013) information structures to the noise structure in (5)-(7) in Appendix D.2 and D.3.

variables to both technology and noise disturbances, ε_t^a and ε_t^v respectively, and compare them to our empirical counterparts. Our goal is to ask whether these models can replicate the *qualitative* patterns in response to the noisy news disturbances ε_t^a and ε_t^v. These impulse responses are shown in Figures 10 and 11, while Table 5 reports the parameter calibration. The model with Epstein and Zin (1989) preferences is solved to a third-order approximation, whereas the CRRA models are solved at first order.

We calibrate the parameters of the Blanchard et al. (2013) information structure so as to match the empirical impulse responses of TFP and TFP expectations displayed in Figures 10 and 11 (column 1). Thus, they are disciplined by the behavior of expectations and actual TFP realizations we estimate. We then study how the transmission of the same ε_t^a and ε_t^v disturbances differs across the three models, comparing the propagation of shocks to the data.

Complete Markets with CRRA Preferences The complete markets model with CRRA preferences yields two main counterfactual implications.

First, expected future changes in TFP do not affect current consumption differentials or the real exchange rate (Figures 10–11, column 2). Under complete markets and time-separable preferences, all shocks are fully insured: marginal utilities are equalized across countries when expressed in common units. Therefore, even if home TFP is expected to rise in the future, this does not translate into a relative welfare or wealth gain for home agents today. As a result, allocations remain unchanged until TFP actually increases. At that point, the social planner allocates a larger share of the expanded endowment to home agents due to home bias in preferences, raising the home consumption differential and depreciating the exchange rate. Prior to the realization of the TFP shock, however, neither consumption nor the real exchange rate responds—contrary to our empirical findings, which show that both variables adjust in anticipation of future TFP improvements.

Second, when TFP actually improves there is both an increase in home consumption and exchange rate depreciation because the usual Backus-Smith condition holds. Yet, in the data we find a clear negative comovement between exchange rates and consumption conditional on the TFP disturbance.

Complete Markets with Epstein and Zin Preferences A leading alternative to CRRA preferences in international finance is the recursive utility model of Epstein and Zin (1989). Under these preferences, marginal utility depends on the entire future consumption path, allowing anticipated TFP changes to affect current outcomes and also breaks the Backus-Smith condition, as shown by Colacito and Croce (2013).

However, while with Epstein-Zin preferences the complete markets model can generate anticipation effects, we find that these effects go *in the opposite direction* of what we find in the data. When expectations of future home TFP improve, the consumption differential falls and the real exchange rate depreciates (Figures 10–11, column 3). This pattern reflects the recursive utility framework, where households derive current utility from favorable news about future consumption. With complete risk-sharing, since the utility of home agents increases purely due to the "good news", they effectively "share" these good news with foreign agents by transferring resources abroad today, which lowers the consumption differential before TFP improves. For the same reason, the exchange rate also depreciates. So while the model produces the (correct) negative correlation between consumption in the real exchange rate, but it does so by implying responses of both variables that are qualitatively the opposite of those we see in the data. In the data, the real exchange rate appreciates and relative consumption rises in response to good news, while in the model both variables fall.

Importantly, these qualitative results are not specific to the structural TFP process we chose in this analysis (outlined in (D.15) in Appendix D.2). In Appendix D.3, we show that similar dynamics arise under a TFP process with long-run risk features, calibrated as in Colacito et al. (2018). In these cases as well, anticipated TFP improvements lead to a decline in the consumption differential and a depreciation of the real exchange rate. The only difference is that with the LRR calibration the excess currency returns are quantitatively much more meaningful. But qualitatively, because foreign currency returns fall in expectation of TFP improvement (as we see in column 2), amplifying these excess returns movements with the LRR process only depreciates the exchange rate even more (while, again, the exchange rate appreciates in anticipation of TFP improvement in the data).

Incomplete Markets with CRRA Preferences Lastly, we consider a model with incomplete international asset markets, where we restrict agents to trading a single international bond and impose a debt-elastic interest rate premium to induce stationarity of equilibrium dynamics (Schmitt-Grohé and Uribe, 2003). We calibrate the elasticity of the interest rate premium to debt to be small enough so that UIP approximately holds, which is consistent with the calibration in Itskhoki and Mukhin (2021) and Kekre and Lenel (2024).

We find that the incomplete markets framework is most successful at generating our empirical findings, although it still misses in some ways. Importantly, the model qualitatively reproduces several empirical features: higher consumption and interest rate differentials, and an appreciated real exchange rate in anticipation of the TFP increases (Figures 10–11, column 4). The reason is that the positive news generate a wealth effect for home households, prompting higher current consumption and external borrowing. In turn, this drives up the interest rate differential and appreciates the real exchange rate. As a result, news about future TFP operates like a demand shock in models with incomplete markets, as also emphasized by Kekre and Lenel (2024).

Nevertheless, the real exchange rate still depreciates strongly when the TFP increase actually occurs (Figure 10, column 4), as the relative price of the now-abundant home good falls. This is counterfactual to our empirical findings, which show that there is no abrupt movement in the real exchange rate up on the realization of the TFP improvement, and the exchange rate depreciates only slowly and smoothly after that point. We conjecture that one ingredient missing from the model might be sufficiently strong movements in expected excess currency returns (currency premia). Specifically, in the data we find that the foreign currency premium increases in the aftermath of TFP improvement, but in the model the expected excess returns are effectively zero. If the model was to similarly imply an increase in the foreign currency premium then this will generate a force that pushes towards exchange rate appreciation, and could thus counteract the depreciation we see at time 0 in our impulse response.

Thus, our analysis suggest that one important ingredient is incomplete markets, and another is potentially appropriately volatile currency premia. However, models that feature both incomplete markets and endogenously volatile currency premia are not well developed and hence we leave the exploration of this research avenue to future work.

6 Conclusions

We document novel empirical evidence that exchange rates are not disconnected from macro aggregates, but that they are indeed tightly linked to fluctuations in noisy expectations of future TFP improvements. Our empirical approach is particularly notable for directly accounting for the possibility that TFP expectations are noisy. Separately identifying the noise in expectations helps sharpen our results, because expectational noise reduces the raw correlation between expectations and *realized* fundamentals in the data.

In addition, we show that the noisy news disturbances we identify appear to generate a number of famous exchange rate puzzles at the same time. Thus, a myriad of FX puzzles share a common and fundamental origin in noisy news about future TFP. Also, the noisy news specifically transmits to exchange rates by causing large fluctuations in expected excess currency returns. Our evidence thus implies that the common thread that ties many FX puzzles together are news-driven UIP wedge fluctuations, providing important guidance for the development of future theoretical models.

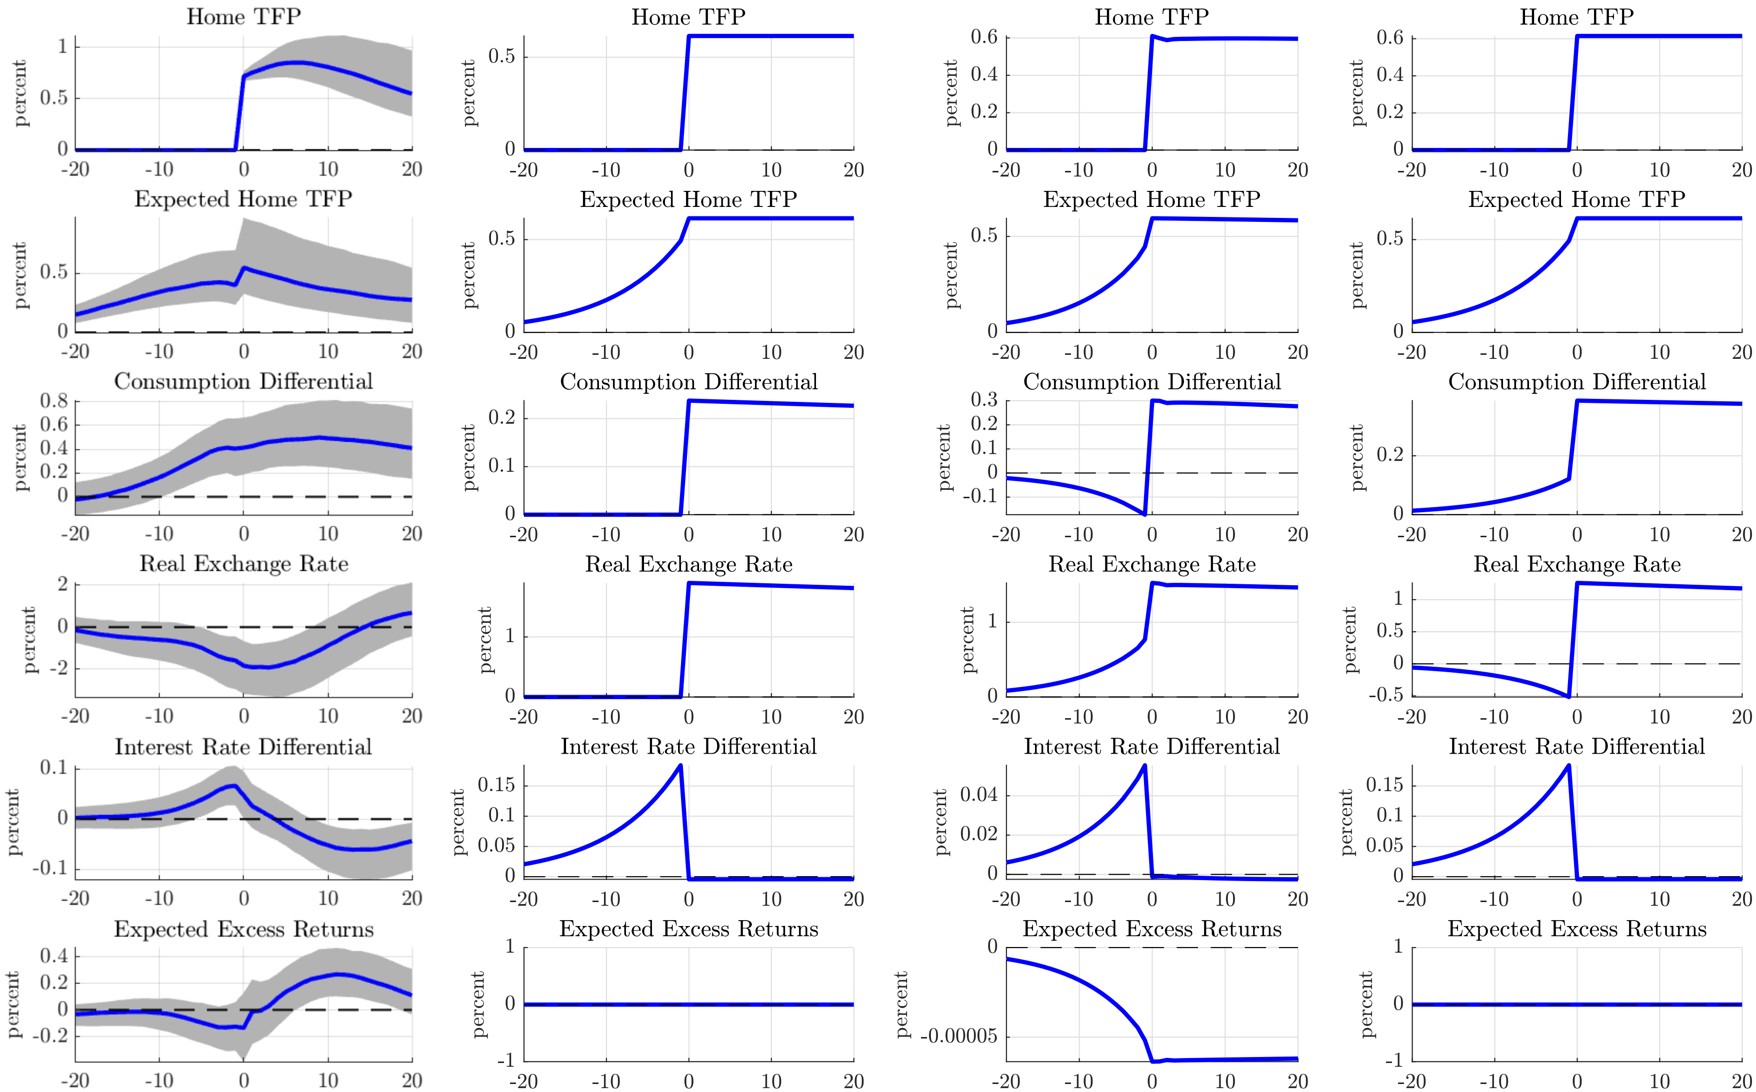

(a) Data

(b) Complete Markets CRRA

(c) Complete Markets EZ

(d) Incomplete Markets CRRA

Figure 10: Impulse Responses to Technology Disturbances: Data vs Models

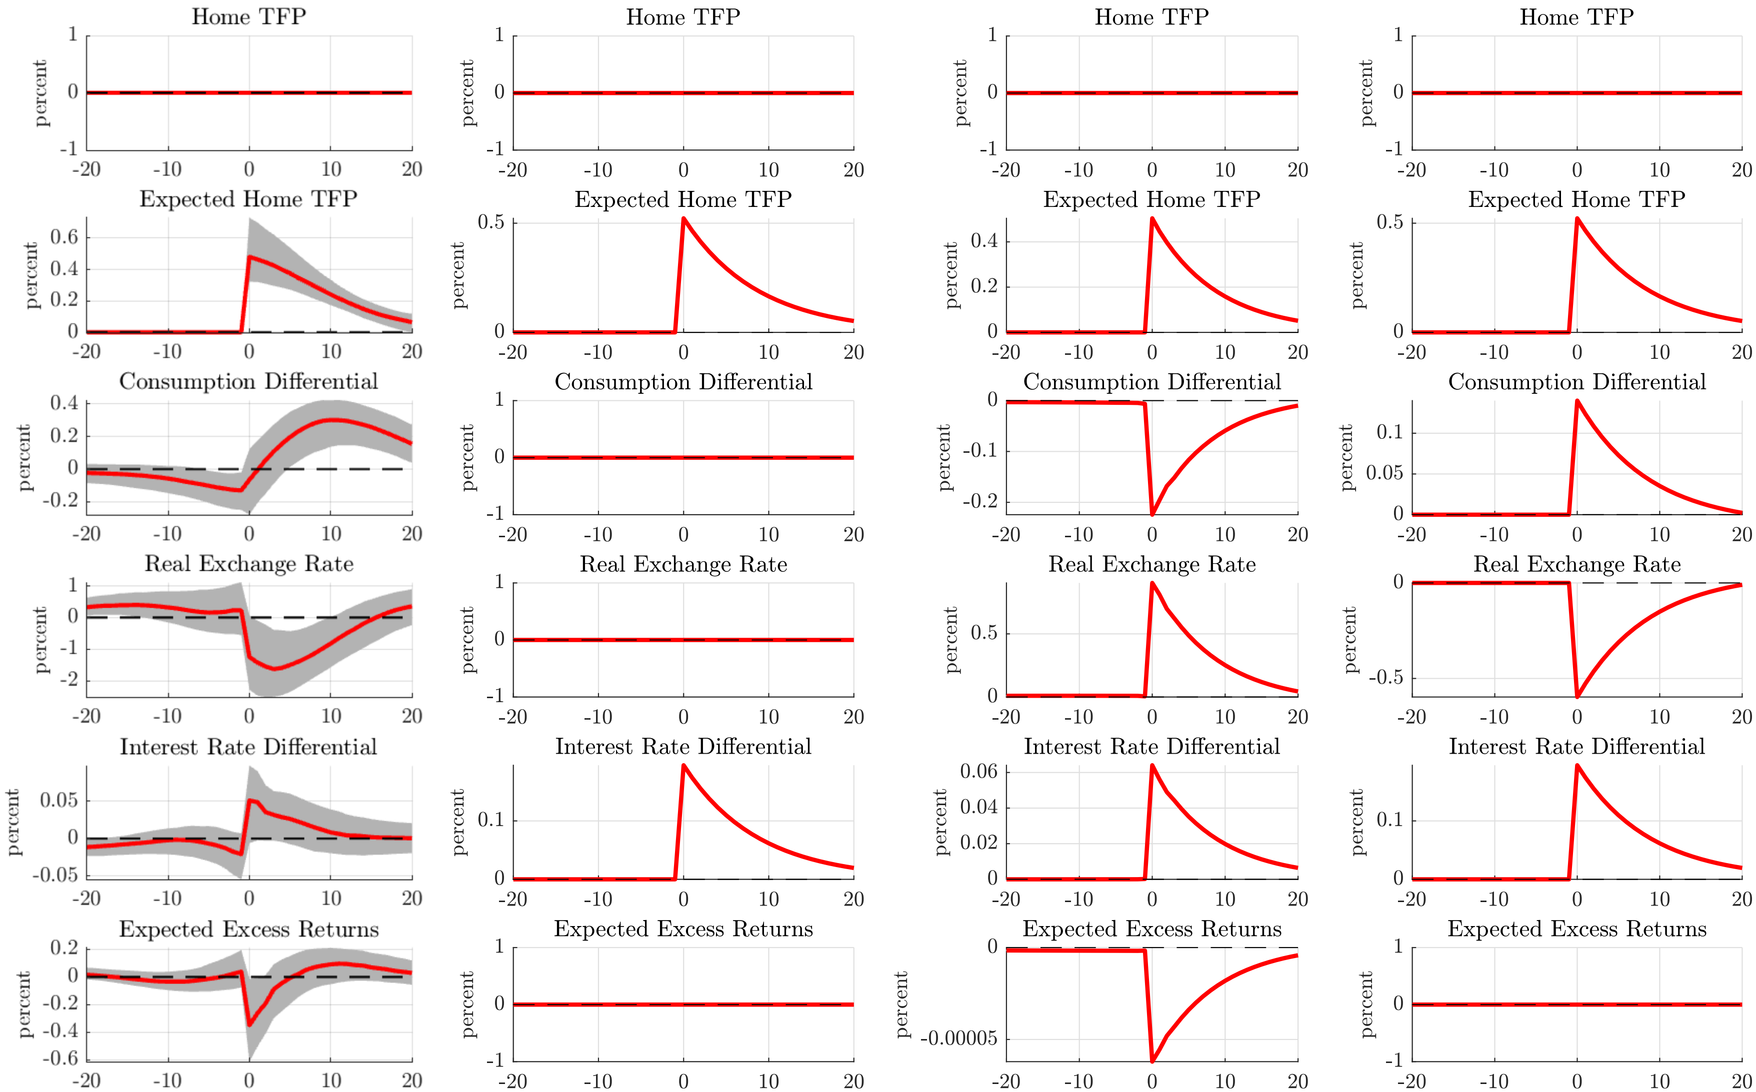

(a) Data

(b) Complete Markets
CRRA

(c) Complete Markets
EZ

(d) Incomplete Markets
CRRA

Figure 11: Impulse Responses to Noise Disturbances: Data vs Models

Table 5: Parameter Calibration

Parameter	Description	Complete Markets CRRA	Complete Markets EZ	Incomplete Markets CRRA
Preferences				
β	Subjective Discount Factor	0.99	0.99	0.99
γ	Relative Risk Aversion	8	8	8
ψ	Elasticity of Intertemporal Subs.	$1/\gamma$	0.9	$1/\gamma$
λ	Home Bias	0.95	0.95	0.95
ϕ_0	Bond Adjustment Costs	-	-	1
ϕ_1	Bond Adjustment Costs	-	-	10^{-5}
TFP processes				
Baseline (Blanchard et al., 2013)				
ρ	Persistence of "permanent shock"	0.89	0.89	0.89
σ_η	Std. dev. of "transitory shock"	0.61	0.61	0.61
σ_ξ	Std. dev. of "noise shock"	5×10^{-5}	5×10^{-5}	5×10^{-5}
σ_η	Std. dev. of "permanent shock"	$\rho\sigma_\mu^2 = (1-\rho)^2\sigma_\eta^2$	$\rho\sigma_\mu^2 = (1-\rho)^2\sigma_\eta^2$	$\rho\sigma_\mu^2 = (1-\rho)^2\sigma_\eta^2$
τ	Co-integration parameter	0.0023	0.0023	0.0023
Appendix D.3 (Long-run risk)				
ρ	Persistence of "long-run shock"	0.9899	0.9899	0.9899
σ_z	Std. dev. of "long-run shock"	0.0049	0.0049	0.0049
σ_a	Std. dev. of "short-run shock"	0.0385	0.0385	0.0385
τ	Co-integration parameter	0.0023	0.0023	0.0023

Notes: Model is calibrated to quarterly frequency. Under the long-run risk process, home and foreign shocks have symmetric standard deviations.

References

Alessandria, G. and H. Choi (2021): "The dynamics of the U.S. trade balance and real exchange rate: The J curve and trade costs?" *Journal of International Economics*, 132.

Alvarez, F., A. Atkeson, and P. J. Kehoe (2009): "Time-varying risk, interest rates, and exchange rates in general equilibrium," *The Review of Economic Studies*, 76, 851–878.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2003): "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review*, 93, 38–62.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-cycle anatomy," *American Economic Review*, 110, 3030–70.

Ayres, J., C. Hevia, and J. P. Nicolini (2020): "Real exchange rates and primary commodity prices," *Journal of International Economics*, 122, 103261.

Bacchetta, P. and E. van Wincoop (2006): "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *The American Economic Review*, 96, 552.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1994): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review*, 84, 84–103.

Backus, D. K. and G. W. Smith (1993): "Consumption and real exchange rates in dynamic economies with non-traded goods," *Journal of International Economics*, 35, 297–316.

Baillie, R. T. and D. Cho (2014): "Time variation in the standard forward premium regression: Some new models and tests," *Journal of Empirical Finance*, 29, 52–63.

Bansal, R. and I. Shaliastovich (2012): "A long-run risks explanation of predictability puzzles in bond and currency markets," *Review of Financial Studies*, hhs108.

Bansal, R. and A. Yaron (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The journal of Finance*, 59, 1481–1509.

Barsky, R. B. and E. R. Sims (2011): "News Shocks and Business Cycles," *Journal of Monetary Economics*, 58, 273 – 289.

Basu, S., G. Candian, R. Chahrour, and R. Valchev (2021): "Risky Business Cycles," NBER Working Papers 28693, National Bureau of Economic Research, Inc.

Beaudry, P. and F. Portier (2006): "Stock Prices, News, and Economic Fluctuations," *The American Economic Review*, 96, pp. 1293–1307.

——— (2014): "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, 52, 993–1074.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013): "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, 103, 3045–70.

Burnside, C. (2011): "Carry trades and risk," *Working Paper*.

Burnside, C., B. Han, D. Hirshleifer, and T. Y. Wang (2011): "Investor overconfidence and the forward premium puzzle," *The Review of Economic Studies*, 78, 523–558.

Candian, G. and P. De Leo (2023): "Imperfect Exchange Rate Expectations," *The Review of Economics and Statistics*, 1–46.

Chahrour, R., S. K. Chugh, and T. Potter (2023): "Anticipated productivity and the labor market," *Quantitative Economics*, 14, 897–934.

Chahrour, R. and K. Jurado (2018): "News or Noise? The Missing Link," *American Economic Review*, 108, 1702–36.

——— (2021): "Recoverability and Expectations-Driven Fluctuations," *The Review of Economic Studies*.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2002): "Can sticky price models generate volatile and persistent real exchange rates?" *The Review of Economic Studies*, 69, 533–563.

Chen, Y.-C., K. S. Rogoff, and B. Rossi (2010): "Can Exchange Rates Forecast Commodity Prices?*," *The Quarterly Journal of Economics*, 125, 1145–1194.

Chernov, M. and D. Creal (2023): "International yield curves and currency puzzles," *The Journal of Finance*, 78, 209–245.

Chernov, M., V. Haddad, and O. Itskhoki (2023): "What do financial markets say about the exchange rate?" *Working Paper*.

Cheung, Y.-W., M. D. Chinn, and A. G. Pascual (2005): "Empirical exchange rate models of the nineties: Are any fit to survive?" *Journal of International Money and Finance*, 24, 1150–1175.

Colacito, R., M. Croce, S. Ho, and P. Howard (2018): "BKK the EZ Way: International Long-Run Growth News and Capital Flows," *American Economic Review*, 108, 3416–49.

Colacito, R. and M. M. Croce (2011): "Risks for the Long Run and the Real Exchange Rate," *Journal of Political Economy*, 119, 153–181.

——— (2013): "International asset pricing with recursive preferences," *The Journal of Finance*, 68, 2651–2686.

Comin, D. A., J. Quintana, T. G. Schmitz, and A. Trigari (2023): "Revisiting Productivity Dynamics in Europe: A New Measure of Utilization-Adjusted TFP Growth," Tech. rep., National Bureau of Economic Research.

Cormun, V. and P. De Leo (2024): "Shocks and exchange rates in small open economies," *Working Paper*.

Corsetti, G., L. Dedola, and S. Leduc (2008a): "High exchange-rate volatility and low pass-through," *Journal of Monetary Economics*, 55, 1113–1128.

——— (2008b): "International Risk Sharing and the Transmission of Productivity Shocks," *The Review of Economic Studies*, 75, 443–473.

——— (2014): "The International Dimension of Productivity and Demand Shocks in the US Economy," *Journal of the European Economic Association*, 12, 153–176.

Devereux, M. B. and C. Engel (2002): "Exchange Rate Pass-Through, Exchange Rate Volatility, and Exchange Rate Disconnect," *Journal of Monetary Economics*, 49, 913 – 940.

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

Eichenbaum, M. and C. L. Evans (1995): "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates," *The Quarterly Journal of Economics*, 975–1009.

Eichenbaum, M. S., B. K. Johannsen, and S. T. Rebelo (2020): "Monetary Policy and the Predictability of Nominal Exchange Rates," *The Review of Economic Studies*, 88, 192–228.

Engel, C. (2014): "Exchange rates and interest parity," in *Handbook of international economics*, Elsevier, vol. 4, 453–522.

——— (2016): "Exchange Rates, Interest Rates, and the Risk Premium," *American Economic Review*, 106.

Engel, C., N. C. Mark, and K. D. West (2008): "Exchange Rate Models Are Not As Bad As You Think," in *NBER Macroeconomics Annual 2007, Volume 22*, National Bureau of Economic Research, Inc, NBER Chapters, 381–441.

Engel, C. and K. D. West (2005): "Exchange rates and fundamentals," *Journal of Political Economy*, 113, 485–517.

Engel, C. and S. P. Y. Wu (2024): "Exchange Rate Models are Better than You Think, and Why They Didn't Work in the Old Days," *Working Paper*.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, pp. 937–969.

Fama, E. F. (1984): "Forward and spot exchange rates," *Journal of Monetary Economics*, 14, 319–338.

Farhi, E. and X. Gabaix (2015): "Rare Disasters and Exchange Rates," *The Quarterly Journal of Economics*.

Farhi, E. and I. Werning (2012): "Dealing with the trilemma: Optimal capital controls with fixed exchange rates," .

Faust, J. (1998): "The robustness of identified VAR conclusions about money," *Carnegie-Rochester Conference Series on Public Policy*.

Faust, J., J. H. Rogers, S.-Y. B. Wang, and J. H. Wright (2007): "The high-frequency response of exchange rates and interest rates to macroeconomic announcements," *Journal of Monetary Economics*, 54, 1051–1068.

Fernald, J. G. (2012): "A quarterly, utilization-adjusted series on total factor productivity," *Working Paper*.

Gabaix, X. and M. Maggiori (2015): "International liquidity and exchange rate dynamics," *The Quarterly Journal of Economics*, 130, 1369–1420.

Gertler, M. and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Gornemann, N., P. Guerron-Quintana, and F. Saffie (2020): "Exchange Rates and Endogenous Productivity," *FRB International Finance Discussion Paper*.

Gourinchas, P.-O., W. Ray, and D. Vayanos (2022): "A Preferred-Habitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers," *Working Paper*.

Gourinchas, P.-O. and H. Rey (2007): "International financial adjustment," *Journal of political economy*, 115, 665–703.

Gourinchas, P.-O. and A. Tornell (2004): "Exchange rate puzzles and distorted beliefs," *Journal of International Economics*, 64, 303–333.

Greenwood, R., S. Hanson, J. C. Stein, and A. Sunderam (2023): "A Quantity-Driven Theory of Term Premia and Exchange Rates," *The Quarterly Journal of Economics*, 138, 2327–2389.

Hassan, T. A. and R. C. Mano (2018): "Forward and spot exchange rates in a multi-currency world," *The Quarterly Journal of Economics*, 134, 397–450.

Hassan, T. A., T. M. Mertens, and T. Zhang (2016): "Not so disconnected: Exchange rates and the capital stock," *Journal of International Economics*, 99, 43–57.

Hau, H. and H. Rey (2005): "Exchange Rates, Equity Prices, and Capital Flows," *The Review of Financial Studies*, 19, 273–317.

Hoffmann, M., M. U. Krause, and T. Laubach (2019): "The Expectations-driven US Current Account," *Economic Journal*, 129, 897–924.

Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2023): "Utilization-adjusted tfp across countries: Measurement and implications for international comovement," *Journal of International Economics*, 103753.

——— (Forthcoming): "International comovement in the global production network," *Review of Economic Studies*.

Ilut, C. (2012): "Ambiguity aversion: Implications for the uncovered interest rate parity puzzle," *American Economic Journal: Macroeconomics*, 4, 33–65.

Itskhoki, O. and D. Mukhin (2021): "Exchange rate disconnect in general equilibrium," *Journal of Political Economy*.

Jarocinski, M. and P. Karadi (2020): "Deconstructing Monetary Policy Surprises: The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

Jeanne, O. and A. K. Rose (2002): "Noise trading and exchange rate regimes," *The Quarterly Journal of Economics*, 117, 537–569.

Karabarbounis, L. (2014): "Home production, labor wedges, and international business cycles," *Journal of Monetary Economics*, 64, 68–84.

Kekre, R. and M. Lenel (2024): "Exchange Rates, Natural Rates, and the Price of Risk," *Working Paper*.

Kim, S.-H., S. Moon, and C. Velasco (2017): "Delayed overshooting: is it an 80s puzzle," *Journal of Political Economy*, 125, 1570–1598.

Kollmann, R. (1995): "Consumption, real exchange rates and the structure of international asset markets," *Journal of International Money and Finance*, 14, 191–211.

——— (2005): "Macroeconomic effects of nominal exchange rate regimes: new insights into the role of price dynamics," *Journal of International Money and Finance*, 24, 275–292.

Kurmann, A. and C. Otrok (2013): "News Shocks and the Slope of the Term Structure of Interest Rates," *American Economic Review*, 103, 2612–32.

LILLEY, A., M. MAGGIORI, B. NEIMAN, AND J. SCHREGER (2020): "Exchange Rate Reconnect," *The Review of Economics and Statistics*, 1–28.

LLOYD, S. AND E. A. MARIN (2020): "Exchange Rate Risk and Business Cycles," .

LUSTIG, H., A. STATHOPOULOS, AND A. VERDELHAN (2019): "The Term Structure of Currency Carry Trade Risk Premia," *American Economic Review*, 109, 4142–77.

MEESE, R. A. AND K. ROGOFF (1983): "Empirical exchange rate models of the seventies: Do they fit out of sample?" *Journal of International Economics*, 14, 3–24.

MIYAMOTO, W., T. L. NGUYEN, AND H. OH (2023): "In Search of Dominant Drivers of the Real Exchange Rate," *The Review of Economics and Statistics*, 1–50.

NAM, D. AND J. WANG (2015): "The effects of surprise and anticipated technology changes on international relative prices and trade," *Journal of International Economics*, 97, 162–177.

OBSTFELD, M. AND K. ROGOFF (2000): "The six major puzzles in international macroeconomics: is there a common cause?" *NBER Macroeconomics Annual*, 15, 339–390.

ROGOFF, K. (1996): "The purchasing power parity puzzle," *Journal of Economic literature*, 34, 647–668.

ROGOFF, K. S. AND V. STAVRAKEVA (2008): "The continuing puzzle of short horizon exchange rate forecasting," *Working Paper*.

SARNO, L. (2005): "Towards a solution to the puzzles in exchange rate economics: Where do we stand?" *Canadian Journal of Economics*, 38, 673–708.

SARNO, L. AND M. SCHMELING (2014): "Which fundamentals drive exchange rates? A cross-sectional perspective," *Journal of Money, Credit and Banking*, 46, 267–292.

SCHMITT-GROHÉ, S. AND M. URIBE (2003): "Closing Small Open Economy Models," *Journal of International Economics*, 61, 163 – 185.

SIENA, D. (2015): "News Shocks and International Business Cycles," *Working Paper*.

SIMS, C. A. (1972): "Money, income, and causality," *American Economic Review*, 62, 540–552.

SIMS, C. A., J. H. STOCK, AND M. W. WATSON (1990): "Inference in linear time series models with some unit roots," *Econometrica*, 113–144.

STAVRAKEVA, V. AND J. TANG (2020): "A Fundamental Connection: Exchange Rates and Macroeconomic Expectations," *Working Paper*.

Steinsson, J. (2008): "The dynamic behavior of the real exchange rate in sticky price models," *American Economic Review*, 98, 519–33.

Uhlig, H. (2003): "What drives GNP?" *Working paper*.

Valchev, R. (2020): "Bond convenience yields and exchange rate dynamics," *American Economic Journal: Macroeconomics*, 12, 124–66.

Verdelhan, A. (2010): "A Habit-Based Explanation of the Exchange Rate Risk Premium," *The Journal of Finance*, 65, 123–146.

A Data Appendix

A.1 Data sources

In the next lines, we describe the data sources used in the paper.

- Nominal exchange rate
  - Daily bilateral exchange rates, Foreign Currency/USD;
  - Source: *Datastream*;
  - Quarterly aggregation: period-average.
- Nominal interest rates
  - Daily Eurodollar deposit rates;
  - Source: *Datastream*;
  - Quarterly aggregation: period-average.
- Consumer Price Indexes
  - CPI Index (Chained 2010)
  - Source: *OECD*, https://stats.oecd.org/index.aspx?DataSetCode=PRICES_CPI.
- Consumption
  - Real consumption;
  - Source: *OECD*, Private final consumption expenditure
- Investment
  - Real Investment;
  - Source: *OECD*, Gross Fixed Capital Formation (`GFCF`), Quarterly growth rates, https://data.oecd.org/gdp/investment-gfcf.htm.
- U.S. TFP:

- U.S. utilization-adjusted TFP as constructed in Fernald (2012);
- Source: John Fernald's website, https://www.johnfernald.net/TFP (latest available vintage, downloaded on January 2, 2022);

- U.S. R&D:
  - Real R&D expenditure
  - Source: *U.S. Bureau of Economic Analysis*, retrieved from *FRED*, https://fred.stlouisfed.org/ser

- U.S. trade balance (% of GDP)
  - Shares of gross domestic product: Net exports of goods and services
  - Source: *U.S. Bureau of Economic Analysis*, retrieved from *FRED*, https://fred.stlouisfed.org/ser

- Equity prices and equity returns
  - MSCI price indexes and total returns indexes
  - Source: retrieved from *Datastream*

- Long-term bond yields
  - Interest Rates: Long-Term Government Bond Yields: 10-Year
  - Source: *Global Financial Data*

B Additional evidence and alternative empirical specifications

Table B.1: Share of forecast error variance explained by the FX shock: target 6-32 quarter periodicities

	Forecast Horizon (Quarter)					
	$Q1$	$Q4$	$Q12$	$Q24$	$Q40$	$Q100$
Home TFP	0.02	0.03	0.13	0.18	0.21	0.21
Home Consumption	0.07	0.10	0.22	0.35	0.32	0.24
Foreign Consumption	0.01	0.01	0.05	0.22	0.27	0.18
Home Investment	0.11	0.17	0.15	0.19	0.20	0.21
Foreign Investment	0.01	0.02	0.05	0.13	0.22	0.18
Interest Rate Differential	0.28	0.22	0.21	0.24	0.24	0.23
Real Exchange Rate	0.89	0.93	0.73	0.58	0.55	0.50
Expected Excess Returns	0.32	0.20	0.41	0.43	0.42	0.41

Notes: The table reports the estimated variance shares at different horizons accounted for by the exchange rate shock that explains most of the forecast error variance of the exchange rate over the [6,32] quarter periodicity range.

Figure B.1: Technology, Noise and the Trade Balance

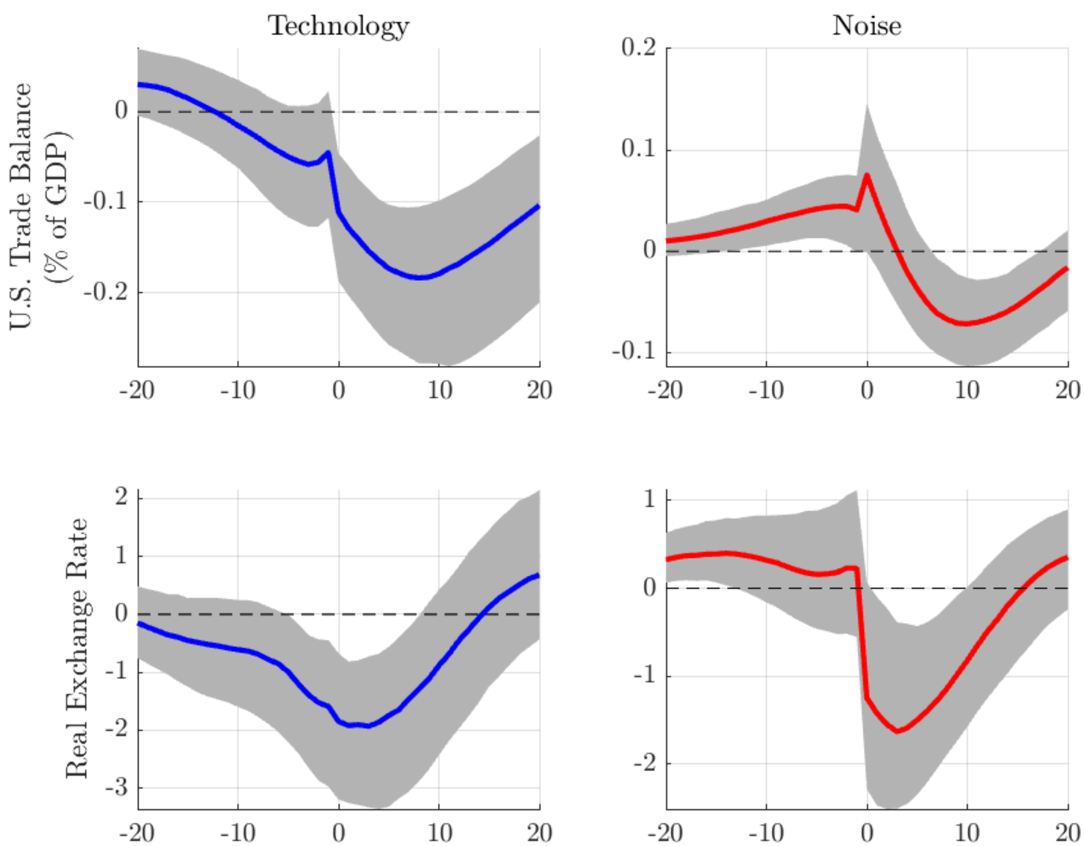

Notes: The figure displays the responses a one standard deviation impulse in the technological disturbance (left column) and the expectational disturbance (right column) at time $t = 0$. The shaded areas are 16-84th percentile bands. Each period is a quarter.

Figure B.2: Impulse responses to Technology disturbances (Alternative specifications)

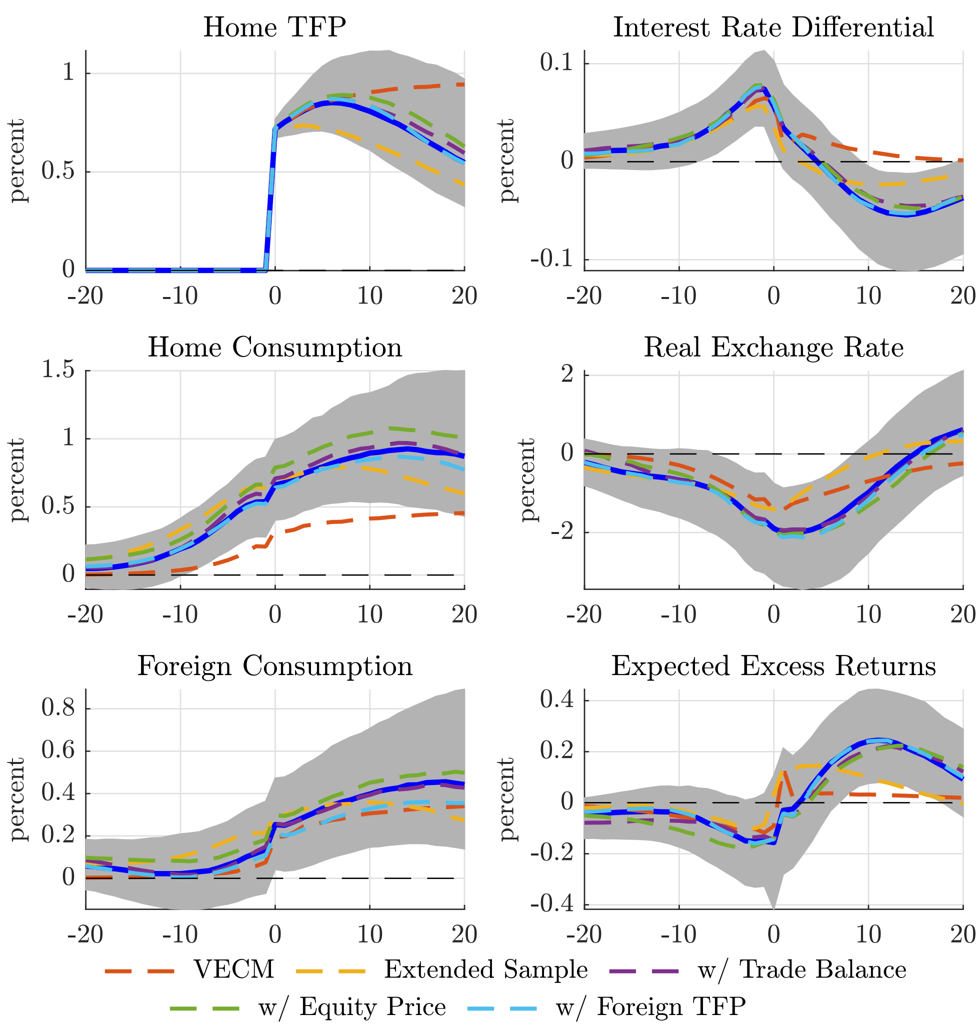

Notes: The figure displays IRFs to a one standard deviation impulse in the technological disturbance at time $t = 0$ for the baseline sample (blue lines), where the shaded areas are 16-84th percentile bands. In addition the figure displays the point estimate of the responses obtained from the VAR applied to the extended sample (1978:Q4-2018:Q4) and from a VECM where the real exchange rate and interest rate differential are assumed stationary. The figure also displays the point estimate of the responses obtained from the baseline VAR with one additional variable. Each period is a quarter.

Figure B.3: Impulse responses to Noise disturbance (Alternative specifications)

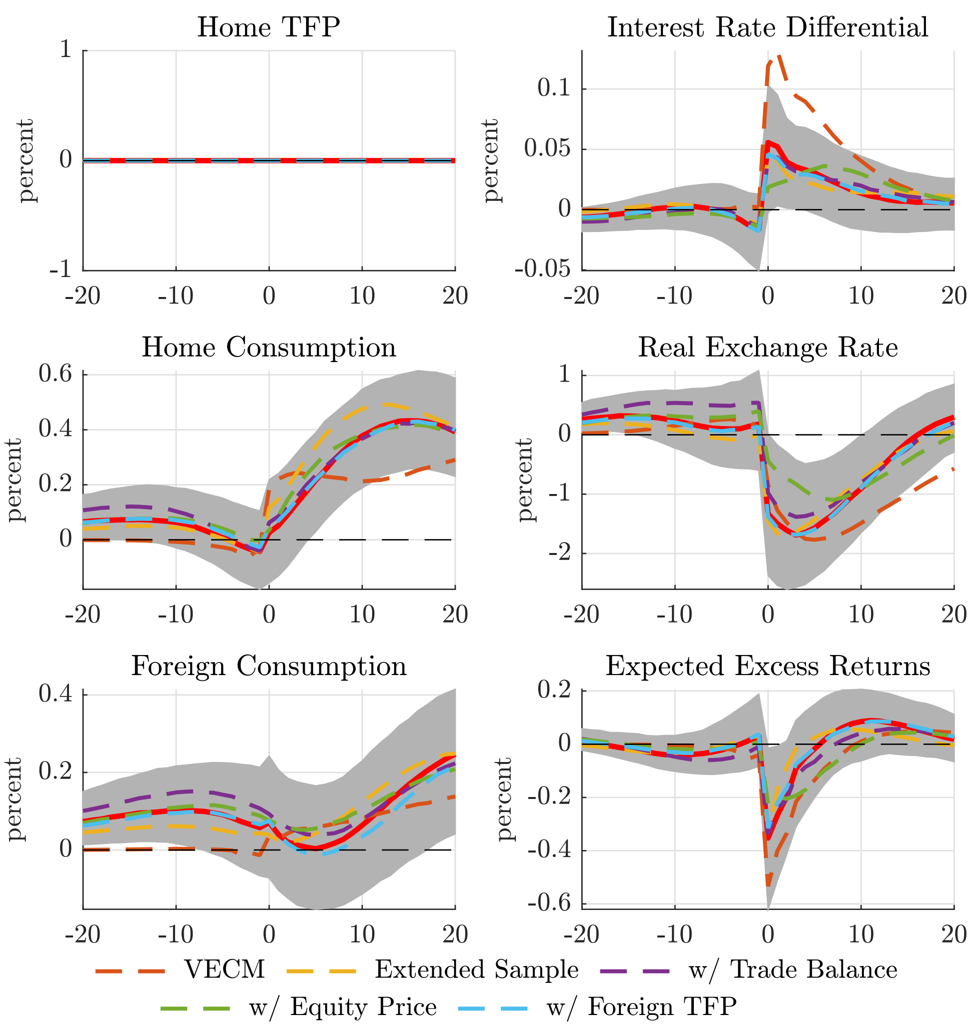

Notes: The figure displays IRFs to a one standard deviation impulse in the expectational disturbance at time $t = 0$ for the baseline sample (blue lines), where the shaded areas are 16-84th percentile bands. In addition the figure displays the point estimate of the responses obtained from the VAR applied to the extended sample (1978:Q4-2018:Q4) and from a VECM where the real exchange rate and interest rate differential are assumed stationary. The figure also displays the point estimate of the responses obtained from the baseline VAR with one additional variable. Each period is a quarter.

Table B.2: Share of variance explained by the Main FX shock (Alternative specifications)

Panel A: Extended sample (1978:Q4-2018:Q4)

	Forecast Horizon (Quarter)					
	$Q1$	$Q4$	$Q12$	$Q24$	$Q40$	$Q100$
Home TFP	0.03	0.03	0.07	0.18	0.28	0.33
Home Consumption	0.04	0.07	0.26	0.45	0.48	0.43
Foreign Consumption	0.01	0.02	0.04	0.18	0.30	0.31
Home Investment	0.20	0.25	0.32	0.38	0.39	0.38
Foreign Investment	0.03	0.03	0.05	0.12	0.22	0.24
Interest Rate Differential	0.41	0.39	0.31	0.30	0.30	0.31
Real Exchange Rate	0.59	0.74	0.83	0.74	0.70	0.67
Expected Excess Returns	0.56	0.29	0.30	0.36	0.36	0.37

Panel B: Vector Error Correction Model

	Forecast Horizon (Quarter)					
	$Q1$	$Q4$	$Q12$	$Q24$	$Q40$	$Q100$
Home TFP	0.01	0.01	0.05	0.10	0.13	0.15
Home Consumption	0.09	0.13	0.25	0.41	0.50	0.58
Foreign Consumption	0.02	0.02	0.04	0.10	0.16	0.20
Home Investment	0.09	0.15	0.14	0.14	0.15	0.18
Foreign Investment	0.02	0.02	0.03	0.06	0.07	0.09
Interest Rate Differential	0.49	0.52	0.45	0.44	0.44	0.45
Real Exchange Rate	0.72	0.87	0.92	0.90	0.90	0.90
Expected Excess Returns	0.55	0.40	0.39	0.43	0.44	0.44

Panel C: Individual countries (Median)

	Forecast Horizon (Quarter)					
	$Q1$	$Q4$	$Q12$	$Q24$	$Q40$	$Q100$
Home TFP	0.03	0.06	0.11	0.26	0.44	0.45
Home Consumption	0.03	0.05	0.16	0.34	0.36	0.40
Foreign Consumption	0.04	0.04	0.15	0.20	0.38	0.38
Home Investment	0.02	0.04	0.20	0.34	0.34	0.36
Foreign Investment	0.05	0.05	0.10	0.19	0.32	0.37
Interest Rate Differential	0.08	0.07	0.11	0.22	0.26	0.32
Real Exchange Rate	0.27	0.42	0.68	0.80	0.76	0.72
Expected Excess Returns	0.21	0.18	0.29	0.32	0.39	0.43

Notes: The table reports the estimated variance shares accounted for by the main exchange rate shock at different horizons.

Table B.3: Variance decomposition (Alternative specifications)

Panel A: Extended sample 1978:Q4-2018:Q4

	Periodicities of 2-100 Quarters			Periodicities of 6-32 Quarters		
	Both	Tech.	Noise	Both	Tech.	Noise
Home TFP	1.00	1.00	0.00	1.00	1.00	0.00
Home Consumption	0.67	0.42	0.25	0.33	0.07	0.26
Foreign Consumption	0.47	0.32	0.15	0.18	0.07	0.12
Home Investment	0.56	0.34	0.21	0.27	0.09	0.19
Foreign Investment	0.50	0.25	0.25	0.35	0.08	0.27
Interest Rate Differential	0.41	0.28	0.13	0.30	0.18	0.11
Real Exchange Rate	0.48	0.26	0.22	0.29	0.08	0.20
Expected Excess Returns	0.38	0.22	0.16	0.32	0.16	0.16

Panel B: Vector Error Correction Model

	Periodicities of 2-100 Quarters			Periodicities of 6-32 Quarters		
	Both	Tech.	Noise	Both	Tech.	Noise
Home TFP	1.00	1.00	0.00	1.00	1.00	0.00
Home Consumption	0.28	0.12	0.16	0.25	0.05	0.20
Foreign Consumption	0.24	0.16	0.09	0.15	0.08	0.07
Home Investment	0.29	0.08	0.22	0.40	0.06	0.34
Foreign Investment	0.22	0.08	0.14	0.16	0.04	0.12
Interest Rate Differential	0.66	0.15	0.51	0.62	0.09	0.53
Real Exchange Rate	0.49	0.14	0.35	0.33	0.06	0.27
Expected Excess Returns	0.49	0.13	0.36	0.47	0.09	0.38

Panel C: Individual countries (Median)

	Periodicities of 2-100 Quarters			Periodicities of 6-32 Quarters		
	Both	Tech.	Noise	Both	Tech.	Noise
Home TFP	1.00	1.00	0.00	1.00	1.00	0.00
Home Consumption	0.69	0.515	0.17	0.34	0.095	0.24
Foreign Consumption	0.54	0.41	0.12	0.29	0.12	0.145
Home Investment	0.62	0.455	0.16	0.375	0.2	0.18
Foreign Investment	0.56	0.43	0.145	0.355	0.105	0.205
Interest Rate Differential	0.395	0.27	0.125	0.245	0.11	0.13
Real Exchange Rate	0.59	0.395	0.195	0.37	0.11	0.245
Expected Excess Returns	0.415	0.25	0.145	0.31	0.13	0.18

Notes: The table reports the estimated variance shares (at periodicities between 2 and 100 quarters and between 6 and 32 quarters) explained by technological disturbances ("Tech."), expectational disturbances ("Noise"), and the combination of both.

Figure B.4: Impulse responses to Technology disturbances (Individual countries)

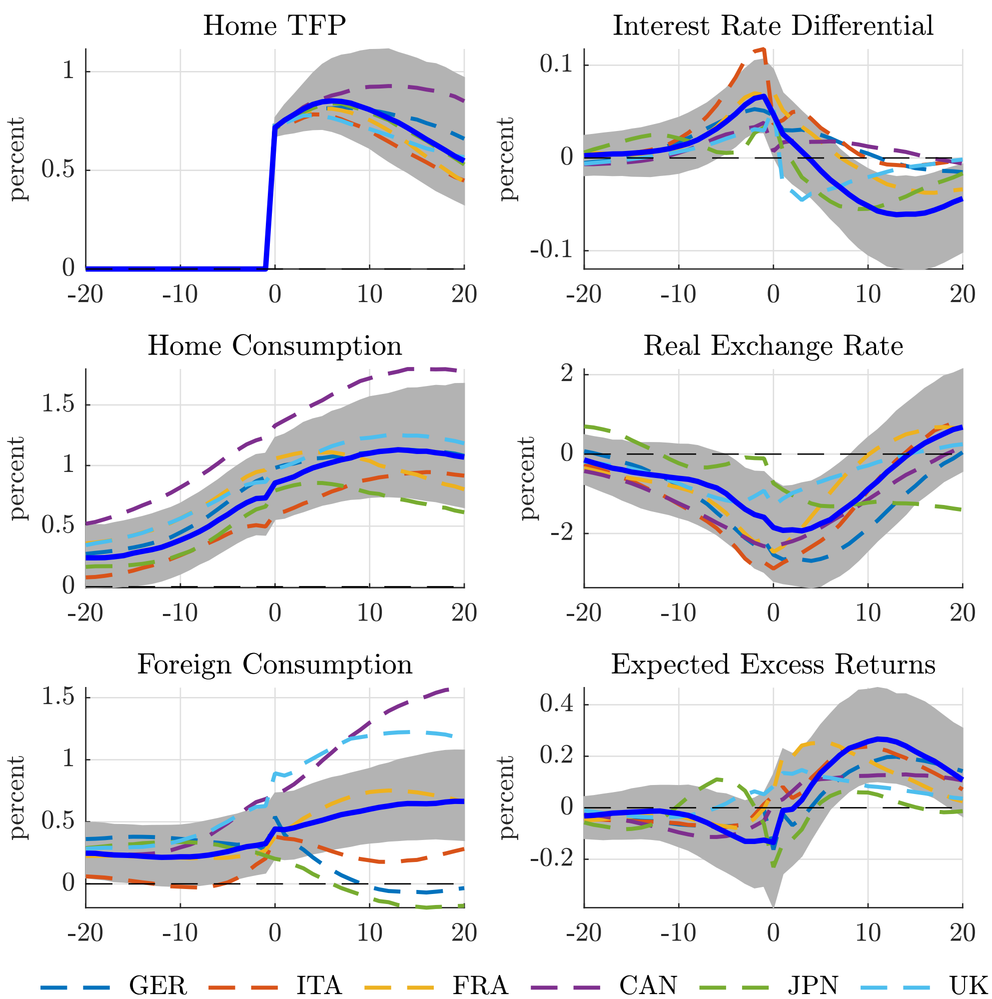

Notes: The figure displays IRFs to a one standard deviation impulse in the technological disturbance at time $t = 0$ for the baseline sample (blue lines), where the shaded areas are 16-84th percentile bands. In addition the figure displays the point estimate of the responses obtained from six different bilateral VARs. Each period is a quarter.

Figure B.5: Impulse responses to Noise disturbance (Individual countries)

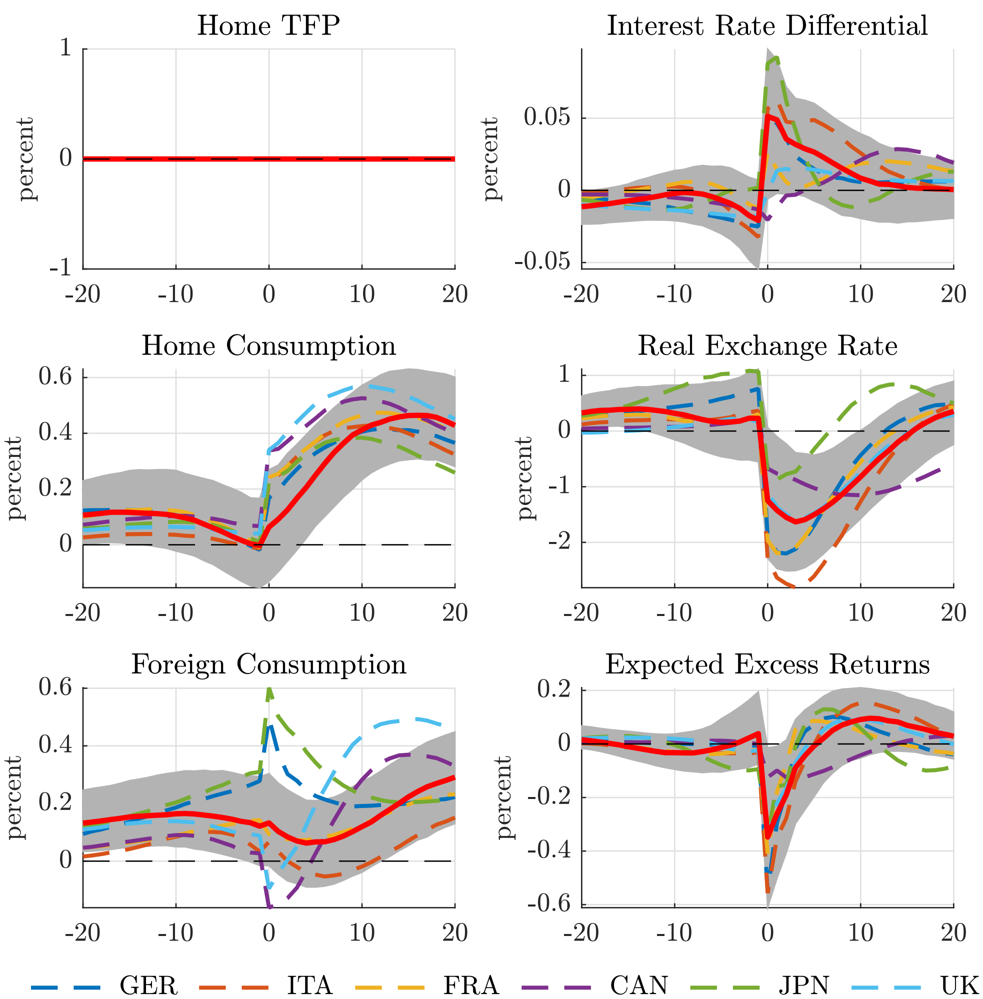

Notes: The figure displays IRFs to a one standard deviation impulse in the expectational disturbance at time $t = 0$ for the baseline sample (blue lines), where the shaded areas are 16-84th percentile bands. In addition the figure displays the point estimate of the responses obtained from six different bilateral VARs. Each period is a quarter.

Figure B.6: Series of estimated technology and noise disturbances

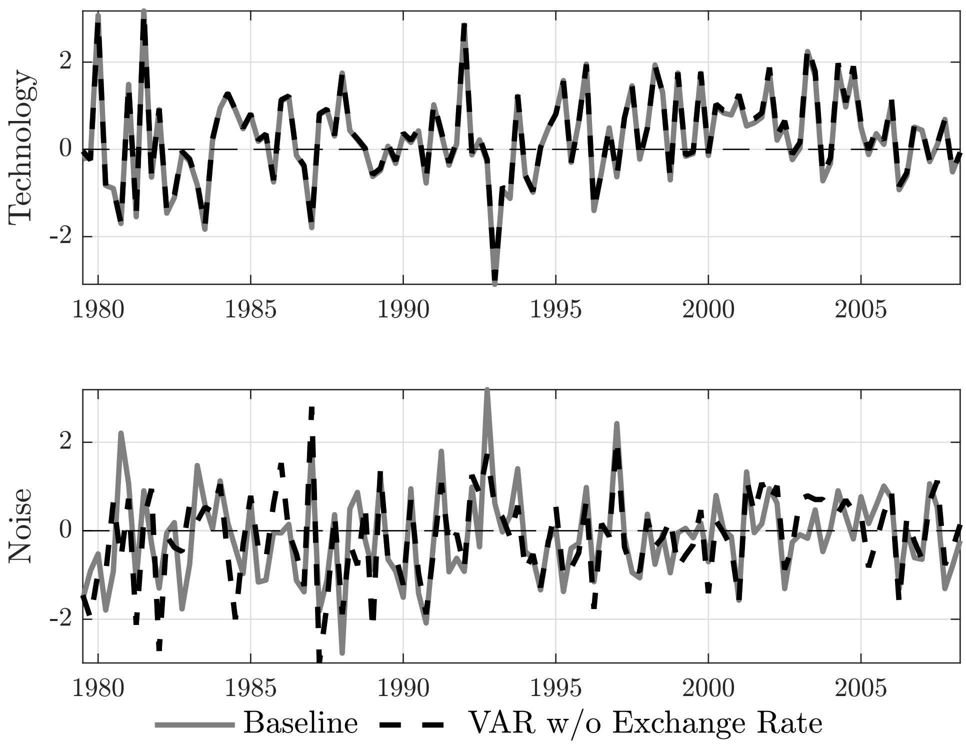

Notes: The figure displays the series of technology and noise disturbances estimated by the baseline VAR as well as the VAR without the exchange rate.

Figure B.7: The exchange rate response to technology and noise disturbances

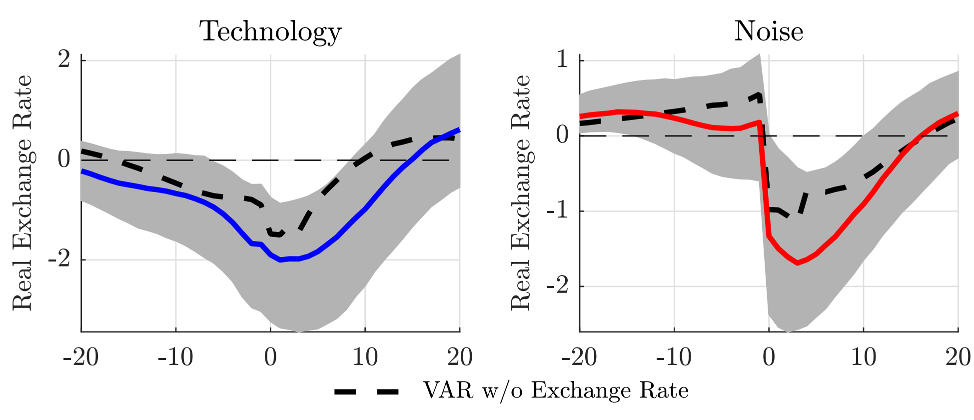

Notes: The figure displays the IRF of the real exchange rate a one standard deviation impulse in the technological disturbance (left column) and the expectational disturbance (right column) at time $t = 0$. The dashed black line is the exchange rate response when the disturbances are estimated using a VAR without the exchange rate. The shaded area is the the 16-84th. Each period is a quarter.

Table B.4: Granger causality: exchange rates and TFP

Panel A: P-values testing whether exchange rates Granger-cause TFP

h	Canada	France	Germany	Italy	Japan	UK	G6
1	0.92	0.78	0.61	0.97	0.65	0.64	0.95
2	0.11	0.86	0.57	0.67	0.80	0.80	0.99
3	0.20	0.04**	0.09*	0.00***	0.17	0.00***	0.09*
4	0.17	0.05**	0.10*	0.02**	0.09*	0.04**	0.21
5	0.11	0.00***	0.00***	0.03**	0.00***	0.04**	0.00***

Panel B: P-values testing whether TFP Granger-cause exchange rates

h	Canada	France	Germany	Italy	Japan	UK	G6
1	0.67	0.76	0.82	0.89	0.99	0.96	0.90
2	0.39	0.80	0.76	0.63	0.99	0.20	0.75
3	0.39	1.00	1.00	0.85	0.81	0.75	0.95
4	0.07*	0.62	0.71	0.44	0.82	0.83	0.44
5	0.00***	0.83	0.88	0.49	0.70	0.97	0.76

Notes: P-values from a Wald test of Granger causality as in Sims (1972). The test is based on the equation: $x_t = \alpha + \beta_0 y_t + \sum_{k=1}^{h} \beta_{-k}^{\text{lag}} y_{t-k} + \sum_{k=1}^{h} \beta_k^{\text{lead}} y_{t+k} + \varepsilon_t$, where x is the quarterly observation of the annual change in the log of the real exchange rate and y is the quarterly observation of the annual change of the log of TFP in Panel A; variables are reversed in Panel B. The null hypothesis is that $\beta_k^{\text{lead}} = 0$ for all $k = 1, \ldots, h$. The sample period is 1976:Q1–2008:Q2. Newey-West standard errors are computed using a data-driven bandwidth, a Bartlett kernel, and prewhitening.

Figure B.8: Bivariate Granger Causality Tests, Different Macro Aggregates (Sample: 1974:Q1-2018:Q4)

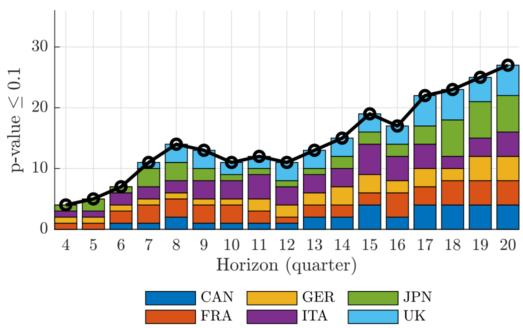

(a) P-values of Granger Causality tests, by country

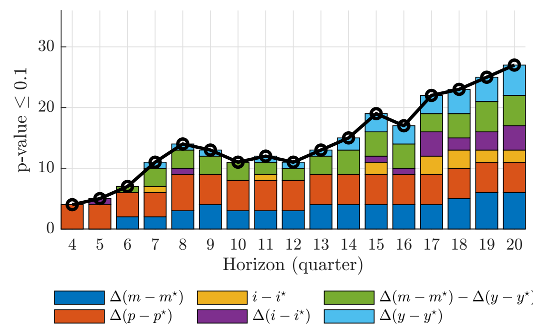

(b) P-values of Granger Causality tests, by variable

Notes: Number of p-values below 0.1 from a Wald test of Granger causality using the procedure of Engel and West (2005) and extending their data through 2018:Q4. We compute Newey-West standard errors using a data-driven bandwidth, a Bartlett kernel, and prewhitening. Each test is based on a bivariate VAR estimated using the quarterly change in the nominal exchange rate, Δs, and a given fundamental, Δf. The horizon on the x-axis denotes the number of lags included in the VAR. The p-value corresponds to the Wald test of the joint significance of the lags of Δs in predicting Δf. Specifically, Δm is the percentage change in M1 (M2 for the United Kingdom); Δy is the percentage change in real GDP; Δp is the percentage change in consumer prices; and i is the short-term interest rate on government debt.

C Interpreting the Granger Causality Evidence

This appendix section presents a simple exchange rate model that helps interpret the Granger-causality evidence reported in Section B.5.

Under UIP and stationary real exchange rate, the equilibrium real exchange rate follows, to a first-order:

$$q_t = -\sum_{k=0}^{\infty} \mathbb{E}_t(\tilde{r}_{t+k}) \tag{C.1}$$

where $\tilde{r}_{t+k} \equiv r_{t+k} - r^{\star}_{t+k}$ denotes the home-foreign real interest rate differential. An equation similar to (C.1) motivated Engel and West (2005) to examine whether exchange rate changes Granger cause interest differentials (or changes in interest differentials). The underlying logic is that if agents have advance information about future innovations in interest differentials and this information is embedded in exchange rates, then exchange rate changes should forecast interest differentials *above and beyond* their own lags.

However, whether exchange rates contain independent information about future interest differentials ultimately depends on the joint *general equilibrium* process of q_t and $\tilde{r}_t$. Consider, for instance, the interest differential in a standard incomplete-market model with CRRA preferences over consumption, to a first-order:

$$r_{t+k} - r^{\star}_{t+k} = \sigma \left[(\mathbb{E}_t \Delta \tilde{c}_{t+k+1}) - (\mathbb{E}_t \Delta \tilde{c}_{t+k}) \right] \tag{C.2}$$

Consider case of financial autarky, where the consumption differential $\tilde{c}_t$ equals the exogenous endowment (or TFP) differential in every period, i.e., $\tilde{c}_t = \tilde{a}_t$ for all t. Let the TFP differential follow the process:

$$\tilde{a}_t = \varepsilon^a_t + \varepsilon^n_{t-1}$$

where ε^a_t and ε^n_{t-1} are observed at times t and $t-1$, respectively. Here, ε^n_t represents a "news shock" to TFP. That is, at time t, agents partially observe the level of next-period's TFP differential.

In this setting, one can show that:

$$\mathbb{E}_t(\tilde{a}_{t+k} \mid \tilde{q}_t, \tilde{a}_t) = \begin{cases} \varepsilon^n_t & \text{for } k = 1 \\ 0 & \text{for } k \geq 2 \end{cases} \qquad \mathbb{E}_t(\tilde{a}_{t+k} \mid \tilde{a}_t) = \begin{cases} 0 & \text{for } k = 1 \\ 0 & \text{for } k \geq 2 \end{cases}$$

The exchange rate contains information about future TFP differentials that is not already contained in current TFP levels. This implies that the exchange rate Granger causes the TFP differential (at horizon 1, in this simple example).

In this setting, one can also show that:

$$\mathbb{E}_t(\tilde{r}_{t+k} \mid q_t, \tilde{r}_t) = \begin{cases} -\sigma\varepsilon_t^n & \text{for } k = 1 \\ 0 & \text{for } k \geq 2 \end{cases} \qquad \mathbb{E}_t(\tilde{r}_{t+k} \mid \tilde{r}_t) = \begin{cases} -\sigma\varepsilon_t^n & \text{for } k = 1 \\ 0 & \text{for } k \geq 2 \end{cases}$$

In this case, the exchange rate does *not* contain information about future interest differentials that is independent of current interest differentials. Hence, exchange rate changes do *not* Granger cause interest differentials, despite the presence of advance information about future TFP.

The intuition is straightforward: interest differentials are themselves forward-looking, like exchange rates, and thus already incorporate similar future information. This is consistent with our empirical finding that interest differentials comove contemporaneously with exchange rates (see, e.g., Figure 1). More broadly, the exchange rate-fundamentals relationship should be stronger for slow-moving fundamentals, such as TFP or output differentials.

D Model Appendix

D.1 Model Environments

All three models we consider are two-country endowment economies with unitary trade elasticity. The formulation of recursive preferences and solution approach follows closely Colacito et al. (2018). When the intertemporal elasticity of substitution equals the inverse of the risk-aversion parameter, the preferences specialize to the CRRA case. Under incomplete markets agents only trade one bond internationally.

Utility function Let C_t denote consumption of the home country. The home-country preferences over an infinite horizon are

$$U_t = \left[(1-\beta)\cdot C_t^{1-1/\psi} + \beta\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{1-1/\psi}{1-\gamma}}\right]^{\frac{1}{1-1/\psi}}. \tag{D.3}$$

With these preferences, agents are risk averse in future utility as well as future consumption. The extent of such utility risk aversion depends on the preference for early resolution of uncertainty, measured by $\gamma - 1/\psi > 0$.

When $\gamma = 1/\psi$, the agent is utility-risk neutral and preferences collapse to the standard time-additive case. When the agent prefers early resolution of uncertainty, that is, $\gamma > 1/\psi$, uncertainty about continuation utility reduces welfare and generates an incentive to trade off future expected utility, $\mathbb{E}_t[V_{t+1}]$, for future utility risk, $\text{var}_t[V_{t+1}]$. This trade-off drives international consumption flows, and it represents one of the most important elements of Colacito and Croce (2013) analysis.

Furthermore, this trade-off is also present in the home stochastic discount factor,

$$M_{t+1} = \beta \left(\frac{\tilde{C}_{t+1}}{\tilde{C}_t} \right)^{-1/\psi} \left(\frac{U_{t+1}}{\mathbb{E}_t \left[U_{t+1}^{1-\gamma} \right]^{\frac{1}{1-\gamma}}} \right)^{1-\gamma}, \tag{8}$$

as the second term captures aversion to continuation utility risk and is extremely sensitive to growth news (Bansal and Yaron, 2004).[19]

The utility of foreign households is symmetric.

Technology Let $\{X_t, Y_t\}$ and $\{X_t^\star, Y_t^\star\}$ denote the time t consumption of goods X and Y in the home and foreign countries, respectively. The consumption aggregates in the two countries are

$$C_t = X_t^{\lambda} \cdot Y_t^{(1-\lambda)}, \quad C_t^\star = X_t^{*(1-\lambda)} \cdot Y_t^{*\lambda}. \tag{D.4}$$

The home (foreign) country produces good X (Y) so that $\lambda > 1/2$ introduces consumption home bias.

In this endowment economy, the following resource constraints apply:

$$A_t \geq X_t + X_t^\star \qquad A_t^\star \geq Y_t + Y_t^\star$$

[19] When $\psi = 1$, preferences take the following form:

$$V_t = (1 - \beta) \cdot \log C_t + \frac{\beta}{1-\gamma} \log \mathbb{E}_t \left[\exp \left\{ V_{t+1} \cdot (1-\gamma) \right\} \right]. \tag{9}$$

When $\psi = 1$, if $\log V_t$ is normally distributed, equation (7) has the following exact counterpart:

$$V_t = (1 - \beta) \cdot \log C_t + \beta \mathbb{E}_t[V_{t+1}] - \frac{(\gamma - 1)}{2} \beta \text{var}_t[V_{t+1}].$$

where A_t and $A_t^\star$ are the exogenous stochastic productivity processes described in the main text.

D.1.1 Complete Markets

To find the solution under complete markets, we follow Colacito et al. (2018) and solve the Pareto problem. We solve the Pareto problem under the general utility function (D.3), which specializes to the CRRA case for $\gamma = 1/\psi$.

Pareto Problem This appendix suppresses notation denoting state and histories and retain only subscripts for time. We represent the Epstein and Zin (1989) utility preference in the following compact way:

$$U_t = W(C_t, U_{t+1}),$$

so that the dependence of current utility on j-step-ahead consumption can easily be denoted as follows:

$$\frac{\partial U_t}{\partial C_{t+j}} = W_{2,t+1} \cdot W_{2,t+2} \cdots W_{2,t+j} W_{1,t+j}, \tag{D.5}$$

where $W_{2,t+j} \equiv \frac{\partial U_{t+j-1}}{\partial U_{t+j}}$ and $W_{1,t+j} \equiv \frac{\partial U_{t+j}}{\partial C_{t+j}}$. Given this notation, the intertemporal marginal rate of substitution between C_t and C_{t+1} is:

$$\text{IMRS}_{C,t+1} = \frac{W_{2,t+1} W_{1,t+1}}{W_{1,t}} = M_{t+1}\pi_{t+1}, \tag{D.6}$$

where π_{t+1} is the probability of a specific state, and M_{t+1} is the stochastic discount factor in C units with the following form:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\psi}} \left(\frac{U_{t+1}}{\mathbb{E}_t[U_{t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}}}\right)^{\frac{1}{\psi}-\gamma}. \tag{D.7}$$

The consumption aggregate combines two goods, X and Y:

$$C_t = C(X_t, Y_t).$$

The planner faces the following constraints:

$$A_t \geq X_t + X_t^\star \tag{D.8}$$

$$A_t^\star \geq Y_t + Y_t^\star \tag{D.9}$$

where A_t and $A_t^\star$ are the exogenous stochastic productivity processes.

The social planner chooses sequences of $\{X_t, X_t^\star, Y_t, Y_t^\star\}_t$ to maximize:

$$\mu_0 W_0 + (1-\mu_0) W_0^\star,$$

subject to (D.8)-(D.9). Let $\lambda_{i,t}$ be the Lagrangian multiplier for the respective constraint i. The Lagrangian is:

$$\Omega = \mu_0 W_0 + (1-\mu_0) W_0^\star + \lambda_{1,t}(A_t - X_t - X_t^\star) + \lambda_{2,t}(A_t^\star - Y_t - Y_t^\star) + \ldots$$

The optimality condition for good X_t for all t is:

$$\mu_0 \left(\prod_{j=1}^{t} W_{2,j} \right) W_{1,t} C_{x,t} = \lambda_{1,t} = C^\star_{x^\star,t} W^\star_{1,t} \left(\prod_{j=1}^{t} W^\star_{2,j} \right) (1-\mu_0). \tag{D.10}$$

Let μ_t be the time-t Pareto weight for the home country:

$$\mu_t = \mu_0 \left(\prod_{j=1}^{t} W_{2,j} \right) W_{1,t} C_t = \mu_{t-1} W_{2,t} \cdot \frac{W_{1,t}}{W_{1,t-1}} \cdot \frac{C_t}{C_{t-1}} = \mu_{t-1} M_t \cdot \frac{C_t}{C_{t-1}}.$$

Thus, equation (D.10) becomes:

$$S_t C_{x,t} \cdot \frac{1}{C_t} = C^\star_{x^\star,t} \cdot \frac{1}{C_t^\star}.$$

where $S_t = \mu_t / \mu_t^\star$. Similarly:

$$S_t C_{y,t} \cdot \frac{1}{C_t} = C^\star_{y^\star,t} \cdot \frac{1}{C_t^\star}.$$

Real Exchange Rate and Real Interest Rates Under complete markets, the real exchange rate satisfies the Backus and Smith (1993) condition:

$$\frac{Q_{t+1}}{Q_t} = \frac{M^\star_{t+1}}{M_{t+1}}$$

where the home stochastic discount factor, M_{t+1}, is reported in equation (D.7), and the foreign stochastic discount factor has an analogous structure.

The home and foreign risk-free rates are:

$$E_t(M_{t+1})R_t = 1; \qquad E_t(M^\star_{t+1})R^\star_t = 1.$$

D.1.2 Incomplete Markets

The budget constraint for the home economy is:

$$\frac{B_{t+1}}{R_t} + \frac{Q_t B^\star_{t+1}}{R^\star_t \Phi\left(Q_t B^\star_{t+1}\right)} = A_t - X_t - P_t Y_t + B_t + Q_t B^\star_t \tag{D.11}$$

where B_{t+1} and $B^\star_{t+1}$ denote domestic- and foreign-good-denominated bonds. The function $\Phi(\cdot)$ represents the cost from international borrowings and it is increasing in the aggregate level of foreign debt: $\Phi'(\cdot) < 0$. Households take the overall cost of international borrowing as given. We further assume a zero steady-state intermediation cost by setting $\Phi(\bar{B}^\star) = 1$. Foreign households only trade foreign-good-denominated bonds, while domestic-good-denominated bonds are in zero net supply. That is, in reality only foreign-good-denominated bonds are traded in equilibrium. As a result, defining the intermediation costs over the foreign bonds only is sufficient to pin down the overall steady-state net foreign asset position. We consider an economy with zero steady-state net foreign asset position. P_t denote the terms of trade.

The home household maximizes utility (D.3) subject to (D.11). The first order condition with respect to home and foreign bonds implies:

$$E_t(M_{t+1})R_t = 1; \qquad E_t\left(M_{t+1}\frac{Q_{t+1}}{Q_t}\right) R^\star_t \Phi\left(Q_t B^\star_{t+1}\right) = 1. \tag{D.12}$$

The budget constraint for the foreign economy is:

$$-\frac{B^\star_{t+1}}{R^\star_t} = A^\star_t - \frac{X^\star_t}{P_t} - Y_t - B^\star_t + \Pi_t \tag{D.13}$$

where Π_t are the intermediation profits, rebated lump-sum to the foreign household. The foreign households' first-order condition for bonds is:

$$E_t(M^\star_{t+1})R^\star_t = 1 \tag{D.14}$$

Combining (D.12) and (D.14):

$$E_t\left(M_{t+1}\frac{Q_{t+1}}{Q_t}\right)\Phi\left(Q_t B^{\star}_{t+1}\right) = E_t(M^{\star}_{t+1})$$

We further assume that:

$$\Phi\left(Q_t B^{\star}_{t+1}\right) = \phi_0\left[1 + Q_t(B^{\star}_{t+1} - \bar{B}^{\star})\right]^{-\phi_1}$$

We calibrate ϕ_0 and ϕ_1 to guarantee a well-defined net-foreign asset position, but small enough so that UIP approximately holds.

D.2 Blanchard et al. (2013) process

With the Blanchard et al. (2013) process, log productivity, a_t, is the sum of a permanent component μ_t (with autocorrelated growth rates) and a transitory component η_t, but agents only observe the realizations of productivity a_t and not the two components separately. In addition, agents observe a noisy signal of the current value of the persistent component, $s_t = \mu_t + \xi_t$, which helps them forecast future TFP growth. Specifically, all stochastic processes $\{a_t\}$ and $\{s_t\}$ evolve according to the following system:

$$\begin{aligned} a_t &= \mu_t + \eta_t, \\ s_t &= \mu_t + \xi_t, \\ \Delta\mu_t &= \rho\Delta\mu_{t-1} + \varepsilon^{\mu}_t, \\ \eta_t &= \rho\eta_{t-1} + \varepsilon^{\eta}_t, \end{aligned} \qquad \begin{bmatrix} \varepsilon^{\mu}_t \\ \varepsilon^{\eta}_t \\ \xi_t \end{bmatrix} \sim \text{iid}\,\mathcal{N}\left(0, \begin{bmatrix} \sigma^2_{\mu} & 0 & 0 \\ 0 & \sigma^2_{\eta} & 0 \\ 0 & 0 & \sigma^2_{\xi} \end{bmatrix}\right). \tag{D.15}$$

Lastly, it is common to impose the parameter restriction $\rho\sigma^2_{\mu} = (1-\rho)^2\sigma^2_{\eta}$, which ensures that the univariate Wold representation of TFP is a random walk as is true in the data.[20]

To properly isolate the independent contributions of beliefs (or the independent effects of technology and noise disturbances), we follow Chahrour and Jurado (2018) to construct a noise representation that is observationally equivalent to representation (D.15).

Lemma 1. *The representation of fundamentals and beliefs in system (D.15) is observation-*

[20] Blanchard et al. (2013) refer to ε^{μ}_t as a permanent productivity shock, ε^{η}_t as a transitory productivity shock, and ξ_t as an information noise shock.

ally equivalent to the noise representation:

$$\Delta a_t = \frac{1-\alpha L}{1-\rho L}\varepsilon_t^m$$
$$\Delta \tilde{s}_t = \sum_{j=0}^{\infty} \alpha^j \varepsilon_{t+j}^m + \tilde{v}_t,$$
$$\tilde{v}_t = \frac{(1-L)(1-\bar{\delta}_1 L)(1-\bar{\delta}_2 L)}{(1-\alpha L)(1-\rho L)}\varepsilon_t^v,$$
$$\begin{bmatrix} \varepsilon_t^m \\ \varepsilon_t^v \end{bmatrix} \overset{iid}{\sim} \mathcal{N}\left(0, \begin{bmatrix} \sigma_\eta^2/\alpha & 0 \\ 0 & \frac{\alpha\sigma_\xi^2}{\bar{\delta}_1\bar{\delta}_2} \end{bmatrix}\right),$$

where $\alpha = 1-\frac{\sigma_\mu}{\sigma_\eta}$*, and* $\bar{\delta}_1, \bar{\delta}_2$ *are the two stable (i.e., inside the unit circle) roots of the quartic polynomial:*

$$D(L) = (1-\alpha L)(1-\alpha L^{-1})(1-\rho L)(1-\rho L^{-1}) - c(1-L)(1-L^{-1})$$

with $c = \frac{\sigma_\eta^2}{\sigma_\mu^2}\alpha$.

Proof. First, let's derive the univariate Wold representation of Δa_t. The process can be written as:

$$\Delta a_t = \frac{1}{1-\rho L}\left(\varepsilon_t^\mu + (1-L)\varepsilon_t^\eta\right)$$

The autocovariance generating function of Δa is:

$$\Gamma_{\Delta a}(L) = \frac{\sigma_\mu^2 + \sigma_\eta^2(2-L-L^{-1})}{(1-\rho L)(1-\rho L^{-1})}.$$

After factorizing the numerator, we obtain:

$$\Gamma_{\Delta a}(L) = \frac{\sigma_\eta^2}{\alpha}\cdot\frac{(1-\alpha L)(1-\alpha L^{-1})}{(1-\rho L)(1-\rho L^{-1})}, \quad \text{with } \alpha = 1-\frac{\sigma_\mu}{\sigma_\eta}.$$

Taking the square root gives the Wold representation:

$$\Delta a_t = \frac{1-\alpha L}{1-\rho L}\varepsilon_t^m, \quad \varepsilon_t^m \overset{iid}{\sim} \mathcal{N}(0, \sigma_\eta^2/\alpha). \tag{D.16}$$

Next, write the signal as:

$$\Delta s_t = \frac{1}{1-\rho L}\varepsilon_t^\mu + (1-L)\xi_t$$

We seek to represent the signal as a function of ε_t^μ and a noise component, v_t, orthogonal to all leads and lags of ε_t^μ, that is,

$$\Delta s_t = B(L)\Delta a_t + v_t.$$

We project Δs_t onto Δa_t:

$$B(L) = \frac{\Gamma_{\Delta a,\Delta s}(L)}{\Gamma_{\Delta a}(L)} = \frac{\sigma_\mu^2}{\Gamma_{\Delta a}(L)} = \frac{\alpha\sigma_\mu^2}{\sigma_\eta^2} \cdot \frac{1}{(1-\alpha L)(1-\alpha L^{-1})}$$

substituting the Wold form of Δa_t we get:

$$B(L)\Delta a_t = \frac{\alpha\sigma_\mu^2}{\sigma_\eta^2(1-\alpha L^{-1})(1-\rho L)}\varepsilon_t^m.$$

Therefore:

$$\Delta s_t = \frac{\alpha\sigma_\mu^2}{\sigma_\eta^2(1-\alpha L^{-1})(1-\rho L)}\varepsilon_t^m + v_t.$$

Scaling both sides by $\frac{\sigma_\eta^2(1-\rho L)}{\alpha\sigma_\mu^2}$ yields:

$$\Delta\tilde{s}_t := \frac{\sigma_\eta^2(1-\rho L)}{\alpha\sigma_\mu^2}\Delta s_t = \sum_{j=0}^{\infty}\alpha^j\varepsilon_{t+j}^m + \tilde{v}_t.$$

To characterize $\tilde{v}_t$, compute:

$$\Gamma_{\tilde{v}}(L) = \left(\frac{\sigma_\eta^2(1-\rho L)}{\alpha\sigma_\mu^2}\right)^2 \left[\Gamma_{\Delta s}(L) - B(L)\Gamma_{\Delta a}(L)B(L^{-1})\right]$$

$$\Gamma_{\tilde{v}}(L) = \left(\frac{\sigma_\eta^2(1-\rho L)}{\alpha\sigma_\mu^2}\right)^2 \left[\frac{\sigma_\mu^2}{(1-\rho L)(1-\rho L^{-1})} + \sigma_\xi^2(1-L)(1-L^{-1}) - \frac{\alpha^2\sigma_\mu^4}{\sigma_\eta^4(1-\alpha L)^2(1-\alpha L^{-1})^2}\right].$$

To express $\Gamma_{\tilde{v}}(L)$ in a rational spectral form, define the quartic polynomial:

$$D(L) = (1-\alpha L)(1-\alpha L^{-1})(1-\rho L)(1-\rho L^{-1}) - c(1-L)(1-L^{-1}),$$

where $c = \frac{\sigma_\eta^2}{\sigma_\mu^2}\alpha$. Then $\Gamma_{\tilde{v}}(L)$ becomes:

$$\Gamma_{\tilde{v}}(L) = \sigma_\xi^2 \left(\frac{\sigma_\eta^2(1-\rho L)}{\alpha\sigma_\mu^2}\right)^2 \left[\frac{D(L)}{(1-\alpha L)^2(1-\alpha L^{-1})^2(1-\rho L)(1-\rho L^{-1})}\right].$$

We factor $D(L) = (1-\bar{\delta}_1 L)(1-\bar{\delta}_2 L)(1-\bar{\delta}_1 L^{-1})(1-\bar{\delta}_2 L^{-1})$, so the square root gives:

$$\tilde{v}_t = \frac{(1-L)(1-\bar{\delta}_1 L)(1-\bar{\delta}_2 L)}{(1-\alpha L)(1-\rho L)}\varepsilon_t^v, \quad \varepsilon_t^v \overset{iid}{\sim} \mathcal{N}(0, \frac{\alpha\sigma_\xi^2}{\bar{\delta}_1\bar{\delta}_2}) \tag{D.17}$$

which is an ARMA(2,3). ■

D.3 Long-run risk process

In this appendix, we consider a TFP process and information structure alternative to the one in (D.15). This alternative reflects the typical process in the long-run risk literature, studied in international finance and macroeconomics by Colacito and Croce (2011, 2013, 2018).

Letting Δa_t denote the growth rate of productivity (in deviations from its mean), the process $\{\Delta a_t\}$ is assumed to follow the law of motion:

$$\begin{aligned} \Delta a_t &= z_{t-1} + \varepsilon_t^a, \\ z_t &= \rho z_{t-1} + \varepsilon_t^z, \end{aligned} \qquad \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^a \end{bmatrix} \overset{\text{iid}}{\sim} \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_z^2 & 0 \\ 0 & \sigma_a^2 \end{bmatrix}\right), \tag{D.18}$$

where $0 < \rho < 1$. Colacito and Croce (2013) refer to ε_t^z as a long-run shock (or news shock), and ε_t^a as a short-run shock.

To properly isolate the independent contributions of beliefs (or the independent effects of technology and noise disturbances), we follow Chahrour and Jurado (2018) and construct a noise representation that is observationally equivalent to representation (D.18). The following lemma presents one such noise representation.

Lemma 2. *The representation of fundamentals and beliefs in system (D.18) is observationally equivalent to the noise representation*

$$\Delta a_t = \rho \Delta a_{t-1} + \varepsilon_t^w - \alpha \varepsilon_{t-1}^w$$

$$\tilde{z}_t = \sum_{j=0}^{\infty} \alpha^j \varepsilon^w_{t+1+j} + \tilde{v}_t$$

$$\tilde{v}_t = \alpha \tilde{v}_{t-1} + \varepsilon^v_t - \rho \varepsilon^v_{t-1}$$

$$\begin{bmatrix} \varepsilon^w_t \\ \varepsilon^v_t \end{bmatrix} \overset{iid}{\sim} \mathcal{N}\left(0, \begin{bmatrix} \kappa & 0 \\ 0 & \sigma_a^2/\sigma_z^2 \end{bmatrix}\right),$$

where $\alpha = \rho\sigma_a^2/\kappa$ *and* $\kappa = \sigma_z^2 + \sigma_a^2(1+\rho^2)$.

Proof. The proof follows the same logic as in the BLL case. First, let's derive the univariate Wold representation of Δa_t. The ACGF of Δa is:

$$\Gamma_{\Delta a}(L) = \frac{\sigma_z^2}{(1-\rho L)(1-\rho L^{-1})} + \sigma_a^2 = \frac{\sigma_z^2 + \sigma_a^2(1-\rho L)(1-\rho L^{-1})}{(1-\rho L)(1-\rho L^{-1})}.$$

We can factor the numerator as a symmetric second-order polynomial:

$$\Gamma_{\Delta a}(L) = \frac{\kappa(1-\alpha L)(1-\alpha L^{-1})}{(1-\rho L)(1-\rho L^{-1})} \tag{D.19}$$

with:

$$\kappa = \sigma_z^2 + \sigma_a^2(1+\rho^2), \quad \alpha = \frac{\rho\sigma_a^2}{\kappa}.$$

Taking the square root yields the Wold representation:

$$\Delta a_t = \frac{1-\alpha L}{1-\rho L} \varepsilon^w_t \tag{D.20}$$

where $\varepsilon^w_t \overset{iid}{\sim} \mathcal{N}(0,\kappa)$ is the Wold innovation.

Second, let's represent the signal process as a function of ε^w_t and noise. Define the linear projection:

$$z_t = B(L)\Delta a_t + v_t \tag{D.21}$$

where $v_t \perp \Delta a_{t-j}$ $\forall j \in \mathbb{Z}$, i.e., the noise component is orthogonal to all leads and lags of technology.

Using standard projection theory and the same factorization as in (D.19):

$$B(L) = \frac{\Gamma_{z,\Delta a}(L)}{\Gamma_{\Delta a}(L)} = \frac{\sigma_z^2 L^{-1}}{\sigma_z^2 + \sigma_a^2(1-\rho L)(1-\rho L^{-1})} = \frac{\sigma_z^2 L^{-1}}{\kappa(1-\alpha L)(1-\alpha L^{-1})}.$$

The noise process v_t has autocovariance generating function:

$$\Gamma_v(L) = \Gamma_z(L) - B(L)\Gamma_{\Delta a}(L)B(L^{-1}) = \frac{\sigma_z^2}{\sigma_a^2\kappa(1-\alpha L)(1-\alpha L^{-1})}$$

taking the square root yields:

$$v_t = \frac{1}{1-\alpha L}\frac{\sigma_z^2}{\sigma_a^2\sqrt{\kappa}}\varepsilon_t^v \tag{D.22}$$

with $\varepsilon_t^v \overset{iid}{\sim} \mathcal{N}(0, \sigma_a^2/\sigma_z^2)$.

Finally, plug equation (D.20) and (D.22) into (D.21), and rearrange to obtain:

$$\underbrace{(1-\rho L)\frac{\sigma_a^2\sqrt{\kappa}}{\sigma_z^2}z_t}_{\tilde{z}_t} = \frac{L^{-1}}{1-\alpha L^{-1}}\varepsilon_t^w + \underbrace{\frac{1-\rho L}{1-\alpha L}\varepsilon_t^v}_{\tilde{v}_t} \tag{D.23}$$

■

Foreign TFP process follows an analogous process, and it is cointegrated with home TFP:

$$\begin{aligned} \Delta a_t^\star &= z_{t-1}^\star + \tau(a_{t-1} - a_{t-1}^\star) + \varepsilon_{a,t}^\star, \\ z_t^\star &= \rho z_{t-1}^\star + \varepsilon_{z,t}^\star, \end{aligned} \qquad \begin{bmatrix} \varepsilon_{z,t}^\star \\ \varepsilon_{a,t}^\star \end{bmatrix} \overset{\text{iid}}{\sim} \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_z^2 & 0 \\ 0 & \sigma_a^2 \end{bmatrix} \right), \tag{D.24}$$

We calibrate the parameter $\tau \in (0,1)$ to a small number to generate moderate cointegration.

Figure D.9 characterizes the impulse responses of the EZ-complete market model both under the baseline process and the long-run risk process that we describe in the main text.

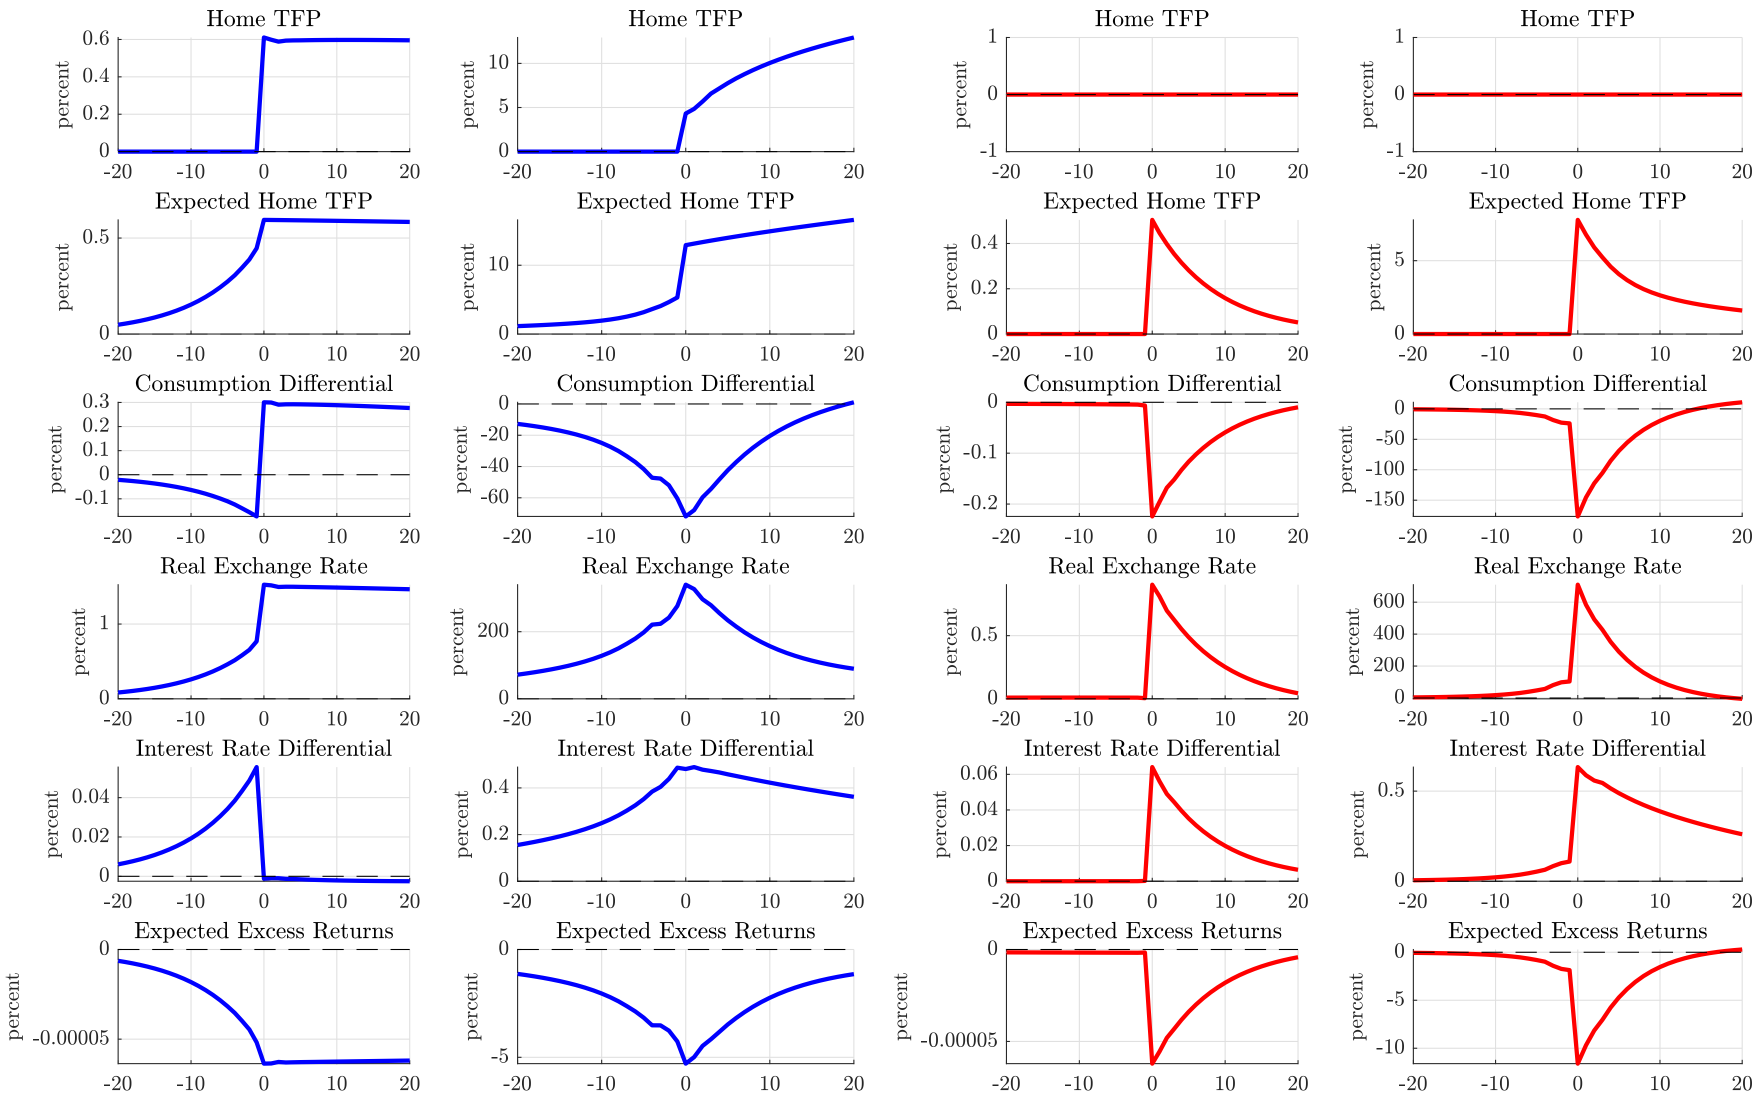

(a) Technology dist.
Baseline process

(b) Technology dist.
Long-run risk process

(c) Noise dist.
Baseline process

(d) Noise dist.
Long-run risk process

Figure D.9: Impulse response under Complete Markets and Epstein and Zin preferences, Alternative Processes