Response to referee report #1

We thank the editor for the chance to respond to this report. We are encouraged to note that the referee did not find major shortcomings with either our methodology or our results, and we hope that the following will provide the clarifications he/she requested.

Comment on page 5

We apologize for not being more explicit on why the *F*-test is a necessary but not a sufficient test of the model. As the referee rightly points out, the *F*-test does test equation (2) in his/her report. However, while equation (1) which defines the current account in the model implies equation (2), the reverse is not true. The reason for this was already noted by Campbell and Shiller (1987, p. 1065), whom we quote:
«Second, while (1) implies (2), the reverse is not true. Equation (2) is consistent with a more general form of (1) that includes a "rational bubble", a random variable b_t satisfying $b_t = \delta E_t b_{t+1}$.»

Comment on page 9

The procedure we followed in the simulations was not exactly as described by the referee, proving his/her point that we were not as clear in our description as we should have been. Here's what we hope is a better description of our simulation procedure, noting that we could easily substitute this description into the paper if it is found to be more informative:

Step 1: we generate data for income changes ΔY and for a variable arbitrarily called z (see below for more details on this second variable) using a bivariate VAR. The coefficients used in this bivariate VAR come from a bivariate VAR in income changes and the current account estimated on our quarterly samples of actual data. As Campbell and Shiller point out, if the model is true then the current account contains all relevant information to forecast ΔY, which is why the literature assumes that the DGP for income changes consists of lagged values of itself and of the current account.

Note, and this is the key point, that we do not use z as our generated current account[1], even though z should look very similar to observed quarterly current account data given the procedure used in step 1. The reason we do not use z as our simulated current account is simple: we want our simulated current accound data to be 100% model-consistent, and there is no reason *a priori* that the simulated z should be. And the reason we want our simulated

[1] This is why we arbitrarily called it z. We could have called it something else, but the key is that z is not our simulated current account. The only reason we generate values of z is to be able to generate values of income changes following a DGP similar to what is assumed in the literature.

current account data to be fully model-consistent is to insure that the simulated current account and the model-predicted current account estimated from simulated data will track each other very closely. If they don't track each other, then there is a clear problem with the methodology to estimate model-predicted series (whereas if the simulated current account is not fully model-consistent, then the two series can diverge for no particular methodological reason).

Step 2: we generate our simulated current account CA_t by plugging our simulated data for ΔY and z together with the VAR coefficients used in step 1 into equation 2 in the paper. This ensures that the simulated current account equals the present discounted value of all expected income declines, where expectations are consistent with the DGP used to generate the income declines. In other words, our simulated current account is fully model-consistent.

Step 3: we estimate a bivariate VAR on the simulated data on income changes and current account from steps 1 and 2 and then plug these simulated data together with the just estimated VAR coefficients onto equation 2 in the paper, to obtain the model-predicted current account, CA_{pt}.

Step 4: we compute the correlation coefficient and variance ratio between CA_{pt} and CA_t. As previously mentioned, CA_{pt} and CA_t should track each other closely (and they will in large samples), implying that the correlation coefficient and variance ratio should be close to one. But the point we want to make in the paper is that estimation uncertainty is so large that the two series may not track each other at all in realistically short samples.
We also test whether the model is a good representation of the simulated current account CA_t using the F-test, as well as the linear and non-linear Wald tests. We the econometricians know that the model is a good (in fact, perfect) representation of CA_t, so the tests should reject the model with probability α, where 1- α is the confidence level of the test.

Step 5: repeat steps 1 – 4 ten thousand times.

Minor questions related to the simulations

- We did not subtract the mean in the calculation of the correlation coefficient and variance ratio (note that many papers in the literature don't either). This is not a problem in our simulations because both series should have similar means in the absence of the small sample issues we are trying to highlight. In other words, if the means are very different then this also proves our point.

- Yes, we follow the literature by calling CA_{pt} and not CA_t the optimal series, because in the literature it is not a given that the observed series is model-consistent. It just happens that we make them to be in our simulations.

- We can easily change the text to denote the rejection probability by α rather than 1 - α .

Comment on page 10

As the referee points out, it is true that correlation coefficients are meaningless if the underlying data is non-stationary. However, our simulation procedure guarantees that the data is stationary (because there is agreement that income *changes* are stationary, and because our simulated current account is by construction the present discount value of all future income changes). Thus, what we said by "correlations are not being very demanding tests" makes sense in a context where correlations are guaranteed to have statistical meaning.

Comment on page 15

In the paper, we wanted to stress that correlation is a poor indicator of the model's fit, and our empirical distributions of correlation coefficients show this point. Our discussion on the relative value of $\hat{k}_1$ and $\hat{k}_2$ was only meant to rationalize why correlations will tend to cluster around -1 and 1 under near-singularity. That said, we have computed empirical values of $\frac{\hat{k}_2}{\hat{k}_1}$, as requested by the referee. The figure attached at the end of this note plots the value of the $\frac{\hat{k}_2}{\hat{k}_1}$ ratio obtained over 10,000 draws of the empirical distribution for each country. As can be easily seen, the graphs fully support the point in the paper that the ratio will tend to be large in absolute value under near-singularity (note that the graphs were censored at -1000 and +1000 for presentation purposes, but the ratio is sometimes higher than this in the data).

Comments on Belgian data

On the data plots, differences in annual and quarterly units arise for various reasons. First, quarterly data are typically not annualized. Second, we remove the mean from our series (as discussed in the data appendix). Third, time coverage differs for quarterly and annual data.

We do not think that quarterly Belgian data contain a trend. To start, economic theory suggests that current accounts cannot trend. A, say, positive trend in current account implies that the country will end up saving infinitely relative to its GDP. Such an outcome is not possible. Moreover, annual data for Belgium, which spans 30 more years than quarterly data, does not look like it contains a trend, suggesting that the apparent trend in the quarterly data is an artifice of the short period span. Because of this, and because our aim is to critique the literature which does not allow for trends for the reasons mentioned, it seems natural not to model a trend for quarterly Belgian data.

Comments on Swedish data

The referee is right to point out the apparent seasonal variation in Swedish data, which after checking comes from seasonality in Y and ΔY_t. We have reestimated the model using de-seasonalized data. The resulting "optimal" current account plotted below (indicated by CAHAT) is again very imprecisely estimated, with the confidence bands even wider than when using non seasonally-adjusted data. In other words, it is not seasonality that is driving our results. We can add a footnote pointing this out if the editor thinks it would be useful.

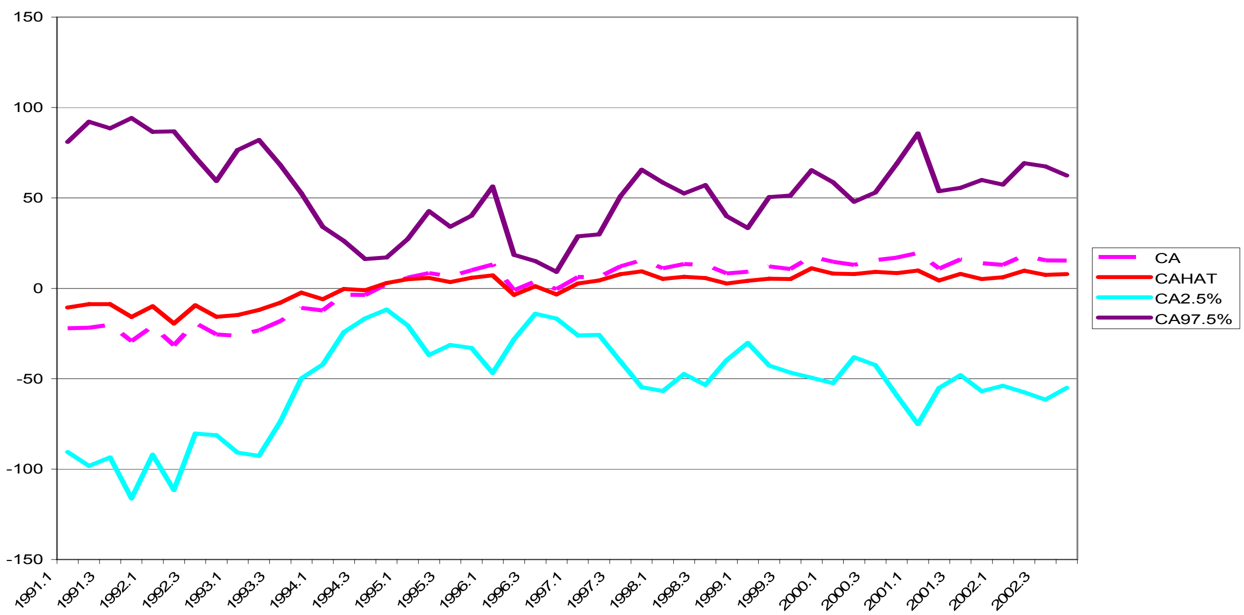

Minor comments

- We can certainly change notation and replace $C\hat{A}_t$ by $\overset{\wedge}{CA}_t$ if it helps the exposition.

- Note that in our paper T does not denote the number of observations in the sample, but denotes a vector of dimension 2**l* (where *l* is the number of lags in the estimated VAR) whose *l*+1 element is equal to 1 and all other elements are zero. Both *K* and *T* appear in the formula for the test statistic because, under the null that the model is a true representation of the data, *K*=*T*, and that's what the Wald test is testing. If the editor so desires, we can easily chose a less confusing notation for this vector.

- Similarly, our choice of "*x*" to define the number of simulated observations may be unconventional, and we can easily change the notation to *n* or *N* as suggested by the referee.

Empirical Values of the $\hat{K}_2 / \hat{K}_1$ Ratio