Online Technical Appendix

A Posterior estimates

Table 1: Estimation Results

	Prior distribution	Prior mean	Prior std.dev.	Posterior Mode	Posterior Mean	Confidence Interval	
AUD/USD							
α_i	γ	1	0.5	1.190	1.292	[0.66 ,	1.872]
γ	γ	1	0.5	0.275	0.233	[0.14 ,	0.323]
ρ^π	β	0.8	0.1	0.896	0.891	[0.86 ,	0.923]
σ^R	γ^{-1}	0.02	0.5	0.019	0.018	[0.015 ,	0.021]
σ^{CF}	γ^{-1}	0.02	0.5	0.038	0.040	[0.033 ,	0.046]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[0.00 ,	0.001]
Marginal Density				2295		Nobs	221

	Prior distribution	Prior mean	Prior std.dev.	Posterior Mode	Posterior Mean	Confidence Interval	
CAD/USD							
α_i	γ	1	0.5	0.764	0.893	[0.35 ,	1.419]
γ	γ	1	0.5	0.357	0.327	[0.24 ,	0.406]
ρ^π	β	0.8	0.1	0.915	0.908	[0.87 ,	0.944]
σ^R	γ^{-1}	0.02	0.5	0.015	0.015	[0.012 ,	0.017]
σ^{CF}	γ^{-1}	0.02	0.5	0.028	0.030	[0.024 ,	0.036]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[0.00 ,	0.000]
Marginal Density				2528		Nobs	221

	Prior distribution	Prior mean	Prior std.dev.	Posterior Mode	Posterior Mean	Confidence Interval	
CHF/USD							
α_i	γ	1	0.5	0.975	1.108	[0.56 ,	1.666]
γ	γ	1	0.5	0.361	0.318	[0.20 ,	0.420]
ρ^π	β	0.8	0.1	0.899	0.893	[0.86 ,	0.931]
σ^R	γ^{-1}	0.02	0.5	0.021	0.020	[0.016 ,	0.024]
σ^{CF}	γ^{-1}	0.02	0.5	0.035	0.037	[0.030 ,	0.045]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[0.00 ,	0.000]
Marginal Density				2417		Nobs	222

	Prior distribution	Prior mean	Prior std.dev.	Posterior Mode	Posterior Mean	Confidence Interval	
EUR/USD							
α_i	γ	1	0.5	1.111	1.246	[0.69 ,	1.824]
γ	γ	1	0.5	0.341	0.292	[0.18 ,	0.402]
ρ^π	β	0.8	0.1	0.886	0.883	[0.84 ,	0.926]
σ^R	γ^{-1}	0.02	0.5	0.019	0.019	[0.015 ,	0.022]
σ^{CF}	γ^{-1}	0.02	0.5	0.032	0.035	[0.028 ,	0.041]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[0.00 ,	0.000]
Marginal Density				1856		Nobs	170

Table 1: Estimation Results cont...

	Prior distribution	mean	std.dev.	Posterior Mode	Posterior Mean	Confidence Interval				
GBP/USD										
α_i	γ	1	0.5	0.848	0.956	[	0.47	,	1.467	]
γ	γ	1	0.5	0.446	0.399	[	0.26	,	0.517	]
ρ^π	β	0.8	0.1	0.935	0.930	[	0.91	,	0.950	]
σ^R	γ^{-1}	0.02	0.5	0.019	0.018	[	0.014	,	0.022	]
σ^{CF}	γ^{-1}	0.02	0.5	0.026	0.029	[	0.023	,	0.035	]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[	0.00	,	0.000	]
Marginal Density				2466			Nobs		222	

	Prior distribution	mean	std.dev.	Posterior Mode	Posterior Mean	Confidence Interval				
JPY/USD										
α_i	γ	1	0.5	0.829	0.977	[	0.44	,	1.501	]
γ	γ	1	0.5	0.434	0.385	[	0.26	,	0.497	]
ρ^π	β	0.8	0.1	0.916	0.912	[	0.88	,	0.940	]
σ^R	γ^{-1}	0.02	0.5	0.025	0.023	[	0.018	,	0.029	]
σ^{CF}	γ^{-1}	0.02	0.5	0.036	0.039	[	0.031	,	0.048	]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[	0.00	,	0.000	]
Marginal Density				2934			Nobs		280	

	Prior distribution	mean	std.dev.	Posterior Mode	Posterior Mean	Confidence Interval				
NZD/USD										
α_i	γ	1	0.5	1.014	1.132	[	0.55	,	1.733	]
γ	γ	1	0.5	0.285	0.249	[	0.16	,	0.334	]
ρ^π	β	0.8	0.1	0.885	0.881	[	0.85	,	0.915	]
σ^R	γ^{-1}	0.02	0.5	0.020	0.019	[	0.016	,	0.022	]
σ^{CF}	γ^{-1}	0.02	0.5	0.040	0.042	[	0.035	,	0.049	]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[	0.00	,	0.000	]
Marginal Density				2299			Nobs		222	

	Prior distribution	mean	std.dev.	Posterior Mode	Posterior Mean	Confidence Interval				
SEK/USD										
α_i	γ	1	0.5	0.937	1.075	[	0.52	,	1.593	]
γ	γ	1	0.5	0.375	0.330	[	0.21	,	0.434	]
ρ^π	β	0.8	0.1	0.916	0.912	[	0.89	,	0.938	]
σ^R	γ^{-1}	0.02	0.5	0.022	0.021	[	0.017	,	0.025	]
σ^{CF}	γ^{-1}	0.02	0.5	0.035	0.038	[	0.031	,	0.046	]
σ^π	γ^{-1}	0.0003	0.005	0.000	0.000	[	0.00	,	0.000	]
Marginal Density				2358			Nobs		222	

B Posterior densities by currency pair

Figure 1: Posterior Densities
a) AUD/USD

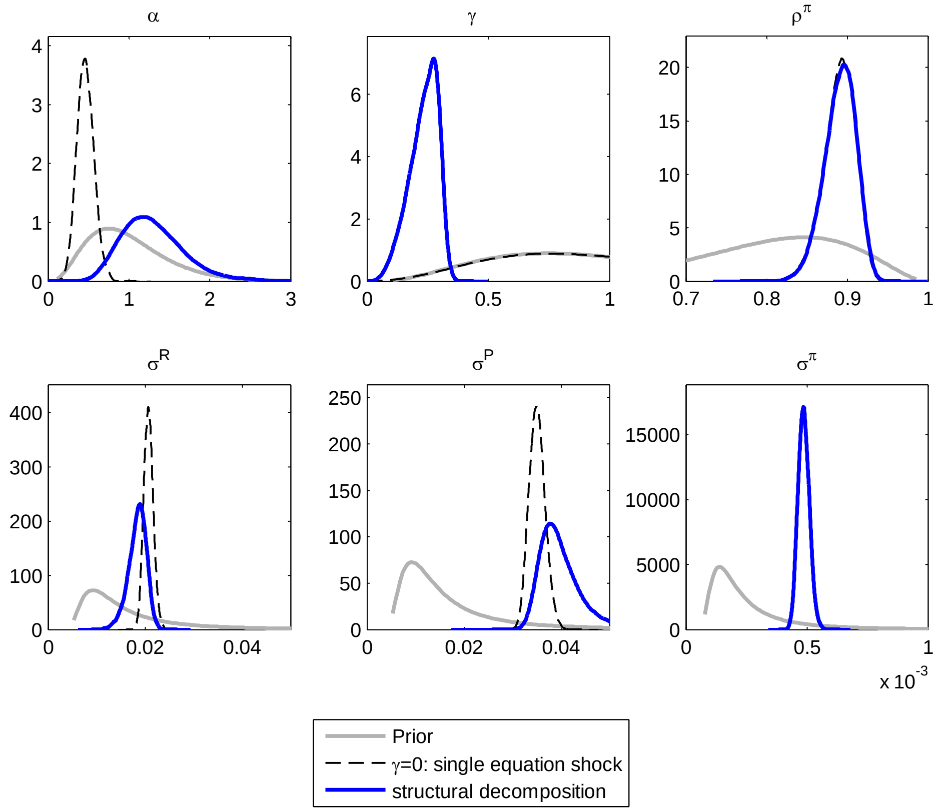

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the risk factor is a single equation exchange rate residual that does not affect the fixed-income market. That setup is equivalent to a single equation estimate with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural decomposition in which the risk factor is identified through its common effect on both the fixed income and foreign exchange markets.

Figure 2: Posterior Densities
b) CAD/USD

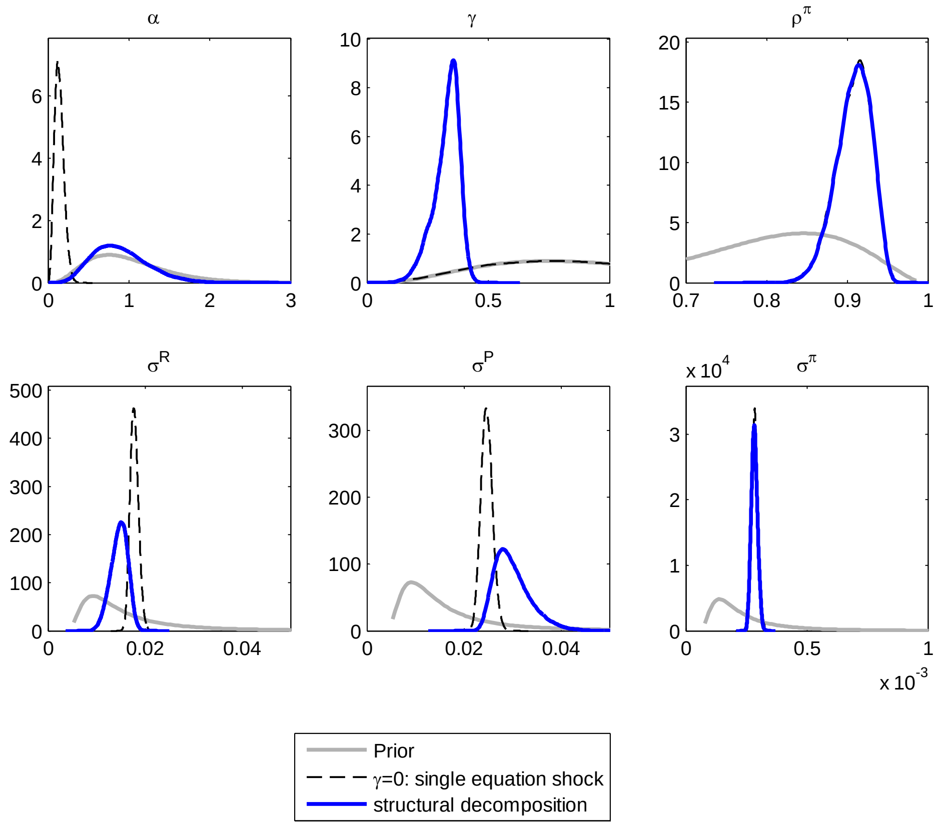

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: Posterior Densities
c) CHF/USD

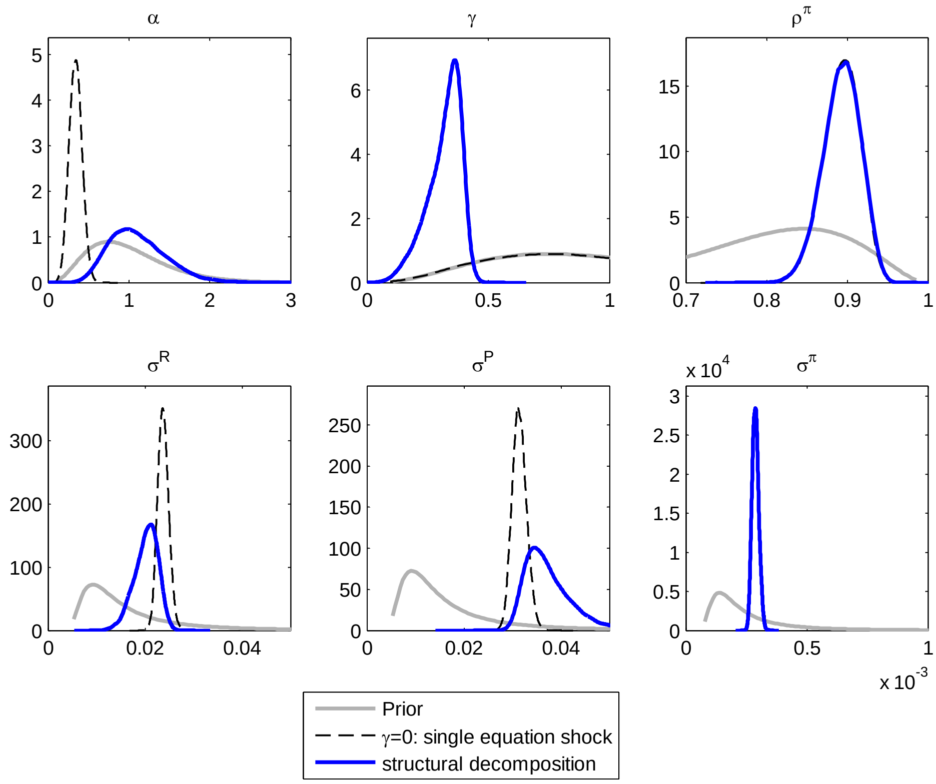

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: Posterior Densities
d) EUR/USD

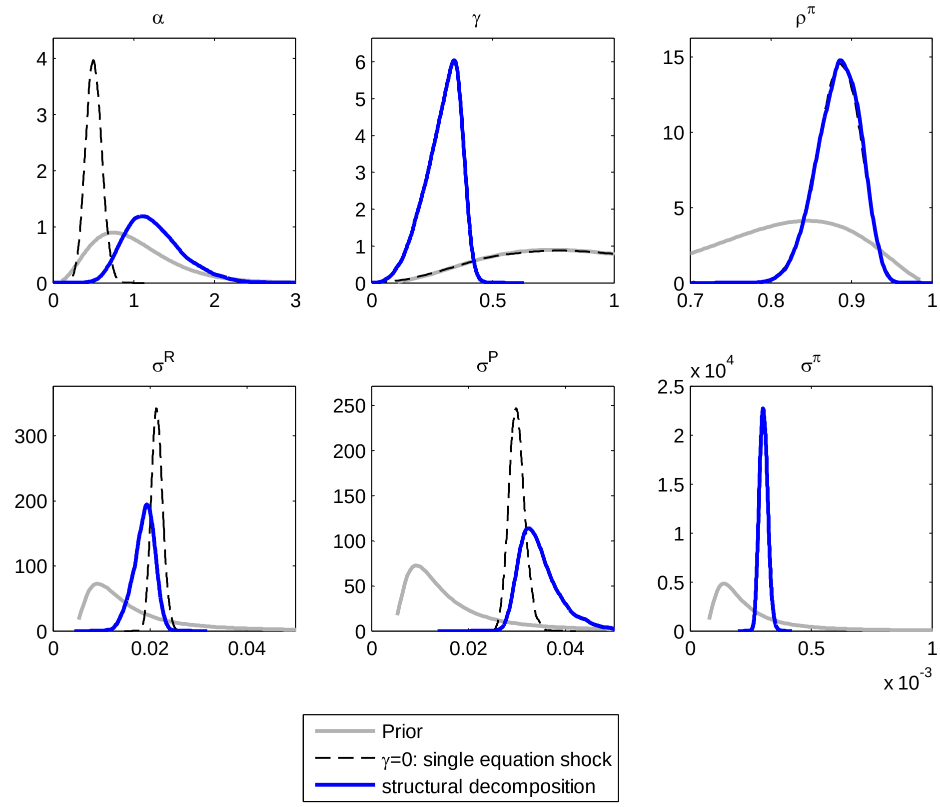

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: Posterior Densities
e) GBP/USD

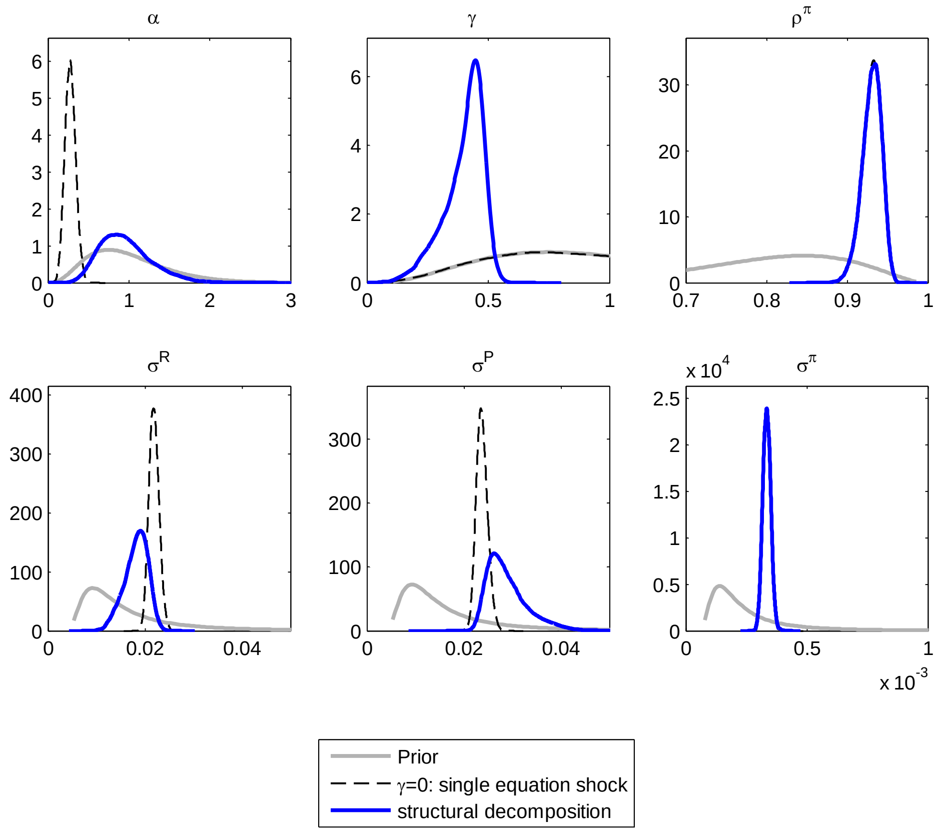

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: Posterior Densities
f) JPY/USD

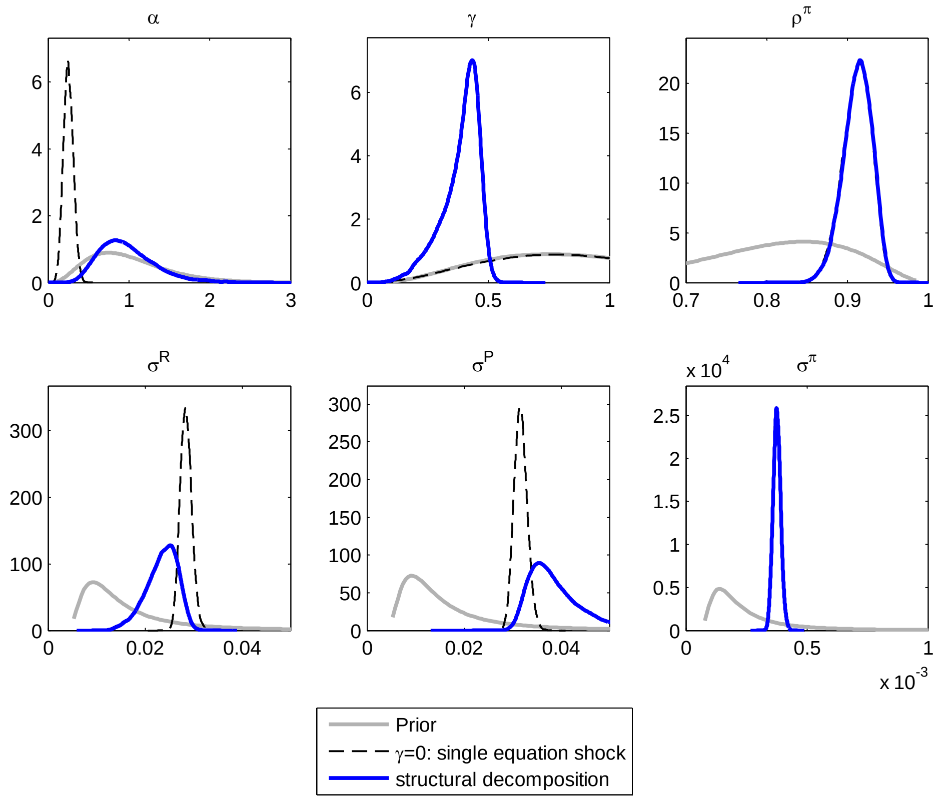

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: Posterior Densities
g) NZD/USD

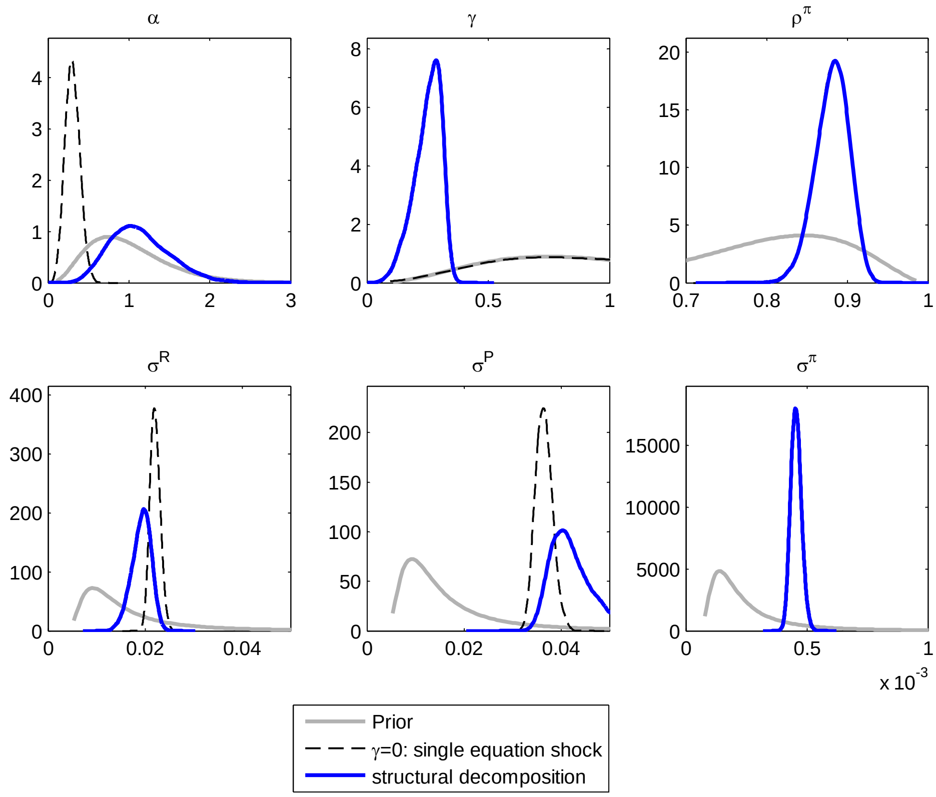

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

Figure 2: h) Posterior Densities: SEK/USD

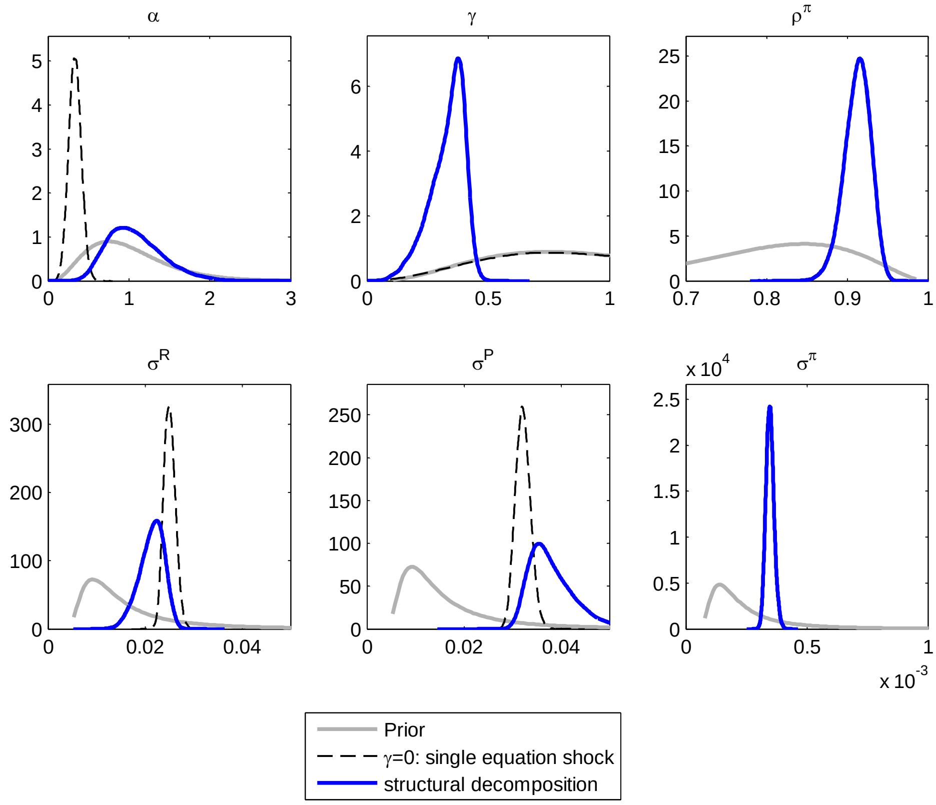

The grey line is the prior density. The thin solid black line is the model for which $\gamma = 0$ so that the shock $\eta^C F$ is a single equation shock to the exchange rate equation that does not affect the fixed-income market. That is model is equivalent to a single equation model with a sign restriction on the prior for α. The solid blue line is the posterior density for the structural model.

C Convergence statistics

Table 2: Geweke (1992) Convergence Tests

Currency	Parameter α	γ	ρ_π	$\sigma^{\hat{R}^f}$	σ^Λ	σ^π
	a) No taper (Assumption of iid draws)					
AUD	0.00	0.00	0.00	0.00	0.00	0.00
CAD	0.00	0.26	0.00	0.00	0.00	0.00
CHF	0.00	0.00	0.00	0.00	0.00	0.00
EUR	0.45	0.02	0.00	0.01	0.09	0.00
GBP	0.00	0.00	0.00	0.00	0.00	0.00
JPY	0.00	0.00	0.31	0.00	0.00	0.00
NZD	0.00	0.00	0.03	0.00	0.00	0.02
SEK	0.00	0.00	0.00	0.00	0.00	0.00
	b) 4% taper					
AUD	0.14	0.19	0.33	0.09	0.14	0.37
CAD	0.15	0.79	0.50	0.05	0.06	0.47
CHF	0.42	0.57	0.01	0.38	0.42	0.65
EUR	0.94	0.64	0.48	0.77	0.87	0.20
GBP	0.67	0.33	0.17	0.61	0.65	0.35
JPY	0.73	0.10	0.82	0.76	0.64	0.59
NZD	0.33	0.52	0.61	0.31	0.51	0.75
SEK	0.61	0.30	0.21	0.59	0.65	0.27
	c) 8% taper					
AUD	0.17	0.20	0.33	0.12	0.16	0.36
CAD	0.13	0.78	0.50	0.04	0.04	0.43
CHF	0.43	0.56	0.00	0.39	0.43	0.63
EUR	0.93	0.62	0.47	0.75	0.86	0.14
GBP	0.67	0.32	0.15	0.60	0.65	0.34
JPY	0.74	0.12	0.81	0.76	0.64	0.55
NZD	0.35	0.52	0.61	0.34	0.53	0.73
SEK	0.61	0.30	0.17	0.59	0.64	0.29

Notes: Figures are p-value of a chi squared test for equality of means in the beginning and the end of the retained section of the MCMC chain (50000 to 125000 vs 125000 to 200000). A value above 0.05 indicates that the null hypothesis of equal means and thus convergence cannot be rejected at the 5 percent level. Differing values over the taper step signals the presence of significant autocorrelation in draws. The estimates using a higher tapering are usually more reliable (Dynare Reference Manual, version 4.4.2) because the posterior draws are serially dependent.

D Robustness

This appendix shows the posterior densities for α, and other parameters where relevant, for variations in the model estimation (use of plain vanilla swaps, 5 or 15-year zero-coupon swaps, and inflation-indexed bond forecasts of inflation), or model formulation (separation of nominal and inflation components of real returns, changes in the error correction structure, and inclusion of the terms of trade):

D.1 Plain vanilla swaps: a discounted sum of returns

First, the model was re-estimated using plain vanilla swaps, rather than zero-coupon swaps. Plain vanilla swaps provide a discounted sum of expected returns. They have the advantage that they are generally available for a longer-period than zero- coupon swap rates for most currency pairs.

$$\begin{aligned} \sum_{k=1}^{N} \beta^k i_t^S &= E_t \sum_{k=1}^{N} \beta^k i_{t+k-1} \\ \frac{1-\beta^N}{1-\beta} i_t^S &\simeq E_t \sum_{k=1}^{N} \beta^k i_{t+k-1} \end{aligned}$$

Time series data for 10-year plain vanilla swaps are visually similar to the zero-coupon time series for most currency pairs; they are visually materially different for the AUD/USD, CAD/USD and NZD/USD. When $\mathbf{R}_t$ was constructed using 10-year plain vanilla swaps (posterior distributions are shown in Figure 3 below), the average posterior mode across the eight USD currency pairs for (α, γ) was (1.01,0.35), compared to (0.96,0.36) in the baseline estimation.

Figure 3: Posterior Densities for α: 10-year vanilla swaps

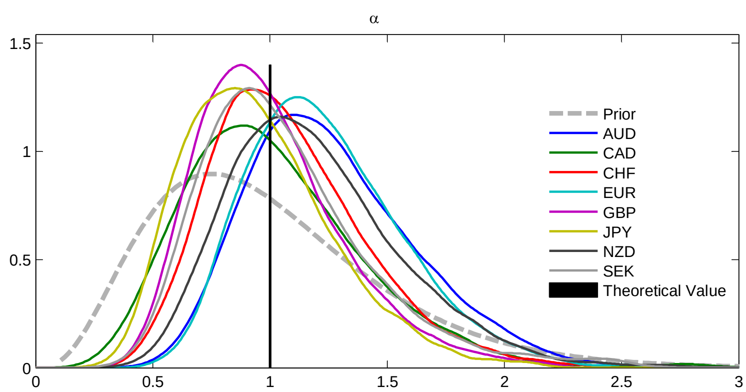

The grey dashed line is the prior density.

D.2 Nominal and inflation components of R_t

To further assess the effect of the simple AR(1) forecast of inflation on the estimated value of α, the system was also estimated with the exchange rate response to $\mathbf{R}_t$ separated into a nominal and inflation component:

$$\Delta q_t \quad = \quad -\alpha_i \Delta \mathbf{R}_{i,t} + \alpha_\pi \Delta \mathbf{R}_{\pi,t} + \eta_t^\Lambda \tag{1}$$

Across the eight USD currency pairs, the average posterior mode for $(\alpha_i, \alpha_\pi, \gamma)$ was (0.97, 0.81, 0.35). The variance of $\mathbf{R}_t$ is dominated by the nominal component, so perhaps unsurprisingly, α_i is very close to α. The response to the inflation component is a little weaker than one. That result could reflect measurement error in view of the simple forecast approach. Another explanation is that the Taylor rule response to inflation[1] requires a rise in future real interest rates to restore balance. Although the nominal equilibrium must still eventually fall to restore PPP, the higher real returns in the short term offset that effect over the short to medium term. The average posterior mode for γ was little changed at 0.35.

Figure 4: Posterior Densities for nominal and inflation components of α

a) Real exchange rate response to ***nominal*** expected returns:

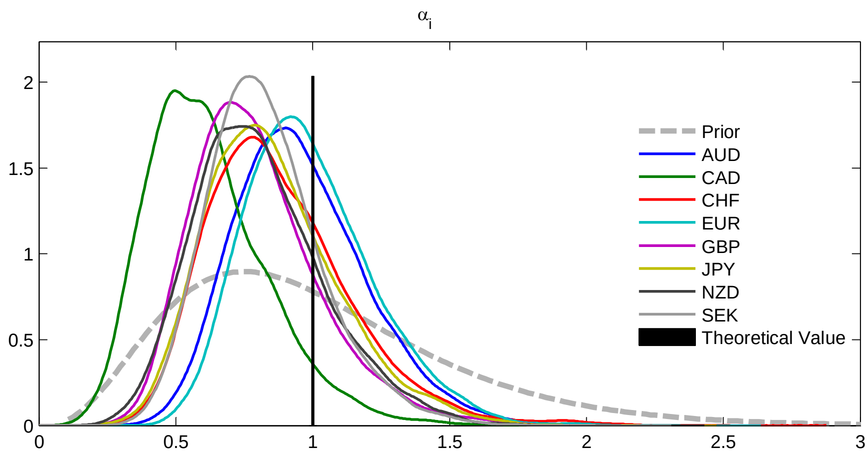

[1] **?**) show that empirically, exchange rates appreciate in response to news about inflation.

b) Real exchange rate response to expected relative inflation component of nominal returns:

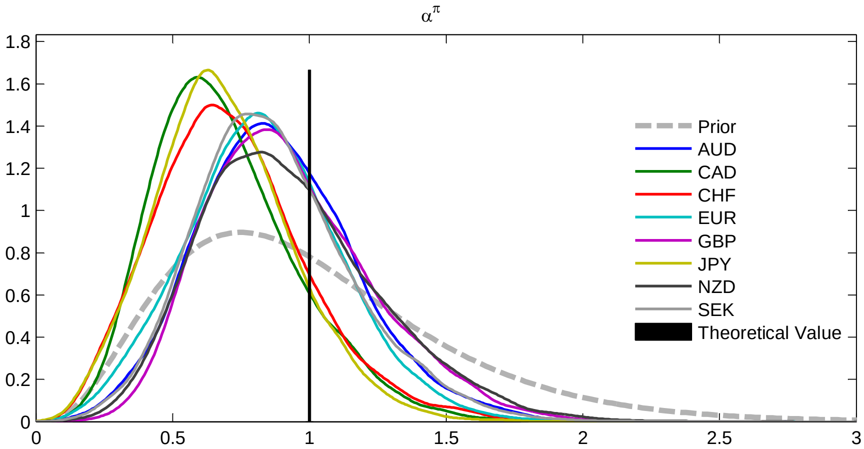

The grey dashed line is the prior density. The term $-\alpha \Delta R_t$ in the model is replaced with $-\alpha^i \Delta R_t^i + \alpha^\pi \Delta R_t^\pi$.

D.3 Inflation forecasts from inflation-indexed bonds

Rather than using a naive AR(1) inflation forecast, inflation forecasts could be derived from long-term inflation-indexed bonds or informed by surveys of inflation expectations. In principle, inflation forecasts derived from inflation-indexed bonds should be more accurate market-based forecasts. Where they exist, however, inflation indexed bond markets are often relatively illiquid and data series are considerably shorter than those for nominal interest rate swaps.

The longest series available is for the GBP/USD pair. Data is available from July 2004 to July 2013, giving only 108 observations. The appropriate forecast contract is the zero-coupon inflation swap which has a structure similar to the zero-coupon swap. The Figure below plots the expected sum of inflation differentials based on annual inflation and the measure using 10-year inflation indexed bonds. The correlation between the two measures is 0.72. The sum of expected inflation differentials from zero-coupon inflation swaps has a standard deviation more than twice that of the measure based on observed inflation.

Figure 5: GBP/USD: relative inflation forecast
(% deviation from mean)

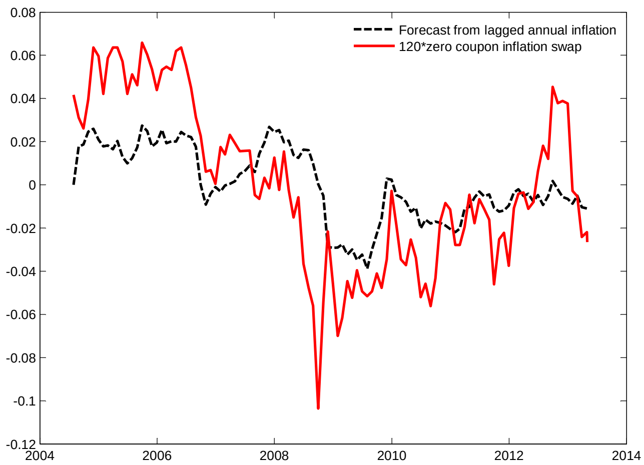

The average estimated posterior mode for α, γ was (0.84,0.38) compared to (0.85, 0.29) for the benchmark formulation estimated over the same short period. Posterior distributions are shown in Figure **??**. Thus the simple AR(1) forecast based on observed annual inflation doesn't appear to have a material effect on the estimated value of α. The estimated weight of the currency risk premium in the fixed income market is larger with the inflation swap forecast. As longer data series for inflation indexed bonds become available, the relative performance of AR(1) forecasts can be examined in more detail. An alternative approach would be to employ survey measures of expected inflation to inform inflation forecasts. An advantage of inflation swaps is that they reflect transactions that require good forecasts rather than survey responses that are not based on actual transactions and can lag observed data.

Figure 6: Prior and posterior densities: GBP/USD 2004-2013 period

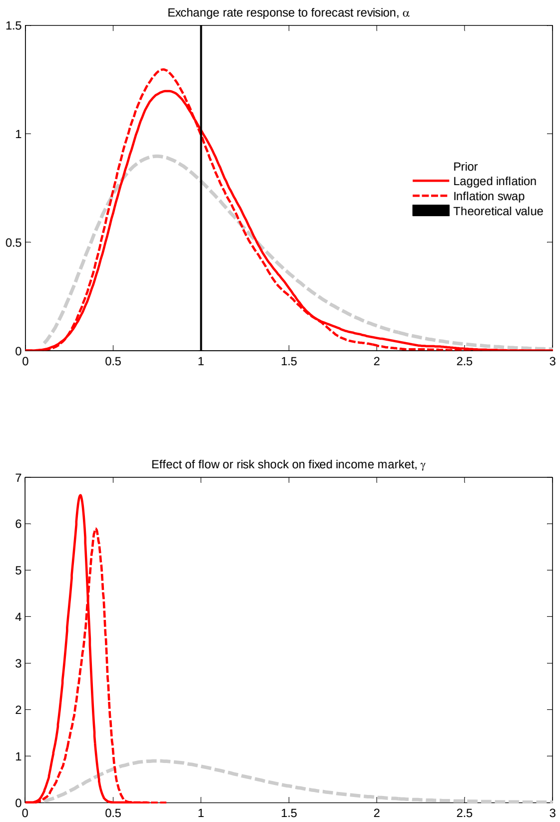

D.4 Flat priors for α and γ

To asses the bias from the priors, the model was re-estimated with flat (uniform distribution) priors for α. The average posterior mode estimate for α was a little higher at 1.12, suggesting that the prior biases the estimated value of α downward slightly. In all cases, unity is still well inside the 90% confidence interval.

Figure 7: Posterior Densities for flat priors

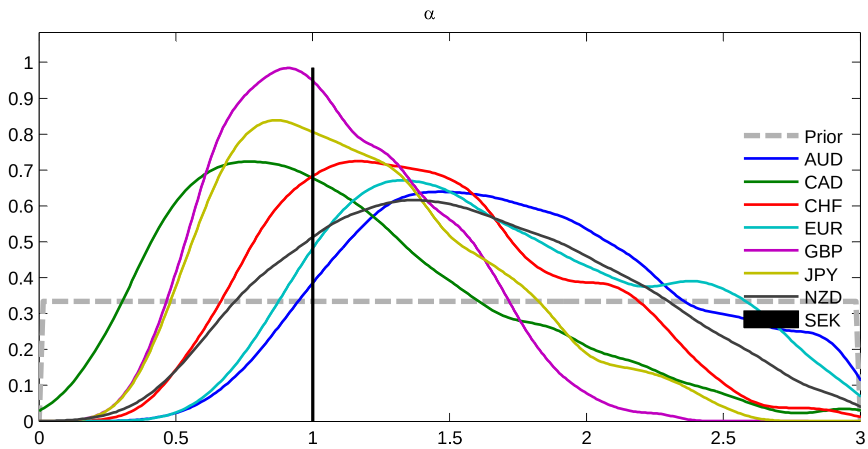

The grey dashed line is the prior density.

D.5 5- and 15-year zero-coupon swaps

A key question for the estimation is, What maturity of swap contract is long enough to proxy the infinite undiscounted sum? If the variables are assumed to have reached steady state within 10 years, then 10 years should be long enough to capture exchange rate fluctuations, with the remaining horizon from ten years to ∞ captured in the steady state constants. In the baseline estimates, the exchange rate response to a 10-year sum of expected returns is near one, suggesting that a 10-year horizon is long enough to proxy the infinite sum.

A useful benchmark for comparison is an AR(1) forecast of future returns. In that framework, the exchange response to the 1-period return should be $1/(1-\rho)$ where ρ is the AR(1) coefficient for $(r_t - r_t^*)$. The expected value of the coefficient α falls from $(1/(1-\rho)$ to $(1-\rho^N)/(1-\rho)$ as the forecast horizon increases. In the limit, α is expected to converge, from above, towards one. Therefore, the estimated response to the 5-year return is expected to be greater than one, while the estimated response tot he 15-year sum of returns is expected to be near one.

When $\mathbf{R}_t$ was constructed as a 5-year forward sum using 5-year zero-coupon swaps, the average posterior mode for (α, γ) was (1.80,0.13) across the eight currency pairs (see Figure **??**). Thus, 5-year swaps appear to be too short a horizon to adequately capture the infinite sum.

Figure 8: Posterior Densities for α: 5-year un-discounted sum

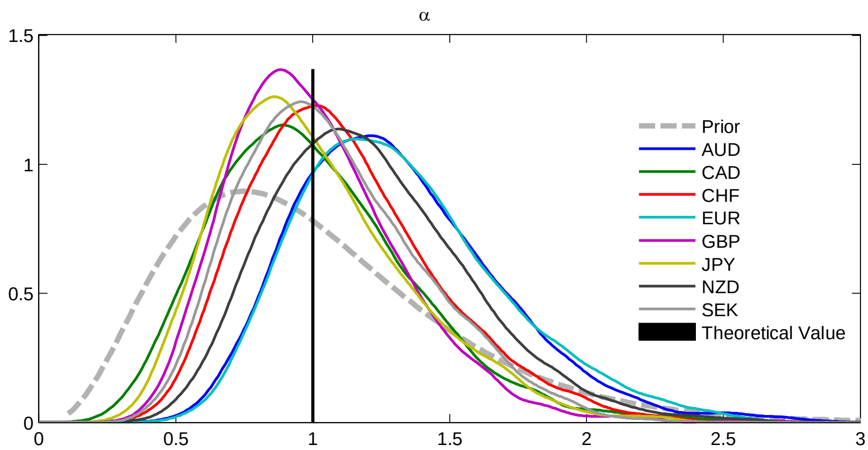

The grey dashed line is the prior density.

When $\mathbf{R}_t$ was constructed as a 15-year forward sum using 15-year swaps, the average posterior mode for (α, γ) was (0.78,0.49). Relative to an AR(1) benchmark, those estimates of $\alpha < 1$ are low.

The higher estimates for γ for longer-horizon sums suggest that a growing role for risk may be part of the explanation. Less certain, longer horiozn payoffs may be subject to higher risk premia.

Figure 9: Posterior Densities for α: 15-year un-discounted sum

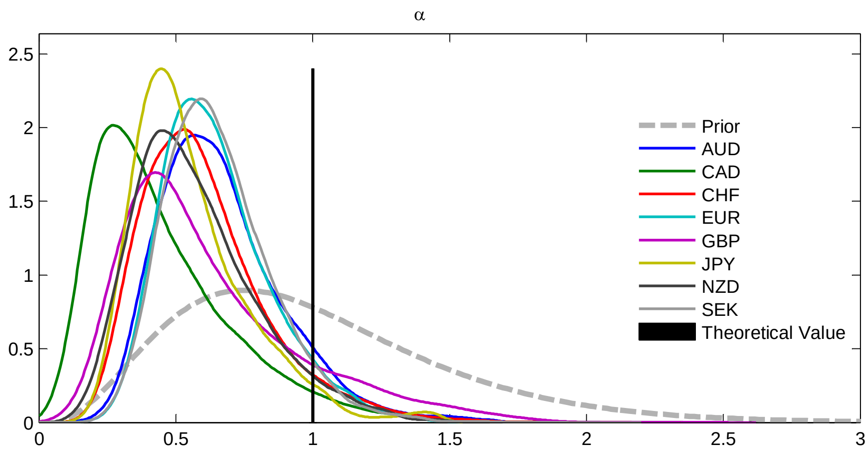

The grey dashed line is the prior density.

Table 3: Data sample for 15-year zero-coupon swaps

Currency	start	end
AUD/USD	01/1995	05/2013
CAD/USD	01/1995	05/2013
CHF/USD	12/1994	05/2013
EUR/USD	03/1999	04/2013
GBP/USD	12/1994	05/2013
JPY/USD	02/1990	05/2013
NZD/USD	12/1994	05/2013
SEK/USD	12/1994	05/2013

D.6 Response to five-year returns at zero, five and ten-year horizons

To explore reasons why the estimated response to longer-horizon sums is less than one, the 15-year sum was divided into three 5-year components: 0-5 year returns, 5-10 year returns and 10-15 year returns. The following graphs show the estimated exchange rate response to expected returns, α, and the weight of the risk premium in gross returns, γ, for the five-year sums of expected returns at those horizons. The estimation is over the same period for each of the three horizons below, but typically for a shorter time period than the five-year sums above because of the available sample period for 15-year zero-coupon swaps. The estimates for (α, γ) for the first five years of real returns is (1.53,0.10), compared to (0.82,0.11) for returns at the five to ten year horizon and (0.71,0.14) for returns at the ten to fifteen years horizon.

Figure 10: Estimates of α and γ at one to five, five to ten and ten to 15 year horizons

a) One to five-year horizon

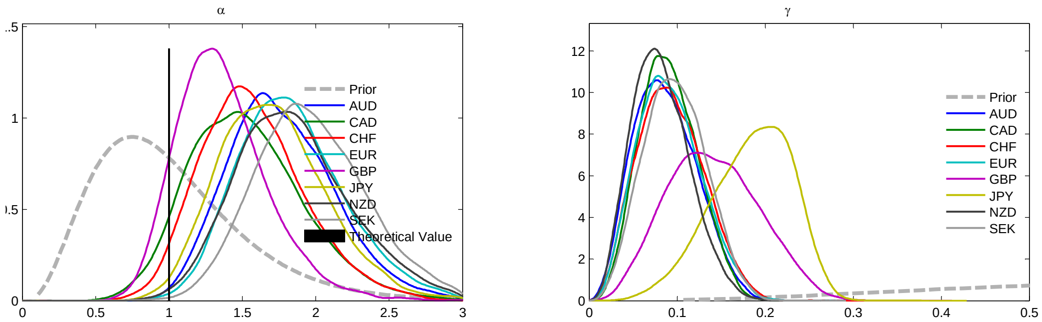

b) Five to ten-year horizon

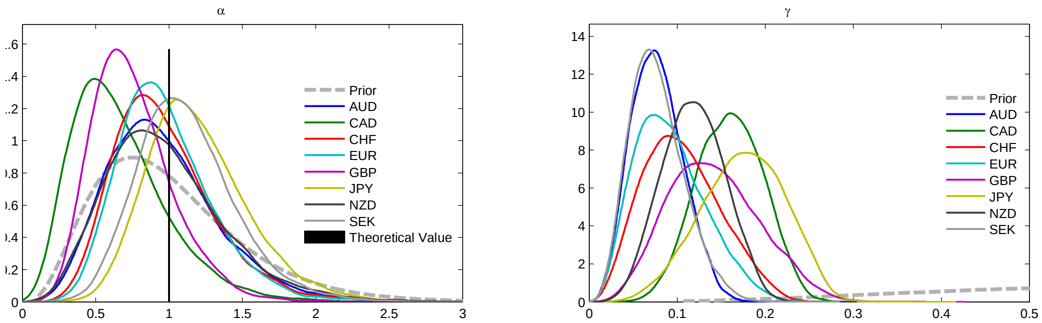

a) Ten to 15-year horizon year horizon

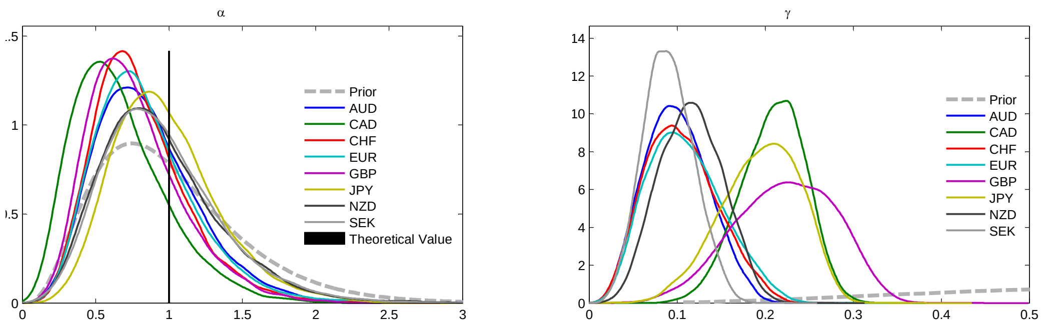

The exchange rate response to subsequent returns was considerably lower. In contrast, a simple AR(1) model would imply a considerably larger coefficient as the 5-year period shifts into the future (the sum becomes smaller, so needs a large coefficient to account for earlier returns). Therefore, it appears that investors put a lower weight on more distant future returns. One explanation is the rate of time-preference. Although the relative asset price is an undiscounted relative sum, it is based on Euler equations for domestic and foreign bonds that are discounted sums. Discounting alone, however, cannot reconcile the low estimate. The coefficients on the 15-year discounted sum derived from plain vanilla swaps were also less than one. Although not significantly so, we know that the prior has biased the estimate slightly upwards.

Another potential explanation is uncertainty. Uncertainty about the future may lead investors to place a lower weight on less certain, longer-horizon outcomes.

The weight of the currency risk premium in the relative yield curve (γ) is estimated to be a little higher for longer-horizon returns, but not significantly so.

D.7 Lag structure

Two variations were considered in terms of the lag structure of the model. First, innovations in Λ_t and R_t^f were modeled as AR(1) shocks rather than iid shocks. The average posterior estimates for (α, γ) were (0.93, 0.36). The estimated AR(1) coefficients were small, averaging 0.021 and 0.023 respectively, consistent with the idea that the innovations are news.

Figure 11: Posterior Densities for α: random walk model

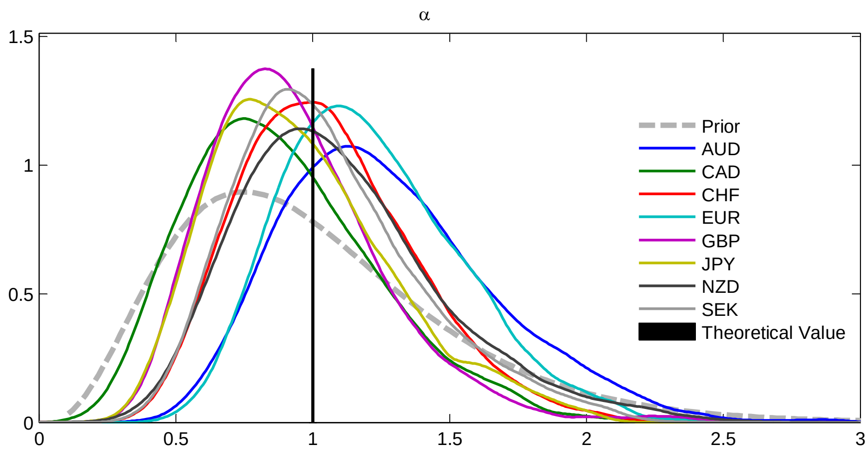

The grey dashed line is the prior density. The error-correction coefficients are set to zero ($\rho_R = \rho_\Lambda = 1$)

Second, the sum of relative risk-free returns, $\hat{R}_t^f$ and the level risk premium, Λ_t, were modeled as AR(1) processes. In that case, the model includes error correction terms. The posterior estimates for α were lower, averaging 0.76, while the estimates for γ were little a little higher averaging 0.41. Posterior distributions are shown in Figure 12.

In this partial equilibrium framework, assumptions about the lag structure are by nature, ad hoc. For example, it may make more sense, in terms of the stabilising role of monetary policy, for R_t to be modeled as adjusting to stabilise the economy and offset the effect of movements in risk premia. Only the gross interest rate, rather than the risk-free rate, is controllable by monetary policy. Implementing the framework in a general equilibrium model is a key area for further work.

D.8 Construction of level risk premium

The long-run relationship that defines the level risk premium was changed to include an estimated parameter, α_L where $q = -\alpha_L \mathbf{R} - \Lambda$. In that case, average posterior estimates for (α, γ) were little changed at (0.97, 0.36). The posterior mode for the long-run exchange rate response to expected returns averaged 0.77. Figure 14 shows the posterior distributions for α and α_L.

Finally, the same model was estimated with the restriction $alpha = \alpha_L$. In that case, average posterior estimates for $(\alpha = \alpha_L, \gamma)$ were (0.95, 0.36). Figure 14 shows the posterior distributions.

Figure 12: Posterior Densities for α: $\mathbf{R}_t$ is an AR(1) process)

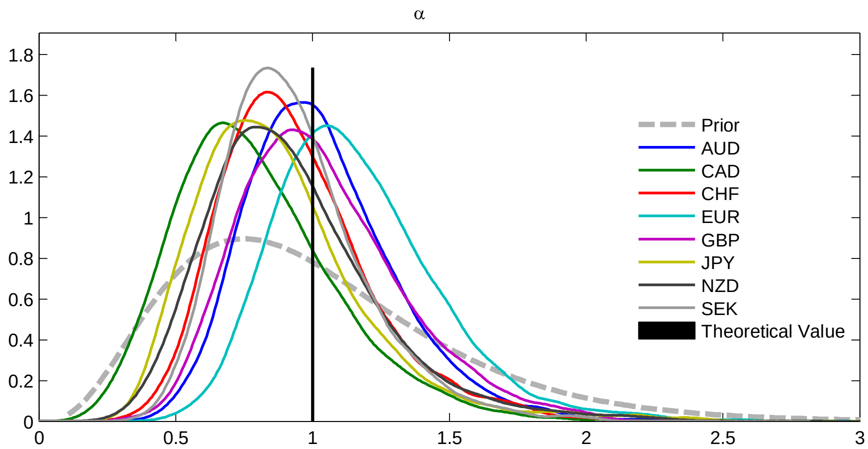

The grey dashed line is the prior density. The error-correction term $-(1-\rho_R)R^f_{t-1}$ is replaced with $-(1-\rho_R)\mathbf{R}_{t-1}$.

Figure 13: Posterior Densities for α: estimated long run relationship

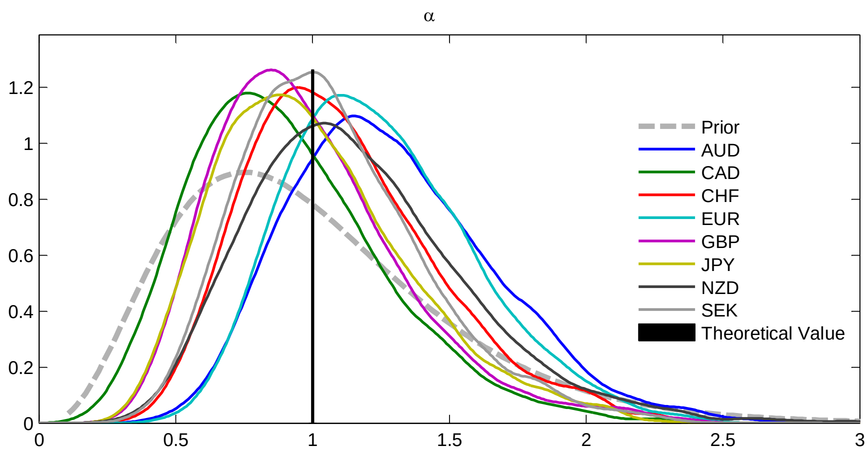

The grey dashed line is the prior density. The level premium, Λ_t is defined as $\Lambda = -(q_t + \alpha_L \mathbf{R}_t)$.

Figure 14: Posterior Densities for α: constrained long-run relationship

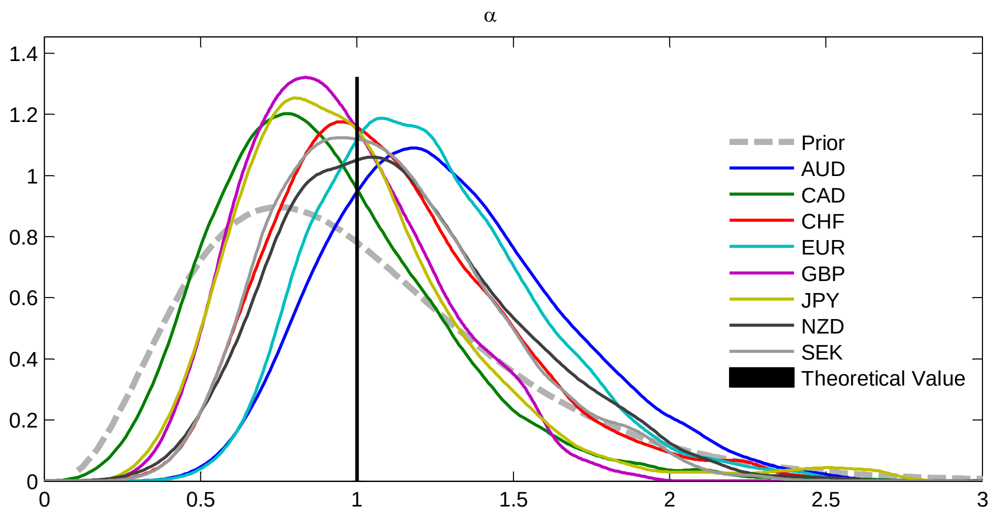

The grey dashed line is the prior density. The error-correction term $-(1-\rho_\Lambda)(q_{t-1}+\mathbf{R}_{t-1}$ is replaced with $-(1-\rho_\Lambda)(q_{t-1}+\alpha\mathbf{R}_{t-1}$.