

Source Matters: Demand vs. Supply Uncertainty in a Small Open Economy

Matthew DeHaven *

October 28, 2026

[Link to latest available version.](#)

Abstract

Using UK newspaper text, I separate economic uncertainty into demand and supply uncertainty using a large language model to label article topics. I find that demand uncertainty depreciates the pound and raises the currency risk premium on the pound relative to the dollar, while supply uncertainty appreciates the pound and lowers the currency risk premium. I rationalize the results using a New Keynesian small open economy model with stochastic volatility, recursive preferences, and incomplete international risk sharing. Under demand shocks a foreign-currency bond shares risk, but under supply shocks it raises risk, so rising uncertainty about each moves the domestic currency risk premium in opposite directions. Incomplete international risk sharing lets the sign depend upon the source, while pricing frictions link the nominal and real responses, which move together in the data. Level shocks cannot generate the observed changes in risk premia, so the data point to time-varying risk. An aggregate measure of uncertainty masks these opposite responses.

Keywords: Time-varying uncertainty, exchange rates, currency risk premium, small open economy, New Keynesian, text analysis

JEL Codes: E32, E44, F41, G15

*Brown University. Email: matthew.dehaven@brown.edu. Website: matthewdehaven.com. I am grateful for the advice and support of my committee, Şebnem Kalemli-Özcan, Gauti Eggertsson, and Amy Handlan. I would also like to thank Fernando Duarte, Pablo Ottonello, Fabrizio Perri, Paula Donaldson, Alexandre Gaillard, Spencer Kwon, Olivia Lattus, Can Soyly, Finn Schüle and the entire Brown Macro Workshop group for comments and feedback. Finally, I wish to thank the participants and organizers at the Federal Reserve Board for their feedback during the Dissertation Fellowship, particularly Dario Caldara, Nils Gornemann, and my mentors, Sai Ma and Francesco Ferrante. All errors remain my own.

1 Introduction

Uncertainty shocks are known to have recessionary effects on the macroeconomy (Bloom 2009, Fernández-Villaverde et al. 2011, Leduc and Liu 2016, Basu and Bundick 2017). Global uncertainty appreciates the US dollar and depreciates risky currencies as investors move assets into the US for safety (Akinci, Kalemli-Özcan, and Queralto 2022, Kekre and Lenel 2024). However, some small open economy currencies, like the pound, are neither inherently risky nor safe. In addition, these economies face not only global shocks, but a substantial amount of domestic uncertainty to which the global ranking of risk does not apply. How does domestic uncertainty affect the currency of a small open economy?

In this paper, I show that the *source* of domestic uncertainty matters in a small open economy: demand uncertainty leads to a depreciation and a rise in the currency risk premium while supply uncertainty leads to an appreciation and a fall in the currency risk premium. I estimate this empirically by developing a novel uncertainty measure from UK newspaper text that decomposes uncertainty into demand and supply sources using a large language model. Then, in a small open economy model, I show how the exchange rate response is determined by households’ pricing of currency risk. What matters is whether the foreign bond pays off well in domestic currency in periods when domestic households most value that payoff. Under positive demand shocks, the domestic currency depreciates in exactly those moments, so the foreign bond pays more when they need it—it shares risk internationally. Under negative supply shocks, the domestic currency appreciates in exactly those moments, so the foreign bond pays less when they need it—it adds to their risk rather than sharing it. When demand uncertainty increases, households value the foreign bond even more for its risk-sharing properties, so they sell their currency to buy it, depreciating the exchange rate today and raising the currency risk premium. The reverse happens for supply uncertainty. Without decomposing uncertainty into demand and supply uncertainty, this distinction would be lost, and the aggregate response would be a mixture of the two effects.

I present empirical evidence of the different response to demand and supply uncertainty in the UK. The UK is an ideal setting as a developed small open economy that is neither the global reserve currency nor an inherently risky currency and has a recent history of large domestic uncertainty shocks. I construct the measure using text analysis of UK newspapers from 2000 to 2025 to classify articles as discussing: (1) the economy and uncertainty and (2) demand vs. supply topics. In addition to the demand and supply topics, I label some articles as relating to policy and financial uncertainty topics, which I use as controls. The topics are labelled by a large language model (LLM), which captures the full context of the article better than keyword searches. Against hand-labelled articles, the LLM’s precision is 96%

for identifying economic uncertainty, and 62% and 88% for the demand and supply topics, compared to 28% and 17% using keyword searches. My measure of overall uncertainty in the UK economy is similar to the estimates of UK economic policy uncertainty by [Baker, Bloom, and Davis \(2016\)](#), but my measure allows me to distinguish the source of economic uncertainty.

I show that a one-standard-deviation innovation in demand uncertainty leads to an increase in the currency risk premium of 59 basis points (annualized) and a depreciation of the exchange rate of 0.80% at a horizon of eight months. A one-standard-deviation innovation in supply uncertainty leads to a decrease in the currency risk premium of 66 basis points (annualized) and an appreciation of the exchange rate of 0.69% at a horizon of eight months, which continues to build. I measure the currency risk premium using the interest differential between the UK and the US minus the depreciation expected by survey forecasters. In my estimations, I control for global uncertainty (the VIX), UK and US macroeconomic conditions, and UK financial and policy uncertainty, and I interpret the remaining innovations as causal under the assumption that, given these controls, they are uncorrelated with other structural shocks. In robustness checks, I show that the results are stable to controlling for sentiment from the same newspaper articles, excluding Covid, and using alternative currency premium measures.

I develop a small open economy New Keynesian model along the lines of [Galí and Monacelli \(2005\)](#), with three additions: (1) stochastic volatility shocks to preferences and productivity, (2) recursive preferences, and (3) incomplete international financial markets, where households can only purchase a single risk-free foreign bond. In the model there are four shocks: a level shock and a volatility shock to each of preferences and productivity. I map demand uncertainty in the data to the volatility of preference shocks, which shift how much households want to consume today. Also, I map supply uncertainty in the data to the volatility of productivity shocks, which shift how much output is created for each unit of labor. Households price the risk of holding foreign currency by how that currency co-moves with their stochastic discount factor (SDF)—a high SDF means that the marginal utility of consumption tomorrow is high relative to their marginal utility today. Negative preference shocks decrease the households' SDF because a negative preference shock lowers their desire to consume, while negative productivity shocks increase the SDF because goods are scarce. A negative preference shock and a negative productivity shock both appreciate the currency in the model but move the SDF in opposite directions, so their respective volatility shocks have opposite-signed currency risk premia.

I take a risk-adjusted log-linear approximation of the model to preserve the variance terms that drive the three risk wedges in the model: a precautionary savings wedge, a currency

risk wedge, and an inflation risk wedge. The inflation risk wedge is small and not important to the exchange rate response. In the calibration, the precautionary savings wedge leads to an appreciation of the currency for both shocks, regardless of the source, so it cannot explain the exchange rate’s behavior in the data. However, the currency risk wedge has a different sign for the two volatility shocks, which allows it to explain the different sign seen in the exchange rate. If markets were instead complete, the currency risk wedge would have a single sign and be tied to the price of risk, regardless of the source of uncertainty. Because the foreign bond cannot pay off differently across states, the risk wedge is compensation for risk that the domestic households cannot share, not an unexploited arbitrage.

Each model ingredient plays a distinct role in my results. The stochastic volatility processes are what actually shift the variance of shocks in the model and move the risk wedges; without them there would be no time-varying risk. The recursive preferences raise risk aversion without affecting the elasticity of intertemporal substitution, which allows me to make the wedges large enough to matter. The incomplete international financial markets are what make the currency risk wedge’s sign depend upon the source of uncertainty. The pricing frictions are not needed for the currency risk wedge to have opposite signs across the two shocks, but they are needed to link the real and nominal exchange rate responses. In the data, these two move very similarly, but without pricing frictions in the model the nominal exchange rate and real exchange rate would have opposite-signed responses for supply volatility shocks. In addition, the pricing frictions are what make the uncertainty shocks recessionary not just for consumption, but also for output and hours worked, as shown in [Basu and Bundick \(2017\)](#). Incomplete risk sharing sets the sign of the currency response and pricing frictions link it to the nominal exchange rate and output.

First-moment shocks cannot explain the currency risk premium responses. In the model, the currency risk premium depends only upon the volatility states in the economy. Level (first-moment) shocks cannot affect the volatility states, and so cannot move the risk wedges, including the currency risk premium. Thus, through the lens of the model a change in the currency risk premium necessitates a change in the perceived risk of one or both of the shocks. While a first-moment productivity shock can match the non-risk premium empirical responses, the first-moment preference shock does not generate the observed depreciation.

The model clarifies that an unconditional correlation of zero between the SDF and exchange rate can hide strong correlations conditional on each uncertainty shock. Under incomplete markets and when only volatility shocks hit the economy, the correlation runs from negative one under only productivity volatility to positive one under only preference volatility. A mix of the two shocks, however, can generate an unconditional correlation at any value in between, including zero. If markets were complete instead, the correlation would always

be negative one. This decomposition of uncertainty can thus help reconcile the observed weak unconditional comovement between the exchange rate and marginal utility ([Backus and Smith 1993](#)) with strong comovement conditional on the source of uncertainty.

The rest of the paper proceeds as follows. Section 2 presents the empirical measure, its validation, and estimates. Section 3 introduces the model. Section 4 discusses the risk wedges in the model. Section 5 presents the calibrated model, its fit to the data, and a comparison of risk sharing. Section 6 concludes.

Literature Review This paper contributes to three strands of the literature: the effects of uncertainty shocks in closed economies, the effects of uncertainty shocks in open economies, and measuring uncertainty shocks.

First, my paper relates to the study of uncertainty shocks in closed economies ([Bloom 2009](#), [Christiano, Motto, and Rostagno 2014](#), [Fernández-Villaverde et al. 2015](#), [Leduc and Liu 2016](#), [Bloom et al. 2018](#), [Ludvigson, Ma, and Ng 2021](#)). In particular, my paper is complementary to the work of [Bianchi, Kung, and Tirsikh \(2023\)](#), who study the different effects of demand and supply uncertainty in a closed economy. They leverage their model and the term structure to back out demand and supply uncertainty from the data, while I introduce new empirical estimates of demand and supply uncertainty separately from the model. They study the response of firm investment, real activity, and the term premium in a closed economy, and find different responses in terms of magnitude, while I study the response of the exchange rate in an open economy and find a different response in sign. I build on their decomposition of the model into analytical expressions for the endogenous risk wedges. As in [Basu and Bundick \(2017\)](#), Rotemberg pricing frictions make uncertainty shocks recessionary for output and hours; in the open economy, because only domestic prices are sticky, they also govern the terms of trade and thus how closely the nominal exchange rate follows the real exchange rate.

Second, my paper relates to the study of uncertainty shocks in open economies ([Fernández-Villaverde et al. 2011](#), [Carrière-Swallow and Céspedes 2013](#), [Kalemli-Özcan, Soylu, and Yildirim 2026](#)). The earlier work of [Benigno, Benigno, and Nisticò \(2012\)](#) shows that real vs. nominal uncertainty can have different impacts in a two-country open economy model through risk sharing; my mechanism builds on the same intuition. They find that real productivity volatility, measured with the VIX, leads to a depreciation of the dollar and little impact on the currency risk premium. In their two-country model, complete markets still let the currency risk premium change sign because it reflects a gap between foreign and domestic prices of risk. In a small open economy with a foreign SDF unaffected by domestic shocks, complete markets tie the currency risk premium to the domestic price of risk alone, and

incomplete markets are what let its sign depend upon the source of the shock. My empirical estimates control for the VIX—which is also a measure of risk aversion (Bekaert, Hoerova, and Lo Duca 2013) and the global financial cycle (Miranda-Agrippino and Rey 2020). Also, I measure supply uncertainty with newspaper text data and find an appreciation of the pound and a significant fall in the currency risk premium. The work of Akinci, Kalemli-Özcan, and Queralto (2022) introduces a model whereby the US dollar appreciates in response to uncertainty shocks due to the role played by US financial intermediaries in global asset markets. In my model no such agent exists—the premium is priced by domestic households’ own conditional covariance. I show that different sources of domestic uncertainty have opposite impacts on a currency that is not inherently safe or risky. A large literature finds that exchange rates and interest differentials are largely driven by time-varying currency risk premia (Alvarez, Atkeson, and Kehoe 2009, Engel 2016). Without relying on segmented markets, I provide a measured source of that variation, domestic demand and supply uncertainty. In addition, my results relate to some of the facts about the UIP premium shown in Kalemli-Özcan and Varela (2021). They show that local risk drives much of the currency risk premium, and link that risk to country-specific policy shocks. I contribute to this work by showing how additional sources of local uncertainty matter, and how the sign response of the currency risk premium depends upon that source.

Third, my paper relates to the literature on the measurement of uncertainty (Bachmann, Elstner, and Sims 2013, Bekaert, Hoerova, and Lo Duca 2013, Jurado, Ludvigson, and Ng 2015, Scotti 2016, Cascaldi-Garcia et al. 2023). The seminal paper of Baker, Bloom, and Davis (2016) introduced newspaper text-based measures of economic policy uncertainty by searching articles for keywords. I follow their approach by using newspaper text for the UK, but I rely upon a large language model for classifying articles instead of a keyword search. I compare my estimates of uncertainty to theirs, and find that my overall economic uncertainty index has a correlation of 0.89 with their economic policy uncertainty, but that masks a lot of variation in the economic topics of the uncertainty. Separating the demand and supply uncertainty is what shows the clear currency risk premium and exchange rate responses. There is a large literature that estimates different types of uncertainty shocks, usually separately, such as trade policy uncertainty (Caldara et al. 2020) and monetary policy uncertainty (Husted, Rogers, and Sun 2020). I extend this literature by jointly classifying the topics of uncertainty, which I will present in the next section on the empirical evidence from UK newspaper text.

2 Empirical Evidence: UK Newspaper Text

I build new newspaper text-based measures of uncertainty for the United Kingdom by classifying articles into economic topics with a large language model. I find that in response to an innovation in overall economic uncertainty, the pound depreciates while the premium slightly falls. However, when I disaggregate uncertainty into demand and supply uncertainty, I find in response to an innovation in demand uncertainty, the pound depreciates and the currency risk premium rises, while in response to an innovation in supply uncertainty the pound appreciates and the currency risk premium falls.

2.1 Measuring Uncertainty: Overall, Demand, and Supply

UK uncertainty is measured by labelling articles that discuss the economy and uncertainty with a large language model (LLM) into different economic topics, which I group into demand, supply, financial, policy, trade, labor, politics, and other uncertainty. The resulting overall index tracks the estimate of economic policy uncertainty from [Baker, Bloom, and Davis \(2016\)](#) and the demand and supply indexes peak at different, recognizable events.

The UK is a prime candidate for this research because it is a small open economy with a floating exchange rate that is widely traded and forecasted and has a large English-language newspaper archive. Importantly, the pound is not the world’s reserve currency and so we cannot predict its behavior based off of a global flight-to-safety. Existing uncertainty measures in the UK do not capture the distinction between supply and demand as sources of uncertainty. The VFTSE, the option-implied expected volatility of the UK FTSE returns, is a common measure of financial market uncertainty. It does not lend itself to be split into different sources of uncertainty and comoves largely with the US VIX (correlation of 0.93 over my sample). The economic policy uncertainty (EPU) index of [Baker, Bloom, and Davis \(2016\)](#) is also constructed for the UK, but this was specifically built to capture policy uncertainty. This index, however, does lend itself to being split into different sources, as it is built using newspaper articles, and the construction of my indexes will build upon their work. The newspaper text is also ideal for this decomposition into uncertainty sources, because articles not only describe the economy and uncertainty, but they also discuss *about what* people are uncertain.

The newspaper text is accessed via ProQuest TDM Studio. I make use of twelve UK publications, some of which publish daily and some weekly, and construct the monthly index from January 2000 to December of 2025.¹ Several publications enter or exit over the sample.

1. The UK papers are the Daily Telegraph, Daily Mail, Guardian, Independent (print + online), Observer, Daily Mirror, Daily Express, Sunday Telegraph, Mail on Sunday, Independent on Sunday, Express on Sunday,

I will account for this when aggregating across papers. For a newspaper article to be included in my candidate set, I require that it contain the terms “uncert*” and “econom*” where the * matches any possible ending to those terms.² This is a very similar initial requirement to Baker, Bloom, and Davis (2016) and leaves me with 52,512 articles over the sample.

I then take each candidate article and feed it into a large language model (GPT-4o mini) through the API on ProQuest TDM Studio to label the following features of the article:

1. a score of 0, 1 or 2 for how much the article discusses the economy, where 2 is high,
2. a score of 0, 1 or 2 for how much the article is about uncertainty, where 2 is high,
3. the topic of the economic uncertainty in the article from the following topics: consumer, housing, output, living costs, energy, inflation, supply chains, finance, investment, monetary policy, fiscal policy, regulation, trade, jobs, wages, strikes, politics, geopolitics, other.

See the Appendix B.1 for the full prompt sent to the LLM. I find that 34,533 (66%) of the articles receive a high score on both uncertainty and the economy and I keep these articles for the next steps.

I use an LLM for this classification step because the LLM is able to take in the context of the article when determining its topic. I attempted keyword searches and found that articles about the economy often mention key terms from both demand and supply sets of terms, making it difficult to separate the two concepts. In addition, keyword searches cannot reliably distinguish an economic use of a term, like “demand”, and other linguistic uses, like “The UK demands the end of the war”. To make this point clear, I present the following example headlines from articles that were in the initial candidate set, but that the LLM deemed off-topic: “Mourinho’s new weapon: he travels economy and lets players go first class” from the Daily Mail, August 2013 and “Obituary: Michelangelo Antonioni... explored the uncertainties between time and space” from the Guardian, August 2007. The first article clearly uses “economy” to describe the class of airplane seats and the second discusses “uncertainty” explored by a film director. The LLM correctly labels each article a zero on the economy and on economic uncertainty, which leads to a final index with more signal and less noise.

I then group economic topics into various sources of uncertainty. I combine housing, consumer, output, and living costs into “demand” (20% of articles) and energy, inflation, and Sunday Mirror. I attempted to include the Sun, but had to drop it because I did not have access to the full article text. For comparison, Baker, Bloom, and Davis (2016) used the Times of London and the Financial Times accessed through a different service to construct their UK economic policy uncertainty measure.

2. This means that “uncert*” would match all of “uncertain”, “uncertainty”, and “uncertainly”.

and supply chains into “supply” (7% of articles). Investment and finance make up the “financial” (25% of articles) group and fiscal, monetary, and regulation make up “policy” (13% of articles). I leave trade as its own topic (9% of articles). Politics and geopolitics get grouped as “politics” (11%). Finally, jobs, wages, and labor strikes get grouped as “labor” (15%), the rest are left ungrouped, as “other” ($< 1\%$ of articles).³ These groups are determined in part based off of topics that generate similar responses. In addition, in Appendix B.4, Figure B11 shows that the results estimated in the next sections are robust to narrower and broader definitions of the demand and supply groups. In particular, it is difficult to determine if uncertainty about labor should be included in the demand or supply groups; it is possible that further disaggregation is needed for this topic. In Appendix B.4, Figure B12 shows that the results estimated in the next sections are robust to assigning all labor articles to either demand or supply. The goal is that demand articles are about consumer spending, confidence, the housing market, firms’ sales and orders, and household specific pricing concerns. The supply group captures input costs for firms, most notably energy costs and supply chain disruptions, and uncertainty about prices.

The newspaper articles are then aggregated into a monthly index. To do so, I construct a separate count of the total articles released by each newspaper in each month, dropping days where twenty or fewer total articles are published for that newspaper.⁴ I then take the ratio of articles labelled by the LLM to a specific uncertainty topic by newspaper for each economic topic group, g , standardize it for each newspaper, average across the newspapers, and then standardize the final indexes:

$$s_{g,n,t} = \frac{\text{Matched Counts}_{g,n,t}}{\text{Total Counts}_{n,t}}, \quad (1)$$


$$I_{g,t} = \frac{\frac{1}{N_t} \sum_{n \in N_t} \frac{s_{g,n,t} - \bar{s}_{g,n}}{\sigma_{g,n}} - \bar{I}_g}{\sigma_g^I}. \quad (2)$$

where $g \in \{\text{Overall, demand, supply, financial, policy, trade, labor, politics, other}\}$. This method accounts for the fact that some newspapers may discuss certain topics with a higher frequency and that some newspapers have many more articles than the others. In addition, averaging across the newspapers available in each period accounts for the fact that sometimes a publication will drop in or out of the sample. The final units of the indexes are normalized, so they have a mean of zero and a standard deviation of one. The resulting overall uncertainty index is shown in Figure 1 with the share of its articles assigned to each economic group shown below.

3. Other includes articles the LLM assigned to other or gave no listed topic.

4. This total count measure captures all types of articles, not just economic or uncertain ones.

Figure 1: Overall Economic Uncertainty and Its Decomposition by Topic



Note: Panel (a) shows the overall economic uncertainty index standardized to a mean of zero and standard deviation of one in yellow, with the six-month moving average of that series in black. The dashed line shows the average value from 2004 to 2007, which was a period of low uncertainty. Each vertical dashed line shows an event, starting from the left: the 9/11 terrorist attack in the US, the beginning of the Iraq war, the Global Financial Crisis (GFC), the European sovereign debt crisis, the vote for “Brexit” (Britain leaving their EU relationship), a delay in a key Brexit vote, the Covid-19 pandemic, the finalized implementation of a Brexit deal, the “mini-budget” crisis in Liz Truss’ administration, and the “Liberation Day” US tariff announcements. Panel (b) shows the share of articles in the overall uncertainty index that are assigned to each topic. The demand topic is in blue and supply in orange. After that, from the bottom, the topics are: financial, policy, trade, labor, and politics. A small number of “other” articles remain ungrouped. Full topic groups can be found in Table B1. Each topic is assigned by an LLM reading articles that mention economic uncertainty.

Figure 1 shows that overall uncertainty spikes during many well-known events in global and UK history. The earliest spikes are during 9/11 and the Iraq war invasion, which the UK participated in. Next, there is a spike for the GFC and then the ensuing Euro sovereign debt crisis towards the end of 2011. The Brexit vote is the largest single uncertainty event in the overall measure, and is followed by additional spikes as Brexit votes are debated and delayed, until the end of 2020, when the trade deal is finalized. The mini-budget crisis, when Liz Truss’ administration introduced a surprising government budget and tax plan and ended up losing her the office in less than two months, shows a large spike in uncertainty. The second largest increase occurs in April of 2025, when Trump announces “Liberation Day” tariffs.

Perhaps surprisingly, Covid is a spike but not massive in this figure. If we follow the line down to the shares of uncertainty, however, we can see that there are shifts between the topics of uncertainty that are not captured in the overall index. For Covid, there is a large increase immediately in uncertainty about labor and jobs (and a fall in the discussion of politics). In this share decomposition, it is also clear that the spike for Trump’s tariffs is captured by the trade uncertainty topic. Finally, notice that a year after Covid, there is an increase in supply uncertainty—this is the beginning of the supply chain disruptions and then the energy crisis in Europe when Russia invaded Ukraine in early 2022.

Table 1: Correlation of Text-Based Uncertainty Indexes with UK EPU and VFTSE


Text index	UK EPU	VFTSE
Overall uncertainty	0.89	0.25
Demand	0.64	0.30
Supply	0.26	-0.01
Financial	0.61	0.46
Policy	0.73	0.03
Trade	0.64	-0.10
Politics	0.57	-0.05
Labor	0.52	0.39
UK EPU		0.04

Note: Overall economic uncertainty and its decomposed topic uncertainty indexes are compared to the UK economic policy uncertainty (EPU) index of [Baker, Bloom, and Davis \(2016\)](#) and the stock market implied volatility index for the FTSE, the VFTSE. The sample period is 2000 to 2025.

Each of the economic uncertainty sources has their own uncertainty indexes. I compare these measures with measures of uncertainty that already exist for the UK: the economic policy uncertainty (EPU) index and the implied volatility of the FTSE stock market (VFTSE). Table 1 shows the correlation of the overall uncertainty index I constructed and each of the

group indexes with these two measures. Overall uncertainty is highly correlated (0.89) with the estimate of UK EPU. This is in part because the EPU is not able to fully distinguish the topic of economic uncertainty through its methods (see Section 2.2). The second most correlated series with the EPU is the policy group (0.73), so my measure of that topic aligns with the EPU. On the right side, the VFTSE is less correlated with these textual measures—in fact it has almost zero correlation with the EPU. Out of the groups of topics, the VFTSE is most correlated with the financial group, which corroborates that topic is picking up financial uncertainty. The supply uncertainty measure has the lowest correlation out of all the topics with either the EPU or the VFTSE index and seems to be picking up a new source of information.

Figure 2: Supply and Demand Uncertainty Indexes



Note: Demand uncertainty index (in blue) and supply uncertainty index (in orange) are constructed from UK newspaper text from articles that are labelled by an LLM. This figure shows the six-month moving average for each index over time. The horizontal dashed line shows the average value of both series from 2004 to 2007 when uncertainty was low. The monthly indexes have a correlation of 0.22 and the six-month moving averages have a correlation of 0.3. See Figure 1 for the full names of the events.

Figure 2 shows the demand and supply uncertainty indexes over time, both standardized to have a mean of zero and standard deviation of one. The two indexes spike at many of the same events discussed for the overall uncertainty index. In 2016, there is a large spike around the Brexit vote and the fallout in the demand uncertainty measure. The second highest demand spike is during the global financial crisis, with little response of supply in this period. The demand index spikes at the start of the Covid crisis, while the supply index rises with the 2021 supply-chain disruption and peaks with the energy crisis in 2022. The two series have a correlation of 0.22 with each other as raw measures (and 0.3 as six-month moving averages, as in Figure 2), which shows they capture different information. This will

be leveraged in the next section to discuss the effect of an innovation in either demand or supply uncertainty.

2.2 Validation of Uncertainty Indexes

I perform a hand-checked validation of a subset of 121 articles that I label without knowing the LLM’s response.⁵ I chose 5 random articles per topic and 26 articles over the LLM’s different economic and uncertainty labels. I then compare how the LLM performs at matching my labels (validation Table B2 is in the Appendix). Out of the 105 articles the LLM labelled as about economic uncertainty, I agreed on 101 (96%). On the topic groups, there was more disagreement: out of the 95 articles the LLM labelled with specific topics, I agreed on 61 (64%). Specifically on the supply topic, I agreed on 12 out of the 15 articles (80%). And on the demand group, I agreed on 12 out of the 20 articles that the LLM chose (60%).

To give context to the performance of the LLM, I want to compare it to a common prior approach: keyword searching. Before running the LLM, I generated a list of demand and supply topic keywords to cover the concepts I was trying to capture. In Table B3 I compare the precision—of the articles each method flags, the share my hand labels agree with—of the LLM and the keyword searches.⁶ On labelling an article about economic uncertainty, the LLM has a precision of 98% while using “econom*” and “uncert*” only gives a precision of 87%. On topics, I find the LLM has a precision of 62% on the demand topic, while a keyword search has a precision of 28%. For the supply topic, the LLM has a precision of 88%, and keywords only 17%. Comparing the LLM to the policy keyword searches of [Baker, Bloom, and Davis \(2016\)](#), I find the LLM’s precision is 62% while their keywords are 15%. The large gap in the keyword searches’ precision comes from the vast number of articles that they incorrectly label as matching the search (false positives). The policy search labels 71% of all articles as being about economic policy, while the LLM labels 9%, and in my sample I find 12%. An example of a misidentified article: “China’s zero-Covid policy threatens global supply chain, says Huawei exec” matched the demand search because it mentioned “consumer”, even though the LLM and I both correctly labelled it as a supply topic.

2.3 Isolating Demand and Supply Uncertainty Innovations

The innovations to the demand and supply uncertainty indexes are the variation that cannot be explained by lagged values of all of the uncertainty indexes or by current and lagged values

5. I plan to extend this pilot sample.

6. For the precision numbers I weight each stratified topic sample to the overall number of articles in that pool.


of US and UK macro conditions and the US VIX. Under the assumption that given these controls the remaining variation is uncorrelated with other structural shocks, the responses to the innovations can be interpreted as causal estimates.

To isolate innovations in topic uncertainty from endogenous responses to UK macroeconomic conditions and US macroeconomic conditions and global risk, I regress the index $I_{g,t}$ on the lags of all four indexes and on current and lagged controls X_t :

$$I_{g,t} = \alpha_g + \sum_{i=0}^p \beta_{g,i} X_{t-i} + \sum_{k \in \{d,s,f,p\}} \sum_{i=1}^p \gamma_{g,k,i} I_{k,t-i} + \epsilon_{g,t}, \quad (3)$$

with $p = 3$ lags, for each topic $g \in \{\text{demand, supply, financial, policy}\}$. The controls include UK real GDP, UK CPI, UK three-month rate, US real GDP, US CPI, the US three-month rate, and the VIX.⁷ The sample is April 2000 to December 2025; the first months are lost to the lags. The residuals $\epsilon_{g,t}$ are the uncertainty innovations, and the demand and supply uncertainty innovations are shown in their raw series in Figure 3.

Figure 3: Supply and Demand Uncertainty Innovations



Note: Demand uncertainty innovations (in blue) and supply uncertainty innovations (in orange) are the remainders from estimating equation (3). These two series have a mean of 0 and are in units of the original uncertainty index, which had a standard deviation of one. The innovations have a correlation of 0.09. See Figure 1 note for the descriptions of the events.

The innovations are in standardized units of the original uncertainty indexes. The correlation between the demand and supply innovations drops to 0.09 after controlling for the lags and other variables. The VIX is included to account for a common global uncertainty factor. These innovations can now be used to estimate local projections.

7. UK has monthly GDP estimates, so I interpolate US GDP to a monthly frequency.

Causal Interpretation Under the assumption that these innovations are uncorrelated with any leads or lags of other structural shocks, they can be interpreted as shocks to each of the uncertainty indexes and they can be used to identify the causal response to demand and supply uncertainty. This follows the framework of [Jordà \(2005\)](#), [Stock and Watson \(2018\)](#), [Plagborg-Møller and Wolf \(2021\)](#), and the inference assumptions of [Montiel Olea and Plagborg-Møller \(2021\)](#). The innovations are serially uncorrelated and the local projection will control for lags of each innovation and outcome to capture any remaining serial correlation. The current controls account for current month co-movement between uncertainty shocks and these macroeconomic variables; the specification conservatively attributes any current month movements to the controls, rather than to the uncertainty measures. In addition, the controls include a global factor, the VIX, and other local uncertainty factors, the UK financial and policy uncertainty indexes, which are likely to comove with the demand and supply uncertainty. There could be possible omitted variables that would invalidate this interpretation: one possible concern could be news about the economy (i.e. not reflecting uncertainty). In a robustness check, in [Appendix B.4](#), I control for the sentiment of the news and find the resulting innovations have a correlation of 0.99 with the baseline ones. If the results are not interpreted causally, they can serve as conditional correlations that will be used to discipline the model.

2.4 Local Projections

Demand, supply, financial and policy uncertainty innovations are included jointly in a local projection regression that traces out the responses of the exchange rate, the currency risk premium, the UK-US short interest rate differential, and the UK-US price level differential over a 24-month horizon. The currency risk premium is measured as the interest differential between a UK and US government bond minus the depreciation of the pound expected in market surveys.

I estimate impulse responses with the local projection method of [Jordà \(2005\)](#) separately for each horizon h :

$$y_{t+h} = \alpha_h + \sum_{g \in \{d,s,f,p\}} \beta_{g,h} \epsilon_{g,t} + \sum_{g \in \{d,s,f,p\}} \sum_{i=1}^{q-1} \eta_{g,h,i} \epsilon_{g,t-i} + \sum_{i=0}^{q-1} \gamma_{h,i} Z_{t-i} + v_{t+h}, \quad (4)$$

where y_{t+h} is the outcome h months ahead, with $q = 3$ lags of the innovations and all of the dependent variables, Z_t , over horizons $h = 1, \dots, 24$. The coefficients $\beta_{g,h}$ can then be read off from the regressions and plotted with their confidence intervals, with Newey-West standard errors. These can be interpreted as one-unit shifts in the residuals of the standardized

uncertainty series, so they are a one standard deviation shift of each uncertainty series.

One local projection is estimated with all four uncertainty topic innovations included at once. While each innovation has low correlation with one another, the Appendix B.4 checks if there are any differences if the demand and supply uncertainty are estimated separately or in a Cholesky ordering. There is little difference in the estimated responses from these alternative specifications.

The dependent variables, $y_t \in Z_t$ are: GBP/USD exchange rate, the UK-US three-month rate differential, UK-US real GDP gap, UK-US CPI gap, and the 12-month currency risk premium. The exchange rate is pounds per dollar (so an increase is a depreciation of the pound) measured at the end of the month. It enters the regression as a logged value, multiplied by 100, so the responses can be interpreted as percent changes in the exchange rate. The 3-month rates are the yields on the government risk-free bonds of that maturity measured in percentage points, so the interest rate differential is in the same unit. The UK-US real GDP gap and UK-US CPI gap enter as log difference of each level, multiplied by 100, so a change in each can be interpreted as a percentage gap. The currency premium is measured in percentage points per year. The sample runs from January 2000 to December 2025 and all data are monthly. As a benchmark, I also estimate equations (3) and (4) with a single aggregate uncertainty measure in place of the topic uncertainty indexes: my overall economic uncertainty index, the UK EPU, and the VFTSE, each standardized and regressed one at a time.

Currency Risk Premium Measurement The currency risk premium is measured following Kalemli-Özcan and Varela (2021) using survey expectations of the exchange rate. The expected excess return, λ_t , of holding the pound vs. the dollar can be expressed in logs as:

$$\lambda_t = (i_t^{UK} - i_t^{US}) - 100 * (\log \mathbb{E}_t \mathcal{E}_{t+12} - \log \mathcal{E}_t), \quad (5)$$

where i_t^c is the one-year government yield in country c , and $\mathcal{E}_t$ is the spot exchange rate between those two countries, in pounds per dollar (an increase is a depreciation of the pound). The survey expectations are from Thomson Reuters (LSEG) and are the median of the survey responses for the expected exchange rate in 12 months. I use end-of-the-month values for the exchange rate as is common for financial market prices. To interpret λ_t as a risk premium, I assume that the surveys capture the market expectation for the exchange rate. When $\lambda_t > 0$, investors expect an excess return from holding the pound, which implies that there must be additional risk associated with holding the pound, and thus a positive

currency risk premium. When $\lambda_t < 0$, investors expect a loss from holding the pound, which implies that there is a negative currency risk premium. The 12-month horizon is the same horizon used in [Kalemli-Özcan and Varela \(2021\)](#).

I compare my estimates of the expected excess return of holding the pound vs. the dollar with the measures used by [Kalemli-Özcan and Soylu \(forthcoming\)](#), who make use of Consensus forecasts for the exchange rate instead of Reuters at the same horizon.⁸ The correlation between the two estimates is 0.72, showing that they are picking up much of the same movement. In Appendix B.2, in Figure B4, I show that my estimated local projections are robust to using the Consensus-based expected excess returns instead.

2.5 Empirical Results

The overall uncertainty measure is similar to the estimated responses of the UK EPU and VFTSE and shows a slight depreciation and negative currency risk premium. However, disaggregating uncertainty into demand and supply measures yields opposite-signed responses on the exchange rate and currency risk premium. The movement in the currency risk premium explains roughly half of the exchange rate movement in the first year. This currency risk premium response is the mechanism that will be formalized and explored in the model.


As a baseline, I estimate the local projection responses to three aggregate measures of uncertainty: my overall economic uncertainty index, the UK EPU, and the VFTSE. The estimated $\beta_{g,h}$ responses for the exchange rate and currency risk premium are shown in Figure 4, with a 68% confidence interval around each. The three sets of responses are very similar, especially the UK EPU and overall economic uncertainty responses. There is no consistent effect on the pound: a small depreciation reverses into an appreciation. The currency risk premium is negative, or insignificant throughout the horizon.

These results on the aggregate uncertainty measures show that there is not a clear answer to the question of the paper. Domestic uncertainty seems to have a small depreciatory impact, but only with a small, mostly negative currency risk premium. I will next show that these aggregate measures of uncertainty mask important exchange rate and currency risk premium responses in the separated uncertainty topics.

Figure 5 shows the responses to demand and supply uncertainty of the exchange rate and currency risk premium, with a 68% confidence interval. A one standard deviation demand uncertainty innovation depreciates the pound by 0.80% and increases the annualized currency risk premium by 59 basis points at an eight-month horizon. A one standard deviation supply uncertainty innovation appreciates the pound by 0.69% and decreases the annualized

8. I thank the authors and the Global Linkages Lab at Brown University for sharing their estimates.


Figure 4: Response of Exchange Rate to Aggregate Uncertainty Measures



Note: Each panel estimates a local projection described in equation (4) but uses only one measure of uncertainty innovations: Overall uncertainty from UK newspaper text, UK economic policy uncertainty from [Baker, Bloom, and Davis \(2016\)](#), or the UK implied stock market volatility (VFTSE). The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Confidence intervals are 68%.

currency risk premium by 66 basis points at an eight-month horizon. The demand uncertainty response depreciates over the first year, peaking at 1.05% at seven months, before drifting slightly past zero in the second year. The supply uncertainty response appreciates over a longer period, peaking at -1.25% at twenty months, before returning to zero at the end of the horizon.

Figure 5: Response of Exchange Rate to Demand and Supply Uncertainty



Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Confidence intervals are 68%.

These two disaggregated uncertainty measures produce clearly opposite-signed responses in the currency risk premium and suggestive opposite-signed responses in the exchange rate in the first year. This distinction is lost in the aggregate measures. The currency premium is clearly moving in Figure 5 and a natural question is how much of the exchange rate response does it explain?


The exchange rate response can be decomposed into three components: (1) shifts in the interest rate differential, (2) shifts in the currency risk premium, and (3) shifts in the exchange rate expectations—see equation (5). Over the first year for the demand uncertainty response, the interest rate differential accounts for 5%, the currency risk premium accounts for 51%, and the rest is changes in expectations. At the same horizon for the supply uncertainty response, the interest rate differential accounts for 9%, the currency risk premium accounts for 61%, and the rest is expectations.⁹ The currency risk premium response is playing an important role in determining the exchange rate response. When in response to a demand uncertainty shock, the pound depreciates by 0.8% at an eight-month horizon, the

9. These contributions are measured as the mean absolute contribution of each piece over the first twelve months.

one-year ahead survey forecast moves by only 0.2%, since the interest differential does not move much (0.01%), the gap must be that forecasters expect an appreciation in the future, and that expected appreciation is the compensation investors require for holding a riskier currency. This will be the key mechanism that is explored in more detail in the model.

Figure 6 shows the estimated $\beta_{g,h}$ responses for the UK-US short term rate differential and the UK-US CPI gap. The demand uncertainty shock is followed by a slight cut in the UK-US rate differential, followed by a return to zero. The price level has an insignificant response at first, but eventually leads to a fall in the price level after one year. On the other hand, a supply uncertainty shock is followed by an increase in the interest rate differential, and a small increase in the price level in the first year that is followed by a return to zero.

Figure 6: Local Projections: Policy Rate and CPI Responses



Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Confidence intervals are 68%.

These additional interest rate and price level responses will help to discipline the model responses. In particular, we see that neither shock generates a large policy or price level response. Since the price level gap between the UK and the US barely shifts, the real exchange rate has a very similar response to the nominal exchange rate, as is normal in the data. This will be an additional fact that the model will be able to reproduce with its pricing frictions.

The Appendix B.3 shows the response of output and industrial production in Figure B6, which are noisy but the GDP gap falls for some periods in response to demand uncertainty, and rises in response to supply uncertainty. The industrial production gap, on the other hand, increases for demand uncertainty, and is insignificant for supply uncertainty.

2.6 Robustness

The results on the currency risk premia are robust across every specification in the appendix, including controlling for newspaper article sentiment. The results on the exchange rate are sensitive to using the broad sterling index, which dampens the depreciation seen in response to demand uncertainty and to dropping the VIX as a control, which loses the appreciation of supply uncertainty, and is robust to the rest.

An objection to the results above is that there is not sufficient control for first-moment level shocks to the economy from demand or supply news. To check if this is the case, I construct sentiment indexes from the same newspaper articles that create the uncertainty indexes, but I create a sentiment score of positive minus negative terms over positive plus negative terms. These two indexes have a slight negative correlation with their respective uncertainty indexes (-0.23, -0.22), which shows that when an article discusses demand uncertainty, it also tends to use language with a negative sentiment. I include these two new measures as controls in the estimation of the demand and supply shocks and find the new shocks have a correlation of 0.99 with the baseline shocks—the current controls already capture the information carried in the sentiment indexes. Thus, including the sentiment indexes as controls has an insignificant impact on the estimated local projection responses.

Another concern is that the findings could be reliant on the exact topics assigned to the demand and supply groups. To rule this out, I run the same identification exercise with different sets of topics assigned to demand and supply. The first, I run the demand group as only the consumer topic, and the supply group as only energy; Figure B11 shows that the results are robust to this specification. The same figure shows that expanding the demand group to include investment (which is currently in the financial group) and the supply group to include wages, labor strikes, and living costs also creates similar results, with a smaller currency risk premium in response to demand uncertainty. A related concern was the omission of the labor category as neither demand nor supply. I show in Figure B12 that including the labor group in either the demand or supply group does not overturn the exchange rate, currency risk premium, or interest rate differential responses. Including labor in the supply category does change the price level response and actually strengthens the respective exchange rate responses for both uncertainty measures. Finally, I check if limiting only to articles that are mainly about economic uncertainty was too strict of a requirement. Instead, Figure B13 shows that using the “partly” about economic uncertainty response from the LLM (which keeps 91% of the original articles) generates very similar results. These efforts show that the signs of both estimated currency risk premia hold in every variant, but the sizes of the demand uncertainty responses are sensitive to the topics included.

A threat to identification is that there are other foreign uncertainty measures that are being captured, rather than UK uncertainty. In an attempt to control for more measures of foreign uncertainty, I include the US EPU and US trade policy uncertainty indexes as further controls. I find these additional controls have little impact on the estimated responses. Which of the controls are necessary for the results to hold? I drop each control one at a time and show the results in Figure B10. The appreciation in response to supply uncertainty is sensitive to controlling for the VIX and for UK GDP; without those it shows no exchange rate response. Finally, estimating the local projection in levels instead of gaps shows similar responses, but larger policy and price level responses. I find that controlling for global uncertainty and current macroeconomic conditions are important for the results, but that additional controls have little additional impact.

An additional concern is that the result could come from one way of measuring the exchange rate or currency risk premium. To address this, I try estimating the local projections with a trade-weighted pound exchange rate and find a similar currency risk premium response, but no depreciation in response to the demand uncertainty. Ideally, to run this comparison, I would want to have a trade-weighted currency risk premium as well, but I do not have access to survey data on all of the needed bilateral exchange rates. In addition, I check if the use of the Reuters surveys for constructing the currency risk premium is driving the results by replacing my premium estimates with ones built from Consensus forecasts and find similar currency risk premium estimates.

Another objection could be about the choices in the estimation of the local projection. To test the number of lags, I try re-estimating the results with six or twelve lags to capture more persistence in the data. Figure B15 shows that responses barely change, except for the UK-US CPI gap response to the supply uncertainty shock, which becomes insignificant. In addition, I check if estimating the demand and supply uncertainty jointly with the financial and policy responses affects the results. The impulse responses, shown in Figure B16 are very similar, though the currency risk premium for demand uncertainty is smaller and less significant when financial uncertainty is not included in the local projection. This is not too surprising, as there is a higher amount of correlation between demand and financial uncertainty indexes and so controlling for the other should be important for identifying the effect of demand uncertainty.

A final concern could be that the results come from the uncertainty shifts around or after Covid. To show this is not the case, I take two approaches: (1) drop the March–June 2020 months where GDP plummeted and then rebounded, (2) stop the sample at December 2019. I find that dropping the Covid months has little impact, while ending the sample in 2019 actually creates stronger, more significant exchange rate and currency risk premium results.

This shows that these findings are not simply a result of the uncertainty and volatility surrounding Covid and the following years.

2.7 Empirical Takeaways

The empirical results have shown that disaggregating uncertainty into demand and supply uncertainty shows important variation in the currency responses to each. The movement of the currency risk premium points to the importance of a risk channel. This is what will be further explored in the model.

3 Model

The model builds on the classic New Keynesian small open economy model of [Galí and Monacelli \(2005\)](#) by introducing incomplete international financial markets and recursive preferences.

3.1 New Keynesian Small Open Economy

The classic model is a small open economy with pricing frictions, home bias, and differentiated goods. I include incomplete international financial markets, which lead to incomplete international risk sharing, a foreign bond Euler equation, and a foreign bond holding equation (the net international investment position).

Setup There is a home country that is infinitesimally small (i.e. the UK) and a single large foreign country (i.e. the rest of the world). The home country takes all values determined by the foreign country as exogenous and given. Foreign country variables are denoted with an asterisk (X^*). Time is infinite and discrete.

Households A representative household maximizes lifetime utility—their value function. Given a discount factor, β , and a flow utility function, U , their value function is defined recursively as:

$$V_t \equiv Z_t U(C_t, N_t) + \beta \cdot CE(V_{t+1}), \quad (6)$$

where $CE(V_{t+1}) = \left(\mathbb{E}_t [V_{t+1}^{1-\Sigma}] \right)^{\frac{1}{1-\Sigma}}$ is a certainty equivalence operator and Z_t is an exogenous preference (demand) shifter that scales the value of current consumption relative to the

future.¹⁰ Larger values of Σ increase the effective risk aversion of the household, while not impacting the intertemporal elasticity of substitution. Flow utility is a standard constant relative risk aversion function, with linear separability between household consumption, C , and labor, N :

$$U(C_t, N_t) \equiv \frac{C_t^{1-\sigma}}{1-\sigma} - b \frac{N_t^{1+\eta}}{1+\eta}, \quad (7)$$

where σ is the inverse of the intertemporal elasticity of substitution and η is the inverse Frisch elasticity of labor supply.

Incomplete International Financial Markets The household can purchase any domestic asset or a foreign bond, B_f , which pays off foreign interest rate R^* . Domestic assets, D_t , are priced using the Arrow-Debreu pricing kernel $Q_{t,t+1}$.¹¹ The nominal exchange rate, $\mathcal{E}_t$, converts the foreign payoff into domestic currency and is written as home currency per unit of foreign currency (so an increase is a depreciation of the home currency). In addition, there is a quadratic cost of holding foreign bonds, $\Psi_{b_f,t} = \frac{\psi_{b_f}}{2} b_{f,t}^2$, which is included to ensure foreign bond holdings are stationary in the model.¹² This cost is quadratic in the real holdings of foreign bonds, $b_{f,t} = \frac{B_{f,t}}{P_{f,t}^*}$.

Household Budget Constraint Households can purchase a consumption bundle, C , at price, P , in domestic currency. They receive income from wages, W , government transfers, T , and firm profits, Pr_t . Bonds purchased in the last period pay off interest, which was determined in that prior period. Their budget constraint, in nominal terms, is:

$$P_t C_t + \mathbb{E}_t [Q_{t,t+1} D_{t+1}] + \mathcal{E}_t (B_{f,t} + P_{f,t}^* \Psi_{b_f,t}) = W_t N_t + D_t + \mathcal{E}_t R_{t-1}^* B_{f,t-1} + Pr_t + T_t.$$

Consumption Bundles Final consumption is a constant elasticity of substitution (CES) bundle over consumption of a home goods bundle, C_h , at price P_h in domestic currency, and a foreign goods bundle, C_f , at price P_f in domestic currency. Allocations are determined by the import share parameter (openness), ω , and elasticity, γ :

$$C_{h,t} = (1 - \omega) \left(\frac{P_{h,t}}{P_t} \right)^{-\gamma} C_t, \text{ and} \quad (8)$$

10. This preference shifter affects all flow utility, from either consumption or labor. This is purposeful as it will then drop out of the labor-leisure tradeoff condition and the preference shock will not have a direct impact on wages.

11. This budget constraint could be defined with a single domestic bond and that would have equivalent results for the model. I define the complete domestic markets to make clear that it is breaking international risk sharing that is important for the results.

12. This is a common technique in the international finance literature, see [Schmitt-Grohé and Uribe \(2003\)](#).

$$C_{f,t} = \omega \left(\frac{P_{f,t}}{P_t} \right)^{-\gamma} C_t. \quad (9)$$

where $P_t \equiv [(1 - \omega)P_{h,t}^{1-\gamma} + \omega P_{f,t}^{1-\gamma}]^{\frac{1}{1-\gamma}}$ is the price index for the consumption bundle, given the CES aggregator. Domestic consumption is itself a CES bundle over different varieties of home goods, $C_{h,i}$, whose allocations are given by the elasticity parameter, ϵ . Since all domestic firms will end up producing the same quantity and price, I do not track the individual varieties demanded.

Relative Prices In this model, only relative prices will matter in the equilibrium (this is important as it allows dropping the price levels—which are nonstationary). I define the following relative prices: the terms of trade $\mathcal{S} = \frac{P_{f,t}}{P_{h,t}}$, the real exchange rate, $\mathcal{Q} = \frac{\mathcal{E}_t P_t^*}{P_t}$, and the nominal exchange rate, $\mathcal{E} = \frac{S_t P_{h,t}}{P_{f,t}^*}$. In addition, I will define inflation rates $\Pi_{j,t} = \frac{P_{j,t}}{P_{j,t-1}}$ for each of the $P_{j,t}$ prices, including the aggregate price bundles, P and P^* . Note that the real exchange rate can be written as a function of the terms of trade:

$$\mathcal{Q}_t = \frac{\mathcal{S}_t}{[(1 - \omega) + \omega \mathcal{S}_t^{1-\gamma}]^{\frac{1}{1-\gamma}}}. \quad (10)$$

Household First Order Conditions Households optimize their value function given their budget constraint (see Appendix), resulting in the following equilibrium conditions. First, the labor-consumption tradeoff,

$$w_t = -\frac{U_{N,t}}{U_{C,t}}, \quad (11)$$

where w is the real wage, which determines how much households value consumption versus additional labor. Second, while the household prices all domestic assets D_t , only the price of domestic risk-free return, R_t , will matter for the equilibrium,

$$1 = \mathbb{E}_t [\text{SDF}_{t,t+1}] R_t. \quad (12)$$

Third, a foreign Euler equation that prices the foreign bond return, in terms of domestic currency:

$$1 = \mathbb{E}_t \left[\text{SDF}_{t,t+1} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right] R_t^* \frac{1}{(1 + \Psi'_{b_f,t})}. \quad (13)$$

Both Euler equations rely upon the definition of the stochastic discount factor (SDF), which is:

$$\text{SDF}_{t,t+1} = \beta \frac{Z_{t+1}}{Z_t} \frac{U_{C,t+1}}{U_{C,t}} \frac{1}{\Pi_{t+1}} \left(\frac{\mathbb{E}_t [V_{t+1}^{1-\Sigma}]^{\frac{1}{1-\Sigma}}}{V_{t+1}} \right)^{\Sigma}, \quad (14)$$

where parameter Σ from the certainty equivalence function provides an additional term that depends upon the continuation value of the households' value function; this makes future volatility important for pricing returns today. This formulation of the SDF follows that of [Rudebusch and Swanson \(2012\)](#). The $\text{SDF}_{t,t+1} = Q_{t,t+1}$, the Arrow-Debreu pricing kernel for the domestic assets. In addition, the household's budget constraint must hold in equilibrium, which will be combined later with the market clearing conditions.

Uncovered Interest Parity (UIP) Condition The UIP condition holds by combining the domestic and foreign Euler equations (equations 12 and 13), which leads to the following relationship between the expected change in the nominal exchange rate and the relative interest rates:

$$\mathbb{E}_t \left[\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right] = \frac{R_t}{R_t^*} (1 + \Psi'_{b_f,t}) - R_t \text{Cov}_t \left[\text{SDF}_{t,t+1}, \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right]. \quad (15)$$

There are two wedges in this equation that lead to deviations from the standard, linear UIP condition: (1) the cost of holding foreign bonds, $\Psi'_{b_f,t}$ which is small and included to induce stationarity in the model, and (2) the endogenous currency risk premium (sometimes referred to as the UIP premium), which is the covariance between the domestic SDF and the nominal exchange rate and would be zero in a linearized model. ¹³

Firms Firms produce differentiated output, Y_t , according to a decreasing returns to scale production function with parameter α . However, in equilibrium, all firms will end up producing the same goods, so aggregate production is simply:

$$Y_t = A_t N_t^{1-\alpha}, \quad (16)$$

where A is the productivity process. This means firms also face the same real marginal costs, mc , which are given by:

$$mc_t = \frac{w_t N_t}{Y_t (1 - \alpha)}. \quad (17)$$

Domestic firms face demand for their output from the foreign country, determined by the foreign country's consumption aggregator (which is symmetric to the home). This demand

13. The UIP equation is not an additional equilibrium condition for the model, but is instead derived from two of the Euler equations which must hold each period.

can be expressed using the terms of trade as:

$$X_t = \omega \mathcal{S}_t^\gamma C_t^*. \quad (18)$$

Firm First Order Condition Firms optimize profits, using the household stochastic discount factor, and face quadratic Rotemberg adjustment costs (see Appendix). This leads to an otherwise textbook nonlinear Phillips curve in terms of the domestic good, with the small addition of the terms of trade and real exchange rate in the marginal cost term:

$$\psi_p (\Pi_{h,t} - 1) \Pi_{h,t} = \epsilon \left(\frac{\mathcal{S}_t}{\mathcal{Q}_t} mc_t - \frac{\epsilon - 1}{\epsilon} \right) + \psi_p \mathbb{E}_t \left[\text{SDF}_{t,t+1} (\Pi_{h,t+1} - 1) \Pi_{h,t+1}^2 \frac{Y_{t+1}}{Y_t} \right], \quad (19)$$

where ψ_p is the price adjustment cost parameter.

Goods Market Clearing Both domestic and foreign goods markets must clear. Foreign goods are simple: $Y^* = C^*$. Domestic good markets have to account for the Rotemberg price adjustment costs and exports:

$$Y_t(1 - \Psi_{p,t}) = C_{h,t} + X_t. \quad (20)$$

Notice that while foreign consumption is constant, exports to the foreign country will adjust endogenously as the prices of home and foreign goods change. This ties the output of the domestic economy to the terms of trade.

Net International Investment Position The household budget constraint collapses to an aggregate net international investment position, which determines the holdings of foreign bonds. In particular, new purchases of real foreign bonds have to equal the net exports of the domestic economy, plus the real return on already held foreign bonds:

$$NX_t = b_{f,t} + \Psi_{b_f,t} - R_{t-1}^* \frac{1}{\Pi_{f,t}^*} b_{f,t-1}, \quad (21)$$

where $NX_t = \omega \mathcal{S}_t^{\gamma-1} C_t^* - C_{f,t}$ is the net exports of the domestic economy. This relationship will keep track of the debt that the domestic economy either claims or owes the foreign country.

Monetary Policy To close the economy, the risk-free return is set by the central bank. Following [Galí and Monacelli \(2005\)](#), the central bank follows a Taylor rule, raising the

risk-free return in response to domestic inflation:

$$R_t = R_{t-1}^{\rho_R} \left(\bar{R} \cdot \Pi_{h,t}^{\phi_\pi} \right)^{1-\rho_R}, \quad (22)$$

where ϕ_π determines the strength of the response, and ρ_R determines the persistence of the interest rate. A value $\phi_\pi > 1$ is needed for determinacy, as in textbook New Keynesian models. Alternative monetary policy rules could target the consumption bundle inflation, Π , or even changes in the nominal exchange rate, but [Galí and Monacelli \(2005\)](#) show that targeting the domestic inflation rate is optimal for the central bank.

Exogenous Stochastic Volatility Shock Process Domestic productivity is an AR(1) process with stochastic volatility:

$$A_t = (1 - \rho_A)\bar{A} + \rho_A A_{t-1} + \exp(\sigma_{A,t}/2)\bar{\sigma}_A \varepsilon_{A,t}, \quad (23)$$

where $\bar{A}$ is the steady state level of productivity, ρ_A is the persistence of the productivity process, $\bar{\sigma}_A$ is the steady state volatility of productivity, and $\varepsilon_{A,t}$ is a standard normal shock. The log-volatility of productivity follows its own AR(1) process:

$$\sigma_{A,t} = \rho_{\sigma_A}\sigma_{A,t-1} + \sigma_{\sigma_A}\varepsilon_{\sigma_A,t}, \quad (24)$$

where ρ_{σ_A} is the persistence of the volatility process, σ_{σ_A} is the volatility of the volatility process, and $\varepsilon_{\sigma_A,t}$ is a standard normal shock. The preference (demand) shifter, Z_t , follows the same structure, an AR(1) process with stochastic volatility:

$$Z_t = (1 - \rho_Z)\bar{Z} + \rho_Z Z_{t-1} + \exp(\sigma_{Z,t}/2)\bar{\sigma}_Z \varepsilon_{Z,t}, \quad (25)$$

$$\sigma_{Z,t} = \rho_{\sigma_Z}\sigma_{Z,t-1} + \sigma_{\sigma_Z}\varepsilon_{\sigma_Z,t}, \quad (26)$$

with the same definitions for the parameters as in the productivity process. I make use of log-volatility processes for the time-varying volatility to ensure that the volatility is always positive. The four shocks, $\varepsilon_{A,t}$, $\varepsilon_{\sigma_A,t}$, $\varepsilon_{Z,t}$, and $\varepsilon_{\sigma_Z,t}$ are independent. I interpret the volatility shock to productivity as the supply uncertainty shock and the volatility shock to preferences as the demand uncertainty shock.

In this economy, domestic productivity and the domestic preference shifter are the only stochastic exogenous variables. Foreign variables: Y^* , Π^* , R^* , are all exogenous, but constant values. Thus the variables for the domestic economy can be interpreted as relative values to

the rest of the world.¹⁴

Equilibrium Definition Given exogenous processes for the shocks, $\{A, \sigma_A, Z, \sigma_Z\}$, and exogenous values for the foreign variables, $\{Y^*, \Pi^*, R^*\}$, an equilibrium is an allocation, $\{V, U, C_h, C_f, Q, S, \mathcal{E}, w, R, \text{SDF}, b_f, \Pi, \Pi_h, N, Y, X, mc\}$, that satisfies the equations (6) through (14) and (16) through (22).

4 Mechanisms

Uncertainty works through two main channels in the model: (1) precautionary savings and (2) currency risk. The precautionary savings channel is common to both shocks, while the currency risk channel depends upon the source of the uncertainty.

4.1 Risk-Adjusted Log-Linearization

The standard approach to log-linearizing models assumes certainty equivalence, such that $\log \mathbb{E}_t[X_{t+1}] = \mathbb{E}_t[\log X_{t+1}]$, where X_{t+1} is a random variable. This removes any role for risk to have an effect in the model. Instead, I will take a risk-adjusted log-linearization of the model that preserves key second moment wedges that will show the mechanisms at play in the model and how they are driven by the uncertainty shocks. I will keep higher order terms for equilibrium conditions that have an expectation and log-linearize everything else. This is a common approach in the finance and macro-finance literature (Jermann 1998, Malkhozov 2014, Bianchi, Kung, and Tirsikh (2023)).¹⁵ A risk-adjusted log-linearization uses the following approximation for expectations:

$$\log \mathbb{E}_t[X_{t+1}] \approx \mathbb{E}_t[\log X_{t+1}] + \frac{1}{2} \text{Var}_t[\log X_{t+1}]. \quad (27)$$

This approximation is exact if X_{t+1} is a log-normal variable, otherwise it is simply a second-order approximation. This approach creates a block-recursive linear model. The risk wedges from the risk-adjustment are linear in the log-volatility states and the rest of the model is linear in the level states and the risk wedges.

Example: Risk-Adjusted Euler Equation I will apply the risk-adjusted log-linearization to the domestic Euler equation, pricing the returns of the domestic risk-free asset. Log-

14. This is why I match the model to UK-US gap variables in Section 5.

15. In Bianchi, Kung, and Tirsikh (2023) they use Markov chain changes in volatility, while I use stochastic volatility processes.

linearized variables are denoted lowercase with a hat, $\hat{x}_t$. The log-linearization is around the non-stochastic steady state. Taking logs of both sides of the Euler equation:

$$0 = \mathbb{E}_t \left[\widehat{\text{sdf}}_{t,t+1} \right] + \underbrace{\frac{1}{2} \text{Var}_t \left[\widehat{\text{sdf}}_{t,t+1} \right]}_{\text{risk wedge}} + \hat{r}_t \quad (28)$$

The variance of the SDF plays a direct role in determining the level of the risk-free interest rate, $\hat{r}_t$.

This approach, applied to the expectations in a few key equations of the model, generates three risk wedges: (1) the price of risk, λ_t , (2) the currency risk premium, ψ_t , and (3) an inflation risk wedge, ξ_t . Importantly, these are the only higher-order moments in the model; everything else is first order. This allows us to decompose the risk wedges and, in the calibrated results, to see the impact of each wedge in isolation. While there are three wedges, the inflation risk wedge plays a small role, so I will now present only the other two wedges in more detail.

4.2 Price of Risk and Precautionary Savings

Precautionary savings is a risk channel that captures that households fear uncertainty in the future and decide to save more today when the future is more uncertain. This causes households to spend less today. The precautionary savings channel, driven by the price of risk, λ_t , is found in the risk-adjusted IS curve:

$$\hat{c}_t = \mathbb{E}_t [\hat{c}_{t+1}] - \frac{1}{\sigma} (\hat{r}_t - \mathbb{E}_t [\hat{\pi}_{t+1}] + \mathbb{E}_t [\Delta \hat{z}_{t+1}] + \lambda_t + \Sigma(1 - \Sigma)\mu_t). \quad (29)$$

This is a textbook New Keynesian IS curve with one additional term, the price of risk:

$$\lambda_t = \frac{1}{2} \text{Var}_t [\widehat{\text{sdf}}_{t,t+1}], \quad (30)$$

$$\mu_t = \frac{1}{2} \text{Var}_t [\hat{v}_{t+1}]. \quad (31)$$

The price of risk is driven by the variance of the log-SDF. The μ_t is an additional risk term that comes from the recursive preferences and increases the volatility of λ_t . Importantly, there are no first-moment terms in the price of risk; it only exists as a higher moment object.

An increase in the price of risk, λ_t , causes households to reduce consumption today, holding everything else constant. This has the exact same impact in the IS curve as raising the current interest rate in the economy, $\hat{r}_t$. In this model, λ_t can only ever be positive. When the volatility in the economy increases, from either preference volatility or productivity

volatility sources, the price of risk will increase as well. Thus this channel cannot have separate effects depending upon the source of uncertainty; it can vary in magnitude, but not in sign. The currency risk channel, however, can have a positive or negative sign.

4.3 Currency Risk

When a domestic household decides to hold the foreign bond for a period, they face risk from the exchange rate changing in value when they receive the return tomorrow. This is the currency risk channel. This behavior is determined by the risk-adjusted log-linearized UIP condition:

$$\mathbb{E}_t [\Delta \hat{e}_{t+1}] = \hat{r}_t - \hat{r}_t^* + \psi_b \tilde{b}_{f,t} - \psi_t. \quad (32)$$

The $\psi_b \tilde{b}_{f,t}$ are the small bond holding costs ensuring stationarity; the risk wedge is ψ_t , which is the endogenous currency risk premium:

$$\psi_t = \frac{1}{2} \text{Var}_t(\Delta \hat{e}_{t+1}) + \text{Cov}_t(\widehat{\text{sdf}}_{t,t+1}, \Delta \hat{e}_{t+1}). \quad (33)$$

Now we can see that when ψ_t rises this causes a fall in $\Delta \hat{e}_{t+1}$, which, holding all else constant, is an increase in the exchange rate today (a depreciation). Since the exchange rate has depreciated today, households expect that it will appreciate tomorrow, which will be compensation for the risk of holding the foreign currency. Conversely, when ψ_t falls this would cause an appreciation in the current nominal exchange rate through the UIP equation, and an expected depreciation tomorrow.

The movements of the currency risk premium depend upon the variance of the exchange rate—which is always positive—and the covariance between the SDF and the exchange rate—which may be positive or negative. This covariance term determines the risk-sharing role of the foreign bond. The SDF measures how valuable a pound is for the household tomorrow versus today. When their SDF is high, a household values the pound today; when their SDF is low, a household values the pound tomorrow. If the exchange rate is high (depreciation) at the same time their SDF is high, then the foreign bond pays off more and risk is shared abroad, so ψ_t rises. If instead the exchange rate is low (appreciation) at the same time their SDF is high, the foreign bond will pay off in fewer units of domestic currency, and the foreign bond becomes a risky asset and provides no risk sharing, so ψ_t falls. This is in contrast to the price of risk, λ_t , which can never switch signs.

This currency risk premium depends upon the source of uncertainty because of the incomplete international financial markets. The UK households have two Euler equations that must hold in all periods, one for their domestic bond and one for the foreign bond. There are

no pricing decisions from the rest-of-the-world because their values are assumed to be fixed and exogenous. They cannot trade any state-contingent foreign asset—that is what makes the international financial markets incomplete and the international risk sharing incomplete, which means all movements of the risk premium in this model come from the domestic side. With incomplete international markets, the exchange rate does not always respond the same way as the domestic SDF. When they price the foreign bond they use their domestic SDF, and they face the currency risk premium, ψ_t , which accounts for the risk of exchanging their savings to foreign currency and back. The currency risk premium, ψ_t , is a compensation for holding risk, not a failure to arbitrage. In a small open economy model with complete markets, which I will compare my results to in Section 5.8, the UK household would price the currency risk as exactly $\psi_t = -\lambda_t$. In that case, the currency risk premium could only have one sign: the opposite sign of the price of risk. The currency risk premium would then have the same sign, regardless of the source of uncertainty.

As shown in Section 2.4, this currency risk premium can be directly estimated in the data using this UIP equation. Given a measure for the expected exchange rate and the relative interest differential, the remaining variation must be the currency risk premium or the small bond holding costs. This allows me to calibrate the currency risk channel of the model to the local projection responses estimated earlier.

This UIP condition is directly comparable to the log-linearized UIP condition presented in [Itskhoki and Mukhin \(2021\)](#).¹⁶ In their model, exogenous financial noise shocks drive changes in the currency risk premium. In this model, the currency risk, ψ_t , is endogenously determined by the representative household.

In the next subsection, I will show how the source of uncertainty shocks determines the sign of this response.

4.4 Covariance as a Function of Uncertainty Shocks

The sign of the covariance is determined by the first-order responses to the level shocks. Because the stochastic volatility process is a log volatility, the covariance is exactly linear in the volatility states:

$$\text{Cov}_t(\widehat{\text{sdf}}_{t,t+1}, \Delta \hat{e}_{t+1}) = a_A b_A \rho_{\sigma_A} \hat{\sigma}_{A,t} + a_Z b_Z \rho_{\sigma_Z} \hat{\sigma}_{Z,t}, \quad (34)$$

where $a_k = \partial \widehat{\text{sdf}} / \partial \epsilon_k$ and $b_k = \partial \Delta \hat{e} / \partial \epsilon_k$. The first takeaway is that more volatility and a higher persistence both increase the movement in the covariance term. Since the volatilities

16. Their condition is written as: $\mathbb{E}_t[\Delta \hat{e}_{t+1}] = \hat{i}_t - \hat{i}_t^* + \chi_2 \tilde{b}_{t+1} - \chi_1 \Psi_t$ where χ_1 and χ_2 are parameters that depend upon the variance of the exchange rate and Ψ_t is their financial noise shock.

are always positive, and persistence is between zero and one, the signs of the responses to each shock are determined by a_k and b_k .

Table 2: Sign of the Covariance by Source of Uncertainty

Shock, k	$a_k = \frac{\partial \widehat{\text{sdf}}_{t,t+1}}{\partial \varepsilon_k}$	$b_k = \frac{\partial \Delta \hat{e}_{t+1}}{\partial \varepsilon_k}$	$a_k b_k$
Productivity, A	—	+	—
Preference, Z	+	+	+

Note: Describes the signs of the coefficients in equation (34), which shows that the sign of the covariance term that determines if the foreign bond is a good or bad hedge is determined by the signs of the first-moment responses to each shock, k .

Table 2 shows the signs at the model’s calibrated values and how they imply an opposite-signed response in the covariance term. At this parameterization, a positive productivity shock causes a depreciation, and a fall in the SDF. A positive preference shock causes a depreciation, and a rise in the SDF. The signs of the SDF responses are fairly direct for this model and not sensitive to the calibration. Thus, for the covariance to have opposite-signed responses, the two first-moment shocks must have the same sign response on the nominal exchange rate.

A positive productivity shock depreciates the pound through the standard channel in Galí and Monacelli (2005): more output, cheaper home goods, and a weaker currency. In some other open-economy models, higher productivity raises demand for home goods and appreciates the currency. However, in this model, the CES demand for home and foreign good varieties creates either scarcity or abundance of the home good in response to productivity shocks. When households consider a surprise of low productivity next period, they know that output will be low, consumption will fall, and so their SDF will be high and the pound will appreciate at the same time. Uncertainty about productivity makes this scenario more likely for the households to face in the future. For UK households facing volatile productivity, pound-denominated assets help them smooth their consumption over the future. Thus, the currency risk premium on the pound falls and the pound appreciates today.

A positive preference shock depreciates the pound and raises the SDF. When households consider a surprise of low preferences next period, they know that their SDF will be low and the pound will appreciate at the same time. When the volatility of preference shocks increases, this scenario becomes more likely for households to face in the future. This is why the households want to save their assets in dollars when volatility of preference shocks is high. Thus, the currency risk premium on the pound rises and the pound depreciates today.

In this model, a positive preference shock leads to a depreciation of the pound, which is

in contrast to other open economy models. This is a feature of the calibration—specifically the persistence of the interest rate. In this model, the interest rate responds very slowly to shocks. A positive preference shock raises spending and inflation, but the nominal rate barely changes, so the real interest rate actually falls. This real interest rate differential then causes the depreciation of the pound as assets shift to the higher real return available abroad. In Appendix E, I show the range of parameter values where this holds and creates the opposite-signed responses in Table 2.

5 Results

The volatility shocks in the model are able to match the uncertainty innovations in the data. The currency risk channel is responsible for the different responses depending upon the source of uncertainty.

5.1 Solution Method

The risk-adjusted log-linear model is solved numerically to a third-order approximation around the non-stochastic steady state. Impulse response functions are then generated by simulating the model till it reaches the stochastic steady state, and then applying the policy functions. A first-order approximation would generate no role for risk and no effect from the volatility shock, while with a second-order approximation, risk only impacts the stochastic steady state. A third-order approximation allows for time-varying risk premia, which are the key mechanism by which the volatility shocks impact the model.

Figures F1 and F2 in the Appendix show that taking a third-order approximation of either the risk-adjusted log-linear approximation or the full nonlinear model gives similar results. I prefer the risk-adjusted log-linearization as it allows for a clean decomposition of the risk wedges and mechanisms by which the volatility shocks impact the model.

5.2 Calibration

The model is fit at a monthly frequency to match the monthly local projection data. Some model parameters are fixed to values that are standard in the literature and then $\{\phi_\pi, \rho_R, \rho_A, \rho_{\sigma_A}, \rho_Z, \rho_{\sigma_Z}\}$ are calibrated to match the local projection responses. The share of imports ω is calibrated to roughly match the UK import share (about 30%). The recursive preferences risk aversion parameter is set to the same value used in Basu and Bundick (2017) and similar to the value used in Rudebusch and Swanson (2012). The full list of parameters and their values is given in Table 3. The following parameters end up at the bounds of their

calibration windows: $\rho_{\sigma_Z}, \rho_R, \rho_A$. In the Appendix, Figure E1 shows the sensitivity of the model's currency risk premium response to the calibration of key parameters.

Table 3: Model Calibration

Parameter	Value	Description
<i>Fixed</i>		
β	$0.99^{1/3}$	Discount factor (monthly)
σ	2	Inverse of intertemporal elasticity of substitution
Σ	80	Recursive risk aversion
η	2	Inverse Frisch elasticity of labor supply
α	0.33	One minus labor share of output
ϵ	6	Elasticity of substitution of domestic goods
ω	0.3	Openness (import share)
γ	1.01	Elasticity of substitution between domestic and foreign goods
ψ_p	500	Price adjustment cost
ψ_{b_f}	0.001	Foreign bond holding cost
$\bar{\sigma}_A$	0.02	Standard deviation of productivity shock
σ_{σ_A}	$2 \log 2$	Volatility of productivity volatility
$\bar{\sigma}_Z$	0.02	Standard deviation of preference shock
σ_{σ_Z}	$2 \log 2$	Volatility of preference volatility
<i>Calibrated to the local projections</i>		
ϕ_π	1.66	Taylor rule response to domestic inflation
ρ_R	0.98	Interest rate smoothing
ρ_A	0.95	Persistence of productivity
ρ_{σ_A}	0.939	Persistence of productivity volatility
ρ_Z	0.969	Persistence of preference shifter
ρ_{σ_Z}	0.96	Persistence of preference volatility

Note: Parameter calibration values are either fixed to standard values in the literature or used to match the empirical local projection responses.

The loss function used to calibrate the parameters is a squared loss on the IRFs compared to the local projection responses. The horizons targeted and the weights on each outcome are in Table 4. I place a higher weight on the exchange rate and UIP deviations as these are the key measures for the model mechanism. The horizons are weighted by the confidence interval for that estimate, thus periods with wide bands are given less weight. I drop the first few periods because the model cannot produce hump-shaped responses without some additional frictions.

Table 4: Loss Function Weights

Outcome	Horizons (months)	Outcome weight	Horizon weight
Exchange rate (GBP/USD)	2–24	1	$1/SE_h^2$
Currency risk premium (12-month)	2–24	1	$1/SE_h^2$
UK-US 3-month interest rate difference	2–24	0.25	$1/SE_h^2$
UK-US CPI gap	2–24	0.25	$1/SE_h^2$

Note: The weights used to target the impulse response function matching of the model and the empirically estimated local projections. The first period is dropped and more weight is placed on the exchange rate and currency risk premium and less weight on horizons with large standard errors.

5.3 Volatility Shocks Match Uncertainty Data


The model is able to match the responses seen in the local projections. A positive shock to productivity volatility causes an appreciation of the pound, a fall in the currency risk premium, an increase in the policy rate, and an increase in the price level in the long run. A positive shock to preference volatility causes a depreciation of the pound, followed by an eventual appreciation, an increase in the currency risk premium, a cutting of the policy rate, and an eventual fall in the price level.

In order to produce changes in the currency risk premium that are of the size seen in the data, the model relied upon high levels of persistence in each of the shocks and in the policy rate. In addition, the model is not able to match the timing of the shocks, as it is not able to produce hump-shaped responses in the exchange rate or UIP deviations. These two outcomes are always monotonic in the model. The model fits a gradual cut in the interest rate in response to demand uncertainty, while the data shows a cut for the first year, and then a raising of the rate in the long run. For supply uncertainty, the model fits an increase in interest rates, but not as large as the data shows. Finally, the model misses on impact for the price level response to both uncertainties, as the model predicts the wrong response at first, and then catches up in the long run.

5.4 First-Moment Shocks Generate No Currency Risk Premium

I re-run the calibration of the model on the same parameters, with the same window of values, and using the same loss function, but instead of using the volatility shocks, I try to match the local projection responses using negative first-moment shocks. As seen in Figure 8, the first-moment shocks cannot match the local projection responses on the currency risk premium or for the exchange rate for demand uncertainty. It is impossible in this model for a first-moment shock to generate a deviation in the UIP condition; the risk wedge, ψ_t ,


Figure 7: Model Evidence: Volatility Shocks and UK Local Projections



Note: The empirical local projection responses from Figure 5 and Figure 6 are repeated here in the dotted lines and 68% confidence intervals. The solid line in blue shows the model's response to a preference volatility shock calibrated using the values in Table 3. The solid line in orange shows the model's response to a productivity volatility shock using the same calibration.

depends only upon the volatility states. The first-moment negative productivity shock does a good job at matching the exchange rate, interest rate gap, and price level gap responses. The first-moment preference shock tries to match the interest gap and price level gap, but mostly estimates very small responses.


Figure 8: Model Evidence: First- vs. Second-Moment Shocks



Note: The empirical local projection responses from Figure 5 and Figure 6 are repeated here in the dotted lines and 68% confidence intervals. The solid line in blue shows the model's response to a preference volatility shock calibrated using the values in Table 3. The solid line in orange shows the model's response to a productivity volatility shock using the same calibration. The dashed lines show a negative level shock for preference and productivity, with the same free parameters re-calibrated as in Table 3.

5.5 Comparing Preference and Productivity Volatility Shocks

I set the persistence of the productivity volatility to be the same as the preference volatility and then I compare the responses to each of the volatility shocks. This allows for a clean comparison between the two sources of uncertainty, holding the size of the fluctuations constant. Figure 9 shows the responses to each of the volatility shocks.




Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric.

Each volatility shock has similar impacts on consumption, but output falls by more in response to productivity volatility. This is due to the fact that preference volatility increases the number of net exports, and so output is supported by the foreign demand for domestic goods. While in response to both shocks households wish to save more, in response to the preference volatility, they see the foreign bond as a good hedge and so they buy much more of it. This leads to the depreciation of the domestic currency. In both cases the policy rate is cut in response to the falling domestic prices (Appendix Figure C1). The final two panels show that the price of risk, λ_t , is positive in both cases and of similar magnitude, while the currency risk premium, ψ_t , is positive in response to preference volatility and negative in response to productivity volatility.

Next, I compare the two shocks with no currency risk premium channel, shown in Figure

10. These responses are classic precautionary savings responses, very similar to the effect of volatility shocks in a closed economy model, like in [Basu and Bundick \(2017\)](#). Households desire to save more, so consumption falls, which causes output to fall. In response to the falling prices, the policy rate is cut ([Appendix Figure C1](#)). With just the precautionary savings channel, the model produces an appreciation of the nominal exchange rate in response to both shocks. This is inconsistent with the local projection estimated responses.




Note: Model calibrated to values in [Table 3](#) except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The currency risk premium channel is manually set to zero.

Without the UIP premium channel, the two volatility shocks produce nearly identical exchange-rate paths—the opposite-signed responses in [Figure 9](#) are generated by the currency risk premium. Next, I compare the two shocks with no precautionary savings channel, isolating the effect of the currency risk channel, shown in [Figure 11](#).

It is clear how the currency risk channel is responsible for the opposite-signed exchange rate responses in the two volatility shocks. For the preference volatility shock, currency risk causes households to wish to hold more of the foreign bond. For the productivity volatility shock, currency risk causes households to wish to hold less of the foreign bond. This leads to opposite-signed responses of net exports, which in turn leads to a drop in consumption for preference volatility and a drop in output for productivity volatility. It is the combination

Figure 11: Model Mechanism: No Precautionary Savings Channel




Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The precautionary savings channel is manually set to zero.

of the two channels that gives the full response of the model.

5.6 The Role of Pricing Frictions

Price stickiness is not needed for the sign of the currency risk premium to change under different sources of uncertainty, but it does carry the effects to the real exchange rate and creates the recessionary response of output.

Without pricing frictions, the currency risk wedge, ψ_t , would barely move, as seen in Figure 12. What the rigidities do is connect the real and nominal sides of the economy. Without them, the model would generate opposite-signed responses to the real and nominal exchange rate, as shown in Figure 12. For productivity volatility, the signs would be opposite on impact and remain so throughout. For preference volatility, the signs would begin matched, but quickly the nominal exchange rate would appreciate while the real exchange rate is still depreciating. This would not match the data, which show there is little difference between the real and nominal responses for either measure of uncertainty. In general, in currency data, the real and nominal exchange rate move together.




Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The dashed lines show the model responses when the nominal rigidity is removed ($\psi_p \rightarrow 0$).

The nominal rigidities apply another force on the model, which is that they ensure that output falls in response to uncertainty. Without nominal rigidities, when households face higher uncertainty, they desire to save more and one way to save more is to work more. When prices are flexible, this leads to higher output in the economy, even as consumption is falling. Nominal rigidities flip this story and the fall in consumption leads firms to cut

production and scale back their demand for work. Thus the nominal rigidities are what generate the recessionary impact of both uncertainty shocks in the model.

5.7 Why Source Matters

In Section 4, I show the sign of the currency risk channel depends upon the sign of the covariance term. In Figure 13, I show the actual calibrated values for the variance of the exchange rate (small) and the covariance between the SDF and the exchange rate (much larger). Thus, the covariance term is the key determinant of the sign of the currency risk channel, and it is what drives the opposite-signed exchange rate response to each of the volatility shocks. The final panel also shows how the two first-moment shocks have the same sign response on the exchange rate on impact, which is the condition for the covariance to have opposite-signed responses.



Note: Model calibrated to values in Table 3. The dashed lines show the model responses to first-moment level shocks.

5.8 International Risk Sharing: Incomplete vs. Complete

In both the baseline model, and in the same model with complete markets, the following condition holds:

$$\lambda_t + \psi_t = \frac{1}{2} \text{Var}_t(\widehat{\text{sdf}}_{t,t+1} + \Delta \hat{e}_{t+1}). \quad (35)$$

This follows naturally from the definitions of the currency risk premium and the price of risk. If markets are complete, then in this small open economy, the domestic SDF will be tied to the fixed foreign SDF, $\widehat{\text{sdf}}^*$, because domestic shocks do not affect the foreign SDF:

$$\widehat{\text{sdf}}_{t,t+1} + \Delta\hat{e}_{t+1} = \widehat{\text{sdf}}^*. \quad (36)$$

This immediately gives that the currency risk premium under complete markets in the small open economy is exactly:

$$\psi_t = -\lambda_t, \quad \text{Complete Markets} \quad (37)$$

because $\text{Var}_t(\widehat{\text{sdf}}_{t,t+1} + \Delta\hat{e}_{t+1}) = \text{Var}_t(\widehat{\text{sdf}}^*) = 0$. In incomplete markets, however, the domestic SDF is not pinned to move one-to-one with the foreign SDF and nominal exchange rate, and the currency risk premium, ψ_t , can have a different sign than the price of risk, λ_t .¹⁷ In this case, the currency risk premium is given by:


$$\psi_t = -\lambda_t + \frac{1}{2}\text{Var}_t(\widehat{\text{sdf}}_{t,t+1} + \Delta\hat{e}_{t+1}). \quad \text{Incomplete Markets} \quad (38)$$

This shows how the unfixed variance of the sum of the SDF and changes in the nominal exchange rate can determine the sign. These shifts in the currency risk premium have important implications for the correlation of the exchange rate with the SDF and with domestic consumption.

I now consider a situation where the only shocks in the model are the volatility shocks to the log-volatility states, $\sigma_{A,t}$ and $\sigma_{Z,t}$. Let $s = \frac{\sigma_{\sigma_Z}^2}{\sigma_{\sigma_Z}^2 + \sigma_{\sigma_A}^2}$ be the share of the total volatility accorded to the demand shock. Figure 14 shows how the correlations between the nominal SDF and nominal exchange rate (and real SDF and real exchange rate) change as the share of the volatility assigned to demand goes from zero to one, under either incomplete markets or complete markets. In panel (a), when the share of demand volatility is zero, the correlation of the exchange rate and SDF is negative one. When the share of demand volatility is one, the correlation is positive one. As the share increases, the correlation becomes a mix, crossing zero at a share of 0.52 (nominal) and 0.5 (real). Where it crosses zero depends upon the loading of the SDF and exchange rate on each of the shocks, so depending upon the calibration it may not always occur at an even mix. In this figure, there is not much of a gap between the nominal and real correlations—this is because sticky prices in the model

17. In Benigno, Benigno, and Nisticò (2012), they have a two-country open economy model with complete markets and so there $\psi_t = \lambda_t - \lambda_t^*$, where λ_t^* is the foreign price of risk. In this setting, ψ_t may still take on positive or negative values depending upon the foreign and domestic prices of risk. In my small open economy setting, $\lambda_t^* = 0$.

Figure 14: Unconditional Correlation of SDF and Exchange Rate



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The model is simulated for 20,000 months with only volatility shocks (no level shocks) at different levels of the demand share of total volatility shock variance, s , with the total variance held fixed. The realized correlation is calculated over the full sample between the real SDF and real exchange rate (solid blue line) and the nominal SDF and nominal exchange rate (dashed orange line). The dotted line shows the relative share of variance in the baseline model.


keep the real and nominal responses close to one another.

Under complete markets, the correlation between the SDF and the exchange rate is always fixed at negative one. This is a feature of the domestic SDF and exchange rate being directly determined by the foreign SDF, which enforces that $\Delta \hat{e}_{t+1} = -\widehat{\text{sdf}}_{t,t+1}$. The takeaway from Figure 14 is that an unconditional correlation of zero between the SDF and nominal exchange rate looks like the classic disconnect story. However, it can mask a strong *conditional* correlation for each shock considered. In this case, I am showing this holds in my model when considering a mixture of only the two volatility shocks.

Next, in Figure 15, I consider the correlation of changes in the real exchange rate and relative consumption growth, which is the subject of the famous Backus–Smith puzzle (Backus and Smith 1993). The puzzle is that under complete markets and CRRA preferences, international risk sharing implies a correlation of one. However, in the data, estimates of this correlation are either negative or close to zero. Figure 15 shows in the golden dashed-dotted line the initial puzzle result: an unconditional correlation of one between changes in the real exchange rate and consumption growth. This correlation is constant, regardless of the demand share of total volatility shocks. Under the baseline for my model, in the solid blue line, the mixture of the two volatility shocks can generate a zero correlation (at $s = 0.68$) and choosing entirely one shock or the other generates a positive one or negative one correlation.

Under only demand uncertainty, the real exchange rate and consumption growth move in opposite directions (correlation of negative one) and under only supply uncertainty they

Figure 15: Backus–Smith Correlation of Exchange Rate and Consumption Growth



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The model is simulated for 20,000 months with only volatility shocks (no level shocks) at different levels of the demand share of total volatility shock variance, s , with the total variance held fixed. The realized correlation is calculated over the full sample between changes in the real exchange rate and relative consumption growth. The solid blue line shows the baseline model with recursive preferences and incomplete markets. The dashed orange line shows the same model with complete markets. The dashed-dotted golden line shows the model with CRRA preferences and complete markets.

move together (correlation of one). With a mixture of the two shocks, the correlation can take any value in between unconditionally. This result holds for the Backus–Smith correlation even under complete markets, as long as there are still recursive preferences. This is because recursive preferences place a wedge between the SDF and domestic consumption. While this correlation does not distinguish incomplete and complete markets, the correlations in Figure 14 show the SDF-exchange rate correlation does separate the behavior of the two models.

6 Conclusion

This paper shows that the source of uncertainty shocks matters. Empirically, overall economic uncertainty in the UK has mixed exchange rate and currency risk premium responses. However, upon disaggregation into demand and supply uncertainty, the role of risk in determining the exchange rate responses is clear. Demand uncertainty depreciates the currency and raises the currency risk premium; supply uncertainty does the opposite. A structural model can account for these results by including stochastic volatility for preference and productivity shocks, recursive preferences, and incomplete international financial markets. The incomplete markets create a role for an endogenous currency risk wedge that drives the opposite exchange rate response across the two uncertainty shocks. Level shocks cannot move

the currency risk premium, so its response in the data points to a change in risk. This research shows the importance of knowing the source of uncertainty for better understanding its macroeconomic responses.

For other countries, this research predicts that their response to domestic uncertainty depends upon which source of uncertainty dominates. Where both sources matter, aggregate uncertainty will have muted effects on the currency and the currency risk premium, and the unconditional correlations between marginal utility and the exchange rate could be close to zero. This would mask underlying comovement and could explain variation across countries' correlations. A natural extension is to estimate the uncertainty topics from newspaper text for other small open economies and test these predictions.

References

- Akinci, Özge, Sebnem Kalemli-Özcan, and Albert Queralto. 2022. *Pricing Risk Globally: Intermediary Constraints, the Dollar, and the Global Financial Cycle*, 30026, NBER, May.
- Alvarez, Fernando, Andrew Atkeson, and Patrick J Kehoe. 2009. “Time-Varying Risk, Interest Rates, and Exchange Rates in General Equilibrium.” *Review of Economic Studies* 76 (3): 851–878.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R Sims. 2013. “Uncertainty and Economic Activity: Evidence from Business Survey Data.” *American Economic Journal: Macroeconomics* 5, no. 2 (April): 217–249.
- Backus, David K., and Gregor W. Smith. 1993. “Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods.” *Journal of International Economics* 35, no. 3 (November): 297–316.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis. 2016. “Measuring Economic Policy Uncertainty.” *The Quarterly Journal of Economics* 131, no. 4 (March): 1593–1636.
- Basu, Susanto, and Brent Bundick. 2017. “Uncertainty Shocks in a Model of Effective Demand.” *Econometrica* 85 (3): 937–958.
- Bekaert, Geert, Marie Hoerova, and Marco Lo Duca. 2013. “Risk, Uncertainty and Monetary Policy.” *Journal of Monetary Economics* 60, no. 7 (October): 771–788.
- Benigno, Gianluca, Pierpaolo Benigno, and Salvatore Nisticò. 2012. “Risk, Monetary Policy, and the Exchange Rate.” *NBER Macroeconomics Annual*.
- Bianchi, Francesco, Howard Kung, and Mikhail Tirsikh. 2023. “The Origins and Effects of Macroeconomic Uncertainty.” *Quantitative Economics* 14 (3): 855–896.
- Bloom, Nicholas. 2009. “The Impact of Uncertainty Shocks.” *Econometrica* 77 (3): 623–685.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. “Really Uncertain Business Cycles.” *Econometrica* 86 (3): 1031–1065.
- Caldara, Dario, Matteo Iacoviello, Patrick Molloy, Andrea Prestipino, and Andrea Raffo. 2020. “The Economic Effects of Trade Policy Uncertainty.” *Journal of Monetary Economics* 109 (January): 38–59.
- Carrière-Swallow, Yan, and Luis Felipe Céspedes. 2013. “The Impact of Uncertainty Shocks in Emerging Economies.” *Journal of International Economics* 90, no. 2 (July): 316–325.
- Cascaldi-Garcia, Danilo, Cisl Sarisoy, Juan M. Londono, Bo Sun, Deepa D. Datta, Thiago Ferreira, Olesya Grishchenko, et al. 2023. “What Is Certain about Uncertainty?” *Journal of Economic Literature* 61, no. 2 (June): 624–654.

- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno.** 2014. “Risk Shocks.” *American Economic Review* 104, no. 1 (January): 27–65.
- Engel, Charles.** 2016. “Exchange Rates, Interest Rates, and the Risk Premium.” *American Economic Review* 106, no. 2 (February): 436–474.
- Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez.** 2015. “Fiscal Volatility Shocks and Economic Activity.” *American Economic Review* 105, no. 11 (November): 3352–3384.
- Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F Rubio-Ramírez, and Martin Uribe.** 2011. “Risk Matters: The Real Effects of Volatility Shocks.” *American Economic Review* 101, no. 6 (October): 2530–2561.
- Galí, Jordi, and Tommaso Monacelli.** 2005. “Monetary Policy and Exchange Rate Volatility in a Small Open Economy.” *The Review of Economic Studies* 72 (3): 707–734.
- Husted, Lucas, John Rogers, and Bo Sun.** 2020. “Monetary Policy Uncertainty.” *Journal of Monetary Economics* 115 (November): 20–36.
- Itskhoki, Oleg, and Dmitry Mukhin.** 2021. “Exchange Rate Disconnect in General Equilibrium.” *Journal of Political Economy* 129 (8): 2183–2232.
- Jermann, Urban J.** 1998. “Asset Pricing in Production Economies.” *Journal of Monetary Economics* 41, no. 2 (February): 257–275.
- Jordà, Òscar.** 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95, no. 1 (February): 161–182.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng.** 2015. “Measuring Uncertainty.” *American Economic Review* 105, no. 3 (March): 1177–1216.
- Kalemli-Özcan, Sebnem, and Can Soylu.** Forthcoming. “Can Policy Uncertainty Drive the Dollar?” *Annual Review of Financial Economics*.
- Kalemli-Özcan, Sebnem, and Liliana Varela.** 2021. *Five Facts about the UIP Premium*. Working Paper, 28923, June.
- Kalemli-Özcan, Sebnem, Can Soylu, and Muhammed A. Yildirim.** 2026. “Tariff Uncertainty and the U.S. Dollar.” NBER, July.
- Kekre, Rohan, and Moritz Lenel.** 2024. “The Flight to Safety and International Risk Sharing.” *American Economic Review* 114, no. 6 (June): 1650–1691.
- Leduc, Sylvain, and Zheng Liu.** 2016. “Uncertainty Shocks Are Aggregate Demand Shocks.” *Journal of Monetary Economics* 82 (September): 20–35.
- Ludvigson, Sydney C., Sai Ma, and Serena Ng.** 2021. “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?” *American Economic Journal: Macroeconomics* 13, no. 4 (October): 369–410.

- Malkhozov, Aytek.** 2014. “Asset Prices in Affine Real Business Cycle Models.” *Journal of Economic Dynamics and Control* 45 (August): 180–193.
- Miranda-Agrippino, Silvia, and Hélène Rey.** 2020. “U.S. Monetary Policy and the Global Financial Cycle.” *The Review of Economic Studies* 87, no. 6 (November): 2754–2776.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller.** 2021. “Local Projection Inference Is Simpler and More Robust Than You Think.” *Econometrica* 89 (4): 1789–1823.
- Plagborg-Møller, Mikkel, and Christian K. Wolf.** 2021. “Local Projections and VARs Estimate the Same Impulse Responses.” *Econometrica* 89 (2): 955–980.
- Rudebusch, Glenn D, and Eric T Swanson.** 2012. “The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks.” *American Economic Journal: Macroeconomics* 4, no. 1 (January): 105–143.
- Schmitt-Grohé, Stephanie, and Martín Uribe.** 2003. “Closing Small Open Economy Models.” *Journal of International Economics* 61, no. 1 (October): 163–185.
- Scotti, Chiara.** 2016. “Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro-Surprises.” *Journal of Monetary Economics* 82 (September): 1–19.
- Stock, James H., and Mark W. Watson.** 2018. “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments.” *The Economic Journal* 128, no. 610 (May): 917–948.

Appendix

A Risk-Adjusted Log-Linear Model Equations

Equilibrium Conditions Here is the full set of equilibrium conditions for the risk-adjusted log-linear model:

$$\hat{v}_t = (1 - \beta)(\hat{z}_t + \hat{u}_t) + \beta \widehat{ce}_t \quad (\text{A1})$$

$$\hat{u}_t = \omega_C(1 - \sigma)\hat{c}_t + \omega_N(1 + \eta)\hat{n}_t \quad (\text{A2})$$

$$\widehat{ce}_t = \mathbb{E}_t[\hat{v}_{t+1}] + (1 - \Sigma)\mu_t \quad (\text{A3})$$

$$\hat{w}_t = \eta\hat{n}_t + \sigma\hat{c}_t \quad (\text{A4})$$

$$\hat{\pi}_t = \hat{\pi}_{h,t} + \Delta\hat{s}_t - \Delta\hat{q}_t \quad (\text{A5})$$

$$\widehat{\text{sdf}}_{t,t+1} = \Delta\hat{z}_{t+1} - \sigma\Delta\hat{c}_{t+1} - \hat{\pi}_{t+1} + \Sigma(\widehat{ce}_t - \hat{v}_{t+1}) \quad (\text{A6})$$

$$\hat{r}_t = -\mathbb{E}_t[\widehat{\text{sdf}}_{t,t+1}] - \lambda_t \quad (\text{A7})$$

$$\mathbb{E}_t[\Delta\hat{e}_{t+1}] = \hat{r}_t + \psi_b\tilde{b}_{f,t} - \psi_t \quad (\text{A8})$$

$$\Delta\hat{e}_t = \hat{\pi}_{h,t} + \Delta\hat{s}_t \quad (\text{A9})$$

$$\tilde{n}x_t = \tilde{b}_{f,t} - \beta^{-1}\tilde{b}_{f,t-1} \quad (\text{A10})$$

$$\tilde{n}x_t = \omega(\hat{x}_t - \hat{s}_t - \hat{c}_{f,t}) \quad (\text{A11})$$

$$\hat{\pi}_{h,t} = \beta\mathbb{E}_t[\hat{\pi}_{h,t+1}] + \beta\xi_t + \frac{\epsilon - 1}{\psi_p}(\widehat{mc}_t + \hat{s}_t - \hat{q}_t) \quad (\text{A12})$$

$$\widehat{mc}_t = \hat{w}_t + \hat{n}_t - \hat{y}_t \quad (\text{A13})$$

$$\hat{y}_t = (1 - \omega)\hat{c}_{h,t} + \omega\hat{x}_t \quad (\text{A14})$$

$$\hat{x}_t = \gamma\hat{s}_t \quad (\text{A15})$$

$$\hat{c}_{h,t} = \hat{c}_t - \gamma(\hat{q}_t - \hat{s}_t) \quad (\text{A16})$$

$$\hat{c}_{f,t} = \hat{c}_t - \gamma\hat{q}_t \quad (\text{A17})$$

$$\hat{y}_t = \hat{a}_t + (1 - \alpha)\hat{n}_t \quad (\text{A18})$$

$$\hat{q}_t = (1 - \omega)\hat{s}_t \quad (\text{A19})$$

$$\hat{r}_t = \rho_R\hat{r}_{t-1} + (1 - \rho_R)\phi_\pi\hat{\pi}_{h,t} \quad (\text{A20})$$

$$\hat{a}_t = \rho_A\hat{a}_{t-1} + \bar{\sigma}_A \exp(\sigma_{A,t}/2)\varepsilon_{A,t} \quad (\text{A21})$$

$$\sigma_{A,t} = \rho_{\sigma_A}\sigma_{A,t-1} + \sigma_{\sigma_A}\varepsilon_{\sigma_A,t} \quad (\text{A22})$$

$$\hat{z}_t = \rho_Z\hat{z}_{t-1} + \bar{\sigma}_Z \exp(\sigma_{Z,t}/2)\varepsilon_{Z,t} \quad (\text{A23})$$

$$\sigma_{Z,t} = \rho_{\sigma_Z}\sigma_{Z,t-1} + \sigma_{\sigma_Z}\varepsilon_{\sigma_Z,t}. \quad (\text{A24})$$

The risk wedges are defined as:

$$\mu_t = \frac{1}{2} \text{Var}_t(\hat{v}_{t+1}) \tag{A25}$$

$$\lambda_t = \frac{1}{2} \text{Var}_t(\widehat{\text{sdf}}_{t,t+1}) \tag{A26}$$

$$\psi_t = \frac{1}{2} \text{Var}_t(\Delta \hat{e}_{t+1}) + \text{Cov}_t(\widehat{\text{sdf}}_{t,t+1}, \Delta \hat{e}_{t+1}) \tag{A27}$$

$$\xi_t = \frac{5}{2} \text{Var}_t(\hat{\pi}_{h,t+1}) + \text{Cov}_t(\widehat{\text{sdf}}_{t,t+1}, \hat{\pi}_{h,t+1}) + \text{Cov}_t(\hat{y}_{t+1}, \hat{\pi}_{h,t+1}). \tag{A28}$$

B Additional Empirical Results

This appendix contains the additional empirical results for the large language model (prompt, topic codes, and validation tables), a comparison of the currency risk premium with an outside estimate, output responses for the local projections, and robustness checks for the local projections.

B.1 LLM Classification and Validation

The full prompt for the LLM is presented in Figure B1. Note that the order in which the topics were listed was shuffled for every article prompted, as there was some bias of the LLM toward picking topics earlier in the list. The two validation tables behind Section 2.2 compare hand labels to the LLM and the LLM to keyword searches. Table B1 maps the 19 topic codes into the topic groups. Finally, Figure B2 shows a decomposition of the overall uncertainty index by each topic of uncertainty.

Figure B1: LLM Prompt

--- system message ---

You label UK newspaper articles for a study of economic uncertainty. Return a JSON object with keys "economy", "uncertainty", "source", "topics".

"economy": how much the article is about the economy (business, markets, prices, jobs, trade or economic policy). 0 = not at all or only in passing; 1 = a secondary theme; 2 = a main theme.

"uncertainty": how much the article discusses uncertainty about the future: outcomes that are unclear, unpredictable or at risk. Reports of events, results, decisions or confident forecasts are 0. 0 = none; 1 = mentioned in passing; 2 = a main theme.

"source": if economy and uncertainty are both at least 1, two to four words naming the specific thing whose future is uncertain. Do not use general words such as conditions, outlook, environment, economy or uncertainty. Otherwise "none".

"topics": if economy and uncertainty are both at least 1, one code from the list below for what the economic uncertainty is about. Add a second code only if it is equally central. Otherwise an empty list.

Codes:

consumer: consumer spending, retail sales, household confidence and finances

housing: house prices, mortgages, housebuilding, property market

investment: business investment, firms' expansion or hiring plans

output: firms' sales, orders, production and profits

jobs: employment, job losses, unemployment

wages: pay, labour costs, staff and skills shortages

strikes: strikes, industrial action, pay disputes

energy: oil, gas, electricity and fuel prices or supply

supply: supply chains, shortages, raw material and other input costs

inflation: the general rate of inflation and price outlook

livingcosts: household bills, food prices, rents

monetary: interest rates, central banks, quantitative easing

fiscal: taxes, public spending, budgets, government borrowing

regulation: regulation, energy and climate policy, competition rules

finance: banks, credit, debt, stock, bond and currency markets

trade: trade, tariffs, trade deals and customs

politics: elections, Brexit negotiations, referendums, government stability

geopolitics: war, terrorism, sanctions, international conflict (not domestic politics)

other: anything else

--- user message ---

Title: <article headline>

Text: <first 400 words of the article body>

Note: The full prompt sent to the LLM for each article. The order of the topic codes is shuffled for every article.

Table B1: LLM Topic Codes and Groups

Group	Topic	Definition in the prompt	Articles	Share
Demand	Housing	house prices, mortgages, housebuilding, property market	2,471	7.2%
	Consumer	consumer spending, retail sales, household confidence and finances	1,786	5.2%
	Output	firms' sales, orders, production and profits	1,747	5.1%
	Living costs	household bills, food prices, rents	993	2.9%
Supply	Energy	oil, gas, electricity and fuel prices or supply	1,226	3.5%
	Inflation	the general rate of inflation and price outlook	925	2.7%
	Supply chains	supply chains, shortages, raw material and other input costs	250	0.7%
Financial	Finance	banks, credit, debt, stock, bond and currency markets	4,281	12.4%
	Investment	business investment, firms' expansion or hiring plans	4,163	12.1%
Policy	Fiscal policy	taxes, public spending, budgets, government borrowing	2,357	6.8%
	Monetary policy	interest rates, central banks, quantitative easing	1,969	5.7%
	Regulation	regulation, energy and climate policy, competition rules	276	0.8%
Trade	Trade	trade, tariffs, trade deals and customs	3,003	8.7%
Labor	Jobs	employment, job losses, unemployment	4,694	13.6%
	Wages	pay, labour costs, staff and skills shortages	324	0.9%
	Strikes	strikes, industrial action, pay disputes	35	0.1%
Politics	Politics	elections, Brexit negotiations, referendums, government stability	3,362	9.7%
	Geopolitics	war, terrorism, sanctions, international conflict (not domestic politics)	502	1.5%
Other	Other	none of the listed codes	169	0.5%
Total			34,533	100%

Note: Topic codes the LLM assigns as the main topic of each article, their definitions in the prompt, and the group each is assigned to. Articles counts the articles labelled as about economic uncertainty with that main topic from 2000 to 2025.

Table B2: Hand Validation of the LLM Labels

	Agree with author
<i>A. About economic uncertainty</i>	
Articles the LLM keeps	101 / 105
<i>B. Same topic group (5 articles per LLM topic)</i>	
Supply (energy, inflation, supply chains)	12 / 15
Demand (consumer, housing, output, living costs)	12 / 20
Demand, excluding output	12 / 15
All 19 topic codes	61 / 95


Note: Hand labels made without seeing the LLM’s labels, compared to the LLM labels. Panel A counts the articles the LLM labels as about economic uncertainty that agree with the hand labels. Panel B samples five articles for each LLM topic and counts the articles hand-labelled to the same topic group.

Table B3: LLM vs. Keyword Classification of the Hand-Labelled Articles

	LLM	Keywords	Keyword rule
Economic uncertainty	98%	87%	econom* and uncert*
Demand	62%	28%	Demand terms (consumer, retail, sales, ...)
Supply	88%	17%	Supply terms (strikes, energy prices, ...)
Policy	62%	15%	EPU policy terms (policy, tax, spending, ...)

Note: Precision is the share of articles flagged by each method that the hand labels agree with, weighting each stratified sample to its number of articles in the full pool. The keyword rule column shows the search terms used for each keyword classification; the policy terms are those of [Baker, Bloom, and Davis \(2016\)](#).

Figure B2: Decomposition of Overall Uncertainty by Subtopic




Note: Same as Figure 1, except that panel (b) shows the share of articles assigned to each topic, shaded by topic group.

B.2 Alternative Currency Risk Premium Measure


The baseline currency risk premium in the paper uses the Reuters forecast of the GBP/USD exchange rate 12 months out. Figure B3 compares my estimate using Reuters forecasts to data provided by [Kalemli-Özcan and Soylu \(forthcoming\)](#), who build estimates of the same currency risk premium from Consensus forecast surveys. The two series are very similar.

Figure B4 shows that replacing the Reuters currency risk premia with the Consensus currency risk premia has no impact on the main result. Figure B5 shows that replacing the currency risk premia estimated at a horizon of 12 months with ones estimated at 6 or 3 months does not change the results.




Note: The currency risk premium built from Reuters forecasts (baseline) and from Consensus forecasts, as estimated by [Kalemli-Özcan and Soylu \(forthcoming\)](#). Both are the UK-US 12-month interest rate differential minus the expected 12-month depreciation of the pound against the dollar, annualized.

Figure B4: Robustness: Consensus Currency Risk Premium



Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline currency risk premium, built from Reuters forecasts, with the Consensus-based estimate of [Kalemli-Özcan and Soylu \(forthcoming\)](#). Confidence intervals are 68% for the baseline.


Figure B5: Robustness: Survey Horizon of the Currency Risk Premium



Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline currency risk premium, from 12-month forecasts and one-year rates, with premia from 6-month and 3-month forecasts. Confidence intervals are 68% for the baseline.

B.3 Output Responses

Figure B6 shows the output responses from the baseline local projection, and on the right-hand side the industrial production responses if that is switched into the local projection for output. These two measures do not agree.




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. The industrial production response replaces the UK-US GDP gap with the UK-US industrial production gap. Confidence intervals are 68%.

B.4 Robustness of the Local Projections


Figures in this section re-estimate the baseline local projection from Figure 5 with one change described in its note. I check the measure of the exchange rate, sentiment and global controls, index construction robustness, sample, and local projection design decisions. Figure B7 compares the estimated sentiment of the newspaper articles for supply and demand to the uncertainty indexes. Figure B9 decomposes the exchange rate response estimated in the local projection into movements in the interest rate differential, the currency risk premium (UIP), and changes in the expected level.

Figure B7: Sentiment and Uncertainty Indexes




Note: The demand and supply uncertainty indexes compared to the sentiment of the same articles, measured as positive minus negative terms over positive plus negative terms. All series are standardized six-month moving averages.

Figure B8: Robustness: Sentiment Controls




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline with adding the demand and supply sentiment indexes (trailing three-month averages) as controls when identifying the uncertainty innovations. Confidence intervals are 68% for the baseline.

Figure B9: Exchange Rate Decomposition




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Decomposes the exchange rate response into the responses of the 12-month survey forecast, the one-year interest rate differential, and the currency risk premium, which sum to the exchange rate response at every horizon.

Figure B10: Robustness: Global Uncertainty Controls




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline controls with adding the US economic policy uncertainty and US trade policy uncertainty indexes, and with dropping either the VIX or UK GDP. Confidence intervals are 68% for the baseline.

Figure B11: Robustness: Topic Grouping




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline topic groups with expanded groups, which add investment to demand and move wages, strikes, and living costs into supply, and with demand as only the consumer topic and supply as only the energy topic. Confidence intervals are 68% for the baseline.

Figure B12: Robustness: Labor Uncertainty in Demand or Supply




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline with adding the labor topics (jobs, wages, and strikes) to the demand group or to the supply group. Confidence intervals are 68% for the baseline.

Figure B13: Robustness: Article Selection Gate




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline, which keeps articles the LLM labels as mainly about the economy and uncertainty, with keeping articles labelled as at least partly about both. Confidence intervals are 68% for the baseline.

Figure B14: Robustness: Covid Sample




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK–US GDP gap, UK–US CPI gap, and UK–US three-month rate gap. Compares the full-sample baseline with a sample that ends in December 2019 for both shock and outcome months, and with dropping observations whose outcome month falls in March–June 2020. Confidence intervals are 68% for the baseline.

Figure B15: Robustness: LP Lag Length



Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline three lags with six and twelve lags in the local projection; the uncertainty innovations are unchanged. Confidence intervals are 68% for the baseline.


Figure B16: Robustness: Joint vs. One-at-a-Time Estimation




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline, estimated jointly, with including only the demand and supply innovations, with estimating each one at a time alongside the financial and policy innovations, and with recursively ordering demand or supply first. Confidence intervals are 68% for the baseline.

Figure B17: Robustness: Exchange Rate Measure

(a) Demand Uncertainty




(b) Supply Uncertainty



— GBP/USD, end of month (baseline) — Broad sterling index, end of month

Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK-US GDP gap, UK-US CPI gap, and UK-US three-month rate gap. Compares the baseline end-of-month pound-dollar exchange rate with the end-of-month broad sterling index. Confidence intervals are 68% for the baseline.

Figure B18: Robustness: UK–US Differentials




Note: Estimates the local projection described in equation (4) with demand, supply, financial, and policy uncertainty innovations. The outcome variables included are the pound-dollar exchange rate, the currency risk premium, UK–US GDP gap, UK–US CPI gap, and UK–US three-month rate gap. Compares the baseline UK–US gaps with UK levels of GDP, CPI, and the three-month rate, with the US three-month rate included separately. Confidence intervals are 68% for the baseline.

C Additional Model Impulse Responses


These complete Figures 9, 10, and 11. Each shows six more variables for the same model: output, domestic and CPI price levels, the interest rate, the real exchange rate, and the inflation wedge.

Figure C1: Model Evidence: Additional Impulse Responses to the Two Volatility Shocks




Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. Additional responses for the model in Figure 9.

Figure C2: Model Mechanism: No UIP Risk Premium Channel, Additional Responses



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The currency risk premium channel is manually set to zero. Additional responses for the model in Figure 10.

Figure C3: Model Mechanism: No Precautionary Savings Channel, Additional Responses



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The precautionary savings channel is manually set to zero. Additional responses for the model in Figure 11.

D Unconditional Correlations at the Fitted Calibration

Figures 14 and 15 use symmetric shocks. Table D1 repeats the exercise at the fitted calibration.


Table D1: Unconditional Correlations at the Fitted Calibration

Correlation	At $s = 0.5$	Crosses zero at $s \approx$
SDF and nominal exchange rate, $\text{corr}(\widehat{\text{sdf}}, \Delta \hat{e})$	0.70	0.09
Backus–Smith, $\text{corr}(\Delta \hat{q}, \Delta \hat{c})$	−0.55	0.19

Note: Model calibrated to values in Table 3. As in Figures 14 and 15, the incomplete markets model is simulated for 20,000 months with only volatility shocks at different levels of the demand share of total volatility shock variance, s , with the total variance held fixed.

E Robustness of the Covariance Signs


The opposite currency risk premium signs require $a_A b_A < 0 < a_Z b_Z$ (Table 2). Figure E1 varies seven parameters one at a time to show where those signs hold.



Note: Model calibrated to values in Table 3. Each panel varies one parameter, holding the others at their calibrated values (dotted line), and shows the products of the first-order loadings, $a_k b_k$, that set the sign of each shock's term in equation (34) for preference (blue) and productivity (orange) shocks.


F Accuracy of the Risk-Adjusted Approximation

Figures F1 and F2 compare the risk-adjusted log-linear model with a third-order perturbation of the full nonlinear model, for each volatility shock. The two give the same results qualitatively, with some quantitative differences.



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The solid line shows the risk-adjusted log-linear model and the dashed line a third-order perturbation of the nonlinear model.

Figure F2: Risk-Adjusted Log-Linear vs. Nonlinear: Preference Volatility



Note: Model calibrated to values in Table 3 except that level shock persistence is 0.97 and the volatility shock persistence is 0.96, so that the two shocks are symmetric. The solid line shows the risk-adjusted log-linear model and the dashed line a third-order perturbation of the nonlinear model.