THE UNIVERSITY OF CHICAGO

MODEL UNCERTAINTY AND FOREIGN EXCHANGE PREDICTABILITY: A ROBUST EXPLANATION OF THE FORWARD PREMIUM PUZZLE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
RYAN HUGHES

CHICAGO, ILLINOIS
JUNE 2018

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I am grateful to Lars Hansen, Harald Uhlig, Lawrence Schmidt, and participants in the Capital Theory working group for providing helpful comments and suggestions. Any errors are my own.

ABSTRACT

I show that concerns for robustness against model uncertainty generate predictable, time-varying violations of uncovered interest parity in a multi-country endowment economy where representative agents average over competing sub-models of growth processes to determine interest rates. I test the theory using consumption and foreign exchange data for several countries, and find that the model's measure of expected currency returns generally explains and predicts actual foreign exchange returns relative to the U.S. Dollar better than the time series of currency forward premiums or interest rate differences.

CHAPTER 1

INTRODUCTION

1.1 Overview

The uncovered interest parity no-arbitrage condition states that cross-currency interest rate differences should correspond to offsetting expected exchange rate movements: the currency of a country whose risk-free bonds pay a higher interest rate will tend to depreciate relative to the currency of a country with a lower interest rate. Contrary to this baseline theory, a wide array of studies have found that the opposite regularly occurs: high interest rate currencies tend to *appreciate* against low interest rate currencies. Agents that borrow in low interest rate currencies to fund lending in high interest rate currencies thus earn predictable excess returns. This feature of the data is known in the international macroeconomics and finance literature as the uncovered interest parity puzzle, the carry trade, or the forward premium anomaly.

This paper offers what is to my knowledge a novel explanation for the forward premium anomaly: agents form their expectations about consumption growth to be robust against various aspects of model misspecification by adopting worst-case beliefs in the style of Hansen and Sargent [12]. I do this by extending the setup of [12] to construct a simple two-country endowment economy with complete markets where the representative agent of each country has preferences over consumption of a country-specific good. With complete markets, differences in interest rates are offset by corresponding expected movements in the exchange rate, so there is no time-varying component to predictable interest parity violations from the perspective of agents within the model. However, when agents adjust their beliefs about future growth to be robust against possible model misspecification/parameter uncertainty, an econometrician standing outside of the model sees that this linkage is broken. Specifically, interest rates for each country are priced according to (dynamically) pessimistically distorted beliefs regarding the persistence of innovations to consumption growth, while the actual joint consumption process - and therefore the evolution of the actual exchange rate - does not feature these distortions. This leads to time-varying expected deviations from uncovered interest parity or predictable currency returns from an econometrician's perspective, although agents within the model remain indifferent between holding the risk-free bonds of either country.

The model offers an easily testable prediction: expected log currency returns are a simple linear function of the differences between ordinary Bayesian and robust (worst-case) estimates of consumption growth for the relevant countries. I test the theory by regressing the quarterly returns of the currencies of seven developed countries relative to the U.S. Dollar against this consumption-derived statistic using over a quarter century of data. I compare the results to benchmark regressions of the same currency returns on the corresponding interest rate differences or log forward discount. For each of the currencies studied, under a variety of metrics, I find that the model's key growth uncertainty statistic - derived from macroeconomic data - explains more of the variation in currency returns, and does so in a more convincing fashion, than the log forward discount. I go on to test the model out-of-sample by comparing the cumulative returns of currency trading strategies based on this paper's consumption growth uncertainty statistic and the ordinary interest rate differential, and I find that the former offers generally better performance across the various currencies for a range of initial learning periods.

The measures of growth uncertainty referenced above are found by applying the Kalman filter to a collection of simple models and comparing the model-averaged expected income growth rate at each point in time with and without a robust probability adjustment. For each country, the endowment growth process takes is a discrete-time version of the hidden Markov model from [12], where observable growth is the sum of a constant unconditional mean and a stationary $AR(1)$ process of unknown persistence. (No relationship is assumed for the persistences of the hidden states across countries.) Neither state variable is visible to the agents inside the model or to the econometrician. This complicates the process of valuing exposure to shocks, since a wide range of persistence parameters with substantially different implications for long-run dynamics have similar likelihoods in the data. To reflect uncertainty regarding the true values of the persistence parameters, agents within the model learn about the data-generating processes over time from observable signals and use this information to update their Bayesian probabilities over a finite collection of models. (I make the assumption that the model set is finite so that the analytical work in Section 2 is congruent with the empirical work in Section 3.) Specifically, agents use the Kalman filter to obtain estimates of the hidden state associated with each growth model they entertain, where all models are constructed so that the sum of the components of the hidden state is identically equal to the expectation of the logarithm of the next period's income growth. Ordinary growth estimates are obtained by weighting the Kalman filter estimates associated with various models by their Bayesian likelihoods. The worst case estimate of growth that helps agents cope with model uncertainty uses the same Kalman filter output, but instead applies a dynamic distortion that magnifies the agents' perceptions of the permanence of bad news while diminishing the agents' perceptions of the permanence of good news. This mirrors what was used in [12] and is developed later in the paper.

The paper is organized as follows: Section 1.2 provides a short overview of related literature. Section 2 describes preferences, endowment dynamics, and the particular variety of uncertainty that agents in the model face. It then solves a planner problem and a competitive economy for agents with the worst-case beliefs described above before presenting the main theoretical result relating predictable interest parity violations to time-varying differences in ordinary and worst-case model-averaged estimates of consumption growth. Section 3 motivates the problem of model selection and tests the implications of the theory. It begins by using bootstrap simulations to generate confidence intervals for the persistence parameters of the eight developed countries in the data. These confidence intervals are then used to form the sets of growth models by maximum likelihood estimation on the remaining parameters. Using these models, I build the model uncertainty statistic described above and compare the results from regressions of currency returns on this statistic against the benchmark regressions of currency returns on the forward premium. Next, I explore the ability of other types of model uncertainty to generate similar results. I then show that the model uncertainty statistic is more than just a proxy for observed interest rates. Finally, I consider the out-of-sample performance of an uncertainty-based trading strategy relative to a baseline forward-premium based strategy. Section 4 concludes. The appendix contains more detailed derivations than are presented in the main paper and additional sensitivity analyses for the empirical work.

1.2 Related Literature

Two foundational papers documenting the forward premium anomaly are Hansen and Hodrick [10] and Fama [9]. Researchers have offered a variety of explanations for the puzzle. These include asset market segmentation (Alvarez, Atkeson, and Kehoe [1]), recursive preferences and long-run risks (Colacito and Croce, [6],[7]), variations in country size as a share of the world economy (Hassan [14]), consumption habits (Verdelhan [23]), and combinations of foreign and domestic Taylor rules (Backus, Gavazzoni, Telmer, and Zin [3]). Scholl and Uhlig [21] demonstrate that the forward premium puzzle is robust in the absence of delayed overshooting reactions to monetary policy shocks, and document that Sharpe ratios for bets on uncovered interest parity violations are markedly larger than those seen in US equity market data.

Another strain of papers examine predictable returns to foreign exchange speculation from a finance perspective. See Backus, Foresi, and Telmer [2] for a look at implications of the forward premium anomaly for affine term structure models. Lustig, Roussanov, and Verdelhan [20] provide an overview of this literature and demonstrate a trading strategy that generates predictable excess returns for US investors by shorting the Dollar when the US price of risk is high. Lustig, Roussanov, and Verdelhan [19] find a slope factor in exchange rates and show that it is related to volatility changes in global equity markets. Recently, Burnside and Graveline [5] have offered a critique of this popular asset market view of exchange rates, advocating instead for fully-specified exchange rate models.

A number of authors have begun to explore other types of uncertainty as an explanation of predictable uncovered interest parity violations. Ismailov and Rossi [17] build an index based on the extent to which exchange rates are more or less difficult to forecast based on their recent past. Berg and Mark [4] consider a series of indices including newspaper-based measures of uncertainty. Husted, Rogers, and Sun [15] consider a variety of sources including option and newspaper data to show that higher uncertainty correlates with larger carry trade excess returns. Isrefi and Mouabbi [18] construct an index of uncertainty using professional forecasts. Ilut [16] features a model with ambiguity averse agents that imagine interest rate shocks follow a process with time-varying volatilities drawn from a finite set.

Of the papers listed above, this paper is most similar in spirit to those described in the last paragraph. The main difference is that our measure of uncertainty is derived entirely from consumption data, and all competing models have constant and equal volatility of shocks to the observable component of the state. Instead, I focus on time-varying beliefs about the persistence of shocks (to a hidden state) to generate predictable violations of uncovered interest parity.

CHAPTER 2

MODEL STRUCTURE AND SOLUTION

I modify the continuous-time, one-country endowment economy structure of Hansen and Sargent [12] by adapting it to a discrete-time framework with two countries and two varieties of consumption good.

2.1 Preferences and Endowment Dynamics

The two countries are labeled by H (Home) and F (Foreign). It will often be helpful to index countries or country-specific variables by $j \in \{H, F\}$. Each country is populated by a continuum of agents. Agents within a country are identical: all have the same preferences, and all receive the same exogenous endowment stream. Agents differ across countries in two ways. The first of these is that the per-capita endowment processes for each country deliver distinct varieties of consumption good. The second manner in which agents differ across countries is that the preferences of all agents are characterized by complete home bias, so that the marginal utility for each agent of consuming the other country's good is always zero. Agents in country j have the following period t utility function over consumption of their country-specific good C_t^j:

$$U_t^j \left(C_t^j \right) = (1 - \beta) \left(\log C_t^j \right) \tag{2.1}$$

Write the per-capita exogenous endowment received by agents in country j of the j-type good at time t as Y_t^j. The growth processes for the two endowment streams each depend on a country-specific vector of two hidden state variables. Let X_t^j be the vector of hidden state variables for country j at time t. The first component of this, $X_t^{j,[1]}$ follows a stationary $AR(1)$ process with an unknown persistence parameter, while the second component, $X_t^{j,[2]}$, is a constant unconditional mean growth rate. The assumption that agents do not know the persistence of the hidden state is central to everything that follows. Log endowment growth of each country's good is the sum of the entries of its state vector plus a temporary shock. The autoregressive component of each country's hidden state vector is also exposed to shocks. Write the shock for country j at time $t+1$ as W_{t+1}^{j}.

Agents do not have full confidence that a single model correctly describes the growth process for either country. Instead, agents believe that for each country j, one of a finite set of K models[1] governs the evolution of the endowment process, and agents update their beliefs concerning the relative likelihoods of these various models whenever new data arrives. Each of these models has the functional form described above with different coefficient matrices. Throughout the paper, I refer to a country-model combination by the pair (j,k), $j \in \{H,F\}$, $k \in \{1,\ldots,K\}$. No relationship is assumed between models (H,k) and (F,k), although at times it will be useful to add subscripts k_h and k_f for additional clarity. The k^{th} model describing income growth for country j is:

$$\begin{aligned} \Delta \log Y_{t+1}^{j} \equiv \log Y_{t+1}^{j} - \log Y_{t}^{j} &= DX_{t}^{j} + G(j)\, W_{t+1}^{j} \\ X_{t+1}^{j} &= A(j,k)\, X_{t}^{j} + B(j,k)\, W_{t+1}^{j} \\ W_{t+1}^{j} &\sim \mathcal{N}(0, I_2), \quad j \in \{H,F\}, \quad k \in \{1,\ldots,K\} \end{aligned} \tag{2.2}$$

The coefficient matrices in (2.2) are given by:

$$\begin{aligned} D &= \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad G(j) = \begin{bmatrix} 0 & g(j) \end{bmatrix} \\ A(j,k) &= \begin{bmatrix} \rho(j,k) & 0 \\ 0 & 1 \end{bmatrix}, \quad B(j,k) = \begin{bmatrix} b(j,k) & 0 \\ 0 & 0 \end{bmatrix} \end{aligned} \tag{2.3}$$

Each country has at most one model with any particular value of ρ, and in what follows it will occasionally be convenient to identify a model by the value of its $AR(1)$ coefficient, so I will refer to model (j,k) and model $\rho(j,k)$ or $\rho_j(k)$ interchangeably.

Agents use the Kalman filter to update their beliefs about the hidden state and the

1. It is not essential that the model sets be the same size, although this is a natural assumption and is maintained in the empirical section of the paper.

Bayesian probability distribution over models. At time t, Associated with each country-model pair, there is a filtered estimate of the hidden state vector $\overline{X}_t^j(k)$, a covariance matrix for this estimate $\Sigma_t(j,k)$, and a Bayesian probability $q_t^j(k)$ of that model generating the history of observed data.

It will be convenient in what follows to be able to compactly describe the joint dynamics of endowment growth for both countries. Toward this goal, write the combined state vector for country j's per-capita endowment growth process as:

$$Z_t^j = \left[\log Y_t^j \;\; \log Y_{t-1}^j \;\; X_t^{j,[1]} \;\; X_t^{j,[2]}\right]' \tag{2.4}$$

This allows us to write the state vector for the entire aggregate real economy at time t as

$$Z_t' = \left[\left(Z_t^H\right)' \;\; \left(Z_t^F\right)'\right] \tag{2.5}$$

Then, stack the vector of shocks together to write:

$$W_{t+1} \equiv \begin{bmatrix} W_{t+1}^H \\ W_{t+1}^F \end{bmatrix}, \quad W_{t+1} \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{I}_4\right) \tag{2.6}$$

From this, the evolution of the real state of the economy follows:

$$Z_{t+1} = \Gamma Z_t + \Phi W_{t+1} \tag{2.7}$$

Here Γ is an 8×8 matrix and Φ is an 8×4 matrix, the entries of which are given by:

$$\Gamma = \begin{bmatrix} \Gamma_H & \mathbf{0}_4 \\ \mathbf{0}_4 & \Gamma_F \end{bmatrix}, \quad \Gamma_j = \begin{bmatrix} \Xi & \Xi' \\ \mathbf{0}_2 & A(j) \end{bmatrix}, \quad \Xi = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \tag{2.8}$$

while Φ is constructed with:

$$\Phi = \begin{bmatrix} \Phi_H \\ \Phi_F \end{bmatrix}$$
$$\Phi_H = \begin{bmatrix} \vec{g}^H & \vec{b}^H & \vec{0}_{4x1} & \vec{0}_{4x1} \end{bmatrix}$$
$$\vec{g}^H = \begin{bmatrix} g(H) & 0 & 0 & 0 \end{bmatrix}', \ \vec{b}^H = \begin{bmatrix} 0 & 0 & b(H) & 0 \end{bmatrix}', \tag{2.9}$$
$$\Phi_F = \begin{bmatrix} \vec{0}_{4x1} & \vec{0}_{4x1} & \vec{g}^F & \vec{b}^F \end{bmatrix}$$
$$\vec{g}^F = \begin{bmatrix} g(F) & 0 & 0 & 0 \end{bmatrix}', \ \vec{b}^F = \begin{bmatrix} 0 & 0 & b(F) & 0 \end{bmatrix}',$$

Each Γ_j above is a 4×4 matrix, $\mathbf{0}_n$ is an $n \times n$ matrix of zeros, Ξ is a 2×2 matrix with scalar entries listed as above, and $A(j)$ has the form given in (2.2) with the (unknown) true value of ρ_j. All entries in $\vec{g}^j$ and $\vec{b}^j$ are scalars, all entries in Φ_j are 4×1 column vectors, each Φ_j is a 4×4 matrix. To describe the evolution of the real state of the economy conditional on any particular pair of models for H and F, replace the appropriate entires of the matrices Γ and Φ.

Some additional notation will help in the next section's study of pricing. Write the observable signal at time t as

$$s_t = \left[\log Y_t^H \quad \log Y_t^F\right]' \tag{2.10}$$

and the signal history up through time t as s^t. As in [11], $\mathcal{S}_t$ is the filtration generated by the signal history through time t, $s_0, \ldots, s_t$, and $\mathcal{X}_t$ is the filtration generated by $Z_0, W_0, \ldots, W_t$. Describe the full state of the economy by collecting the real state of the economy, the current output from the Kalman filter that agents use to construct their beliefs, and the true values of the persistence parameters together as follows:

$$\zeta_t = \left\{ Z_t, \left\{ \overline{X}_t^j(k), \Sigma_t(j,k), q_t^j(k) \right\}_{1 \le k \le K}^{j \in \{H,F\}}, \rho_H, \rho_F \right\} \tag{2.11}$$

Partition this into an $\mathcal{S}_t$-measurable component $\widehat{\zeta}_t$, together with the information that is hidden from the agents and the econometrician. Then we can write:

$$\zeta_t = \left\{ \widehat{\zeta}_t, X_t, \rho_H, \rho_F \right\} \tag{2.12}$$

where $\widehat{\zeta}_t$ includes the current signal and the current collected output from the Kalman filter, and X_t is the collected hidden state vector

$$X_t' = \left[\left(X_t^H \right)' \left(X_t^J \right)' \right] \tag{2.13}$$

Conditional on $\mathcal{S}_t$, the (model-averaged) expected log endowment growth of the numeraire good for agents of country j between times t and $t+1$ is distributed $\mathcal{N}\left(\overline{\kappa}_{j,t}, g^2(j)\right)$, with:

$$\begin{aligned}\overline{\kappa}_t^j &\equiv \mathbb{E}\left[\log Y_{t+1}^j - \log Y_t^j \middle| \mathcal{S}_t\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\log Y_{t+1}^j - \log Y_t^j \middle| \text{ model } (j,k), \mathcal{S}_t\right]\middle| \mathcal{S}_t\right] \\ &= \sum_{k=1}^{K_j} \mathbb{E}\left[\log Y_{t+1}^j - \log Y_t^j \middle| \text{ model } (j,k), \mathcal{S}_t\right] \Pr\left(\text{model } (j,k) | \mathcal{S}_t\right) \\ &= \sum_{k=1}^{K_j} D\overline{X}_t^j(k) q_t^j(k) = \sum_{k=1}^{K_j} \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \overline{X}_t^{j,[1]}(k) \\ \overline{X}_t^{j,[2]}(k) \end{bmatrix} q_t^j(k)\end{aligned} \tag{2.14}$$

2.2 Robust Planner's Problem

We begin by studying the problem of an imaginary robust social planner for each country. Like the agents, the planner does not know the true value of ρ_j and is limited to the information set $\mathcal{S}_t$. In order to cope with model uncertainty, the planner considers playing a dynamic zero-sum game against a minimizing player (malevolent Nature) to construct a sequence of worst-case beliefs concerning the probability distribution over models of growth. Neither player is forced to commit to a particular sequence of actions for all time. Instead, each player is free to choose a new action at each time period as new information arrives.

This means that both players must consider the actions of their contemporaneous opponent and also the actions of future maximizing and minimizing players. The solution concept for this game is a Markov perfect equilibrium as outlined in [11]. Since there is no storage technology and marginal utility is always positive, the maximizing planner will have the single representative agent consume the entire endowment of their variety of good in each period, or

$$C_t^j = Y_t^j \tag{2.15}$$

This problem is computationally simple since the action of the maximizing player does not depend on the behavior of the minimizing player, but studying the panner's problem will help prepare us for analyzing the competitive economy in the next section. The sequence of worst-case beliefs evolves as follows: at each time t, the planner for country j receives a new signal and updates his or her set of Bayesian probabilities associated with the various models of endowment growth. Nature then twists these beliefs[2] in such a way as to choose from among a set of statistically close distributions the one that will lead to the lowest level of expected lifetime utility for the agents within the model, taking into account the future behavior of both maximizing and minimizing players. Specifically, the planner imagines a malevolent agent choosing a relative probability distribution over models of the unobservable hidden state in such a way as to minimize the sum of agents' (model-averaged) lifetime utility together with the relative entropy between the distorted and baseline probability distributions over models. With this decision rule for the maximizing player in the planner's problem, the value function updates as:

$$V^j(\zeta_t) = (1-\beta)\log Y_t^j + \beta \mathbb{E}\left[V^j(\zeta_{t+1}) \middle| \mathcal{X}_t\right] \tag{2.16}$$

To help the maximizing player choose a set of beliefs that are robust against model uncertainty, the minimizing player chooses an $\mathcal{S}_t$-measurable distortion to the probability

2. Sections 3.4 and A.2 explore other varieties of model misspecification.

distribution over models of growth that twists away from the Bayesian probabilities to minimize the sum of the expected conditional continuation value plus the penalized relative entropy term:

$$\min_{\left\{\hat{q}_t^j(k)\right\}_{k=1}^K} \sum_{k=1}^K \hat{q}_t^j(k) \left(\mathbb{E}\left[V(\zeta_t) \middle| \mathcal{S}_t, \rho_j(k) \right] + \theta_2 \log\left(\frac{\hat{q}_t^j(k)}{q_t^j(k)} \right) \right) \tag{2.17}$$

I will solve the minimizing player's problem (2.17) first, then I will guess and verify a solution to the value function recursion (2.16). To make notation easier in what follows, introduce the abbreviation below for the expected[3] level of the country j individual's value function at time t given that the true data-generating process for country j is model (j,k):

$$\widehat{V}^j(\zeta_t, k) \equiv \mathbb{E}\left[V^j(\zeta_t) \middle| \mathcal{S}_t, \rho_j(k) \right] \tag{2.18}$$

Using (2.18), write the Lagrangian for the problem of the minimizing player (2.17) as:

$$\mathcal{L} = \sum_{k=1}^K \hat{q}_t^j(k) \left(\widehat{V}^j(\zeta_t, k) + \theta_2 \log\left(\frac{\hat{q}_t^j(k)}{q_t^j(k)} \right) \right) + \widetilde{\eta} \left(1 - \sum_{k=1}^K \hat{q}_t^j(k) \right) \tag{2.19}$$

The minimizing player's optimality condition with respect to the twisted probability on model k' is:

$$\frac{\partial \mathcal{L}}{\partial \hat{q}_t^j(k')} = \widehat{V}^j(\zeta_t, k') + \theta_2 \left(\log\left(\frac{\hat{q}_t^j(k')}{q_t^j(k')} \right) + \hat{q}_t^j(k') \frac{1}{\frac{\hat{q}_t^j(k')}{q_t^j(k')}} \frac{1}{q_t^j(k')} \right) - \widetilde{\eta} = 0 \tag{2.20}$$

The second term in the parentheses above reduces to one, so we can simplify (2.20) further as

$$\frac{\partial \mathcal{L}}{\partial \hat{q}_t^j(k')} = \widehat{V}^j(\zeta_t, k') + \theta_2 \left(\log\left(\frac{\hat{q}_t^j(k')}{q_t^j(k')} \right) + 1 \right) - \widetilde{\eta} = 0 \tag{2.21}$$

3. Note that this allows the value function to depend on variables which are not measurable with respect to the observed information set $\mathcal{S}_t$, and also limits the malevolent player to making a choice based on observable signals.

Rearranging, (2.21) becomes

$$-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k'\right) = \log\left(\frac{\hat{q}_t^j\left(k'\right)}{q_t^j\left(k'\right)}\right) - \left(\frac{\widetilde{\eta}}{\theta_2} - 1\right) \tag{2.22}$$

To ease notation in what follows, introduce a transformation of the Lagrange multiplier:

$$\eta \equiv \left(\frac{\widetilde{\eta}}{\theta_2} - 1\right) \tag{2.23}$$

Then (2.22) becomes

$$-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k'\right) = \log\left(\frac{\hat{q}_t^j\left(k'\right)}{q_t^j\left(k'\right)}\right) - \eta \tag{2.24}$$

Exponentiate (2.24) and multiply both sides by $q_t^j\left(k'\right)$ to obtain:

$$q_t^j\left(k'\right)\exp\left(-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k'\right)\right) = \hat{q}_t^j\left(k'\right)\exp\left(-\eta\right) \tag{2.25}$$

Now, add (2.25) over all models of growth for country j, $k = 1, \ldots, K$, and use the fact that the sum of the twisted model probabilities must equal one to obtain:

$$\sum_{k=1}^{K} q_t^j\left(k\right)\exp\left(-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k\right)\right) = \exp\left(-\eta\right) \tag{2.26}$$

Finally, by substituting (2.26) in (2.25) and rearranging, we obtain the twisted probability for model k' of growth for country j at time t,

$$\widetilde{q}_t^j\left(k'\right) = \frac{q_t^j\left(k'\right)\exp\left(-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k'\right)\right)}{\sum_{k=1}^{K} q_t^j\left(k\right)\exp\left(-\theta_2^{-1}\widehat{V}^j\left(\zeta_t, k\right)\right)} \tag{2.27}$$

where $\widetilde{q}_t^j\left(k\right)$ is the minimizing value for $\hat{q}_t^j\left(k\right)$.

Now we will solve the value function recursion as promised. Recall from (2.16) that the value function updates via

$$V^{j}\left(\zeta_{t}\right)=\left(1-\beta\right)\log Y_{t}^{j}+\beta\mathbb{E}\left[V^{j}\left(\zeta_{t+1}\right)\Big|\mathcal{X}_{t}\right] \tag{2.28}$$

Guess a solution to (2.28) of the form:

$$V^{j}\left(\zeta_{t}\right)=\log Y_{t}^{j}+\lambda'\left(j\right)X_{t}^{j} \tag{2.29}$$

Then, with dynamics given by (2.2), we have

$$\begin{aligned}\beta\mathbb{E}\left[V^{j}\left(\zeta_{t+1}\right)\Big|\mathcal{X}_{t}\right]&=\beta\mathbb{E}\left[\log Y_{t+1}^{j}+\lambda'\left(j\right)X_{t+1}^{j}\Big|\mathcal{X}_{t}\right]\\&=\beta\left(\log Y_{t}^{j}+DX_{t}^{j}+\lambda'\left(j\right)A\left(j\right)X_{t}^{j}\right)\end{aligned} \tag{2.30}$$

Combining (2.28) and (2.30) gives

$$V^{j}\left(\zeta_{t}\right)=\log Y_{t}^{j}+\beta\left(D+\lambda'\left(j\right)A\left(j\right)\right)X_{t}^{j} \tag{2.31}$$

Equating coefficients in (2.29) and (2.31) gives the system:

$$\lambda'\left(j\right)=\beta D+\beta\lambda'\left(j\right)A(j) \tag{2.32}$$

or

$$\beta=\lambda^{[1]}\left(1-\beta\rho_{j}\right) \tag{2.33}$$

and

$$\beta=\lambda^{[2]}\left(1-\beta\right) \tag{2.34}$$

where $\lambda^{[n]}$ is the n^{th} component of λ.

Combining the results above, we see that the value function does in fact have the form written in (2.29), so that

$$V^{j}\left(\zeta_{t}\right)=\log Y_{t}^{j}+\lambda^{\prime}\left(j\right)X_{t}^{j} \tag{2.35}$$

with

$$\lambda^{\prime}(j)=\left[\begin{array}{cc}\frac{\beta}{1-\beta\rho(j)} & \frac{\beta}{1-\beta}\end{array}\right] \tag{2.36}$$

where X_{t}^{j} is the hidden state governing the growth of country j's endowment at time t, and $\rho(j)$ is the unknown autoregressive parameter for the time-varying component of country j's hidden state vector under the true model.

Given the realized signal history $\mathcal{S}_{t}$, the expected level of the value function (2.35) conditional on model (j,k) being the true data-generating process is given by:

$$\widehat{V}^{j}\left(\zeta_{t},k\right)\equiv\mathbb{E}\left[V^{j}\left(\zeta_{t}\right)\Big|\mathcal{S}_{t},\rho_{j}\left(k\right)\right]=\log Y_{t}^{j}+\lambda^{\prime}(j,k)\overline{X}_{t}^{j}(k) \tag{2.37}$$

where $\overline{X}_{t}^{j}\left(k\right)$ is the estimate of X_{t}^{j} obtained via the Kalman filter supposing that model k generated the data, and $\lambda\left(j,k\right)$ includes the extra argument to highlight that the value of ρ for country j is assumed to be the k^{th} element of the grid of possible values:

$$\lambda^{\prime}(j,k)=\left[\begin{array}{cc}\frac{\beta}{1-\beta\rho_{j}(k)} & \frac{\beta}{1-\beta}\end{array}\right] \tag{2.38}$$

Now we can use (2.37) and (2.38) to compute the values of the worst-case model probabilities. First, observe that

$$\begin{aligned}\exp\left(-\theta_{2}^{-1}\widehat{V}^{j}\left(\zeta_{t},k\right)\right)&=\exp\left(-\theta_{2}^{-1}\log Y_{t}^{j}-\theta_{2}^{-1}\lambda^{\prime}\left(j,k\right)\overline{X}_{t}^{j}\left(k\right)\right)\\&=\left(Y_{t}^{j}\right)^{-\theta_{2}^{-1}}\exp\left(-\theta_{2}^{-1}\lambda^{\prime}\left(j,k\right)\overline{X}_{t}^{j}\left(k\right)\right)\end{aligned} \tag{2.39}$$

Then, substituting every instance of the first term of (2.39) with the third term of (2.39) in (2.27) and cancelling the common power of current income Y_{t}^{j} from both the numerator and the denominator, we obtain:

$$\widetilde{q}_t^j(k) = \frac{q_t^j(k)\exp\left(-\frac{1}{\theta_2}\widehat{V}^j(\zeta_t,k)\right)}{\sum_p q_t^j(p)\exp\left(-\frac{1}{\theta_2}\widehat{V}^j(\zeta_t,p)\right)} = \frac{q_t^j(k)\exp\left(-\theta_2^{-1}\lambda'(j,k)\,\overline{X}_t^j(k)\right)}{\sum_p q_t^j(p)\exp\left(-\theta_2^{-1}\lambda'(j,p)\,\overline{X}_t^j(p)\right)} \tag{2.40}$$

The exponential tilting above places heavier weights on models for which the associated value function is closer to $-\infty$ at time t, and lower weights on models for which the associated value function is closer to $+\infty$. This adjustment takes place anew each time period as agents update their estimates of the hidden state. The nature of the agent's pessimism is not that models with some particular levels of persistence are always weighted more heavily than ordinary Bayesian probabilities would suggest, and some other ranges of persistence are always weighted less heavily. Instead, the models that would lead to lower lifetime utility at any particular instant are regarded to be relatively more likely, with no consideration for prior (or future) adjustments. Mechanically, during good times, if we suppose (counterfactually[4]) that filtered estimates of the hidden state are positive and equal across submodels, the value function will be increasing as a function of the model index k or equivalently in the size of the persistence parameter ρ. Twisted probabilities are then relatively higher for models with low persistence, and relatively lower for models with high persistence. In bad times, we may imagine that the reverse is true, and the value function becomes a decreasing function of the persistence parameter ρ. During bad times, then, agents' pessimism causes them to attach higher probabilities to models with a more persistent hidden state, and lower probabilities to models with a less persistent hidden state. Relative to ordinary Bayesian beliefs, the permanence of good news is diminished, while the permanence of bad news is exaggerated.

Under the distorted probabilities (2.40), the (model-averaged) expected log endowment

4. Of course, estimates of the hidden state will in general vary across competing submodels, but we abstract away from this in order to more clearly illustrate the mechanism.

growth of country j's numeraire good is:

$$\begin{aligned}
\widetilde{\kappa}_t^j &\equiv \widetilde{\mathbb{E}}^j \left[\log Y_{t+1}^j - \log Y_t^j \middle| \mathcal{S}_t \right] \\
&= \widetilde{\mathbb{E}}^j \left[\mathbb{E} \left[\log Y_{t+1}^j - \log Y_t^j \middle| \rho_j(k), \mathcal{S}_t \right] \middle| \mathcal{S}_t \right] \\
&= \sum_{k=1}^{K_j} \mathbb{E} \left[\log Y_{t+1}^j - \log Y_t^j \middle| \rho_j(k), \mathcal{S}_t \right] \widetilde{\Pr} \left(\rho_j(k) \middle| \mathcal{S}_t \right) \\
&= \sum_{k=1}^{K_j} D\overline{X}_t^j(k) \widetilde{q}_t^j(k) = \sum_{k=1}^{K_j} \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \overline{X}_t^{j,[1]}(k) \\ \overline{X}_t^{j,[2]}(k) \end{bmatrix} \widetilde{q}_t^j(k)
\end{aligned} \tag{2.41}$$

where $\widetilde{\mathbb{E}}^j$ indicates the expectation under the distorted or worst-case model probabilities. More generally, as in [12], we can calculate the expectation of any random variable under the worst-case model probabilities by evaluating the expectation under the regular Bayesian probabilities of the product of that random variable with the likelihood ratio martingale for worst-case model-averaged growth relative to baseline model-averaged growth. In the next section this will allow us to compute prices with a concern for robustness against model uncertainty. To see how to do this, decompose the signal for country j at time $t+1$:

$$\Delta \log Y_{t+1}^j = \overline{\kappa}_t^j + g(j) \overline{W}_{t+1}^j \tag{2.42}$$

(Recall that $\overline{\kappa}_t^j$ is expected log growth under the ordinary Bayesian probabilities given by (2.14).) Then we have that country j's likelihood ratio martingale for worst-case model probabilities relative to Bayesian model probabilities is given by:

$$M_{t+1}^j \equiv \exp \left(g_2^{-1}(j) \left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j \right) \overline{W}_{t+1}^j - \frac{1}{2} \left[g_2^{-1}(j) \left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j \right) \right]^2 \right) \tag{2.43}$$

Now, let the payoff of an arbitrary asset at time $t+1$ be given by π_{t+1}. Given the ordinary (no concern for model uncertainty) stochastic discount factor for agents in country j, S_{t+1}^j, form a candidate time t price for the payoff π_{t+1} by computing the expectation of the payoff times the country j robust discount factor, and write this as

$$p_t^j(\pi_{t+1}) = \mathbb{E}\left[M_{t+1}^j \beta \frac{Y_t^j}{Y_{t+1}^j} \pi_{t+1} \middle| \mathcal{S}_t \right] = \widetilde{\mathbb{E}}^j \left[\beta \frac{Y_t^j}{Y_{t+1}^j} \pi_{t+1} \middle| \mathcal{S}_t \right] \tag{2.44}$$

where either of the last two expressions above are equivalent to:

$$p_t^j(\pi_{t+1}) = \sum_{k=1}^{K} \widetilde{q}_t^j(k)\, \mathbb{E}\left[\beta \frac{Y_t^j}{Y_{t+1}^j} \pi_{t+1} \middle| \mathcal{S}_t, \rho_j(k) \right] \tag{2.45}$$

In the next section, I show that these candidate prices are in fact equilibrium prices for Arrow securities when agents adopt worst-case beliefs regarding the probability distribution over models of growth. Later, we will use this to compute how risk-free interest rates for each country vary over time.

2.3 Robust Competitive Economy

Now we are prepared to study a competitive economy where agents in the two countries form their beliefs to be robust against model uncertainty as in the previous section and [12]. Agents in each country j trade a complete set of Arrow securities contingent on the next period's observable signal, and receive an exogenous dividend processes Y_t^j with dynamics given in (2.2). As in the planner's problem, agents imagine themselves playing a game against a malevolent Nature that is able to dynamically distort the probability distribution over models relative to the Bayesian model likelihoods. Importantly, neither the minimizing nor maximizing player are able to observe the true value of the hidden state or the true persistence coefficients, so the distorted probabilities and the maximizing player's actions depend only on the observable signal history $\mathcal{S}_t$.

Given the time-t exogenous state of the economy ζ_t, stochastic discount factor $\widetilde{S}_{t+1}$, and total current wealth $a_t^j \equiv Y_t^j + \pi_t^j$, the agent in country j chooses current consumption C_t^j and Arrow security holdings for each of next period's possible states. The maximizing player works to achieve the highest level of expected utility under the assumption that a minimizing player simultaneously twists the probability distribution over models to lower the lifetime utility of the maximizing player. Write the joint[5] Bayesian probability at time t of model k_h having generated the data for country H and model k_f having generated the data for F as $q_t\left(k_h, k_f\right)$, and write the distorted joint probability for the same models chosen by the malevolent player as $\widetilde{q}_t\left(k_h, k_f\right)$.

The agent's problem at each time t solves:

$$\begin{gathered}\max_{C_t^j, \pi_{t+1}^j(\cdot)} \min_{\{\widetilde{q}_t(k_h,k_f)\}} (1-\beta)\log C_t^j \\ +\sum_{k_h=1}^{K}\sum_{k_f=1}^{K}\widetilde{q}_t\left(k_h,k_f\right)\left(\beta\mathbb{E}\left[V^j\left(a_{t+1}^j,\zeta_{t+1}\right)\middle|\mathcal{S}_t,\rho_H\left(k_h\right),\rho_F\left(k_f\right)\right]\right) \\ +\theta_2\sum_{k_h=1}^{K}\sum_{k_f=1}^{K}\widetilde{q}_t\left(k_h,k_f\right)\log\left(\frac{\widetilde{q}_t\left(k_h,k_f\right)}{q_t\left(k_h,k_f\right)}\right)\end{gathered} \tag{2.46}$$

The minimizing player must choose a valid set of worst-case probabilities, and the maximizing agent faces the budget constraint:

$$\begin{gathered}C_t^j + \mathbb{E}\left[\widetilde{S}_{t+1}a_{t+1}^j\middle|\mathcal{S}_t\right] = Y_t^j + \pi_t^j \equiv a_t^j \\ a_{t+1}^j \equiv \left(Y_{t+1}^j + \pi_{t+1}^j\right)\end{gathered} \tag{2.47}$$

Where $\widetilde{S}_{t+1}$ is a stochastic discount factor that prices future payoffs mentioned above, and asset holdings at the next period are the sum of Arrow security holdings and the exogenous endowment.

(Please note that in (2.47) above, $\widetilde{S}_{t+1}$, a random variable, is not to be confused with $\mathcal{S}_{t+1}$, which is an information set.) After the agent chooses consumption C_t^j and asset

5. Although the ordinary probabilities are independent, I avoid factoring the Bayesian joint distribution for now to keep the presentation of both sets of probabilities consistent. Although the worst-case probability distributions are in fact independent, this is not assumed in the way that (2.46) is written.

holdings to solve (2.46), the value function updates as follows:

$$V^j\left(a_t^j,\zeta_t\right)=(1-\beta)\log C_t^j+\beta\mathbb{E}\left[V^j\left(Y_{t+1}^j+\pi_{t+1}^j,\zeta_{t+1}\right)\middle|\mathcal{X}_t\right] \tag{2.48}$$

A basic form of the Euler equation is immediate:

$$\mathbb{E}\left[\widetilde{S}_{t+1}\middle|\mathcal{S}_t\right]=\sum_{k_h=1}^{K}\sum_{k_f=1}^{K}\widetilde{q}_t\left(k_h,k_f\right)\left(\beta\mathbb{E}\left[\frac{C_t^j}{C_{t+1}^j}\middle|\mathcal{S}_t,\rho_H\left(k_h\right),\rho_F\left(k_f\right)\right]\right) \tag{2.49}$$

We will soon simplify expression on the right hand sides of (2.49). For now, suppose that the stochastic discount factor used to price future payoffs in the agent's budget constraint coincides with the stochastic discount factor from the planner's problem. That is, consider the candidate pricing function:

$$\widetilde{S}_{t+1}^j=\beta M_{t+1}^j\frac{Y_t^j}{Y_{t+1}^j} \tag{2.50}$$

Where M_{t+1}^j is the worst-case likelihood ratio martingale from the planner's problem (2.43):

$$M_{t+1}^j\equiv\exp\left(g_2^{-1}(j)\left(\widetilde{\kappa}_t^j-\overline{\kappa}_t^j\right)\overline{W}_{t+1}^j-\tfrac{1}{2}\left[g_2^{-1}(j)\left(\widetilde{\kappa}_t^j-\overline{\kappa}_t^j\right)\right]^2\right) \tag{2.51}$$

The agent's optimal policy functions are:

$$C_t^j=(1-\beta)\,a_t^j \tag{2.52}$$

$$\pi_{t+1}^j=\frac{\beta}{1-\beta}Y_{t+1}^j \tag{2.53}$$

Then, with prices given by (2.50), and budget constraint (2.47), policy functions (2.52) and (2.53) solve the individual's maximization problem with associated value function

$$V^j\left(a_t^j,\zeta_t\right)=\log\left(a_t^j\right)+\lambda'(j)\,X_t^j+\log(1-\beta) \tag{2.54}$$

While the same worst-case beliefs as in the planner's problem solve the minimization problem (2.46).

Note that market clearing requires that $C_t^j = Y_t^j$, so using the rule (2.52) above we must have

$$a_t^j = \frac{Y_t^j}{1-\beta} \tag{2.55}$$

In particular, recalling the definition $a_0^j \equiv Y_0^j + \pi_0^j$, the starting level of security holdings must be

$$\pi_0^j = \frac{\beta}{1-\beta} Y_0^j \tag{2.56}$$

Section A.1 verifies the form of the value function (2.54) above. To see why the worst-case probabilities are the same, given the value function and prices, observe first that

$$\begin{aligned} \mathbb{E}\left[V^j\left(a_{t+1}^j, \zeta_{t+1}\right)\middle|\mathcal{S}_t, \rho_H\left(k_h\right), \rho_F\left(k_f\right)\right] &= \mathbb{E}\left[V^j\left(a_t^j, \zeta_t\right)\middle|\mathcal{S}_t, \rho_j = \rho_j\left(k\right)\right] \\ &\equiv \widehat{V}_t^j\left(a_t^j, \zeta_t, k\right) \end{aligned} \tag{2.57}$$

The last term has been introduced as an abbreviation to make the notation easier in what follows. In words, the value function (or conditional value function) of a given country does not depend on the persistence of growth shocks to the other country's dividend stream, but only on the persistence of shocks to its own dividend stream. With this observation, first-order conditions in (2.46) can be manipulated in such a way as to reveal that the problem of the country j minimizing player reduces to choosing a worst-case marginal distribution over the models of growth for country j, so that the algebra follows the derivation for the planner's problem very closely.

In this case, the individual's worst-case probabilities are:

$$\begin{aligned}\widetilde{q}_t^j(k) &= \frac{q_t^j(k)\exp\left(-\frac{1}{\theta_2}\widehat{V}^j(a_t^j,\zeta_t,k)\right)}{\sum_p q_t^j(p)\exp\left(-\frac{1}{\theta_2}\widehat{V}^j(a_t^j,\zeta_t,p)\right)}\\ &= \frac{q_t^j(k)\exp\left(-\theta_2^{-1}\left(\log\left(a_t^j\right)+\lambda'(j,k)\overline{X}_t^j(k)+\log(1-\beta)\right)\right)}{\sum_p q_t^j(p)\exp\left(-\theta_2^{-1}\left(\log\left(a_t^j\right)+\lambda'(j,p)\overline{X}_t^j(p)+\log(1-\beta)\right)\right)}\\ &= \frac{q_t^j(k)\exp\left(-\theta_2^{-1}\lambda'(j,k)\overline{X}_t^j(k)\right)}{\sum_p q_t^j(p)\exp\left(-\theta_2^{-1}\lambda'(j,p)\overline{X}_t^j(p)\right)}\end{aligned} \tag{2.58}$$

The first equality above comes from the same sequence of steps as we used to solve for the minimizing probabilities in the planner's problem, since the steps did not rely on the particular form or level of the value function. The second equality follows by substituting the particular form of the individual's value function, and replacing the hidden state with its filtered estimate to reflect conditioning on a particular model. The final equality comes from factoring out common terms and canceling.

With the above, the right side of the Euler equation for Arrow securities reduces to:

$$\begin{aligned}&\sum_{k_h=1}^{K}\sum_{k_f=1}^{K}\widetilde{q}_t(k_h,k_f)\left(\beta\mathbb{E}\left[\frac{C_t^j}{C_{t+1}^j}\middle|\mathcal{S}_t,\rho_H(k_h),\rho_F(k_f)\right]\right)\\ &= \sum_{k_j=1}^{K}\widetilde{q}_t^j(k)\left(\beta\mathbb{E}\left[\frac{C_t^j}{C_{t+1}^j}\middle|\mathcal{S}_t,\rho_j(k)\right]\right) = \sum_{k_j=1}^{K}\widetilde{q}_t^j(k)\left(\beta\mathbb{E}\left[\frac{Y_t^j}{Y_{t+1}^j}\middle|\mathcal{S}_t,\rho_j(k)\right]\right)\\ &= \beta\mathbb{E}\left[M_{t+1}^j\frac{Y_t^j}{Y_{t+1}^j}\middle|\mathcal{S}_t\right]\end{aligned} \tag{2.59}$$

Which is what we wanted to show, since this verifies the equality (2.49) for the proposed discount factor (2.50). (Recall that multiplication by the worst-case likelihood ratio martingale M_{t+1}^j and then taking the expectation of the product is the same as computing the expectation under the twisted model probabilities.)

We obtain from (2.58) that the worst-case beliefs from the country j individual's problem are the same as the worst-case beliefs from the country j planner's problem regardless of the country j individual's starting wealth. Moreover, since in equilibrium the representative agent's total wealth for country j is $a_t^j = \frac{1}{1-\beta} Y_t^j$, substituting this expression in (2.54) shows that the levels of the planner's value function and the individual's value function coincide.

2.4 Predictable Interest Parity Violations

Now, we observe that the actual evolution of the exchange rate is governed by

$$\frac{E_t^{F,H}}{E_{t+1}^{F,H}} = \frac{S_{t+1}^F}{S_{t+1}^H} \tag{2.60}$$

while the absence of arbitrage opportunities requires that the agents inside the model believe it will evolve according to:

$$\frac{E_t^{F,H}}{E_{t+1}^{F,H}} = \frac{M_{t+1}^F S_{t+1}^F}{M_{t+1}^H S_{t+1}^H} \tag{2.61}$$

Equation (2.60) above can be written as:

$$\begin{aligned} \frac{E_t^{F,H}}{E_{t+1}^{F,H}} = \frac{\frac{Y_t^F}{Y_t^H}}{\frac{Y_{t+1}^F}{Y_{t+1}^H}} = \left(\frac{Y_t^F}{Y_{t+1}^F}\right)\left(\frac{Y_{t+1}^H}{Y_t^H}\right) \\ = \exp\left(-\overline{\kappa}_t^F - g^2(F) W_{t+1}^F\right) \exp\left(\overline{\kappa}_t^H + g^2(H) W_{t+1}^H\right). \end{aligned} \tag{2.62}$$

Computing the expectation of the logarithm of the term above under the ordinary Bayesian probabilities yields:

$$\mathbb{E}\left[\log\left(\frac{E_t^{F,H}}{E_{t+1}^{F,H}}\right)\middle|\mathcal{S}_t\right] = \overline{\kappa}_t^H - \overline{\kappa}_t^F \tag{2.63}$$

Moreover, equilibrium risk-free interest rates for claims to each type of good are given by:

$$\left(1 + i_{t+1}^H\right) = \left(\widetilde{\mathbb{E}}^H\left[\beta \frac{Y_t^H}{Y_{t+1}^H}\middle|\mathcal{S}_t\right]\right)^{-1} = \beta^{-1}\exp\left(\widetilde{\kappa}_t^H - \tfrac{1}{2}g^2(H)\right) \tag{2.64}$$

$$\left(1+i^F_{t+1}\right) = \left(\widetilde{\mathbb{E}}^F\left[\beta\frac{Y^F_t}{Y^F_{t+1}}\middle|\mathcal{S}_t\right]\right)^{-1} = \beta^{-1}\exp\left(\widetilde{\kappa}^F_t - \tfrac{1}{2}g^2(F)\right) \tag{2.65}$$

Combining these together, we have:

Proposition 1. *Define the log return for currency F relative to currency H from times t to $t+1$ as:*

$$r^{F,H}_{t+1,FX} \equiv \log\left(\frac{\left(1+i^F_{t+1}\right)E^{F,H}_t}{\left(1+i^H_{t+1}\right)E^{F,H}_{t+1}}\right) \tag{2.66}$$

Then, when equilibrium interest rates are determined under the worst-case probabilities, the expected log currency return under ordinary Bayesian beliefs is:

$$\mathbb{E}\left[r^{F,H}_{t+1,FX}\middle|\mathcal{S}_t\right] = \left(\overline{\kappa}^H_t - \widetilde{\kappa}^H_t\right) - \left(\overline{\kappa}^F_t - \widetilde{\kappa}^F_t\right) + \frac{1}{2}\left(g^2(H) - g^2(F)\right) \tag{2.67}$$

In particular, when the two countries have equal variances for the one-period shocks to each of their respective non-traded goods, the expression above reduces to

$$\mathbb{E}\left[r^{F,H}_{t+1,FX}\middle|\mathcal{S}_t\right] = \left(\overline{\kappa}^H_t - \widetilde{\kappa}^H_t\right) - \left(\overline{\kappa}^F_t - \widetilde{\kappa}^F_t\right) \tag{2.68}$$

The proof is immediate from combining (2.63) - (2.65) above. Proposition 1 says that the agents' pessimistic adjustments to their expectations of consumption growth create predictable[6], time-varying departures from uncovered interest parity. This is because the interest rate for country j is determined[7] by the distorted model-averaged expected log consumption growth term $\widetilde{\kappa}^j_t$, while the actual best guess for the evolution of the exchange rate is governed by the regular expected growth terms $\overline{\kappa}^j_t$. The premium for converting from

6. That is, predictable to the econometrician who uses Bayesian probabilities to compute the expected exchange rate movement. Agents within the model use the twisted probabilities both to price bonds and also to form their expectations regarding future exchange rate movements, and so they are indifferent toward holding either country's risk-free bond.

7. See Section A.4.

currency H to currency F, lending for one time period, and converting back in the next time period is increasing in the extent to which the distorted estimate of consumption growth for the home country falls short of the ordinary estimate of consumption growth for the home country, $\left(\overline{\kappa}_t^H - \widetilde{\kappa}_t^H\right)$, and decreasing in the size of the analogous quantity for the foreign country, $\left(\overline{\kappa}_t^F - \widetilde{\kappa}_t^F\right)$.

CHAPTER 3

EMPIRICAL RESULTS

I use series of per-capita consumption growth for the United States, Australia, Canada, Switzerland, The United Kingdom, Japan, Norway, and New Zealand, together with the spot and forward (three month) foreign exchange rates for each of the currencies of the countries just listed relative to the US Dollar. All data comes from Datastream. Daily spot and forward exchange rate data begins in Q2 1990 and runs through Q3 2016. To relate these to the consumption growth series, I use population and consumption data beginning in Q4 1989 and running through Q2 2016. If monthly population or consumption data was available, I collapsed it to quarterly data by averaging. When only yearly data was available, I interpolated between the sample periods to obtain quarterly data. I treated mid-year estimates as taking place at the end of Q2. I convert the foreign exchange data to quarterly values by taking the last observation in each quarter. Series names are provided in Table A.1 of the appendix.

3.1 Why Model Uncertainty?

To motivate the parameter uncertainty problem facing the agents within the model, I calculate maximum likelihood parameter estimates of (2.2) for each country, and then construct confidence intervals for $\rho(j)$ following the procedure of Stoffer and Wall (1991) [22]. A wide range of values for ρ fit the data for each country about equally well, but ostensibly small variations in the magnitude of ρ can lead to sizable differences in long-run consumption dynamics. Table 3.1 provides maximum likelihood point estimates of ρ together with 95% and 99% confidence intervals obtained from bootstrap simulations. All 95% confidence intervals contain the range $(.5, .88)$, with many wider still, and many 99% intervals contain practically[1] the entire range $(0, 1)$. The top half of Figure 3.1 displays counts for estimates

1. In cases where simulations produced a right endpoint for the confidence interval larger than 0.99, the value was top-coded as 0.99.

Table 3.1: Maximum likelihood estimates and confidence intervals for hidden state persistence parameter ρ. Obtained via bootstrap with $2,000$ samples.

	ρ_{MLE}	95% Confidence Interval	99% Confidence Interval
USD	0.8462	$(0.4934, 0.9251)$	$(0.2678, 0.9561)$
AUD	0.7031	$(0.2548, 0, 8833)$	$(0.1258, 0.9405)$
CAD	0.9220	$(0.1145, 0.9658)$	$(0.0310, 0.9842)$
CHF	0.7928	$(0.1734, 0.9217)$	$(0.0714, 0.9586)$
GBP	0.9075	$(0.4537, 0.9623)$	$(0.2030, 0.9822)$
JPY	0.3797	$(0.0307, 0.9408)$	$(0.0069, 0.9900)$
NOK	0.8878	$(0.0322, 0.9415)$	$(0.0089, 0.9900)$
NZD	0.7921	$(0.0529, 0.9453)$	$(0.0183, 0.9900)$

of ρ for each country using 2,000 resamples, the bottom half displays cumulative distribution functions. Figure 3.2 plots impulse response functions for a one-standard deviation shock to the first component of the hidden state vector for the US under each of the two models associated with the endpoints of the 99% confidence interval for ρ. The initial shock is larger for $\rho = 0.2678$ due to greater values of the first entry of the B matrix for smaller values of ρ, but after three quarters the effects of the shock are barely perceptible. In contrast, one sees a smaller response on impact for the $\rho = 0.9561$ model, but the effect of the shock wears out much more slowly: it takes longer than a decade for the magnitude of the shock in the high ρ model to be comparable to the magnitude of the low ρ model after only two quarters.

3.2 Model Set Construction

For each country j, I form a collection of plausible models of consumption growth by dividing the 99% confidence interval for the persistence parameter ρ into a grid of 20 evenly spaced points to generate a set[2] of coefficients $\{\rho(j,k)\}_{k=1}^{K=20}$. I then iterate over the values in $\{\rho(j,k)\}_{k=1}^{K}$, with $g(j)$ fixed at its maximum likelihood value, to estimate the remaining parameters of (2.2) -(2.3) by maximum likelihood. This produces model-specific[3] values for

2. The sets include the endpoints of the interval.

3. The empirical results to follow do not vary appreciably if instead all non-ρ coefficients are fixed at their maximum likelihood values.

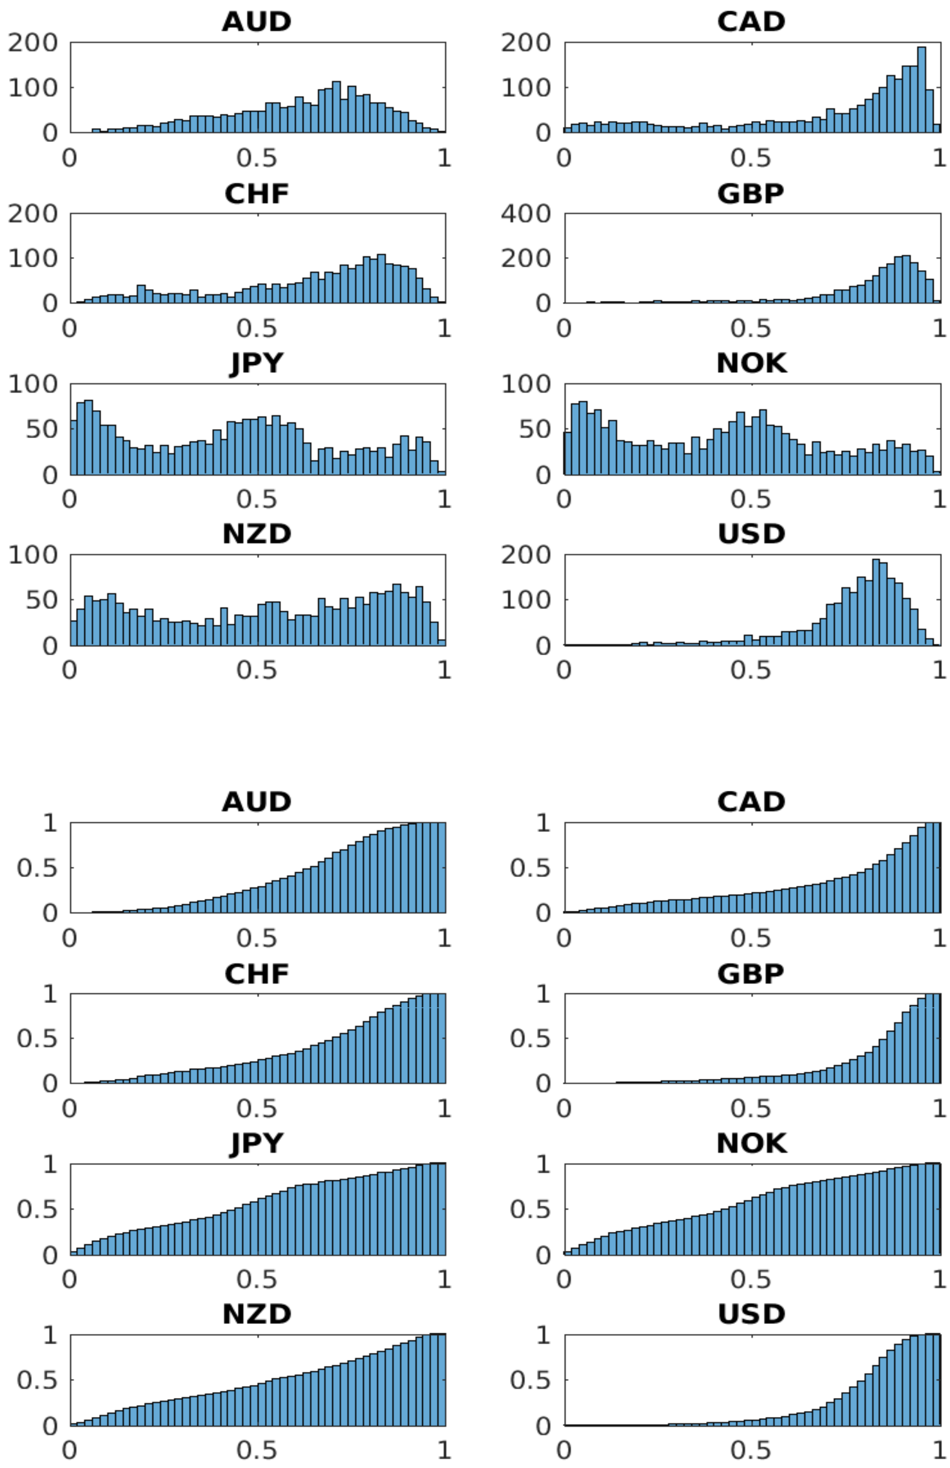

Figure 3.1: Bootstrap counts (top panel) and distributions (bottom panel) for ρ, with 2,000 samples.

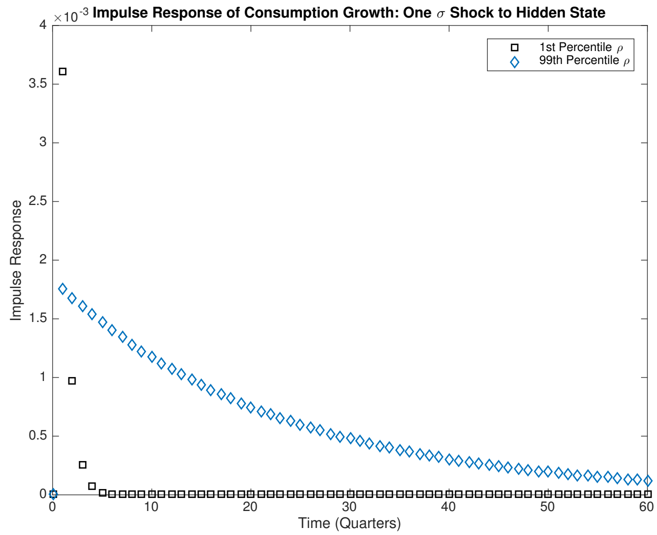

Figure 3.2: Impulse response of one standard deviation shock to hidden state vector $X_t^{US,[1]}$ for first percentile ($\rho = 0.2678$) and ninety-ninth percentile ($\rho = 0.9561$) values of persistence parameter ρ.

the standard deviation of shocks to the time-varying component of the hidden state vector $b(j,k)$, and the constant component of the hidden state vector or unconditional mean of consumption growth[4] $X_t^{j,[2]}$. After calculating the ordinary filtered state estimates $\overline{X}_t^j(k)$ and Bayesian probabilities $q_t^j(k)$, one has the series $\overline{\kappa}_t^j$ by (2.14). The adjustment (2.40) then lets us calculate[5] $\widetilde{\kappa}_t^j$ by (2.41). This produces the series of differences between ordinary and worst-case estimates of consumption growth for each country, $\overline{\kappa}_t^j - \widetilde{\kappa}_t^j$. Figure 3.3 plots Bayesian and worst-case probabilities for the various sub-models of United States consumption growth. The value of ρ increases moving from the bottom of each graph to the top.

4. I omit a model index here since it is not referenced in the remainder of the paper.

5. For tables presented in the main body of the paper, I set $\theta_2^{-1}(j) = 1$ for all countries j, where θ_2 is the penalty parameter in the minimization problems (2.17) and (2.46) as well as the worst-case probabilites (2.40). The appendix explores alternate specifications, and shows that our empirical results are not especially sensitive to moderate changes in the value of θ_2^{-1}. In particular, one obtains similar results by setting $\theta_2^{-1} = 0.5$ and $\theta_2^{-1} = 1.5$. These come very close to spanning the range of country-specific values calibrated according to the detection error probability method of [12]. The case of country-specific θ_2^{-1} values is the last sequence of tables presented in the appendix, and again the results are qualitatively similar to those given in the main paper body.

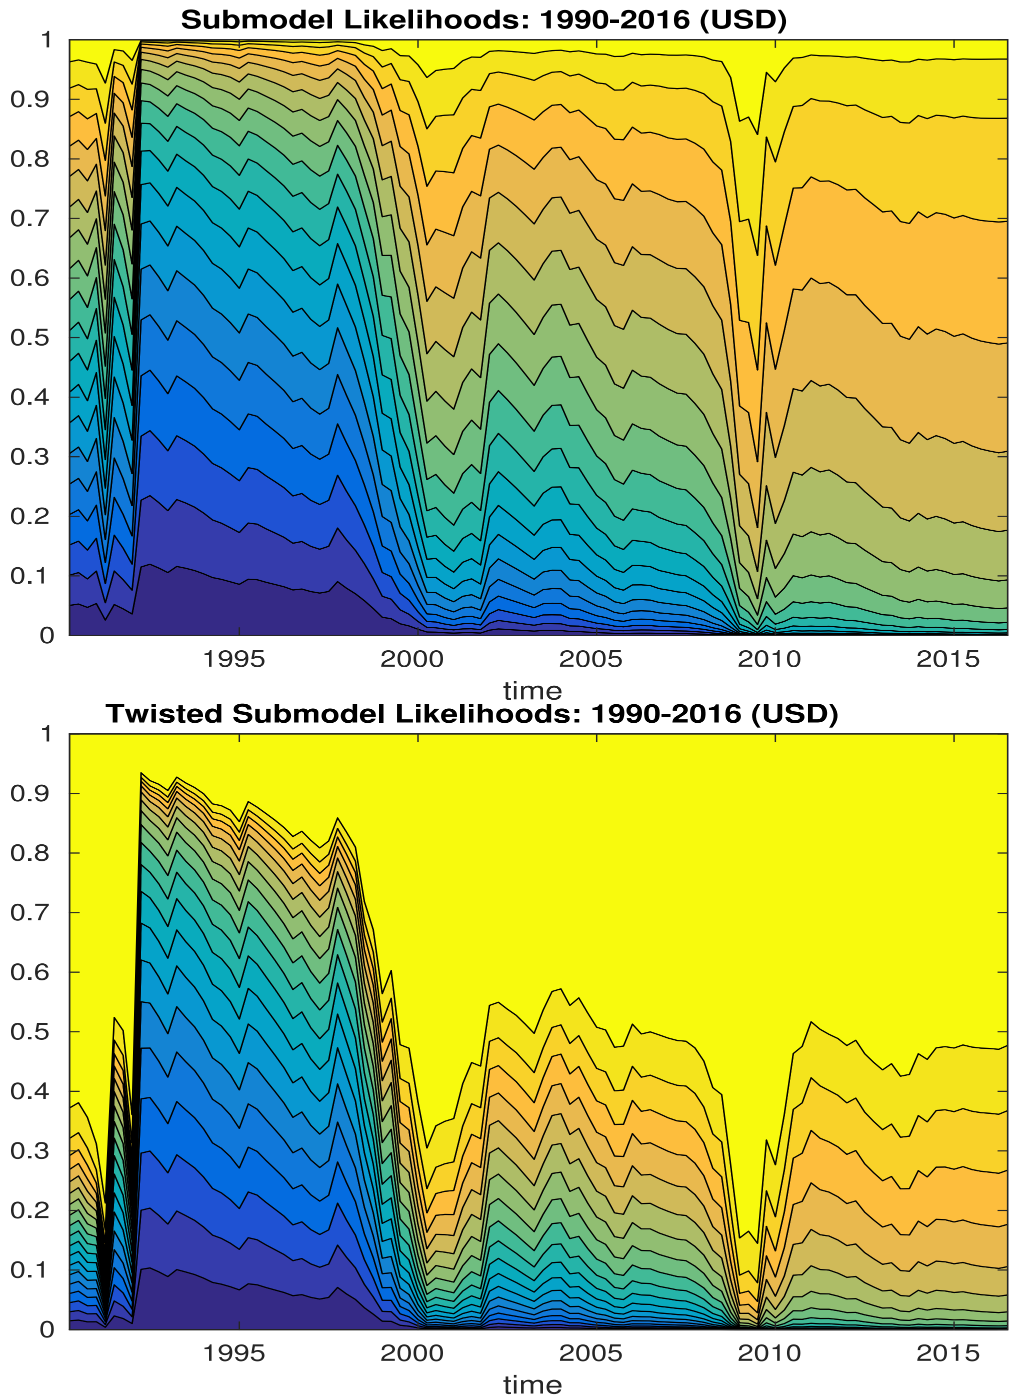

Figure 3.3: Bayesian (top panel) and worst-case (bottom panel) probabilities for sub-models of U.S. consumption growth. Color bands correspond to time-t likelihoods of the various sub-models, and are ordered monotonically so that the value of the persistence parameter ρ increases moving from the bottom to the top of each graph.

3.3 Testing the Theory

Now that we have constructed time series for Bayesian and worst-case consumption growth estimates, we are prepared to test (2.67) from Proposition 1. To do this, form the series of log real currency returns

$$r_{t+1,FX}^{j,USD} \equiv \underbrace{f_t^{j,USD} - e_{t+1}^{j,USD} - \pi_{t+1}^{USD}}_{\text{log real returns to currency } j} \tag{3.1}$$

and estimate the regression

$$r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \underbrace{\left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD}\right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j\right)\right]}_{\substack{\text{predicted return, currency } j \\ \text{(model uncertainty), (2.67)}}} + \varepsilon_{U,t+1}^j \tag{3.2}$$

for each of currencies $j \in \{\text{AUD}, \text{CAD}, \ldots, \text{NZD}\}$, where $f_t^{j,USD}$ is the log of the time t three-month forward rate between currency j and the U.S. Dollar, $e_t^{j,USD}$ is the log of the time-t spot exchange rate between currency j and the U.S. Dollar, and π_{t+1}^{USD} is the U.S. inflation rate between periods t and $t+1$. (Recall from (2.14) and (2.41) that $\overline{\kappa}_t^j$ is the expected (model-averaged) log consumption growth for country j at time $t+1$ under ordinary Bayesian model probabilities, while $\widetilde{\kappa}_t^j$ is the model-averaged expected level of log consumption growth for time $t+1$ computed using the worst-case model probabilities $\widetilde{q}_t$.) As a benchmark for comparison, we estimate a similar regression using the log currency forward discount (which approximates the interest rate difference between the two currencies[6]) $f_t^{j,USD} - e_t^{j,USD} \approx i_{t+1}^j - i_{t+1}^{USD}$:

$$r_{t+1,FX}^{j,USD} \equiv= \alpha_C^j + \beta_C^j \underbrace{\left[f_t^{j,USD} - e_t^{j,USD}\right]}_{\substack{\text{log forward discount,} \\ \text{currency } j \\ \approx i_{t+1}^j - i_{t+1}^{USD}}} + \varepsilon_{C,t+1}^j \tag{3.3}$$

6. By *covered* interest parity, we have that $f_t^{j,USD} \approx i_{t+1}^j - i_{t+1}^{USD} + e_t^{j,USD}$, so that $f_t^{j,USD} - e_t^{j,USD} \approx i_{t+1}^j - i_{t+1}^{USD}$.

Table 3.2 reports results for regressions (3.2) and (3.3). The top portion of the table lists statistics for regressions of (log) foreign exchange returns on our consumption-derived series of expected log returns from (2.67). The bottom portion of the table provides benchmark statistics from regressions of (log) foreign exchange returns on the log currency forward discount (or interest rate difference). The first column gives slope coefficient estimates. The second column provides t-statistics on the various slope coefficients. (In each case, standard errors are computed using the Hansen-Hodrick correction with 12 lags.) The third column lists the correlation coefficient between observed returns and the right-hand-side variable. Column IV reports p-values for bootstrap test of zero correlation[7] between realized returns and each of the two right-hand-side series. Column V reports regression R^2 values.

Comparing the various statistics from the two regressions for each currency, our measure of consumption growth uncertainty appears to have a stronger relationship to realized foreign exchange returns, and to explain more of the variation in returns, than the series of forward discounts (or the difference between foreign and Dollar interest rates). In particular, slope coefficient t-statistics, correlations, and R^2 values are for all currencies uniformly larger in regression specification (3.2) relative to regression specification (3.3). Similarly, p-values for permutation tests of no correlation between the left-hand-side and right-hand-side variable are uniformly lower for the growth uncertainty specification relative to the corresponding p-values relating the same currencies' returns to their forward discount.

Examining the second column across both regressions, one sees that six of the seven slope coefficients on consumption growth uncertainty are significantly different from zero against a one-sided alternative at the 1% level, and all are significantly different from zero at the 10% level. In contrast, for the regressions on log forward discounts, only one currency (JPY) has a slope coefficient significantly different from zero at the 1% level, while two others (AUD) and (CHF) are significantly different than zero at the 10% level but not at the 5% level.

7. The p-value of the test is obtained by computing the fraction of simulations for which an artificial dataset has greater correlation than the actual dataset, where each artificial dataset is formed by replacing one variable sequence with a random permutation of itself. See Davison and Hinkley [8] for further discussion.

Looking at the third column, we have that correlations between foreign exchange returns and consumption growth uncertainty are uniformly higher than the analogous correlations between currency returns and log forward discounts. The differences in magnitudes range from about .05 (AUD, CHF, JPY) to as much as about 0.15 (NOK). In the Column IV bootstrap tests of zero correlation between currency returns and our measure of consumption growth uncertainty, we reject the null hypothesis of zero correlation at the 10% significance level for all currencies, with three of these also rejected at the 5% significance level. In contrast, for the bootstrap tests of zero correlation between currency returns and currency forward discounts, we reject the null of zero correlation at the 10% significance level for only three of the seven currencies: AUD, CHF, and JPY, with the null also rejected at the 1% significance level for JPY.

In the final column, one sees regression R^2 values that are uniformly between one and two percentage points higher for the first regression specification (returns on consumption growth uncertainty) relative to the second (returns on the forward discount). In the cases of CAD, GBP, and NOK, we see an economically significant increase in the R^2 value from about one-half of one percent or less to around 2% or more when moving from regressions on the forward discount to regressions on the model's consumption-based measure of predicted returns.

In light of this evidence, one can conclude that over the time range studied, the measure of expected uncovered interest parity violation from (2.67) based on model uncertainty explains more of the variation in currency returns in the following time period, and does so in a more convincing fashion, than the ordinary currency forward discount.

3.4 Which Model Misspecification Fears Are Essential?

In the previous sections, concerns for robustness against model uncertainty have manifested themselves entirely in the form of agents distrusting their probability distribution over the set of competing models. Implicit within this choice is the decision to ignore other varieties of

Table 3.2: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth (top section) and currency forward discounts (bottom section). Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD}\right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j\right)\right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	33.8614	2.5355	0.1915	0.0248	0.0367
CAD	9.3980	2.5779	0.1583	0.0556	0.0251
CHF	29.2950	5.5940	0.1951	0.0225	0.0381
GBP	6.3112	2.6758	0.1293	0.0966	0.0167
JPY	31.2662	4.8637	0.2046	0.0191	0.0419
NOK	22.2097	2.4815	0.1454	0.0699	0.0211
NZD	22.9692	1.6523	0.1560	0.0579	0.0243
Regress currency return on (log) forward discount: $r_{t+1,FX}^{j,USD} = \alpha_C^j + \beta_C^j \left[f_t^{j,USD} - e_t^{j,USD}\right] + \varepsilon_{C,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	2.3562	1.3899	0.1591	0.0549	0.0253
CAD	0.8095	0.9370	0.0725	0.2312	0.0053
CHF	1.6580	1.4446	0.1480	0.0694	0.0219
GBP	0.3535	0.4186	0.0314	0.3779	0.0010
JPY	1.7730	2.0181	0.1680	0.0439	0.0282
NOK	0.0862	0.0636	0.0086	0.4671	0.0001
NZD	0.8241	0.2625	0.0508	0.3057	0.0026

model uncertainty. Two other varieties of uncertainty that have received considerable attention in the literature are the possibility that, conditional on a particular model being the true data generating process: i) the dynamics of the shock process are misspecified, and ii) the true value of the hidden state vector differs from the estimated value. Both of these varieties are explored in [12]. This section of the paper explores whether these alternate specification doubts can function as a complement or substitute for the distorted probability distribution over models in predicting uncovered interest parity violations. Specifically, how do these other varieties of uncertainty affect the empirical results in Table 3.2? Table 3.3 provides the same output as Table 3.2 for modified regressions of currency returns on modifications of the $\left(\overline{\kappa}_t^H - \widetilde{\kappa}_t^H\right) - \left(\overline{\kappa}_t^F - \widetilde{\kappa}_t^F\right)$ series that omit or include various elements of agent's pessimistic belief adjustments. The first subtable provides results for all three adjustments. The second subtable provides regression results when agents no longer pessimistically adjust their probability distribution over models, but do distrust their specification of model dynamics and their estimation of the hidden state. Adjustments to the value function and model-averaged growth estimates follow the procedures detailed in [12] and are reproduced in Section A.2 of the appendix. In the first case, results are essentially unchanged from those in Table 3.2. In the second case, we see that the model offers a dramatically poorer fit[8] to the data. Slope coefficients change signs for most countries, R^2 values fall from 2-4% to often less than a tenth of a percentage point, and correlations drop by as much as 20%, often changing signs. One now fails to reject the null hypotheses of zero slope coefficients and zero correlation of predicted returns with actual returns at any common significance level for all currencies. In particular, a random permutation of the predicted return series now has greater correlation with observed returns than the original series of predicted returns between one-half and two-thirds of the time, depending on the currency. We conclude that the distortions to the

8. This is not surprising, since here the $\widetilde{\kappa}$ series differ from the $\overline{\kappa}$ series only by the addition of μ (constant model-specific distortions to the mean of the shock) and $\Sigma_t \lambda$ (where Σ_t is the variance-covariance matrix for the estimate of the hidden state) terms, and Σ_t, the non-constant portion of this dies, down monotonically over time as the precision of the hidden state estimate improves See Section A.2 for more details.

probability distributions of the ρ values are both necessary and sufficient to generate our empirical result.

3.5 Comparing Returns and Model-Implied Interest Rates

One might ask if the empirical result is simply a consequence of one or the other κ_t^j series serving as a proxy for observed interest rates. To show that this is not the case, I demonstrate that the series of predicted currency returns from (2.67) has a stronger relationship with observed returns than with the series of interest rate[9] differences calculated under either set of beliefs. I measure this by comparing the results of bootstrap permutation tests[10] for the null hypotheses of zero correlation between the respective series and actual returns.

The null hypotheses for the tests are:

Zero Correlation of Returns with Predicted Returns:

First Column, Table 3.4:

$$\begin{aligned} H_0 : Corr\left(r_{t+1,FX}^{j,USD}, i_{t+1}^{j} - i_{t+1}^{USD} - \mathbb{E}\left[\Delta e_{t+1}^{j,USD}\middle|\mathcal{S}_t\right]\right) \\ = Corr\left(r_{t+1,FX}^{j,USD}, \left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD}\right) - \left(\overline{\kappa}_t^{j} - \widetilde{\kappa}_t^{j}\right)\right) = 0 \end{aligned} \tag{3.4}$$

Zero Correlation of Returns with (Bayesian) Interest Rate Difference:

Second Column, Table 3.4

$$H_0 : Corr\left(r_{t+1,FX}^{j,USD}, i_{t+1}^{j} - i_{t+1}^{USD}\right) = Corr\left(r_{t+1,FX}^{j,USD}, \overline{\kappa}_t^{j} - \overline{\kappa}_t^{USD}\right) = 0 \tag{3.5}$$

---

9. Recall that by taking the logarithm of (2.64) or (2.65), one obtains that the time-varying component of the log interest rate for country j, i_{t+1}^j, is $\widetilde{\kappa}_t^j$, or:

$$\log\left(1 + i_{t+1}^j\right) = \widetilde{\kappa}_t^j - \tfrac{1}{2}g^2\left(j\right) - \log\beta$$

From which we have the approximation:

$$i_{t+1}^j \approx \widetilde{\kappa}_t^j - \tfrac{1}{2}g^2\left(j\right) - \log\beta$$

If the agents were to compute model probabilities using ordinary Bayesian updating, the $\widetilde{\kappa}_t$ term would instead be $\overline{\kappa}_t^j$ from (2.14).

10. See explanation in Footnote 7.

Table 3.3: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth under alternate misspecification fears. Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. The first subtable includes distortions to dynamics, state estimation, and sub-model probabilities. The second subtable includes distortions to dynamics and estimation of the hidden state, but not sub-model probabilities.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
Distortions to dynamics, hidden state, and model probabilities					
AUD	33.5310	2.5604	0.1917	0.0244	0.0367
CAD	9.4258	2.6203	0.1602	0.0520	0.0257
CHF	29.1710	5.5100	0.1949	0.0245	0.0380
GBP	6.3122	2.6748	0.1294	0.0928	0.0167
JPY	31.1969	4.9644	0.2049	0.0202	0.0420
NOK	22.1966	2.4888	0.1458	0.0693	0.0213
NZD	22.7992	1.6486	0.1553	0.0574	0.0241
Distortions to dynamics and hidden state, but *not* model probabilities:					
AUD	-207.4552	-1.0412	-0.0261	0.6046	0.0007
CAD	-136.6356	-0.4566	-0.0312	0.6251	0.0010
CHF	-268.0356	-1.5744	-0.0470	0.6818	0.0022
GBP	19.4683	0.1087	0.0088	0.4598	0.0001
JPY	101.2136	0.1782	0.0190	0.4262	0.0004
NOK	-26.7035	-0.0988	-0.0049	0.5243	0.0000
NZD	-170.1591	-0.4575	-0.0337	0.6310	0.0011

Zero Correlation of Returns with (Robust Beliefs) Interest Rate Difference

Third Column, Table 3.4

$$H_0 : Corr\left(r^{j,USD}_{t+1,FX}, i^j_{t+1} - i^{USD}_{t+1}\right) = Corr\left(r^{j,USD}_{t+1,FX}, \widetilde{\kappa}^j_t - \widetilde{\kappa}^{USD}_t\right) = 0 \tag{3.6}$$

Table 3.4 reports the p-values for the tests in (3.4)-(3.6). Column I is repeated from Table 3.2 and reports p-values for tests of zero correlation between actual returns and the series of predicted returns from (2.67). Column II provides p-values for tests of zero correlation between actual returns and model-implied differences in interest rates under ordinary Bayesian model averaging, $\overline{\kappa}^j_t - \overline{\kappa}^{USD}_t$. Column III lists p-values for tests of zero correlation between returns and the model-implied difference in interest rates under the worst case probabilities, $\widetilde{\kappa}^j_t - \widetilde{\kappa}^{USD}_t$. For all currencies in our sample, we reject the null hypothesis of no correlation between actual and predicted currency returns at the 10% significance level or less, while for all but two of the currencies (AUD and NZD), we fail to reject the null at the 10% significance level in either of the other two series. For the five currencies for which we reject the null of zero correlation at the 10% significance level in the first test but not the second or third, we see differences in p-values ranging from about 10% to as much as 35%. In this sense, the model of uncovered interest parity violations being driven by the *differences* in the $\overline{\kappa} - \widetilde{\kappa}$ series presents a more convincing explanation of the data than one could obtain via either the model-implied Bayesian or worst-case interest rate difference series alone. This highlights that the measure of model uncertainty in Proposition 1 provides useful information beyond simply acting as a proxy for interest rates.

3.6 Comparison of Trading Strategy Returns

How does the performance of a currency trading strategy based on using (2.67) to predict returns compare to the performance of a naive carry trade strategy? To explore this, I examine the out-of-sample cumulative returns of each using the following method, starting

Table 3.4: Bootstrap permutation test p-values of zero correlation between observed foreign exchange returns and various series derived from consumption data listed in (3.4). The left column lists results for predicted currency returns given by (2.67). The middle column provides the analogous values for the series of interest rate differences implied by the model under ordinary Bayesian beliefs. The final column lists p-values for the null hypothesis of zero correlation between actual currency returns and the model-implied difference in interest rates under worst-case beliefs.

p-values for zero correlation of actual returns with:			
	Predicted	Consumption-Implied Interest Rate Difference	
	Returns	Bayesian Beliefs	Worst-Case Beliefs
AUD	0.0244	0.0434	0.0286
CAD	0.0504	0.2316	0.1675
CHF	0.0225	0.1598	0.1256
GBP	0.0965	0.2489	0.1424
JPY	0.0171	0.3700	0.2809
NOK	0.0686	0.1627	0.1490
NZD	0.0599	0.0515	0.0423

from common initial learning periods of varying length:

1. Estimate regression coefficients in (3.2) and (3.3) using the first $\overline{T}$ observations in the data set, $\left\{r^{j,USD}_{s+1,FX}\right\}_{s=0}^{\overline{T}-1}$ and either $\left\{\left(\overline{\kappa}^{USD}_s - \widetilde{\kappa}^{USD}_s\right) - \left(\overline{\kappa}^j_s - \widetilde{\kappa}^j_s\right)\right\}_{s=0}^{\overline{T}-1}$ or $\left\{f^{j,USD}_s - e^{j,USD}_s\right\}_{s=0}^{\overline{T}-1}$.

2. For both right hand side variables, form predicted returns at time $T+1$ using coefficients estimated above and observed data at time T.

3. Cumulative return rises (falls) by the size of the actual return when the signs of the actual and predicted returns at time $T+1$ agree (disagree).

4. Add the next pair of observations to the training set in the first point above, and repeat the process.

Figure 3.4 shows cumulative returns to each strategy starting from an initial \$1 exposure, for starting training set sizes of 10, 20, and 40 quarters. (Training set size increases as one moves down the page.) Returns to the trading strategy based on consumption growth uncertainty (from Equation (2.67)) are drawn in a solid line with a black ciricle, $\circ$, where returns to the carry trade strategy are plotted with a red x.

For several of the currencies (CHF, GBP, and JPY) the uncertainty-based strategy has better performance (here, higher cumulative returns) than the carry trade strategy at virtually all time periods where the payoffs from following the two strategies diverge, and this result is not sensitive to the size of the training dataset. For CAD and NZD, the uncertainty-based strategy typically has similar to better performance by the end of the data series, although the carry trade strategy sometimes has higher returns in the beginning or middle of the sample, depending on the size of the initial learning period. AUD appears to be the only currency for which the baseline carry trade strategy consistently provides more accurate predictions regardless of training set size. In general, though, when the returns to the two strategies show a noticeable divergence from one another, the regression (3.2) based on our measure of consumption growth uncertainty typically predicts the sign of currency returns with greater accuracy than the carry trade regression (3.3).

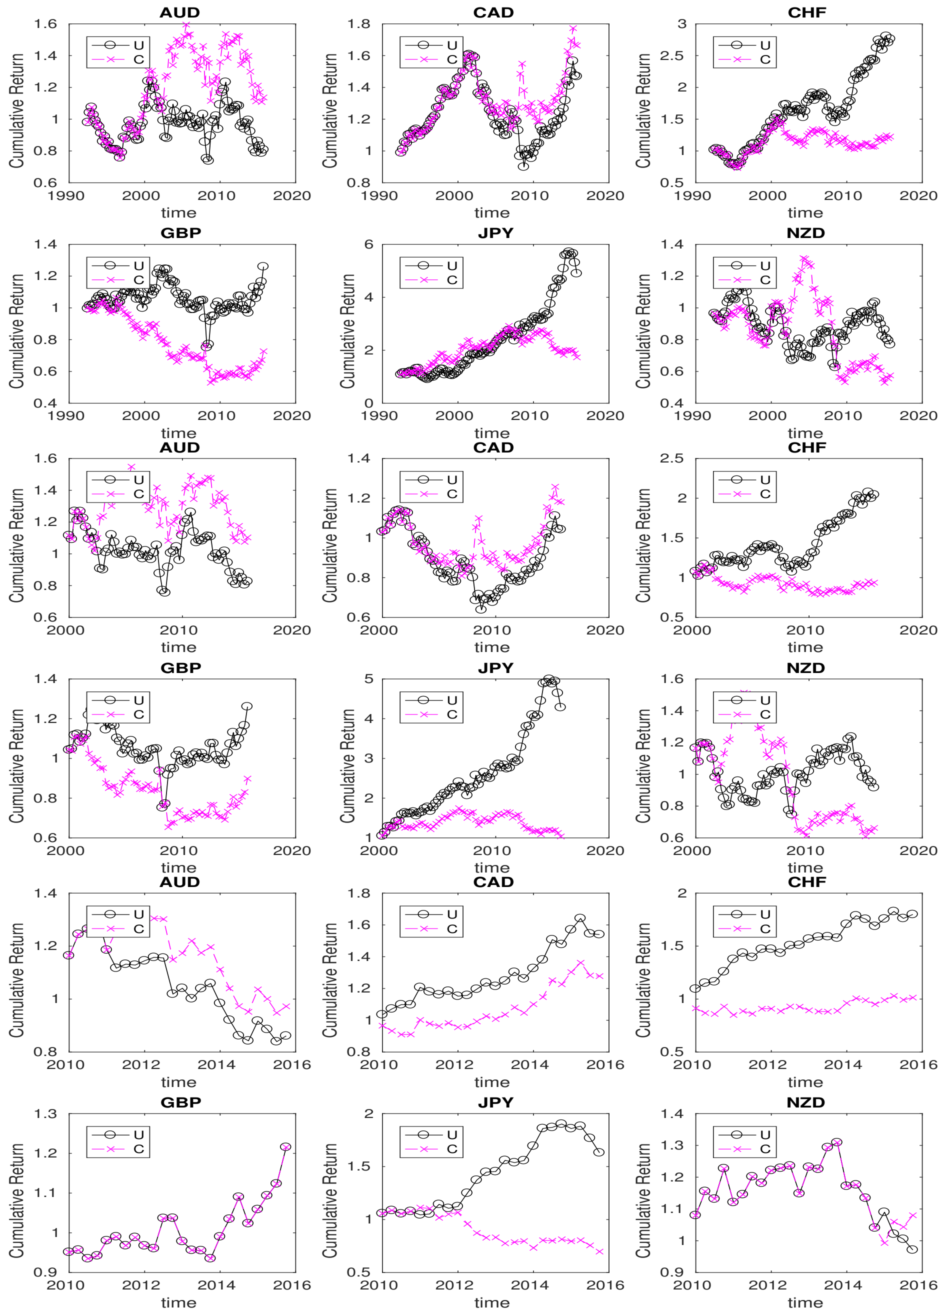

Figure 3.4: Comparison of cumulative returns to trading strategies based on model uncertainty (labelled with letter U, black circles $\circ$), from equation (2.67) and ordinary carry trade (letter C, red x) for initial learning periods of 10, 40, and 20 quarters. (Learning period increases moving from the top of the page to the bottom of the page.)

CHAPTER 4

CONCLUSION

I have illustrated a mechanism where robustness against model uncertainty creates predictable uncovered interest parity violations (or returns to foreign currency speculation) that evolve over time as a simple linear function of the differences in ordinary and worst-case estimates of consumption growth for the relevant countries. Although these deviations from uncovered interest parity are apparent to an outside observer that uses ordinary Bayesian model averaging to form their expectations about exchange rate movements, agents within the model are indifferent between holding risk-free bonds denominated in either currency. I test the model and find that its consumption-based time series of expected currency returns explains the variation in actual currency returns as well or better than the benchmark time series of currency forward premiums. I motivate the agents' misspecification fears by examining a collection of models where an unobservable, time-varying component to expected consumption growth follows an autoregressive process of order one with unknown persistence. Competing models (indexed by persistence coefficient values) are observationally similar but have very different implications for long-run dynamics. Agents inside the model cope with this uncertainty by dynamically applying a pessimistic distortion to their ordinary Bayesian probability distribution over the various sub-models that are candidates for the true data-generating processes. I go on to explore the ability of other types of uncertainty to generate the result, and also to show that the paper's measure of expected currency returns offers a more convincing explanation of the foreign exchange data than either consumption-implied interest rate difference series[1] is able to provide on its own. Finally, I document that the out-of-sample predictive performance of a trading strategy based on the expected currency return series generated by the model typically exceeds that of one based on the ordinary currency forward discount or carry trade.

1. Calculated under Bayesian or worst-case probabilities.

In addition to providing a new explanation of predictable currency returns or the forward premium puzzle, this paper is the first (to my knowledge) to empirically test and find support for the effects of robust model probability adjustments in the context of unknown consumption growth persistence as in [12]. Further research in this area might examine the extent to which a similar mechanism can explain other puzzles in macroeconomics and finance.

APPENDIX A

APPENDICES

A.1 More on the Individual's Value Function

In this section, I will guess and verify[1] the policy and value functions for the individual presented in (2.52) and (2.54). In Section 2.3, I claimed that these were:

$$C_t^j = (1-\beta)\, a_t^j \tag{A.1}$$

$$\pi_{t+1}^j = \frac{\beta}{1-\beta} Y_{t+1}^j \tag{A.2}$$

With the form of the value function being:

$$V^j\left(a_t^j, \zeta_t\right) = \log\left(a_t^j\right) + \lambda'(j)\, X_t^j + \log(1-\beta) \tag{A.3}$$

For total wealth

$$a_t^j = \frac{Y_t^j}{1-\beta} \tag{A.4}$$

To see that the form of the value function stated above satisfies the update equation (2.48), compute:

$$\begin{aligned} V\left(a_t^j, \zeta_t\right) &= \log\left(a_t^j\right) + \Psi^j(\zeta_t) + \log(1-\beta) \\ &= (1-\beta)\log C_t^j + \beta\mathbb{E}\left[V\left(a_{t+1}, \zeta_{t+1}\right)|\mathcal{X}_t\right] \\ &= (1-\beta)\log\left((1-\beta)\, a_t^j\right) + \beta\mathbb{E}\left[\log\left(Y_{t+1}^j + \frac{\beta}{1-\beta}Y_{t+1}^j\right)\middle|\mathcal{X}_t\right] \\ &\quad + \beta\mathbb{E}\left[\Psi^j(\zeta_{t+1})\middle|\mathcal{X}_t\right] + \beta\log(1-\beta) \end{aligned} \tag{A.5}$$

For some function of state variables $\Psi^j(\zeta_t)$ to be determined.

1. See [13] for a similar derivation in continuous time.

Notice that the second term in the final expression can be written as

$$\beta\mathbb{E}\left[\log\left(\frac{Y_{t+1}^{j}}{1-\beta}\right)\middle|\mathcal{X}_t\right]=\beta\mathbb{E}\left[\log\left(\frac{Y_{t+1}^{j}}{Y_t^{j}}a_t^{j}\right)\middle|\mathcal{X}_t\right]=\beta\log a_t^{j}+\beta DX_t^{j} \tag{A.6}$$

Combining (A.6) above with (A.5), we have:

$$\begin{aligned}V\left(a_t^{j},\zeta_t\right)&=\log\left(a_t^{j}\right)+\Psi^{j}\left(\zeta_t\right)+\log\left(1-\beta\right)\\&=\left(1-\beta\right)\log\left(\left(1-\beta\right)a_t^{j}\right)+\beta\log a_t^{j}+\beta DX_t^{j}\\&\quad+\beta\mathbb{E}\left[\Psi^{j}\left(\zeta_{t+1}\right)\middle|\mathcal{X}_t\right]+\beta\log\left(1-\beta\right)\end{aligned} \tag{A.7}$$

Equating coefficients in (A.7) has the $\log\left(a_t^{j}\right)$ and $\log\left(1-\beta\right)$ terms vanish, which leaves the following equality to be solved:

$$\Psi^{j}\left(\zeta_t\right)=\beta DX_t^{j}+\beta\mathbb{E}\left[\Psi^{j}\left(\zeta_{t+1}\right)\middle|\mathcal{X}_t\right] \tag{A.8}$$

Now we recognize that the function Ψ^{j} that solves (A.8) is given by the same function of state variables from the planner's value function, namely

$$\Psi^{j}\left(\zeta_t\right)=\lambda'(j)X_t^{j} \tag{A.9}$$

with the vector $\lambda\left(j\right)$ defined as in (2.36):

$$\lambda'(j)=\begin{bmatrix}\frac{\beta}{1-\beta\rho(j)} & \frac{\beta}{1-\beta}\end{bmatrix} \tag{A.10}$$

To see that the choices of C_t^{j} and π_{t+1}^{j} in (2.52) and (2.53) are optimal when agents have worst-case beliefs about the future represented by multiplication with the likelihood ratio martingale

$$M_{t+1}^{j} \equiv \exp\left(g_2^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\overline{W}_{t+1}^j - \tfrac{1}{2}\left[g_2^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right) \tag{A.11}$$

and prices are given by the stochastic discount factor from the planner's problem, write the agent's maximization problem as

$$\max_{C_t^j, \pi_{t+1}^j} (1-\beta)\log C_t^j + \beta\mathbb{E}\left[M_{t+1}^j V\left(Y_{t+1}^j + \pi_{t+1}^j, \zeta_{t+1}\right)\middle| \mathcal{S}_t\right] \tag{A.12}$$

subject to the budget constraint

$$C_t + \beta\mathbb{E}\left[M_{t+1}^j \frac{Y_t^j}{Y_{t+1}^j} a_{t+1}^j \middle| \mathcal{S}_t\right] = Y_t^j + \pi_t^j \equiv a_t^j \tag{A.13}$$

where next period's asset holdings are

$$a_{t+1}^j = \pi_{t+1}^j + Y_{t+1}^j \tag{A.14}$$

The first-order condition for π_{t+1}^j is

$$\frac{1-\beta}{C_t^j}\beta\mathbb{E}\left[M_{t+1}^j \frac{Y_t^j}{Y_{t+1}^j}\middle| \mathcal{S}_t\right] = \beta\mathbb{E}\left[M_{t+1}^j \frac{\partial V\left(a_{t+1}^j, \zeta_{t+1}\right)}{\partial a_{t+1}}\middle| \mathcal{S}_t\right] \tag{A.15}$$

The envelope condition is:

$$\frac{\partial V\left(a_t^j, \zeta_t\right)}{\partial a_t^j} = \frac{1-\beta}{C_{t+1}^j} \tag{A.16}$$

Combining (A.15) with (A.16) and rearranging we gives the Euler equation for Arrow securities π_{t+1}^j:

$$\mathbb{E}\left[M_{t+1}^j \frac{Y_t^j}{Y_{t+1}^j}\middle| \mathcal{S}_t\right] = \mathbb{E}\left[M_{t+1}^j \frac{C_t^j}{C_{t+1}^j}\middle| \mathcal{S}_t\right] \tag{A.17}$$

We know this will be satisfied because agents consume a constant fraction of their total wealth each period (owing to their unitary intertemporal elasticity of substitution) and wealth is proportional to the total endowment, so we have the equality (A.17).

A.2 More Types of Uncertainty, Distorted Value Function

This section provides a short review of the other types of model uncertainty used in constructing Table 3.3. For a more detailed derivation and discussion, see [12]. One can consider a broader set of specification doubts, specifically that agents may distrust their understanding of model dynamics and/or their filtered estimate of the hidden state vector associated with a particular model. As before, agents form their beliefs by imagining that they are playing a game against a malevolent Nature that has twisted some aspect(s) of their approximating model. Each of these distortions is the solution to a problem that minimizes the sum of agents' lifetime expected utility (under the worst-case distortion) together with a penalty term times the relative entropy of the distorted distribution relative to the original. For an example of this, consider a change to the model structure where instead $W_{t+1}^{j} \sim \mathcal{N}\left(\mu, g^2(j)\right)$, with $\mu \neq \vec{0}$. Then the value function for the representative agent of country j becomes[2]

$$\log Y_t^j + \lambda'(j) X_t^j + \frac{\beta}{1-\beta}\left[G(j) + \lambda'(j)B(j)\right]\mu \tag{A.18}$$

In this case, the relative entropy term is $\mu'\mu$. For penalty term[3] $\frac{\beta}{1-\beta}\frac{\theta_1(j)}{2}$, the solution is

$$\mu(j) = -\frac{1}{\theta_1(j)}\left[G'(j) + B'(j)\,\lambda(j)\right] \tag{A.19}$$

So that the value function with worst-case dynamics becomes

$$\log Y_t^j + \lambda'(j) X_t^j - \left(\frac{1}{\theta_1(j)}\right)\left(\frac{\beta}{1-\beta}\right)\left|G(j) + \lambda'(j)B(j)\right|^2 \tag{A.20}$$

2. Derivation provided later in this section.

3. Here, we follow the notation of [12] and write the formulas to allow for variation in relative entropy penalty parameters across both countries and types of distortion.

A similar problem finds the worst-case estimate of the hidden state vector. For penalty parameter $\theta_2(j)$, the robust estimate at time t conditional on model k is

$$\widetilde{X}_t^j(k) = \overline{X}_t^j(k) - \left(\frac{1}{\theta_2(j)}\right) \Sigma_t(j,k)\lambda\,(j,k) \tag{A.21}$$

One can then calculate the expected level of the value function conditional on model (j,k) and signal history $\mathcal{S}_t$ under fears of both misspecified dynamics and incorrect estimation of the hidden state by replacing $X_t^j(k)$ with $\widetilde{X}_t^j(k)$ in (A.20). Write this as $\widetilde{V}^j\,(\zeta_t, k)$:

$$\begin{aligned} \widetilde{V}_t^j(k) &= \log Y_t^j + \lambda'(j,k)\overline{X}_t^j\,(k) \\ &- \left(\frac{1}{\theta_1(j)}\right)\left(\frac{\beta}{1-\beta}\right)|G(j) + \lambda'(j,k)B(j,k)|^2 - \left(\frac{1}{\theta_2(j)}\right)\lambda'\,(j,k)\,\Sigma_t(j,k)\lambda\,(j,k) \end{aligned} \tag{A.22}$$

The worst-case model probabilities found under the distortions above are obtained by replacing $\widehat{V}^j\,(\zeta_t, k)$ with $\widetilde{V}^j\,(\zeta_t, k)$ in (2.40). In this case (2.67) holds with different values of the $\widetilde{\kappa}$ terms. One might suspect that since the mean distortions are constant, and distortions to the hidden state die out monotonically over time, that the misspecification fears above are not crucial for our result. Table 3.3 shows that this is in fact the case.

To verify the results above, for a modification to the mean of the shock vector so that $\mathbb{E}\left[W_{t+1}^j \middle| \mathcal{X}_t\right] = \mu$, guess a form for the value function for agents in country j:

$$V^j\,(\zeta_t) = \log Y_t^j + \widetilde{\lambda}'(j)X_t^j + \psi'\mu \tag{A.23}$$

where $\widetilde{\lambda}$ and ψ are vectors to be determined. The value function must satisfy

$$\log Y_t^j + \widetilde{\lambda}'(j)X_t^j + \psi'\mu = (1-\beta)\left(\log Y_t^j\right) + \beta\mathbb{E}\left[V^j\,(\zeta_{t+1})\middle|\mathcal{X}_t\right] \tag{A.24}$$

but by assumption

$$\beta\mathbb{E}\left[V^{j}\left(\zeta_{t+1}\right)\middle|\mathcal{X}_{t}\right]=\beta\left(\log Y_{t}^{j}+DX_{t}^{j}+G(j)\mu+\widetilde{\lambda}'(j)A(j)X_{t}^{j}+\widetilde{\lambda}'(j)B(j)\mu+\psi'\mu\right) \tag{A.25}$$

Equating coefficients on X_t^j and ψ gives the system of equations

$$\begin{aligned}\widetilde{\lambda}'(j)&=\beta D+\beta\widetilde{\lambda}'(j)A(j)\\ \psi'\mu&=\beta G(j)\mu+\beta\widetilde{\lambda}'(j)B(j)\mu+\beta\psi'\mu\end{aligned} \tag{A.26}$$

The first row of (A.26) is the same as (2.36) , so the vector $\widetilde{\lambda}$ is unchanged from the no-distortion case. From the second row, we obtain

$$\psi'(j)=\left(\frac{\beta}{1-\beta}\right)\left[G(j)+\lambda'(j)B(j)\right] \tag{A.27}$$

A.3 Likelihood Ratio Martingale: Combined Worst-Case Beliefs

As in (2.41), write the model-averaged worst-case mean for numeraire endowment growth as

$$\widetilde{\kappa}_{t}^{j}\equiv\sum_{k=1}^{n}\widetilde{q}_{t}^{j}(k)\left(D\widetilde{X}_{t}^{j}(k)+G(j)\mu(j,k)\right) \tag{A.28}$$

Where the expression above allows for worst-case distortions to each of model probabilities, model dynamics, and state estimation. From (2.14) the ordinary Bayesian mean is

$$\overline{\kappa}_{t}^{j}\equiv D\sum_{k=1}^{n}q_{t}^{j}(k)\overline{X}_{t}^{j}(k) \tag{A.29}$$

Decompose the signal observed for country j at time $t+1$ into

$$y_{t+1}^{j}\equiv\Delta\log Y_{N,t+1}^{j}=\overline{\kappa}_{t}^{j}+g(j)\overline{W}_{t+1}^{j} \tag{A.30}$$

Then, since $G(j)$ does not vary across models, the likelihood ratio for the worst-case consumption growth process relative to the regular filtered estimate of the consumption growth process reduces to

$$M_{t+1}^{j} \equiv \frac{\exp\left[-\frac{1}{2}g^{-2}(j)\left(y_{t+1}^{j}-\widetilde{\kappa}_{t}^{j}\right)^{2}\right]}{\exp\left[-\frac{1}{2}g^{-2}(j)\left(y_{t+1}^{j}-\overline{\kappa}_{t}^{j}\right)^{2}\right]} = \frac{\exp\left[-\frac{1}{2}g^{-2}(j)\left(g(j)\overline{W}_{t+1}^{j}+\left(\overline{\kappa}_{t}^{j}-\widetilde{\kappa}_{t}^{j}\right)\right)^{2}\right]}{\exp\left[-\frac{1}{2}g^{-2}(j)\left(g(j)\overline{W}_{t+1}^{j}\right)^{2}\right]}$$

$$= \frac{\exp\left[-\frac{1}{2}g^{-2}(j)\left(g^{2}(j)\left(\overline{W}_{t+1}^{j}\right)^{2}+2g(j)\overline{W}_{t+1}^{j}\left(\overline{\kappa}_{t}^{j}-\widetilde{\kappa}_{t}^{j}\right)+\left(\overline{\kappa}_{t}^{j}-\widetilde{\kappa}_{t}^{j}\right)^{2}\right)\right]}{\exp\left[-\frac{1}{2}g^{-2}(j)\left(g(j)\overline{W}_{t+1}^{j}\right)^{2}\right]}$$

$$= \frac{\exp\left[-\frac{1}{2}\left(\overline{W}_{t+1}^{j}\right)^{2}-\frac{1}{2}g^{-2}(j)2g(j)\overline{W}_{t+1}^{j}\left(\overline{\kappa}_{t}^{j}-\widetilde{\kappa}_{t}^{j}\right)-\frac{1}{2}g^{-2}(j)\left(\overline{\kappa}_{t}^{j}-\widetilde{\kappa}_{t}^{j}\right)^{2}\right]}{\exp\left[-\frac{1}{2}g^{-2}(j)g^{2}(j)\left(\overline{W}_{t+1}^{j}\right)^{2}\right]}$$

$$= \exp\left[g^{-1}(j)\left(\widetilde{\kappa}_{t}^{j}-\overline{\kappa}_{t}^{j}\right)\overline{W}_{t+1}^{j}-\frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_{t}^{j}-\overline{\kappa}_{t}^{j}\right)\right]^{2}\right] \tag{A.31}$$

which gives (2.43).

A.4 Determining Interest Rates

The ordinary stochastic discount factor for agents of country j that prices time $t+1$ payouts at time t is

$$S_{t+1}^{j} = \beta\frac{C_{t}^{j}}{C_{t+1}^{j}} \tag{A.32}$$

but the pessimism of the representative agents means that assets are priced so that returns R_{t+1} obey

$$1 = \mathbb{E}\left[M_{t+1}^{j}S_{t+1}^{j}R_{t+1}\middle|\mathcal{S}_{t}\right] \tag{A.33}$$

Now, using (A.31), we can price bonds for each country. We have that

$$1 = \mathbb{E}\left[M_{t+1}^{j}S_{t+1}^{j}\left(1+i_{t+1}^{j}\right)\middle|\mathcal{S}_{t}\right] \rightarrow \left(1+i_{t+1}^{j}\right) = \left(\mathbb{E}\left[M_{t+1}^{j}S_{t+1}^{j}\middle|\mathcal{S}_{t}\right]\right)^{-1} \tag{A.34}$$

Where we can evaluate (A.34) as

$$\begin{aligned}
&\mathbb{E}\left[M_{t+1}^{j}S_{t+1}^{j}\middle|\mathcal{S}_t\right] \\
&= \mathbb{E}\left[\exp\left\{g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\overline{W}_{t+1}^j - \frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right\}\beta\frac{Y_t^j}{Y_{t+1}^j}\middle|\mathcal{S}_t\right]
\end{aligned} \tag{A.35}$$

but from (2.2), we have that

$$\frac{Y_t^j}{Y_{t+1}^j} = \exp\left(-\overline{\kappa}_t^j - g(j)\overline{W}_{t+1}^j\right) \tag{A.36}$$

So the expression above reduces to

$$\begin{aligned}
&\mathbb{E}\left[M_{t+1}^{j}S_{t+1}^{j}\middle|\mathcal{S}_t\right] \\
&= \mathbb{E}\left[\exp\left\{g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\overline{W}_{t+1}^j - \frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right\}\times\right. \\
&\qquad \left.\beta\exp\left\{-\overline{\kappa}_t^j - g(j)\overline{W}_{t+1}^j\right\}\middle|\mathcal{S}_t\right] \\
&= \beta\exp\left\{-\overline{\kappa}_t^j - \frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right\}\mathbb{E}\left[\exp\left\{\left(g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right) - g(j)\right)\overline{W}_{t+1}^j\right\}\middle|\mathcal{S}_t\right] \\
&= \beta\exp\left\{-\overline{\kappa}_t^j - \frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right\}\exp\left\{\frac{1}{2}\left(g^{-1}(j)\left(\widetilde{\kappa}_t^t - \overline{\kappa}_t^j\right) - g(j)\right)^2\right\} \\
&= \beta\exp\left\{-\overline{\kappa}_t^j - \frac{1}{2}\left[g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right]^2\right\} \\
&\quad \times\exp\left\{\frac{1}{2}\left(g^{-1}(j)\left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right)\right)^2 - \left(\widetilde{\kappa}_t^j - \overline{\kappa}_t^j\right) + \frac{1}{2}g^2(j)\right\} \\
&= \beta\exp\left(-\widetilde{\kappa}_t^j + \frac{1}{2}g^2(j)\right)
\end{aligned} \tag{A.37}$$

Taking the inverse of the final term in (A.37) we have that the time-t interest rate for country j is:

$$\left(1 + i_{t+1}^j\right) = \beta^{-1}\exp\left(\widetilde{\kappa}_t^j - \frac{1}{2}g^2(j)\right) \tag{A.38}$$

Table A.1: Datasources: Series Identifiers from Datastream for consumption growth and foreign exchange data.

Currency/Country	Consumption	Population	Spot Rate	3-month Forward Rate
USD/United States	USCNPER.D	USPOPTOTP	*	*
AUD/Australia	AUCNPER.D	AUPOPTOT	TDAUDSP	TDAUD3M
CAD/Canada	CNCNPER.D	CNPOPTOTP	TDCADSP	TDCAD3M
CHF/Switzerland	SWCNPER.D	SWPOPTOT	TDCHFSP	TDCHF3M
GBP/United Kingdom	UKCNPER.D	UKPOPTOT	TDGBPSP	TDGBP3M
JPY/Japan	JPCNPER.D	JPPOPTOT	TDJPYSP	TDJPY3M
NOK/Norway	NWCNPER.D	NWPOPTOT	TDNOKSP	TDNOK3M
NZD/New Zealand	NZCNPER.D	NZPOPTOTP	TDNZDSP	TDNZD3M

A.5 Data Sources

Consumption, population, spot exchange rate, and three-month forward exchange rate data all come from Datastream with series identifiers given in Table A.1.

A.6 Sensitivity Analysis: Alternate values of relative entropy penalty parameter θ

The following subsections show that the results of Tables 3.2 - 3.4 are not qualitatively sensitive to moderate changes in the value of the θ_2 parameter. The first subsection imposes a common θ_2 across countries and penalizes Nature more heavily for an equivalent distortion to the model probabilities. The second subsection similarly imposes a common penalty parameter, but penalizes Nature less heavily for its probability distortions relative to the results presented in the main paper body. The third and final subsection relaxes the assumption of a common θ_2 parameter across countries, and repeats the analysis with individually calibrated values.

Penalty Parameter (Inverse) = 0.5

Tables A.2 - A.4 provide results when $\theta_2 = 2$ for all countries. Malevolent Nature (the minimizing player) now pays a higher penalty for an equivalent probability distortion relative to the baseline case of $\theta_2 = 1$, so one would expect the worst-case estimates of consumption growth to differ less relative to baseline estimates than is the case for our analysis in the main paper body. We see some variation in estimated slope coefficients (as expected) but the same pattern of t-statistics, p-values, and R^2 values.

Table A.2: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth (top section) and currency forward discounts (bottom section). Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. In all computations for this table, $\theta_2^{-1} = 0.5$.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	112.1648	5.2685	0.2319	0.0094	0.0538
CAD	20.5751	2.7379	0.1488	0.0653	0.0221
CHF	78.8327	4.3326	0.1895	0.0282	0.0359
GBP	5.3382	3.1080	0.1021	0.1504	0.0104
JPY	87.7699	4.5503	0.2078	0.0188	0.0432
NOK	65.0155	2.5346	0.1540	0.0572	0.0237
NZD	75.0826	2.7323	0.1841	0.0285	0.0339
Regress currency return on (log) forward discount: $r_{t+1,FX}^{j,USD} = \alpha_C^j + \beta_C^j \left[f_t^{j,USD} - e_t^{j,USD} \right] + \varepsilon_{C,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	2.3562	1.3899	0.1591	0.0555	0.0253
CAD	0.8095	0.9370	0.0725	0.2326	0.0053
CHF	1.6580	1.4446	0.1480	0.0685	0.0219
GBP	0.3535	0.4186	0.0314	0.3733	0.0010
JPY	1.7730	2.0181	0.1680	0.0450	0.0282
NOK	0.0862	0.0636	0.0086	0.4693	0.0001
NZD	0.8241	0.2625	0.0508	0.3026	0.0026

Table A.3: Bootstrap permutation test p-values of zero correlation between observed foreign exchange returns and various series derived from consumption data listed in (3.4). The left column lists results for predicted currency returns given by (2.67). The middle column provides the analogous values for the series of interest rate differences implied by the model under ordinary Bayesian beliefs. The final column lists p-values for the null hypothesis of zero correlation between actual currency returns and the model-implied difference in interest rates under worst-case beliefs. In all computations for this table, $\theta_2^{-1} = 0.5$.

p-values for zero correlation of actual returns with:			
	Predicted Returns	Consumption-Implied Interest Rate Difference	
		Bayesian Beliefs	Worst-Case Beliefs
AUD	0.0092	0.0434	0.0361
CAD	0.0612	0.2316	0.2011
CHF	0.0267	0.1598	0.1482
GBP	0.1506	0.2489	0.1725
JPY	0.0158	0.3700	0.3291
NOK	0.0589	0.1627	0.1637
NZD	0.0322	0.0515	0.0468

Table A.4: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth under alternate misspecification fears. Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. The first subtable includes distortions to dynamics, state estimation, and sub-model probabilities. The second subtable includes distortions to dynamics and estimation of the hidden state, but not sub-model probabilities. In all computations for this table, $\theta_2^{-1} = 0.5$.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
Distortions to dynamics, hidden state, and model probabilities:					
AUD	110.5722	5.6945	0.2310	0.0092	0.0534
CAD	20.8379	2.7747	0.1507	0.0628	0.0227
CHF	78.4091	4.3266	0.1887	0.0277	0.0356
GBP	5.3639	3.1588	0.1024	0.1448	0.0105
JPY	87.8555	4.7029	0.2085	0.0178	0.0435
NOK	65.0438	2.5591	0.1544	0.0575	0.0238
NZD	74.7525	2.8516	0.1835	0.0314	0.0337
Distortions to dynamics and hidden state, but *not* model probabilities:					
AUD	-414.9104	-1.0412	-0.0261	0.6046	0.0007
CAD	-273.2711	-0.4566	-0.0312	0.6251	0.0010
CHF	-536.0711	-1.5744	-0.0470	0.6818	0.0022
GBP	38.9367	0.1087	0.0088	0.4598	0.0001
JPY	202.4273	0.1782	0.0190	0.4262	0.0004
NOK	-53.4070	-0.0988	-0.0049	0.5243	0.0000
NZD	-340.3181	-0.4575	-0.0337	0.6310	0.0011

Penalty Parameter (Inverse) = 1.5

Tables A.5 - A.7 provide results when $\theta_2 = \frac{2}{3}$ for all countries. Malevolent Nature (the minimizing player) now pays a lower penalty for an equivalent probability distortion relative to the baseline case of $\theta_2 = 1$, so one would expect the worst-case estimates of consumption growth to differ by a larger amount (relative to baseline estimates) than is the case for our analysis in the main paper body. Again, we see some variation in estimated slope coefficients (as expected) but the same pattern of t-statistics, p-values, and R^2 values.

Table A.5: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth (top section) and currency forward discounts (bottom section). Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. In all computations for this table, $\theta_2^{-1} = 1.5$.

Regress currency return on consumption growth uncertainty:

$$r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$$

Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	18.1913	1.4224	0.1546	0.0597	0.0239
CAD	6.8451	2.4989	0.1693	0.0426	0.0287
CHF	18.7714	4.6865	0.1950	0.0245	0.0380
GBP	6.9544	2.7309	0.1389	0.0796	0.0193
JPY	20.3168	4.5013	0.2068	0.0184	0.0428
NOK	13.7452	2.5570	0.1400	0.0762	0.0196
NZD	12.5746	1.1425	0.1327	0.0845	0.0176

Regress currency return on (log) forward discount:

$$r_{t+1,FX}^{j,USD} = \alpha_C^j + \beta_C^j \left[f_t^{j,USD} - e_t^{j,USD} \right] + \varepsilon_{C,t+1}^j$$

Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	2.3562	1.3899	0.1591	0.0555	0.0253
CAD	0.8095	0.9370	0.0725	0.2326	0.0053
CHF	1.6580	1.4446	0.1480	0.0685	0.0219
GBP	0.3535	0.4186	0.0314	0.3733	0.0010
JPY	1.7730	2.0181	0.1680	0.0450	0.0282
NOK	0.0862	0.0636	0.0086	0.4693	0.0001
NZD	0.8241	0.2625	0.0508	0.3026	0.0026

Table A.6: Bootstrap permutation test p-values of zero correlation between observed foreign exchange returns and various series derived from consumption data listed in (3.4). The left column lists results for predicted currency returns given by (2.67). The middle column provides the analogous values for the series of interest rate differences implied by the model under ordinary Bayesian beliefs. The final column lists p-values for the null hypothesis of zero correlation between actual currency returns and the model-implied difference in interest rates under worst-case beliefs. In all computations for this table, $\theta_2^{-1} = 1.5$.

p-values for zero correlation of actual returns with:			
	Predicted	Consumption-Implied Interest Rate Difference	
	Returns	Bayesian Beliefs	Worst-Case Beliefs
AUD	0.0576	0.0434	0.0254
CAD	0.0408	0.2316	0.1389
CHF	0.0221	0.1598	0.1080
GBP	0.0809	0.2489	0.1300
JPY	0.0168	0.3700	0.2459
NOK	0.0772	0.1627	0.1365
NZD	0.0921	0.0515	0.0401

Table A.7: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth under alternate misspecification fears. Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. The first subtable includes distortions to dynamics, state estimation, and sub-model probabilities. The second subtable includes distortions to dynamics and estimation of the hidden state, but not sub-model probabilities. In all computations for this table, $\theta_2^{-1} = 1.5$.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
Distortions to dynamics, hidden state, and model probabilities:					
AUD	18.1167	1.4376	0.1554	0.0576	0.0242
CAD	6.8279	2.5144	0.1707	0.0424	0.0291
CHF	18.7495	4.7001	0.1952	0.0246	0.0381
GBP	6.9560	2.7300	0.1390	0.0767	0.0193
JPY	20.2996	4.4826	0.2070	0.0194	0.0428
NOK	13.7698	2.5680	0.1404	0.0760	0.0197
NZD	12.5318	1.1353	0.1324	0.0914	0.0175
Distortions to dynamics and hidden state, but *not* model probabilities:					
AUD	-138.3035	-1.0412	-0.0261	0.6046	0.0007
CAD	-91.0904	-0.4566	-0.0312	0.6251	0.0010
CHF	-178.6904	-1.5744	-0.0470	0.6818	0.0022
GBP	12.9789	0.1087	0.0088	0.4598	0.0001
JPY	67.4758	0.1782	0.0190	0.4262	0.0004
NOK	-17.8023	-0.0988	-0.0049	0.5243	0.0000
NZD	-113.4394	-0.4575	-0.0337	0.6310	0.0011

Table A.8: Estimated/calibrated values for more persistent ($i = 1$) and less persistent hidden state ($i = 0$) models.

Currency	$\rho(1)$	$\rho(0)$	$B(1)$	$B(0)$	$\mu_y(1)$	$\mu_y(0)$	g_2	θ_1^{-1}	θ_2^{-1}
USD	0.9800	0.6522	0.0016	0.0026	0.0039	0.0043	0.0034	14.4000	1.4000
AUD	0.9800	0.4807	0.0008	0.0027	0.0036	0.0046	0.0049	10.0000	1.6000
CAD	0.9800	0.7513	0.0007	0.0018	0.0026	0.0038	0.0058	8.5500	1.4000
CHF	0.9800	0.3673	0.0006	0.0015	0.0007	0.0012	0.0030	16.2000	1.5000
GBP	0.9800	0.6657	0.0017	0.0037	0.0039	0.0043	0.0050	9.8000	1.1000
JPY	0.9800	0.7096	0.0100	0.0138	0.0029	0.0023	0.0089	5.5000	1.2000
NOK	0.9800	0.4554	0.0010	0.0040	0.0050	0.0056	0.0102	4.7500	1.5000
NZD	0.9800	0.5713	0.0000	0.0018	0.0045	0.0039	0.0090	5.4500	1.5000

Individually Calibrated Values

Tables A.9 - A.11 provide results when θ_2 is calibrated individually for all countries using the detection error probability method of [12]. In this case, I estimate model parameters by maximum likelihood, fix the volatility of the shock matrix G at this value, and then find two models (one with a highly persistent hidden state, one with a less persistent hidden state) that have equal empirical likelihoods over the data sample. Each country's value of θ_1 is calculated to give an average detection error probability of 0.4 (supposing the hidden state is known), and then θ_2 is calibrated to set the overall detection error probability equal to 0.2. Parameter values are given in Table A.8. Again, the results are qualitatively similar to those in the rest of the paper.

Table A.9: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth (top section) and currency forward discounts (bottom section). Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. For computations in this table, parameters have the values listed in Table A.8.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	20.6222	1.5445	0.1610	0.0477	0.0259
CAD	7.1440	2.5194	0.1671	0.0475	0.0279
CHF	19.8194	4.7086	0.1939	0.0239	0.0376
GBP	7.0337	2.6929	0.1403	0.0802	0.0197
JPY	21.3473	4.5586	0.2049	0.0194	0.0420
NOK	14.6469	2.5307	0.1406	0.0724	0.0198
NZD	13.7799	1.2177	0.1371	0.0807	0.0188
Regress currency return on (log) forward discount: $r_{t+1,FX}^{j,USD} = \alpha_C^j + \beta_C^j \left[f_t^{j,USD} - e_t^{j,USD} \right] + \varepsilon_{C,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
AUD	2.3562	1.3899	0.1591	0.0543	0.0253
CAD	0.8095	0.9370	0.0725	0.2302	0.0053
CHF	1.6580	1.4446	0.1480	0.0665	0.0219
GBP	0.3535	0.4186	0.0314	0.3756	0.0010
JPY	1.7730	2.0181	0.1680	0.0428	0.0282
NOK	0.0862	0.0636	0.0086	0.4610	0.0001
NZD	0.8241	0.2625	0.0508	0.3041	0.0026

Table A.10: Bootstrap permutation test p-values of zero correlation between observed foreign exchange returns and various series derived from consumption data listed in (3.4). The left column lists results for predicted currency returns given by (2.67). The middle column provides the analogous values for the series of interest rate differences implied by the model under ordinary Bayesian beliefs. The final column lists p-values for the null hypothesis of zero correlation between actual currency returns and the model-implied difference in interest rates under worst-case beliefs.

p-values for zero correlation of actual returns with:			
	Predicted	Consumption-Implied Interest Rate Difference	
	Returns	Bayesian Beliefs	Worst-Case Beliefs
AUD	0.0496	0.0434	0.0260
CAD	0.0422	0.2316	0.1433
CHF	0.0230	0.1598	0.1113
GBP	0.0784	0.2489	0.1306
JPY	0.0182	0.3700	0.2530
NOK	0.0758	0.1627	0.1391
NZD	0.0849	0.0515	0.0406

Table A.11: Output from regression of realized excess currency returns on model-implied deviations from uncovered interest parity generated by differences between Bayesian and worst-case estimates of consumption growth under alternate misspecification fears. Column I reports regression betas. Column II provides the t-statistic for regression betas using the method of Hansen and Hodrick with 12 lags. Column III reports the correlation between the two series in the regression. Column IV lists p-values for a bootstrap permutation test of zero correlation between the two series. The final column reports the regression R^2. The first subtable includes distortions to dynamics, state estimation, and sub-model probabilities. The second subtable includes distortions to dynamics and estimation of the hidden state, but not sub-model probabilities.

Regress currency return on consumption growth uncertainty: $r_{t+1,FX}^{j,USD} = \alpha_U^j + \beta_U^j \left[\left(\overline{\kappa}_t^{USD} - \widetilde{\kappa}_t^{USD} \right) - \left(\overline{\kappa}_t^j - \widetilde{\kappa}_t^j \right) \right] + \varepsilon_{U,t+1}^j$					
Column	I	II	III	IV	V
Statistic	β	$t_{\beta,HH}$	$Corr.$	$Pr\,(\text{corr} = 0)$	R^2
Distortions to dynamics, hidden state, and model probabilities:					
AUD	20.5112	1.5639	0.1618	0.0496	0.0262
CAD	7.1279	2.5397	0.1686	0.0422	0.0284
CHF	19.7856	4.7070	0.1940	0.0230	0.0376
GBP	7.0343	2.6913	0.1404	0.0784	0.0197
JPY	21.3253	4.5460	0.2051	0.0182	0.0421
NOK	14.6698	2.5418	0.1411	0.0758	0.0199
NZD	13.7209	1.2090	0.1367	0.0849	0.0187
Distortions to dynamics and hidden state, but *not* model probabilities:					
AUD	-101.0684	-0.3880	-0.0176	0.5710	0.0003
CAD	-97.5959	-0.4566	-0.0312	0.6302	0.0010
CHF	-190.2410	-1.5488	-0.0464	0.6789	0.0022
GBP	22.6299	0.1560	0.0124	0.4462	0.0002
JPY	72.2930	0.1783	0.0190	0.4285	0.0004
NOK	-19.3317	-0.1001	-0.0050	0.5209	0.0000
NZD	-122.0689	-0.4581	-0.0338	0.6317	0.0011

REFERENCES

[1] Fernando Alvarez, Andrew Atkeson, and Patrick J. Kehoe. Time-Varying Risk, Interest Rates, and Exchange Rates in General Equilibrium. *Review of Economic Studies*, 76(3):851–878, 2009.

[2] David K. Backus, Silverio Foresi, and Chris Telmer. Affine term structure models and the forward premium anomaly. *Journal of Finance*, 56(1):279–304, 02 2001.

[3] David K. Backus, Federico Gavazzoni, Christopher Telmer, and Stanley E. Zin. Monetary policy and the uncovered interest parity puzzle. NBER Working Papers 16218, National Bureau of Economic Research, Inc, July 2010.

[4] Kimberly Berg and Nelson Mark. Measures of global uncertainty and carry-trade excess returns. *Journal of International Money and Finance*, 2017.

[5] A. Craig Burnside and Jeremy J. Graveline. On the asset market view of exchange rates. Working Paper 18646, National Bureau of Economic Research, December 2012.

[6] Riccardo Colacito and Mariano M. Croce. Risks for the long run and the real exchange rate. *Journal of Political Economy*, 119(1):153 – 181, 2011.

[7] Riccardo Colacito and Mariano M. Croce. International Asset Pricing with Recursive Preferences. *Journal of Finance*, 68(6):2651–2686, December 2013.

[8] A.C. Davison and D.V. Hinkley. *Bootstrap Methods and Their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1997.

[9] Eugene F. Fama. Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3):319 – 338, 1984.

[10] Lars Peter Hansen and Robert J Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5):829–53, October 1980.

[11] Lars Peter Hansen and Thomas J. Sargent. Recursive robust estimation and control without commitment. *Journal of Economic Theory*, 136(1):1–27, September 2007.

[12] Lars Peter Hansen and Thomas J. Sargent. Fragile beliefs and the price of uncertainty. *Quantitative Economics*, 1(1):129–162, 2010.

[13] Lars Peter Hansen and Thomas J. Sargent. The Price of Macroeconomic Uncertainty with Tenuous Beliefs. Working papers, December 2017.

[14] Tarek A. Hassan. Country Size, Currency Unions, and International Asset Returns. *Journal of Finance*, 68(6):2269–2308, December 2013.

[15] Lucas Husted, John Rogers, and Bo Sun. Uncertainty, currency excess returns, and risk reversals. *Journal of International Money and Finance*, 2017.

[16] Cosmin Ilut. Ambiguity aversion: Implications for the uncovered interest rate parity puzzle. *American Economic Journal: Macroeconomics*, 4(3):33–65, July 2012.

[17] Adilzhan Ismailov and Barbara Rossi. Uncertainty and deviations from uncovered interest parity. *Journal of International Money and Finance*, 2017.

[18] Klodina Istrefi and Sarah Mouabbi. Subjective interest rate uncertainty and the macroeconomy: A cross-country analysis. *Journal of International Money and Finance*, 2017.

[19] Hanno Lustig, Nikolai Roussanov, and Adrien Verdelhan. Common risk factors in currency markets. *Review of Financial Studies*, 24(11):3731–3777, 2011.

[20] Hanno Lustig, Nikolai Roussanov, and Adrien Verdelhan. Countercyclical currency risk premia. *Journal of Financial Economics*, 111(3):527–553, 2014.

[21] Almuth Scholl and Harald Uhlig. New evidence on the puzzles: Results from agnostic identification on monetary policy and exchange rates. *Journal of International Economics*, 76(1):1–13, September 2008.

[22] David S. Stoffer and Kent D. Wall. Bootstrapping state-space models: Gaussian maximum likelihood estimation and the Kalman filter. *Journal of the American Statistical Association*, 86(416):1024–1033, 1991.

[23] Adrien Verdelhan. A habit-based explanation of the exchange rate risk premium. *The Journal of Finance*, 65(1):123–146, 2010.