Exchange Rate Disconnect Redux*

Ryan Chahrour
Boston College

Vito Cormun
Santa Clara University

Pierre De Leo
University of Maryland

Pablo Guerron-Quintana
Boston College

Rosen Valchev
Boston College

February 15, 2021

PRELIMINARY AND INCOMPLETE
Please do not circulate

Abstract

In a sample of G6 currencies vis-a-vis the USD, we find that the variation in expectations of future US productivity explains more than half of exchange rate fluctuations. We separately identify significant effects of both anticipated changes in productivity and shocks that influence expectations of productivity, but are unrelated to any eventual change in productivity (a.k.a. noise shocks). Moreover, we find that each of these shocks generates all three well-known exchange rate anomalies of predictable excess returns, Backus-Smith puzzle and excess volatility. Thus, our findings suggests exchange rates are indeed connected to fundamentals (but mainly with *future* ones), and that three celebrated exchange rate puzzles have a *common* empirical origin.

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

*Contact: ryan.chahrour@bc.edu, vcormun@scu.edu, deleop@umd.edu, guerron@bc.eud, valchev@bc.edu.

1 Introduction

Exchange rate behavior is notoriously plagued by numerous "puzzles" in the sense that exchange rates in the data violate many of the implications of theoretical models. On the one hand, there is a general lack of connection between a variety of macroeconomic aggregates (e.g. output, investment, money supply) and exchange rates, both contemporaneously and in a forecasting sense, in a set of results the literature broadly refers to as the "exchange rate determination" puzzle.[1] A specific issue that has received outsize attention and is often discussed separately, is the lack of correlation between current interest rate differentials and subsequent exchange rate changes, which results in forecastable excess returns and thus a violation of the Uncovered Interest Parity (UIP) condition.[2] Similarly, the literature has also emphasized the low correlation between real exchange rates and the consumption differentials across countries, which is a violation of the Backus and Smith (1993) risk-sharing condition, another central implication of a large class of models. Overall, the combination of these different phenomena has led to the general observation that exchange rates tend to "live a life of their own" – i.e. there seems to be a basic "disconnect" between exchange rates and other macroeconomic quantities, very much contrary to model implications.

Exchange rates are at the heart of the international transmission mechanism of open economy models, and as a result the challenge of designing models that fit the empirical behavior of exchange rates has received significant attention. To guide this process, the literature has searched extensively for *any* empirical relationship between the exchange rate and the broader macroeconomy. The general finding is a pronounced lack of contemporaneous correlation between exchange rates and macroeconomic aggregates such as output or consumption, but some evidence of a correlation between exchange rates and *future* macroeconomic quantities (e.g. Engel and West (2005), Sarno and Schmeling (2014)).[3] The theoretical literature, on the other hand, has been converging to the conclusion that to explain multiple exchange rate puzzles simultaneously, models need to allow for what is essentially exchange-rate-specific shocks that remain largely disconnected from the broader economy (Rebelo et al. (2019), Itskhoki and Mukhin (2019)).

[1]See for example Meese and Rogoff (1983) and Engel and West (2005) among others.

[2]The UIP puzzle has been central to the exchange rate literature since the seminal work of Fama (1984), see Engel (2014) for an excellent survey

[3]Lilley et al. (2019) find a contemporaneous connection between US purchases of foreign bonds and the dollar, but only in the post-2009 period. Such contemporaneous relationships have proven elusive over a longer time span.

In this paper, we use an agnostic, model-free empirical analysis to uncover the shock(s) that acts as the major driver of exchange rate fluctuations in the data. We find that there are two shocks, both related to expectations of productivity growth, that simultaneously account for the lion's share of the variation in the real exchange rate, drive deviations from both the UIP and the Backus-Smith risk-sharing conditions, and explain a large portion of fluctuations in macroeconomic quantities such as consumption. These two shocks, which we separately identify, consist of (i) a fundamental shock to technology, which people partially anticipate; and (ii) an expectational "noise" shock, capturing expected movements in technology *that never realize*. Such unrealized expectations may correspond, for examples, to blue prints that never reach the adoption phase or to overly optimistic or pessimistic projections for new technologies under development. Overall, the empirical results are consistent with an account in which all of the major stylized facts about exchanges rates, including their relationship with real variables, are driven by agents' noisy information about future productivity. A broad insight is that many major exchange rate puzzles are in a sense one-and-the-same, as they arise conditional on the same shock.

Our analysis proceeds in two steps. First, we seek a purely "agnostic" description of the comovement patterns associated with changes in exchange rates. To do this, we follow the VAR procedure of Uhlig (2003), and recover a set of orthogonal shocks ordered by their respective importance in explaining exchange rate variation. We find that the "first" (i.e. most important to exchange rate dynamics) shock recovered in this way already explains a full three-quarters of exchange rate variation, and roughly half of the variation in macro aggregates. The shock also generates all three celebrated exchange rate puzzles that we described earlier. The key is that while this shock causes an immediate reaction in the exchange rate, its effect on macroeconomic quantities, with the exception of interest rates, are delayed. Thus, it only generates a correlation between exchange rates and *future* macro aggregates, in-line with the exchange rate determination and Backus-Smith puzzles. Yet, it also drives significant variation in forecastable excess returns and interest rate differentials in a way that generates not only the classic UIP puzzle, but also the "reversal" in UIP deviations at longer horizons recently emphasized by Engel (2016) and Valchev (2020).

This first step of our analysis establishes that exchange rate fluctuations are associated with subsequent real fluctuations and, in particular, with changes in future productivity (as measured by the Fernald (2014) TFP series). These results suggest that *anticipated* changes in productivity could play an important role in exchange rate dynamics and the broader empirical patterns we observe. The second part of our analysis analyzes this hypothesis by using the structural identification assumptions of Chahrour and Jurado (2019). For our purposes, the primary advantage of this empirical methodology is that it explicitly distinguishes between shocks that are associated with eventual changes in productivity (a.k.a. fundamental shocks) and shocks that influence expectations of productivity, but are unrelated to any eventual change in productivity (a.k.a. noise shocks). Beyond imposing that fundamental and noise shocks are orthogonal, their approach is quite general (see Chahrour and Jurado, 2019, for a full discussion.)

Implementing this approach in our baseline VAR, we find that both of these identified shocks play an important role in driving exchange rates and in generating the three puzzles summarized above. On the one hand, each accounts for roughly a quarter of the variation in real exchange rates by itself. On the other hand, the Impulse Response Functions (IRFs) to both shocks generate significant fluctuations in expected currency returns, in-line with both the classic UIP puzzle of high interest rates forecasting domestic currency profits and the newly documented "reversal" in this forecastability pattern at longer horizons. Both types of shocks also cause conditional movements in exchange rates and (delayed) movements in aggregates that in turn generate the Backus-Smith and exchange rate determination puzzle more broadly.

Lastly, we come full circle in our analysis and directly relate the identified shocks and the initial principal-components-type of shock we extract. There is no structural restriction between these two empirical analyses that would require that the identified fundamental and noise shocks span the same information as the shock that comes out of the Uhlig (2003) rotation of the variance-covariance matrix of innovations. However, we find that the combination of the fundamental and noise shocks virtually perfectly replicate the impulse responses to the principal-components-type of shock. Thus, the purely statistical "major" driver of exchange rates that we initially recovered is effectively fully spanned by the two structural shocks that we identify, giving it meaning as related to a mechanism of imperfect foresight of future productivity.

Overall, our results strongly suggest that exchange rates *are* connected to fundamentals, and the near-consensus favoring "disconnect" is at best incomplete. Even though one of our identified shocks is a "noise" kind of shock, it is very different from the model of exogenous shocks to the demand for foreign currency bonds, that is typically used to capture "noise" in the exchange rate literature. Instead, the noise shock we identify is a shock to expectations about future productivity and, hence, an information shock about a deep economic fundamental. Thus, our results show that over half of exchange rate variation, and a significant proportion of the specific conditional dynamics that generate a number of well known exchange rate puzzles, are all tightly linked to shocks related to productivity.

In the last part of the paper, we make headway towards a model that can explain these results. Beyond the obvious implication that one needs a model driven by imperfectly anticipated productivity, our empirical results also show that the exchange rate effects of these shocks run almost exclusively through UIP deviations. Decomposing the real exchange rate in two components, one driven by the sum of future expected interest rates, and one by the sum of expected excess currency returns, we find that the noisy-news shocks drive exchange rate dynamics predominantly through the excess returns component. Thus, a model that can explain the data must (i) feature noisy information about future productivity; and (ii) contain a mechanism through which informational shocks can drive endogenous UIP deviations (in a way consistent with the IRFs we estimate).

Related Literature

This paper is related to several different strands of the international and macro literatures. On the empirical side, we speak to the exchange rate determination puzzle which refers to the lack of correlation between exchange rates and macroeconomic fundamentals, both contemporaneously and in terms of forecasting future exchange rates with current macroeconomic fundamentals (e.g. Meese and Rogoff (1983), Cheung et al. (2005), Engel and West (2005)). There is also the related observation that the exchange rate is "excessively" volatile and persistent, as compared to macroeconomic fundamentals – see for example Obstfeld and Rogoff (2000), Chari et al. (2002), Sarno (2005), Steinsson (2008).

Our finding that there is a connection between exchange rates and macroeconomic fundamentals, but one that runs between current exchange rates and future fundamentals, is the opposite of the forecasting relationship between current and past macro variables and exchange rates, that some of the literature has tried to, but failed to find robustly (e.g. Meese and Rogoff (1983), Rogoff and Stavrakeva (2008)). However, it is consistent with previous studies that have documented that exchange rates Granger-cause some macroeconomic quantities (Engel and West (2005), Sarno and Schmeling (2014)). Our results contribute to this discussion, by showing that the link between current exchange rates and future fundamentals runs specifically through imperfect foresight regarding future productivity. Moreover, our results show that this imperfect foresight mechanism is also responsible for the two most famous "pricing" exchange rate puzzles – the UIP and the Backus-Smith puzzles.

A recent related paper is Stavrakeva and Tang (2018), who use survey of expectations to measure the surprises in macroeconomic announcements. They 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, and one that is especially important for the portion of the exchange rate driven by expected future currency returns. Our definition of "news" is different, as we specifically identify shocks to beliefs about *future* US TFP *innovations* (as opposed to revision of beliefs about past endogenous variables such as output), hence we add to their result by documenting the importance of the arrival of information about *future* productivity developments is a significant driver of exchange rates and currency returns.

Relative to the papers discussed above, our results also specifically show a link between the imperfect information about the future and two seminal exchange rate puzzles – the UIP (e.g. Fama (1984), Engel (2014)) and the Backus-Smith puzzles (Backus and Smith (1993)). Both puzzles have received extensive theoretical attention, and numerous potential mechanisms have been proposed as resolution of one or the other.[4] Such models, however, have typically relied on the standard assumption that agents have full information on current and past innovations to the exogenous shocks driving the economy, but no information on their future innovations. As a result, while the models are consistent with the pricing puzzles, they often run counter to the exchange rate "disconnect" that the previously described literature tries to grapple with, since shocks drive contemporaneous moves in both exchange rates and other macroeconomic quantities.

To confront this challenge, a new strand of the literature has emerged that has analyzed mechanisms that can generate the exchange rate pricing puzzles based on exchange-rate-market specific "noise trader" shocks that have only a muted effect on the broader macroeconomy (Rebelo et al. (2019), Itskhoki and Mukhin (2019)). This is a new and more elaborate take on the older idea that, given the exchange rate disconnect fact, UIP-specific or FX-risk shocks are a convenient and powerful way of generating empirically realistic exchange rate dynamics (e.g. Devereux and Engel (2002), Jeanne and Rose (2002), Kollmann (2005), Bacchetta and van Wincoop (2006), Farhi and Werning (2012)).[5]

Relative to this recent literature emphasizing the role of shocks to noise trader FX-

[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), Bacchetta and Van Wincoop (2010), Burnside et al. (2011), Ilut (2012)) and liquidity premia (Valchev (2020) have been proposed as explanations of the UIP Puzzle. On the other hand, Corsetti et al. (2010), Colacito and Croce (2013), Karabarbounis (2014) develop models that explain the Backus-Smith puzzle.

[5] Relatedly, Huo et al. (2020) find that international comovement between macro aggregates is also likely explained by non-fundamental shocks, though they do not speak to correlation with exchange rates

demand, our empirical results suggest that another promising avenue is to examine models with partial information about future productivity. While both paradigms feature some "noise", the two are conceptually different. In the case of the existing literature, the "noise" shock is an exogenous shift in the demand for one currency relative to another, that is orthogonal to macroeconomic fundamentals. Our results, instead, provide evidence of a shock that creates noise in the expectations of future fundamentals. Hence, while ex-post the shock is indeed orthogonal to fundamentals at all leads and lags, agents do not know this in real time and do react to it as if it carries information about future productivity. In that sense, it is both a shock about fundamentals, and one that is perceived as such by the agents.

Thus, our results suggest a mechanism that provides a comprehensive explanation of empirical exchange rate dynamics must be able to generate all major exchange rate puzzles conditional on the *same* shock related to imperfect foresight of future productivity. Models that can generate multiple exchange rate puzzles out of TFP shocks are rare – one such model (albeit with no room for anticipation of future productivity) is Colacito and Croce (2013). Nevertheless, that model also cannot generate the reversal of UIP at longer horizons and the broader exchange rate determination puzzle. The particular details of a "successful" exchange rate model that is consistent with the full details of our empirical results is still an open question.

Lastly, there is a small but growing literature specifically documenting the effects of news shocks in the international data and developing 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 evidence by Beaudry and Portier (2014). Perhaps most closely related to us is the work of Nam and Wang (2015), who use a Barsky and Sims (2011) approach to identifying news-to-TFP shocks, and find that they are indeed an important driver of exchange rates in the data. In contrast to us, however, they do not consider the effect of the shocks on exchange rate puzzles and also do not separately identify the effects of fundamental shocks from those driven by expectations shocks that are orthogonal to fundamentals.

2 Initial Empirical Analysis

In this section we describe our first-step, model-free empirical analysis. Following the methodology in Uhlig (2003), we are interested in examining the properties of the shocks that matter the most to exchange rate variation. Intuitively, this approach is akin to a principal components analysis of the k-step ahead forecast error variance of exchange rates.

Specifically, given a VAR

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

Uhlig (2003) suggests to find the matrix A and orthogonal series of shocks ε_t such that

$$u_t = A\varepsilon_t$$

where ε_t are ordered in terms of their importance in explaining the forecast error variance (as implied by the VAR in equation (1)) of any one variable (or set of variables) $y_{i,t} \in Y_t$.

In our benchmark analysis, the vector Y_t contains data on the US and a trade-weighted aggregate for the other G6 countries.[6] The vector Y_t contains (i) the nominal exchange rate S_t expressed in terms of number of USD per foreign currency, (ii) the Fernald series on US-TFP cleaned out of endogenous components like utilization, (iii) US real consumption and investment, (iv) foreign real consumption, (v) the interest rate differential, (vi) and the CPI price level differential vis-a-vis the US:

$$Y_t = [\ln(S_t), \ln(TFP_t^{US}), \ln(C_t^{US}), \ln(I_t^{US}), \ln(C_t^*), \ln(\frac{1+i_t^{US}}{1+i_t^*}), \ln(\frac{CPI_t^{US}}{CPI_t^*})]$$

In applying the Uhlig (2003) method, we seek a rotation matrix A and ordering of shocks such that the first element of the vector ε_t explains the most of forecast error variance (FEV) of the exchange rate at horizons for 1 to 20 quarters in the future, and then the second element of ε_t has the second highest explanatory power for the 20-step ahead FEV and so on, so that:

$$\sum_{h=1}^{20} Var(E_t(u_{1,t+h}|\varepsilon_{1,t})) > \sum_{h=1}^{20} Var(E_t(u_{1,t+h}|\varepsilon_{2,t})) > \ldots$$

where $h = 1, ..., 20$, time is measured in quarters, $x_{i,t}$ is the $i-$th element of the vector x_t,

[6] In the Appendix we also report separate estimation results for each G7 country, including them in Y_t one at a time; results are consistent across all exchange rate pairs.

Figure 1: Average share of forecast error variance over horizons [2,100] explained by $\varepsilon_{1,t}$

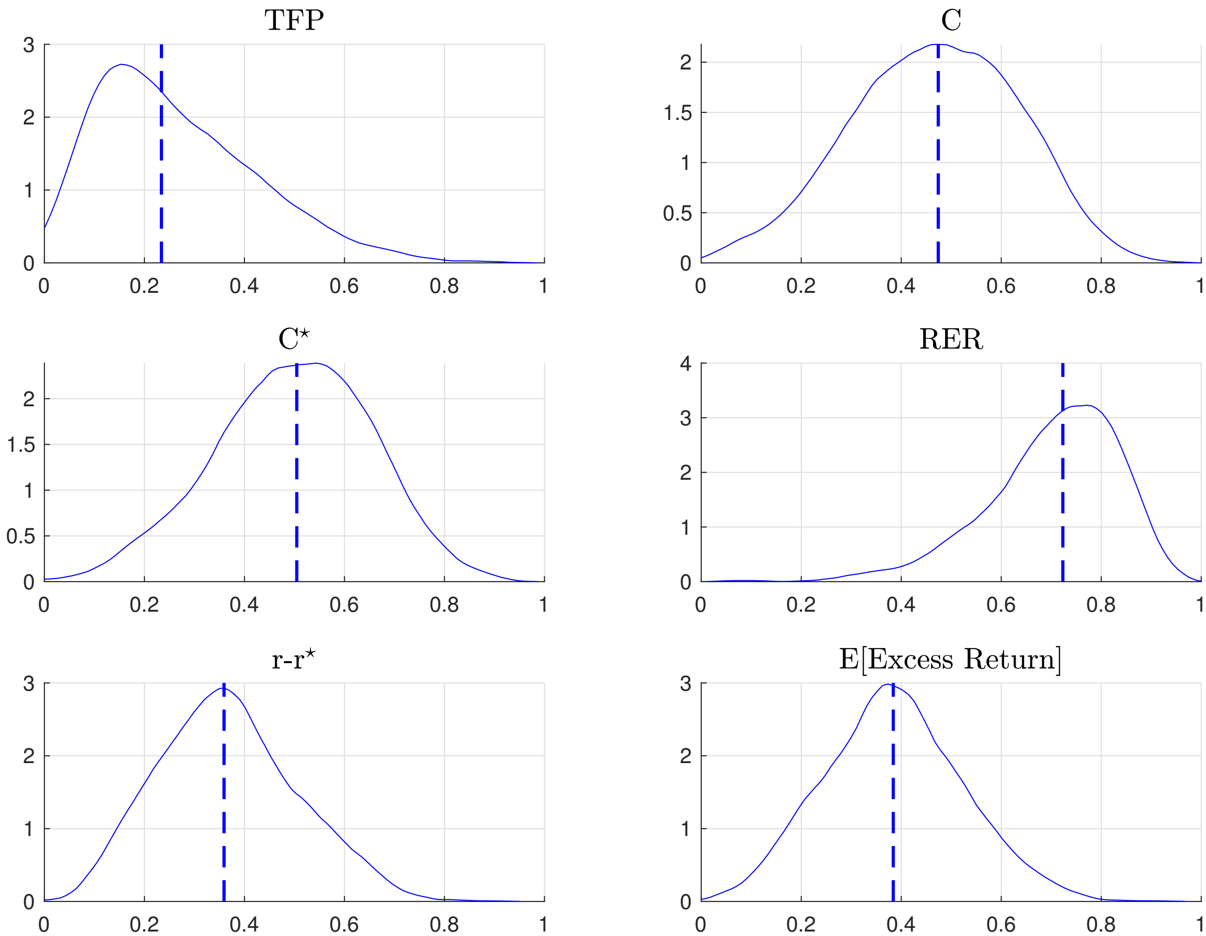

and thus for example $u_{1,t}$ is the residual in the forecasting equation for the exchange rate.

We use quarterly data for the time period 1976:Q1-2008:Q2 for the G7 countries, and in our benchmark analysis the foreign variables in Y_t are trade-weighted G6 averages – e.g. the exchange rate is the trade-weighted exchange rate of the US and so on. We include four lags, and estimate the VAR via Bayesian methods using Minnesota priors.

It turns out that the first (i.e. most important to the exchange rate) shock $\varepsilon_{1,t}$ already explains roughly three-quarters of the variation in the real exchange rate. As a result, we will focus the rest of our analysis on this first shock. Note that for consistency with model equations that we will get to later, we discuss the results in terms of the (log) real exchange rate $q_t = \ln(S_t) - \ln(\frac{CPI_t^{US}}{CPI_t^*})$, which is simply a linear combination of variables already included in the VAR, hence we can directly infer its IRFs, variance decomposition and etc..

We show the posterior distribution of the share of variance explained by $\varepsilon_{1,t}$ for all six-variables in our VAR in Figure 1. Specifically, this figure plots the share of the forecast error variance over horizons of 2 to 100 quarters, which is effectively close to the unconditional variance of the variables given the finite sample. There is a whole distribution of these shares due to our Bayesian analysis. The vertical dashed line denotes the median estimate and is what we use as our benchmark point estimate in the discussion below.

Our first result is that this one shock is responsible for the lion's share of variation in the real exchange rate, with a median estimate of the variance share explained by this shock being 75%. The fact that one shock can be so important to the exchange rate is perhaps intuitive, given the previous literature on the exchange rate disconnect (e.g. Engel (1999), Engel and West (2005)), because it gives the impression that there is indeed an exchange-rate specific shock that accounts for most of the movements. From that point of view, however, one might expect that in turn this shock will not be responsible for significant fluctuations in macro variables.

Yet, that is not true this shock – which is of such central importance to exchange rate fluctuations – is also responsible for roughly half of the fluctuations in US and foreign consumptions. The shock is similarly important for US TFP, being able to explain roughly a quarter of its variance. Lastly, in addition to explaining fluctuations in the level of exchange rates and macro aggregates, the shock also drives a significant proportion of the fluctuations in interest rate differential and expected excess returns. Taking all three facts together, this hints at the possibility that a single shock might be responsible for a significant portion of the fluctuations of (i) exchange rate levels, (ii) macro aggregates and (iii) the UIP puzzle.

In Figure 2 we plot the impulse response functions of the six variables to this FX shock. The median impulse response is plotted with a dashed blue line, and the shaded areas around it are the 68% and 90% confidence intervals respectively. A number of notable results emerge.

First, the real exchange rate shows a significant response on impact, appreciating by about 1.25% after a one standard deviation positive shock. The exchange rate also shows non-monotonic, delayed overshooting dynamics with the peak appreciation coming at around one year after the shock, after which time the exchange rate steadily depreciates back to its long-run mean. The non-monotonic dynamics we recover are similar to the ones previously emphasized by Eichenbaum and Evans (1995) and Steinsson (2008), and this results in a dynamic response that is very persistent – with a half life of three to three-and-a-half years – in line with the "excess persistence" puzzle documented by previous results. Overall, the results suggests that our different empirical procedure is indeed picking up a source of exchange rate variation that is responsible for important and familiar empirical patterns in

Figure 2: Impulse Response Functions to $\varepsilon_{1,t}$

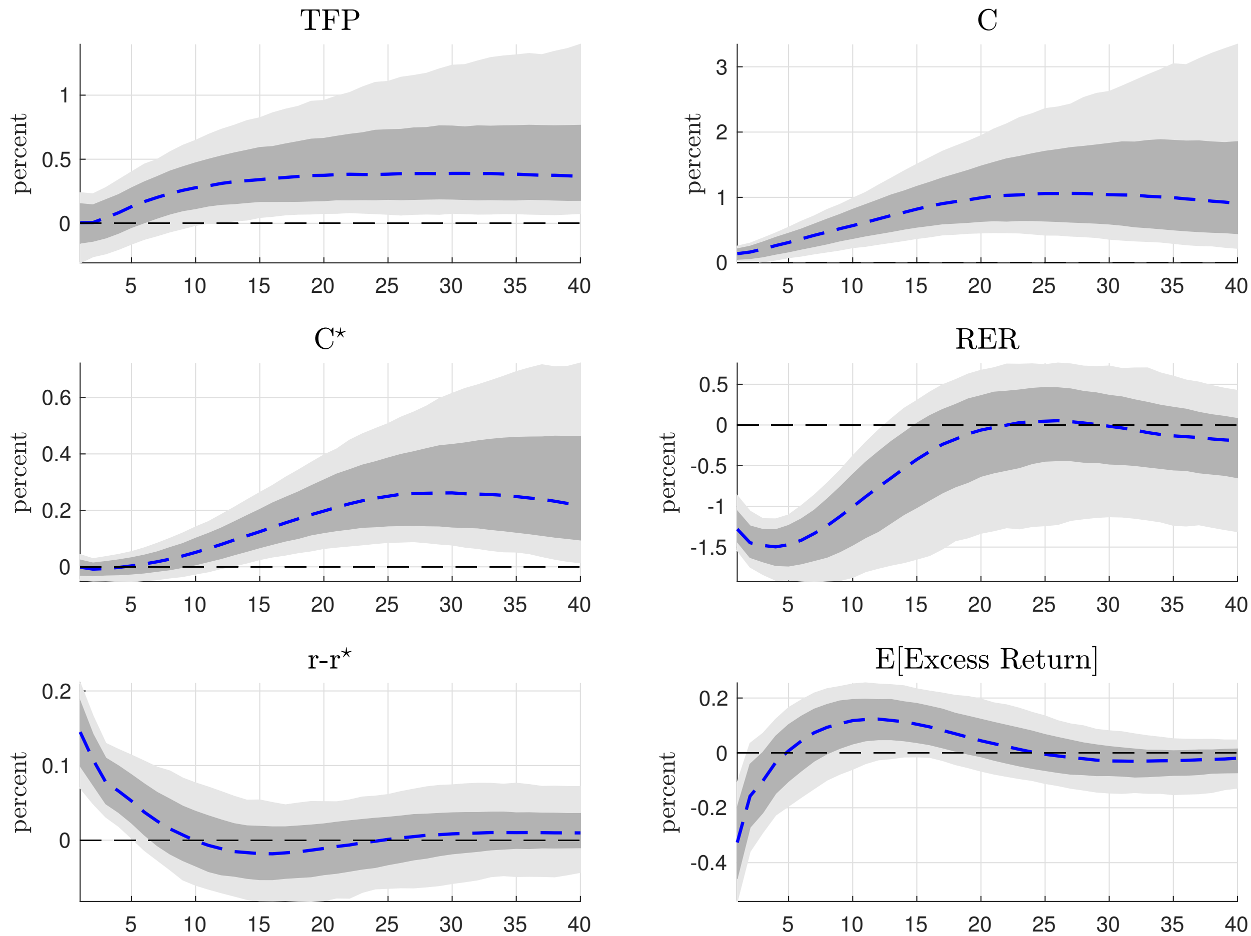

the exchange rate.

Importantly, the non-monotonicity we estimate is driving a "reversing" or "cyclical" pattern in the deviations from uncovered interest parity, in line with the results of Engel (2016). Specifically, while the shock causes delayed overshooting in the exchange rate, it leads to a monotonic impulse response in the interest rate differential, which increases on impact and gradually returns to its long-run mean. In addition, the shock also causes non-monotonic movements in the *expected* excess return, $\mathbb{E}_t(\lambda_{t+1}) \equiv \mathbb{E}_t(\Delta q_{t+1} + r_t^* - r_t)$, with the expectation being defined by the VAR in equation (1). Such predictable variation in the expected excess returns is a violation of the uncovered interest parity (UIP) condition.

In particular, our IRF estimates show that the initial sustained appreciation of the exchange rate leads to an increase in the expected excess return on the dollar – meaning that borrowing in foreign currency and investing in the USD makes money – precisely at times of elevated US interest rates. This is a manifestation of the "classic" UIP Puzzle that currencies are expected to earn higher returns following an increase in their interest rate, or put another way, the observation that exchange rates do not depreciate enough to offset movements in interest rate differentials, leaving potential profit opportunities on the table. We can read this off the bottom two IRFs which show that $r_t - r_t^*$ and $\mathbb{E}_t(\lambda_{t+1})$ move in opposite directions at short horizons.

In addition, the eventual depreciation of the exchange rate, causes a "reversal" in the UIP violation with the USD being expected to loose money against the foreign currency at horizons of 5 to 20 quarters in the future (which manifests in the the IRF of $\mathbb{E}_t(\lambda_{t+1})$ turning significantly *positive* at such medium-term horizons). This is in-line with the recent evidence that the UIP puzzle is more involved than the basic observation that "high interest currencies make money", as there are lower-frequency reversals in that relationship as exemplified by our impulse responses. The fact that the FX shock we identify can explain both the initial increase in excess returns and their eventual drop, supports the hypothesis that there is a common driver behind that pattern (e.g. Valchev (2020)).

While the fact that there is a common factor in both the short and long-horizon UIP violations might not be overly surprising, our result that this shock also causes significant movements in macro aggregates, thus connecting both exchange rates levels and returns to macro fundamentals is an interesting and novel result. In terms of consumption, the FX shock we identify has very little effect on impact and in the short-run overall. Both US and foreign consumption only respond in quantitatively significant terms to the shock after a couple of years. The effect on both US and foreign consumption peaks at around 25 quarters in the future, with US consumption moving by a significantly bigger amount, achieving a peak of roughly 1% increase while foreign consumption increases by 0.3% at its peak.

Interestingly, the effect on US TFP is similarly delayed, with the shock having an insignificant impact on productivity up to 5 quarters in the future, but eventually leads to what looks like a virtually permanent 0.4% increase in US productivity.

The dynamic patterns we uncover are consistent with several different exchange rate puzzles, suggesting that this shock might be a common driver to them all. First, we have already discussed how the shock leads to fluctuations in expected excess returns that are consistent with the literature's latest understanding of the UIP Puzzle. Second, the fact that the shock causes significant movements in the exchange rate on impact, but has no effect on US or foreign consumption implies that this shock is at least partially responsible for the Backus-Smith risk-sharing puzzle – the surprisingly low $corr(\Delta(C_t - C_t^*), \Delta q_t)$. And third, the lack of on-impact response extends to other macro aggregates as well – see for example the TFP IRF, and we find have similar results for output and investment (see Appendix) – suggesting that this shock can also explain the so called exchange rate "determination puzzle", the surprisingly low contemporaneous correlation between exchange rates and both endogenous macroeconomic variables like consumption and output, and commonly agreed-upon macroeconomic drivers such as TFP.

The way our results tie together a number of previous puzzles is through the fact that the estimates point to an immediate response in exchange rates and their expected returns, but only a delayed reaction in macroeconomic aggregates. Thus, the common shock we find is not an exchange rate-specific shock that is somehow disconnected from the macroeconomy – the shock is indeed connected to the broader economic fundamentals, but mainly to *future* fundamentals and not today's fundamentals.

To showcase these results in another way, in Figure 3 we show the share of the FEV for individual quarter-horizons for each of the six variables in the VAR, that can be explained by our shock up to 40 quarters in the future.

As foreshadowed by the shape of the impulse response functions, this shock is the major driver of exchange rate fluctuations regardless of the horizon. The shock is most important for the short-to-medium term variation of about 1 to 2 years, where its explanatory power peaks at 90%, but it is similarly important across all horizons with a minimum explanatory power of roughly 70%. In addition to explaining fluctuations in the level of exchange rates and macro aggregates, the shock also drives a significant proportion of the fluctuations in interest rate differential and expected excess returns.

On the other hand, while the shock is responsible for roughly half of the total variation in US and foreign consumption, this is primarily due to its effect on medium-to-long run fluctuations of 5+ years. The shock has a similar effect on US TFP, being able to explain roughly 40% of medium-to-long run fluctuations in TFP, but very little of the high frequency fluctuations of one year or less.

2.1 Exchange Rate Decomposition

The results so far suggest that, broadly speaking, the exchange rate is a leading indicator or "anticipates" future macroeconomic fundamentals. This is not surprising by itself, given the forward looking nature of the exchange rate. A Campbell-Shiller type of decomposition for the level of the exchange rate q_t shows that it can be expressed as sum of future expected

Figure 3: Share of FEV explained by $\varepsilon_{1,t}$, at different horizons

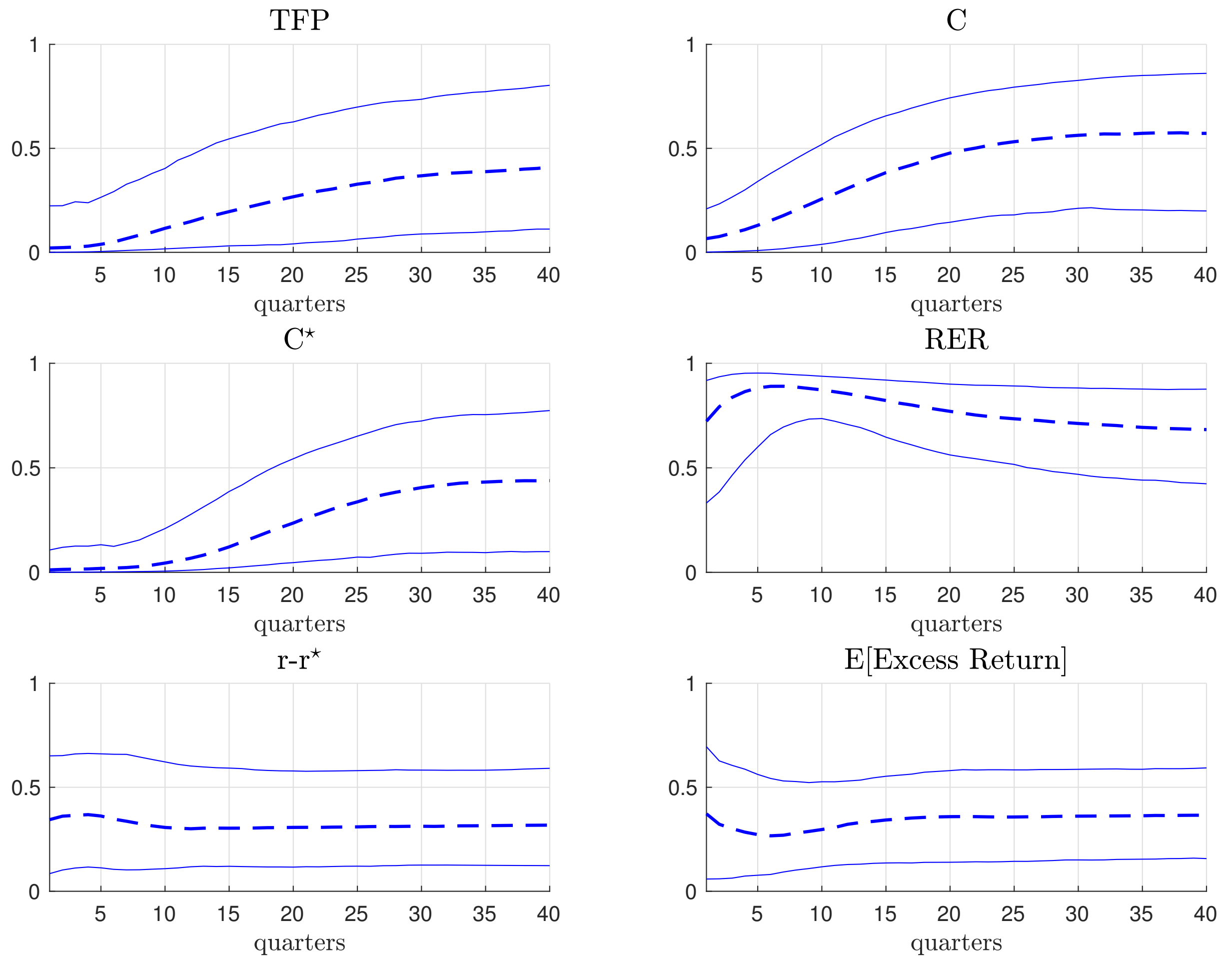

interest rate differentials and excess returns:

$$q_t = \sum_{h=0}^{\infty} \mathbb{E}_t(r_{t+h} - r^*_{t+h}) - \sum_{h=0}^{\infty} \mathbb{E}_t(\lambda_{t+1+h}) \tag{2}$$

In a no-arbitrage, equilibrium setting both interest rates and excess returns $\lambda_{t+1} = \Delta q_{t+1} + r^*_t - r_t$ must be driven by macroeconomic fundamentals, hence, it is natural to expect that whatever drives exchange rates must also be able to significantly influence and forecast future interest rates and excess returns.

This fundamental relation can help us understand and motivate the basic hypothesis behind the so-called "exchange rate determination puzzle" that the exchange rate should be closely related to contemporaneous macroeconomic aggregates such as C_t or TFP_t. That hypothesis is essentially based on the assumption that equilibrium is a linear system with

AR(1) dynamics, so that the best time-t forecast of future fundamentals are todays fundamentals.

Yet, as we found in our results, the data strongly suggests that the type of shock that is important to exchange rate fluctuations is inconsistent with the basic characteristics of a world driven by AR(1) shocks, where agents observe all current and past innovations, but have no information on future innovations. In such a framework, shocks have a strong immediate impact. On the other hand, while the shock we recover is indeed important to future fundamentals, it does not cause much of an immediate response.

3 Anticipated Productivity and Expectations Shocks

Our application of the maximum share procedure provides valuable insights regarding the empirical patterns associated with exchange rate fluctuations. Specifically, our results show that exchange fluctuations today are associated with changes in real economic outcomes in subsequent periods and, after significant delay, with statistically significant changes in productivity. These patterns suggest the possibility, at least, that anticipated changes in future technology may play an important role in driving exchange rates.

Our goal in this section is to provide a more structural assessment of the hypothesis that expectations about technology are a dominant driver of exchange rates. To do this we adopt the structural identification assumptions of Chahrour and Jurado (2019), which those authors show can be used to independently identify the "fundamental" shock driving realized changes in productivity and expectational "noise" shocks, which explain changes in productivity forecasts that are never realized.

The Chahrour and Jurado (2019) approach to identification is especially well-suited for our question for several reasons. First, it separates the effects of actual changes in technology from the effects of "pure beliefs" by construction, a feature that is central to our objective in this paper, given the predominant view in the literature that "noise" matters for exchange rates. This contrasts with the family of News shock identification schemes, such as Barsky and Sims (2011), which necessarily commingle the effects of beliefs with fundamentals. Second, it avoids the assumption that the underlying structural data generating process has an invertible representation, which is often violated in models of economic foresight (Blanchard et al., 2013). Finally, as we discuss below, this procedure allows for an arbitrary structure for the fundamental process and for the signal thereof, so that we need make essentially no assumptions about what aspects of productivity people learn about, or when they do so.

To fix ideas, we present a simplified discussion of the Chahrour and Jurado (2019) procedure here. The main assumption is that agents in the economy receive a noisy signal η_t about future TFP, with the signal being *any* linear linear combination of future innovations to TFP plus an orthogonal noise component v_t:

$$\eta_t = \sum_{k=1}^{\infty} \zeta_k \varepsilon^a_{t+k} + v_t,$$

where ε^a_{t+k} are the Wold representation innovations to the TFP process a_t:

$$a_t = A(L)\varepsilon^a_t. \tag{3}$$

Further assumptions on the particular structure of the TFP process or on the coefficients ζ_k are not necessary. Moreover, the noise component of the signal is also allowed to have an arbitrary lag structure:

$$v_t = \sum_{k=1}^{\infty} \nu_k \varepsilon^v_{t-k}.$$

The assumptions of the Chahrour and Jurado (2019) procedure are that (i) the productivity shocks ε^a_t explain 100% of the variation in TFP (i.e. they are indeed the Wold innovations in equation (3)) and (ii) the signal-noise innovations ε^v_t are orthogonal to TFP at all leads and lags. In the case of a two-variable var in $[a_t, \eta_t]$, the restrictions we impose amount to placing zeros in the MA representation of the data in the following:

$$\begin{bmatrix} a_t \\ \eta_t \end{bmatrix} = \cdots + \begin{bmatrix} 0 & 0 \\ * & 0 \end{bmatrix} \begin{bmatrix} \varepsilon^a_{t+1} \\ \varepsilon^v_{t+1} \end{bmatrix} + \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \varepsilon^a_{t} \\ \varepsilon^v_{t} \end{bmatrix} + \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \varepsilon^a_{t-1} \\ \varepsilon^v_{t-1} \end{bmatrix} + \cdots$$

This structure amounts to assuming that the productivity series a_t is orthogonal to the signal noise shocks ε^v_t at all leads and lags, while the signal η_t can contain information about future productivity ε^a_{t+k}.

The Chahrour and Jurado (2019) has several benefits. First, it separately identifies the fundamental shocks, ε^a_{t+k}, from the "noise" component of expectations, v_t. By examining the responses of economic variables, like the exchange rate, to the "fundamental" shock ε^a_{t+k}, we therefore see an indication of how (and if) fundamental shocks are anticipated. By examining responses to the second type of shock, ε^v_t we learn how much of economic fluctuations are associated with movements in expectations that are completely orthogonal to productivity – e.g. misplaced optimism or pessimism. This is especially useful for deriving insights that can guide model development, as the estimates can help recover the information sets of economic agents, and thus put tight restrictions on the modeling framework and its parameters.

Second, in practice, we do not need to directly observe the agents' expectations or actual signals η_t. Instead, we rely on (i) the assumption of rational expectations, which implies that the equilibrium variables load on agents' information, e.g. for the exchange rate:

$$q_t = \kappa_\eta \eta_t + \dots$$

and (ii) that our VAR empirical specification includes enough forward looking variables (in our case not just the exchange rate, but also consumption, interest rates, and prices, which are all "jump" variables in an an equilibrium framework) so as to be able to fully span the relevant part of the agents' information set (i.e. all information relevant to predicting future TFP). Under these auxiliary assumptions, we can identify the fundamental and noise shocks without making further assumptions about the information structure in the economy, and expectations of any variables in the system can be backed out using the dynamics implied by the VAR.

Applying this procedure to our estimated VAR system in equation (1), we obtain separately responses to the two shocks, ε_t^a and ε_t^v. We then begin by analyzing the respective shares of the variances of the six variables in the VAR that these two separate expectational shocks (fundamental and noise) can explain. The posterior distribution of those variance shares (median estimates denoted by dashed lines) are plotted in Figures 4 and 5. The left column presents the results for the fundamental shock ε_{t+k}^a and the right column the results for the noise shock ε_t^v, with the different variables whose variance we are decomposing presented in separate rows.

The first row of Figure 4 simply shows that our empirical procedure works as expected – the technology shock we estimate ε_t^a indeed accounts for 100% of the variation in TFP, while the noise shock is completely orthogonal to it.

In the second and third rows, we show the results for US and foreign consumption respectively. The estimates indicate that the fundamental technology shock explains about 45% of the variation in US consumptions, and 30% of the variation in foreign consumption, while the noise shock explains 30% of the variation in US consumption and about 20% of the variation in foreign consumption. Thus, consumption is not driven only by the actual productivity shock, but also by shocks to the expectations of future TFP. Intuitively, one would expect this latter expectational effect to also have an impact on asset prices.

Figure 4: Share of FEV (2-100Q) explained by ε_t^a and ε_t^v

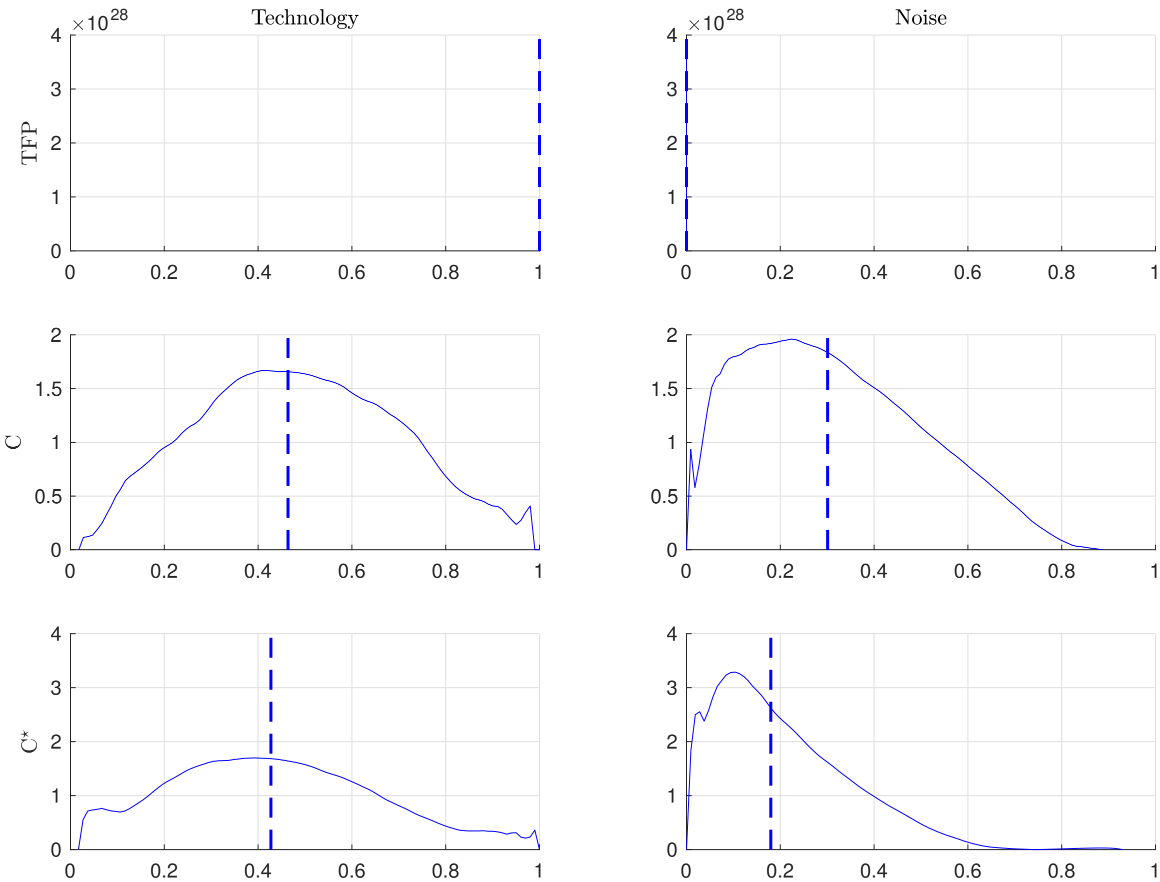

We can see this directly in the first row of Figure 5, which shows the shares of the variation in exchange rates (the international asset price of key interest to this study) that are driven by those two shocks. In the left column, we see that fundamental shocks to productivity are indeed significantly related to the exchange rate and can explain roughly 30% of its fluctuations. In the right column, we see that the noise shock is also quantitatively important and can explain about 20% of the exchange rate variation. We find a similar split in the importance of the two shocks for the interest rate differential, with actual productivity shocks explaining 30% and the expectational noise shocks explaining a fifth of the interest rate differential fluctuations.

Lastly, both shocks explain roughly the same amount of the variations in expected currency returns – about 20% each. This is intuitive and what one would indeed expect if excess returns are generated by a rational expectations, equilibrium model as we would discuss in more detail later. For now, we would just explain intuitively that the empirical procedure

Figure 5: Share of FEV (2-100Q) explained by ε_t^a and ε_t^v

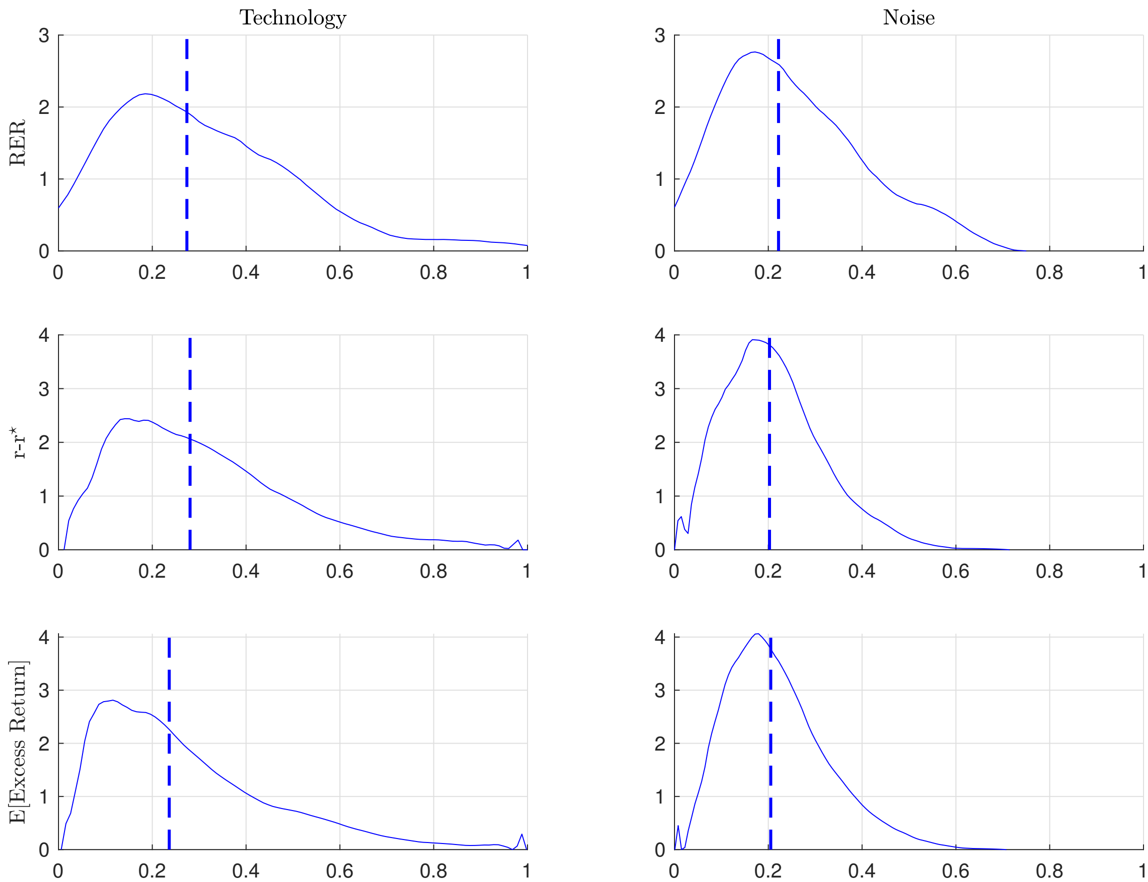

assumes agents only observe η_t but not ε_t^a and ε_t^v separately, hence they face a signal extraction problem. If information is indeed incomplete in this way, then agents would not be able to separately identify ε_t^a and ε_t^v in real time, hence the time-t expectation of the currency return should be influenced by both shocks similarly – there is rational confusion over which shock actually occurred. Note that we as economists are able to extract both shocks separately but only because we look at the data *ex-post*. Intuitively, our estimation procedure leverages the availability of a historical data series, and effectively runs a smoothing filter backwards to disentangle the two shocks from one another in *past* data. In real time, e.g. at the last period of the data sample, both us as econometricians and the agents are unable to disentangle the two shocks.

In Figures 6 and 7 we plot the estimate Impulse Response Functions (IRFs) to technology shocks (left column) and the expectational noise shock (right column). The x-axis is in terms of periods after the shock, with 0 representing the period of the shock itself. As

Figure 6: Impulse Responses to ε_t^a and ε_t^v

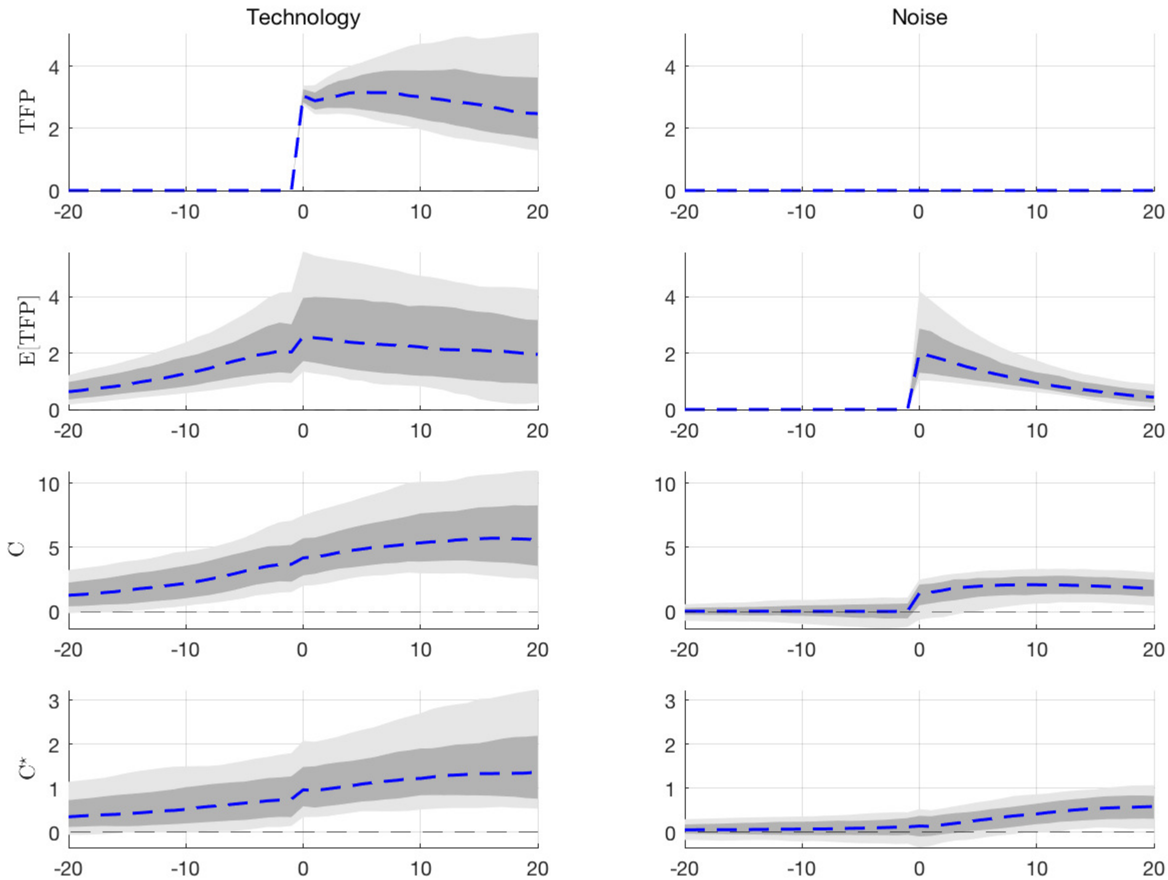

opposed to a standard IRF plot, we also show 20-quarters *before* the shock actually occurs, because anticipation implies that endogenous variables may start reacting to the shock before it actually affects productivity. Whether or not the endogenous variables respond before productivity actually moves is not assumed but is to be estimated – if the estimates show no significant early response of these variables, this would constitute clear evidence against the hypothesis of expectational effects of productivity.

In the first row of Figure 6, we again plot the impulse response of the TFP series itself as a "sanity check" – and we can see that, as expected, the technology shock ε_t^a only affects TFP from period 0 onwards since this is its Wold innovation, while the information noise shock has no effect on actual TFP at any horizon. We also note that the TFP process appears to be highly persistent, with some but weak evidence of mean reversion.

To help with the interpretation of our results and IRFs, in the second row of Figure 6 we also plot the response of the estimated agents' expectation of one-step-ahead productivity $\mathbb{E}_t(a_{t+1})$. In the left-column, we can see that the expectations of TFP are significantly higher than their long-run mean up to 20 quarters before the actual technology shock ε_t^a is realized. This showcases that the data speaks strongly in favor of anticipation effects, and thus supports the idea that agents have at least some advance information of future TFP. This information is, however, noisy, which can be inferred from the fact that the "noise" shock (which is by construction orthogonal to TFP at all leads and lags) also adjusts expectations up. Thus, we are indeed uncovering results consistent with the noisy-information paradigm, where agents do have some advance information and thus *partially* anticipate future movements in TFP, yet that information is noisy hence expectations sometimes move even though there is no actual future increase in productivity.

We can see a similar dynamic play out in the IRF of consumption. There are significant anticipation movements in consumption up to 20 quarters before the actual increase in productivity is realized, which consumption increasing smoothly and staying high after productivity actually ticks up. On the other hand, there are of course no anticipatory movements in respect to noise shocks, but once the noise shock arrives there is similarly a persistent boom in consumption. Thus, optimistic expectations about future TFP lead to an increase in consumption today. The quantitative impact of the noise shock is very similar to the early (i.e. $t-20$) anticipatory effects in response to the technology shock – with consumption increasing by roughly 1% in the US. This is natural, given that no matter how optimistic expectations are today, the actual improvement of technology has not arrived yet, hence the resource constraint of the economy has not been loosened.

In Figure 7 we show the impulse response functions of the real exchange rate and interest rate differential, and the expected currency returns. As expected given the consumption results, the real exchange rate similarly shows anticipatory effects, although those are significant much closer to the actual date of the shock – with the real exchange rate appreciating about 2% 5 quarters before an actual one standard deviation increase in productivity. The peak appreciation of 4% occurs concurrently with the increase in productivity, and then the exchange rate gradually depreciated back to its mean.

The exchange rate similarly appreciates when expectations of future TFP improve due to a noise shock. In that case, however, the peak impact of 3.5% is achieved immediately, and then the exchange rate gradually depreciates back to the its mean as agents learn gradually over time that their optimistic beliefs are in fact incorrect.

Figure 7: Impulse responses to ε_t^a and ε_t^v

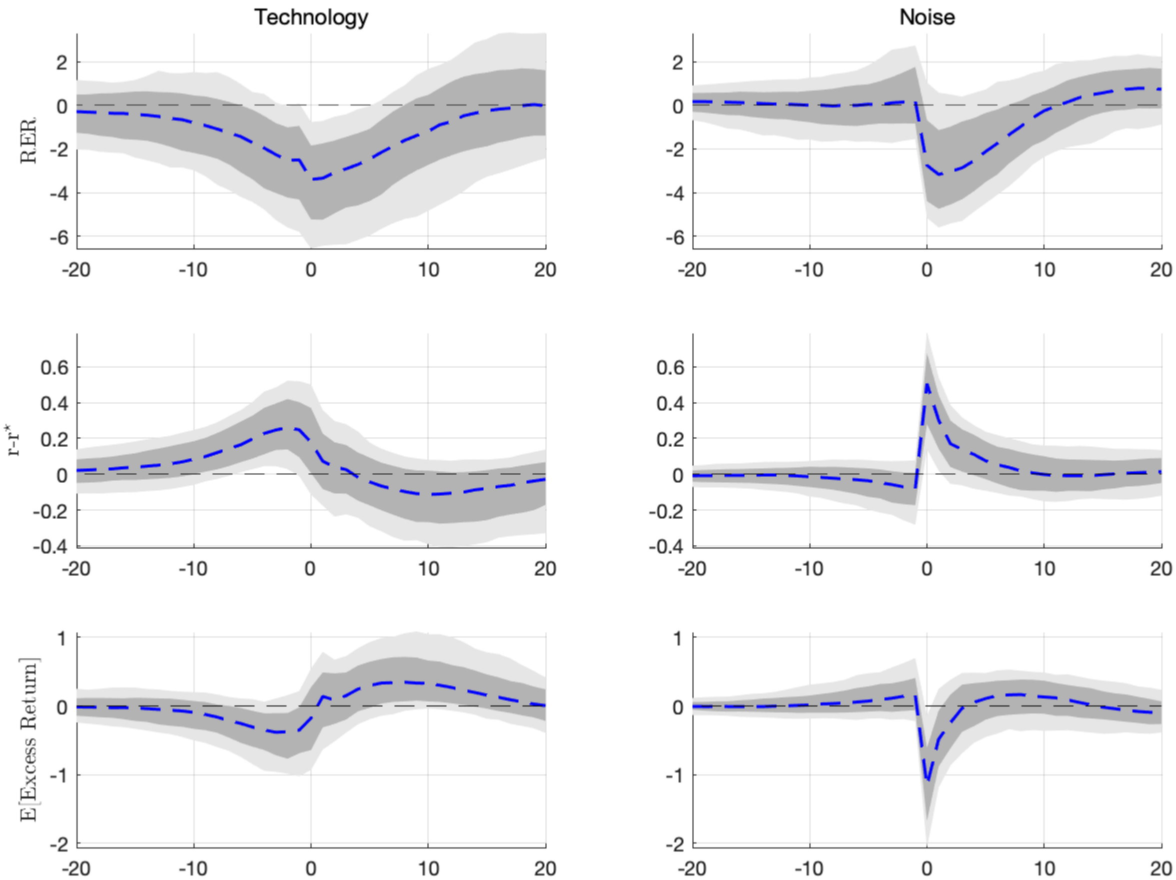

The interest rate differential similarly shows anticipatory effects, gradually increasing up to a peak of 0.4% 2 quarters before the actual increase in productivity, and then quickly falling below its long-run mean shortly after the shock realizes, and then increasing back to its steady state. A noise shock that increases expectations of future productivity similarly increases interest rates temporarily.

Lastly, we find that the expected currency excess return is depressed and significantly below its mean up to 8 quarters before an actual increase in productivity. Once the productivity shock realizes, the expected excess return rises quickly, and is significantly above its long-run mean for horizons of 4- to 12-quarters following the increase in productivity. The expectational noise shock also affects the expected excess returns significantly, with an increase in expected future US productivity leading to sharply lower expected USD returns.

The effect is temporary (again as agents learn that the shock was indeed simply noise), and return to steady state after 4 quarters.

3.1 Connecting back to Uhlig (2003)

Our results so far indicate that both fundamental and expectational shocks about future TFP are indeed closely related to exchange rates. We were originally motivated to explore this hypothesis due to the initial Uhlig (2003) results we obtained, which suggested the recovered FX shock mainly affects macro aggregates with a delay. The two sets of results align well with each other qualitatively, which is a result in itself, since the shocks identified by the Chahrour and Jurado (2019) procedure are not guaranteed to span the same space as the shock we recovered first by the Uhlig (2003) procedure.

To get a direct sense of how much of the original Uhlig (2003)-identified shock the two structurally-identified shocks we recover can account for, we use the estimated VAR to simulate counter-factual data conditional *only* on the Chahrour and Jurado (2019)-recovered technology and expectational shocks. This gives us a counter-factual dataset driven exclusively by the two identified shocks. We then perform the Uhlig (2003) maximum-share procedure on this simulated data, and plot the resulting IRFs against the original Uhlig (2003)-IRFs we recovered in Section 2.

The results are displayed in Figure 8. We reproduce the original IRFs and their associated standard errors with blue dashed-line and shaded regions around it, and with the red dashed line we present the resulting IRFs from the counter-factual, simulated sample. Remarkably, the two impulse responses align almost perfectly, especially in the case of the real exchange rate, the currency excess returns and the interest rate differential.

This suggests that, in terms of exchange rates at least, the identified fundamental and noise shocks we identified separately are essentially perfectly replicating the economic content of $\varepsilon_{1,t}$, the "most important" source of exchange rate variation as per Uhlig (2003) procedure. Hence, we have come full circle in our conclusions that anticipation of future TFP is an important driver of exchange rates.

Lastly, we also examine what percent of the variation in the exchange rate is due to the combination of (i) anticipation of future TFP shocks and (ii) shocks to the expectation of future TFP. This would answer the question of how much of the exchange rate variation is due to shocks about expectations of future productivity. To do so, we simulate fundamental+noise economy and compute the $1 - R^2$ after regressing the change in exchange rate on

Figure 8: Chahrour and Jurado (2019) vs Uhlig (2003)

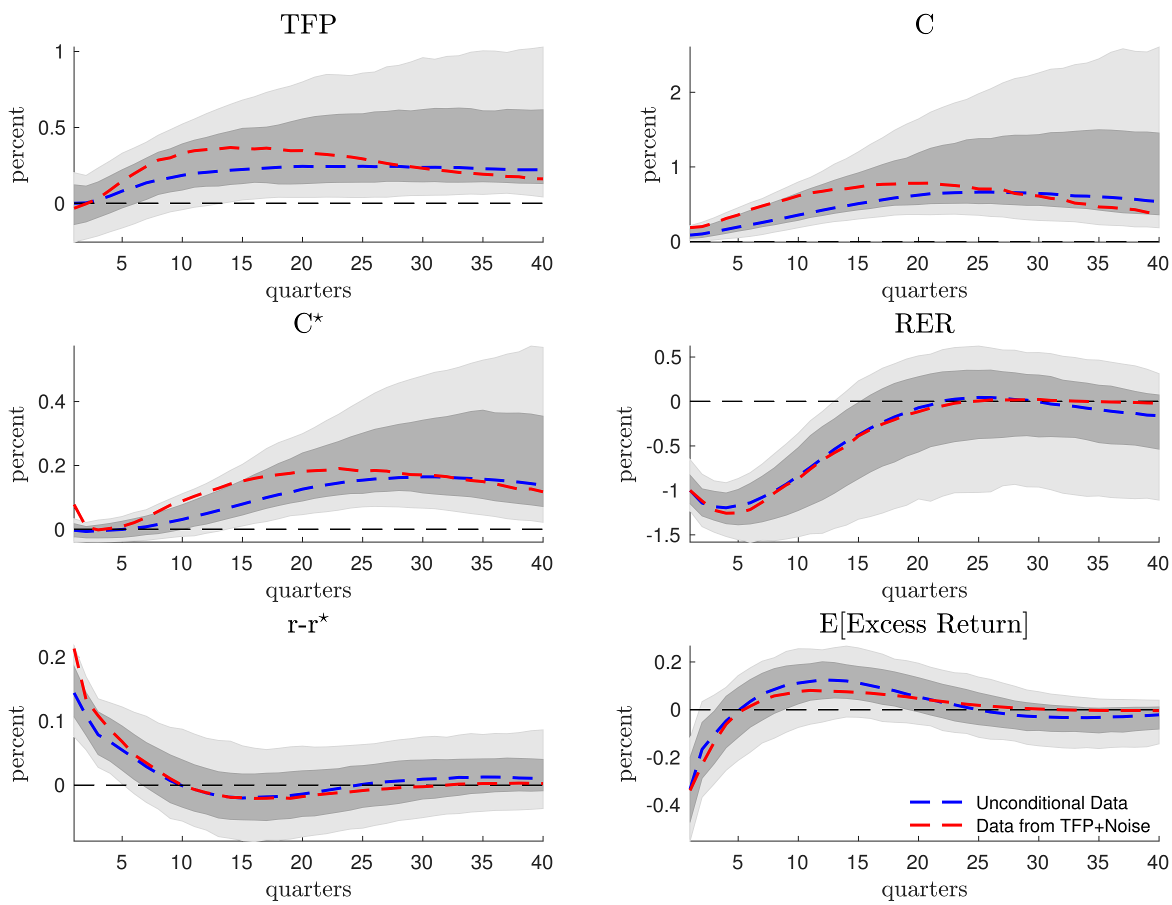

present and past fundamental shocks.

	$1 - R^2$
Germany	0.69
Italy	0.72
France	0.68
Canada	0.47
Japan	0.65
United Kingdom	0.66
G6	0.65

This shows that roughly two-thirds of the exchange rate variation is due to such shocks about future outcomes and expectations, and only a third of the exchange rate variation can be attributed to current and past productivity shocks.

4 Exchange rate puzzles and TFP anticipation

Given the large effect our two identified shocks play in exchange rate dynamics, it is interesting to consider whether the shocks are also driving some or all of the three broad exchange rate puzzles we outlined in the beginning. Namely, (i) the UIP puzzle, (ii) the Backus-Smith puzzle and (iii) the exchange rate determination puzzle.

Deviations from Uncovered Interest Parity

Starting with the UIP puzzle, we can drill down into the exchange rate dynamics further, by decomposing it into its two components – the one driven by future expected real interest rates and the one driven by future expected excess returns.

$$q_t = \underbrace{-\sum_{k=0}^{\infty}(r_{t+k} - r^*_{t+k})}_{=q_t^{UIP}} + \underbrace{\sum_{k=0}^{\infty} -E_t(\lambda_{t+k+1})}_{=q_t^{ExcessReturns}}$$

We call the first component, the one exclusively driven by the interest rate differentials, q_t^{UIP} to signify that this is a counter-factual exchange rate path that would respect the UIP condition. The actual exchange rate is of course different, because UIP is violated and hence $\mathbb{E}_t(\lambda_{t+1}) \neq 0$. To get a sense of how our shocks impact the exchange rate through the interest rate differential and through the cumulated deviations from UIP, in Figure 9 (third row) we plot the portion of the exchange rate driven exclusively by UIP violations:

$$RER(Ex.Ret) = q_t - q_t^{UIP}$$

As we can see, this quantity displays both significant anticipatory effects and effects upon and following a shock to expectations. In fact, the effects of both of these shocks are virtually indistinguishable from the IRFs of the raw exchange rate itself q_t. Thus, we can conclude that the identified we recover are indeed affecting the exchange rate predominantly through the UIP deviations channel.

Next, we focus specifically on the "classic" UIP puzzle that high interest rates predict high currency returns, in the sense that the seminal Fama (1984) UIP regression

$$\lambda_{t+1} = \alpha + \beta(r_t - r^*_t) + u_t$$

recovers an estimated coefficient $\beta < 0$. In our raw data, in the case of the G7 average

Figure 9: Exchange rate puzzles

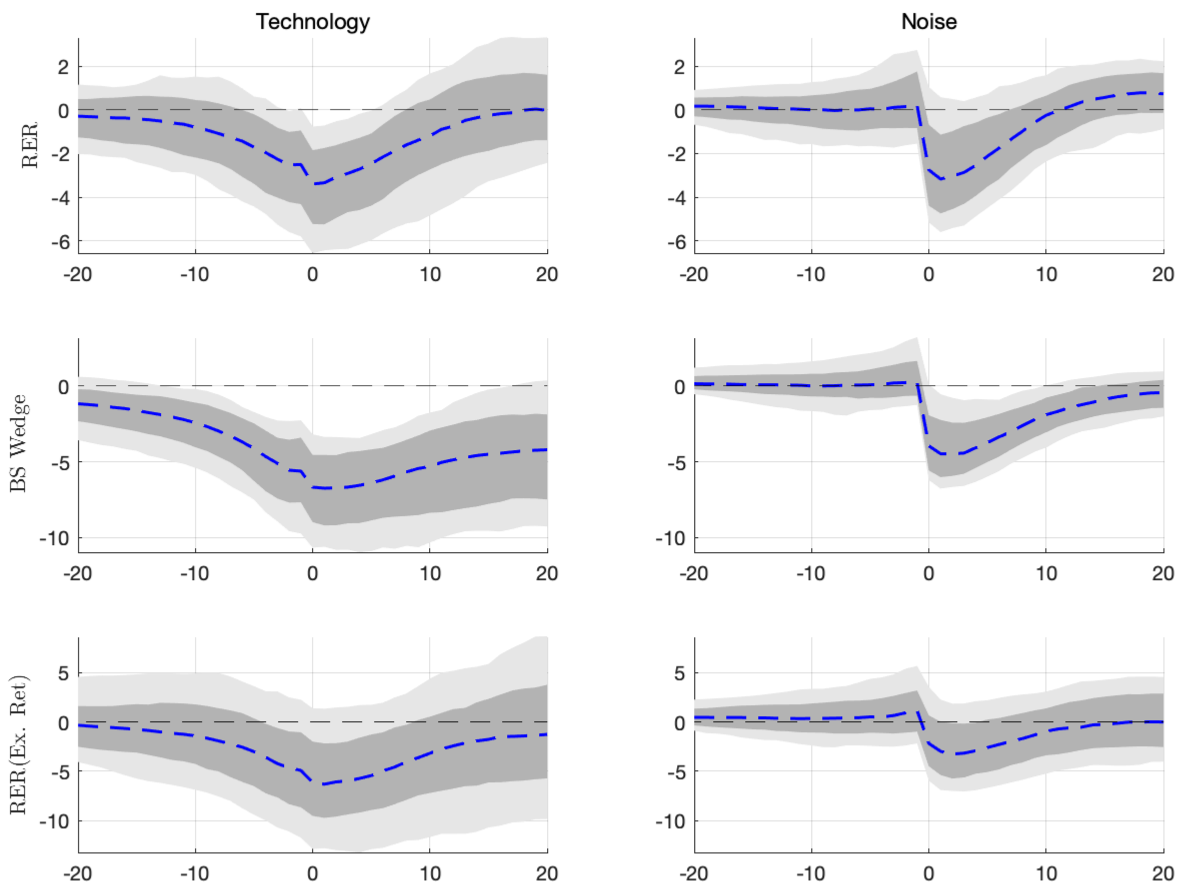

we find a significantly negative β of roughly -1.75 in line with previous findings (e.g. Engel (2014)). This is denoted by the same blue vertical line in all three panels of Figure 10.

Next, we compute the resulting UIP β in our counter-factual dataset simulated based only on the two shocks ε_t^a and ε_t^v. Given our Bayesian methods, we obtain a whole posterior distribution of UIP β estimates, which density we plot in the bottom panel of Figure 10. For clarity, with the solid red line we denote the median of that distribution and the dotted red lines represent the 95% coverage interval.

The median estimate is -1.4, hence a back-of-the-envelope calculation tells us that the combination of shocks to TFP and to expectations about TFP can explain roughly 80% of the classic UIP Puzzle relationship. Drilling down further, in the top panels of the Figure we plot the resulting posterior distributions of the UIP-β based on either only-TFP shocks

Figure 10: UIP regression decomposition

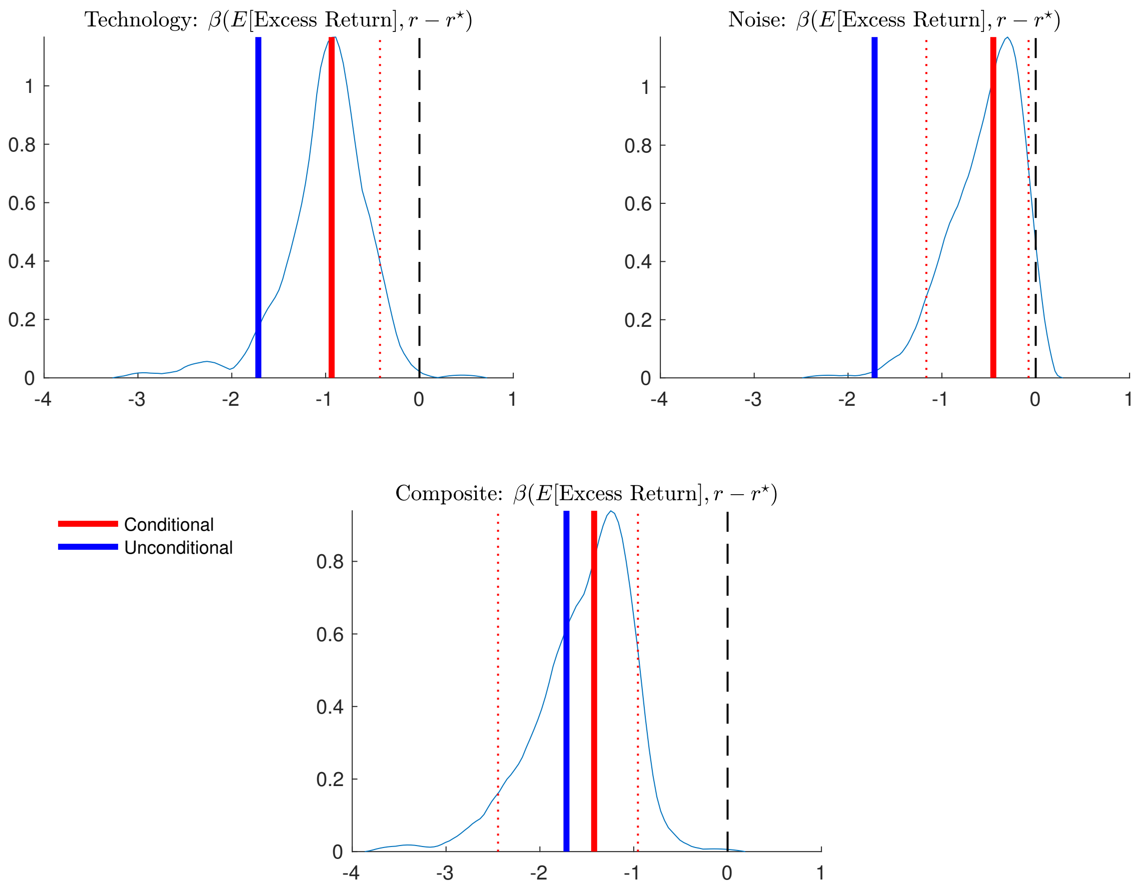

(including anticipation effects) and only expectational noise shocks. From the top left panel, we can see that the TFP shocks by themselves generate a UIP-β of -1, while the UIP-β based on only expectational shocks is -0.4 (by construction, the sum of the two equal the UIP-β based on a simulation where both of these shocks are active, i.e. the one presented in the bottom panel).

Thus, overall we conclude that fundamental and expectational shocks about TFP can explain a significant portion of the classic UIP puzzle, and more generally, those two shocks account for the majority of fluctuations in the real exchange rate *precisely* through their effect on UIP deviations.

Risk-sharing Puzzle

Next we turn to the Backus-Smith risk-sharing puzzle. As a first, step we consider the IRF of the Backus-Smith wedge defined as

$$\text{BS Wedge}_t = \Delta q_t - (c_t - c_t^*)$$

The impulse responses of that variable in respect to the fundamental TFP shock and the expectational shock are plotted in the second row of Figure 9. We can again see a significant anticipation effect in response to the actual TFP shock, with the wedge being significantly negative as early as 20 quarters before the actual shock. The fact that the wedge is negative, means that in anticipation of a positive US TFP shock, the dollar does not depreciate sufficiently to offset the gap in the consumption differential (which is also positive, as we can infer from the consumption IRFs). After the realization of the shock, the wedge adjusts slightly towards zero, but remains significantly different from zero for all 20 quarters after the shock we plot.

The expectational noise shock also causes significant effects on the BS Wedge. On impact of heightened expectations of high future productivity, the wedge also moves sharply negative and then converges back to zero over 20-quarters. Thus again, optimistic expectations of future TFP leads to a situation where the exchange rate does not depreciate sufficiently to offset the resulting boom in domestic consumption.

Overall, this shows that the two shocks we recover with the Chahrour and Jurado (2019) procedure are responsible for significant and volatile deviations from the perfect risk-sharing condition of Backus and Smith (1993). To examine this result from a different angle, we also evaluate what has become the benchmark Backus-Smith Puzzle moment

$$Corr(\Delta q_t, c_t - c_t^*)$$

in the counter-factual simulations based on only two identified shocks, and compare the resulting moment to the Backus-Smith correlation in the raw data. The results are presented in Figure 11, in the same way as in the UIP-β figure. The blue line represents the correlation estimated in the raw data, and with red lines we show the counter-factual estimates based on both shocks together, and on each one of them separately.

The bottom panel again shows posterior distribution of $Corr(\Delta q_t, c_t - c_t^*)$ in the counter-factual simulation based on both the fundamental productivity shocks and the shocks to expectations. The median estimate is -0.5, suggesting that these two shocks together generate a strong form of the Backus-Smith puzzle. The raw correlation in the data itself is roughly zero, weakly negative, while under the null hypothesis of perfect risk-sharing it is exactly one.

Figure 11: Backus-Smith correlation decomposition

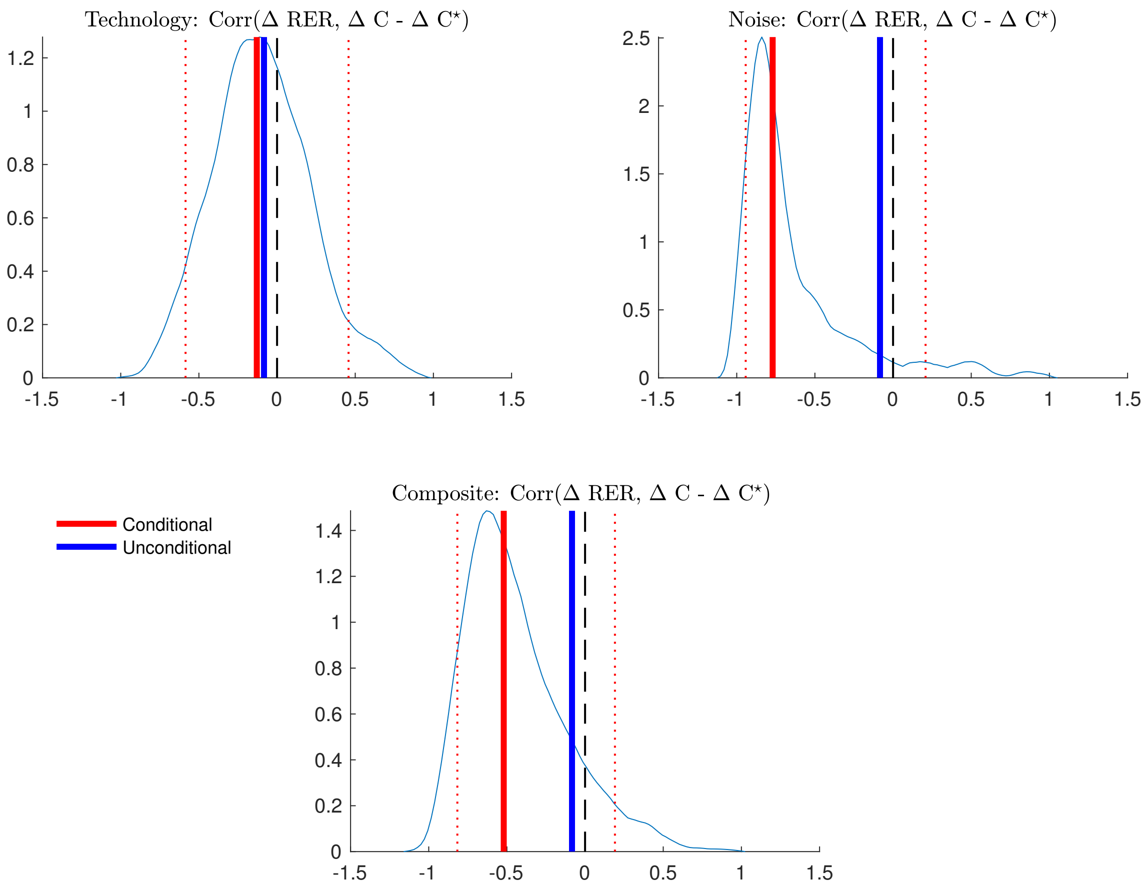

The top two panels look at the effects of the two shocks, ε_t^a and ε_t^v, separately. In the top left panel, we see that the fundamental productivity shocks by themselves generate a correlation that is roughly zero, while the expectational shocks by themselves generate a sharply negative correlation. In any case, the 95% credible intervals contain the raw data estimate in all three cases – both shocks together and each taken separately.

Exchange Rate Determination Puzzle

Lastly, we consider the exchange rate determination puzzle. This refers to the general observation that exchange rates are "disconnected" from the broader economy, as if they are "living a life of their own". There are a myriad of ways the previous literature has quantified this phenomenon in the data, and in here we are going to focus on two particular moments – the ratio between the variance of the quarterly growth rate of US consumption and the real exchange rate and the contemporaneous correlation between those growth rates.

Figure 12: FX-consumption variance ratio

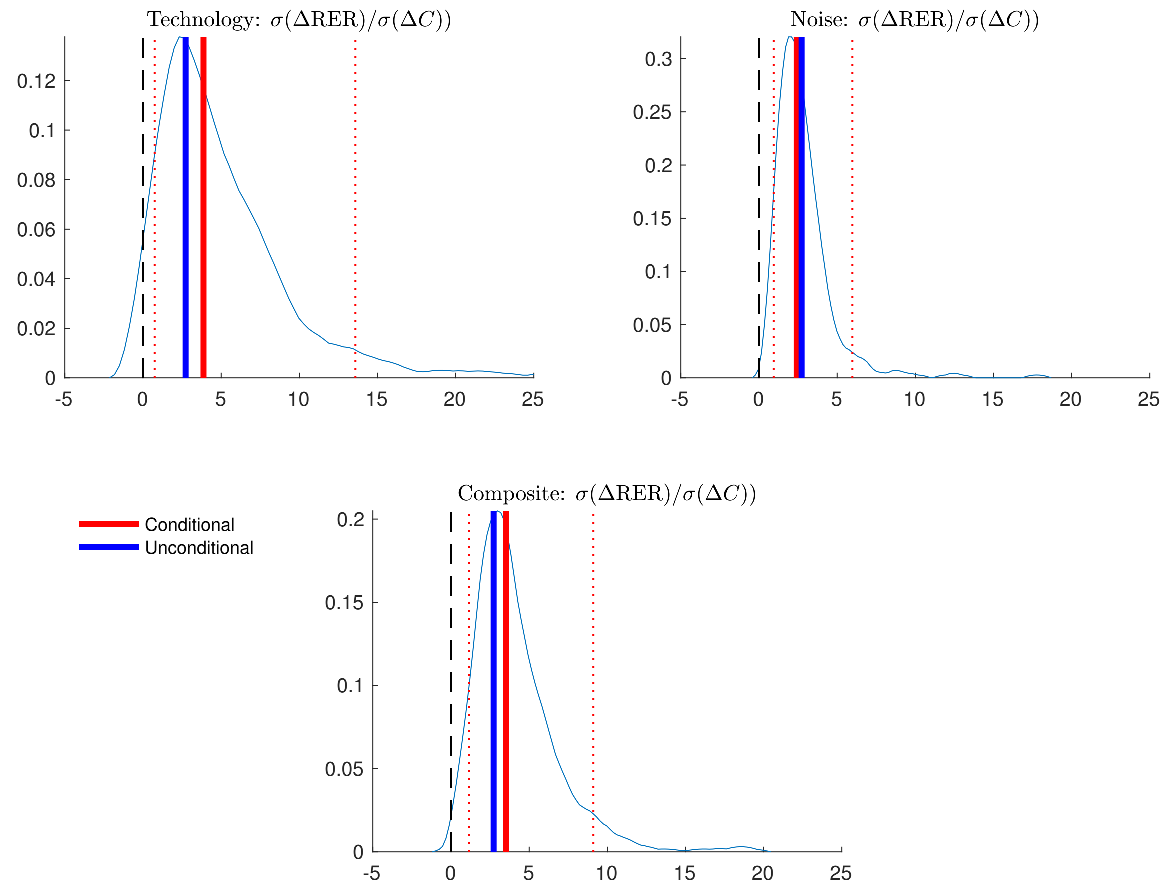

Those two moments speak to the fact that (i) the exchange rate is much more volatile that macro aggregates and (ii) there is virtually no contemporaneous correlation. The variance ratio in the raw data is roughly 3, while the contemporaneous correlation is zero. We obtain similar results if we substitute other macro aggregates, like output, investment or TFP, for consumption, so in the text we only talk about consumption.

We compare these moments to those emerging from the counter-factual simulation based only on the productivity and expectational shocks we recover. In Figure 12, we show the estimated posterior distribution of the variance ratio $\frac{Var(\Delta q_t)}{Var(\Delta c_t)}$. The bottom panel shows that the median estimate of the variance ratio aligns virtually perfectly with the raw data (exchange rate growth rates being three times more volatile than consumption growth rates). There is not much difference in the results when conditioning on only one of the two identified shocks either, as we can see from the top two panels.

The correlation $corr(\Delta q_t, \Delta c_t)$ is roughly similar to the Backus-Smith correlation, al-

Figure 13: FX-Consumption correlation

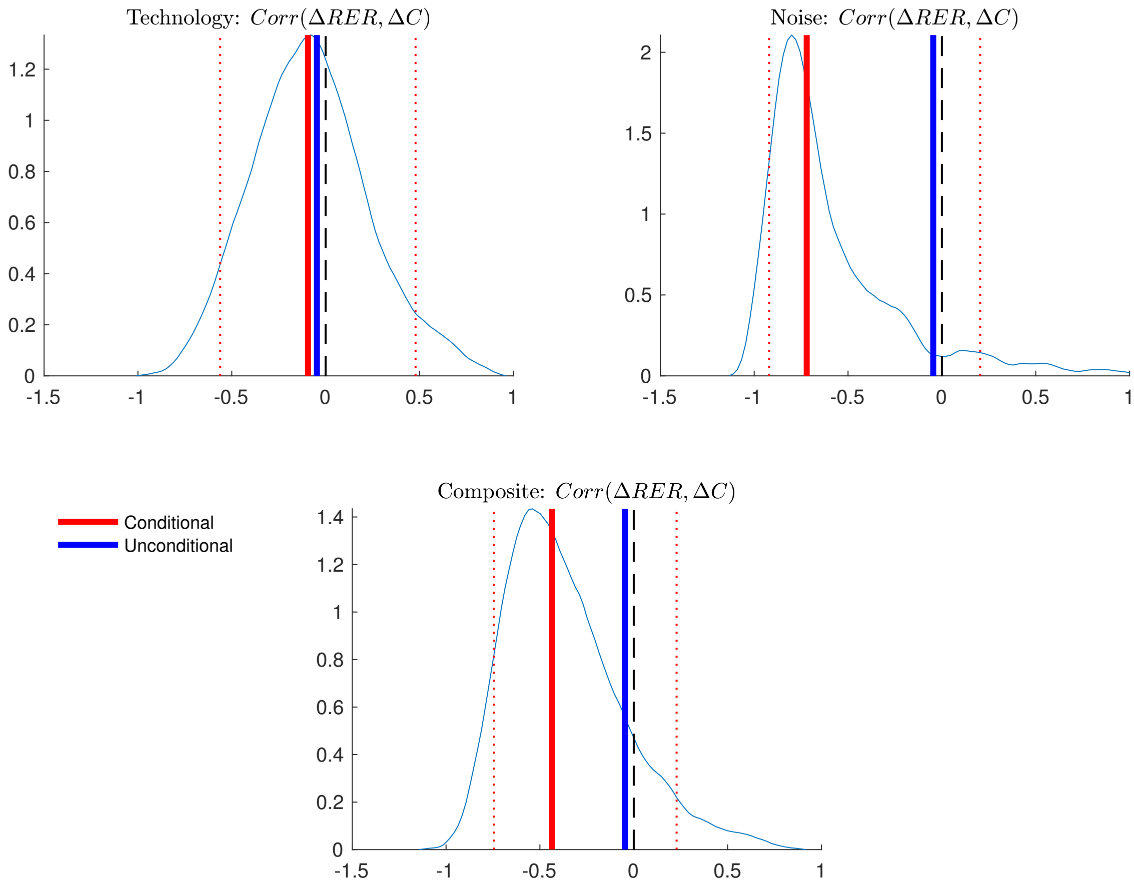

though it is not the same moment as it does not also take into account foreign consumption growth.

Takeaways

Overall, our results indicate that the great majority of exchange rate variation is in fact closely connected to macroeconomic fundamentals, however, the connection is primarily with *future* fundamentals. Moreover, we find that this connection appears to run primarily through a mechanism of imperfect foresight about future TFP.

Interestingly, the link between TFP fundamentals and noise run through UIP deviations, as those two shocks cause significant fluctuations in exchange rates, but mainly through their impact on expected currency returns. Combined with our additional results that the two identified shocks are also significant drivers of the Backus-Smith and exchange rate determination puzzles, this suggests that imperfect foresight might be a common source of both exchange rate fluctuations and puzzles.

Thus, there is promise that a model driven by the single mechanism of noisy expectations regarding TFP, can help explain a number of exchange rate puzzles, and generate empirically realistic exchange rate volatility.

5 Towards a model

The obvious question is if model of noisy expectations can indeed generate the empirical patterns we uncover. The short answer, is that this depends on the specific mechanism for endogenous UIP deviations. Endogenous UIP deviations are crucial, as the empirical results suggest that the effects of both ε_t^a and ε_t^v run through excess currency returns.

[IN PROGRESS]

References

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.

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

Bacchetta, P. and E. Van Wincoop (2010): "Infrequent portfolio decisions: A solution to the forward discount puzzle," *The American Economic Review*, 100, 870–904.

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.

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

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

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

Chahrour, R. and K. Jurado (2019): "Recoverability and Expectations-Driven Fluctuations," Tech. rep., mimeo.

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.

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

Corsetti, G., L. Dedola, and S. Leduc (2010): "Optimal monetary policy in open economies," *Handbook of Monetary Economics*.

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.

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.

Engel, C. (1999): "Accounting for US real exchange rate changes," *Journal of Political economy*, 107, 507–538.

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

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," Tech. rep., National Bureau of Economic Research.

Fernald, J. (2014): "A quarterly, utilization-adjusted series on total factor productivity," Federal Reserve Bank of San Francisco.

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

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

Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2020): "International comovement in the global production network," .

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 (2019): "Exchange rate disconnect in general equilibrium," Tech. rep., National Bureau of Economic Research.

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.

Kollmann, R. (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.

Lilley, A., M. Maggiori, B. Neiman, and J. Schreger (2019): "Exchange rate reconnect," Tech. rep., National Bureau of Economic Research.

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.

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.

Rebelo, S., M. Eichenbaum, and B. Johannsen (2019): "Monetary Policy and the Predictability of Nominal Exchange Rates," .

Rogoff, K. S. and V. Stavrakeva (2008): "The continuing puzzle of short horizon exchange rate forecasting," Tech. rep., National Bureau of Economic Research.

Sarno, L. (2005): "Towards a solution to the puzzles in exchange rate economics: Where do we stand?" *Canadian Journal of Economics/Revue canadienne d'économique*, 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.

Siena, D. (2015): "News Shocks and International Business Cycles," .

Stavrakeva, V. and J. Tang (2018): "Survey-based exchange rate decomposition: New methodology and new facts," Tech. rep., 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.