The Return of the Half-Life: A Note on the Persistence of the Real Exchange Rate

Julian di Giovanni*
University of California, Berkeley

March 30, 2003
All Comments Welcome

1 Introduction

A very interesting recent paper by Imbs, Mumtaz, Ravn and Rey (2002) argues that the estimated persistence of the real exchange rate estimated in previous research is biased upwards because past work has not taken into account the heterogeneity in the dynamics of disaggregated relative prices. This is a very interesting and insightful point. The purpose of this note is not to argue against this fact, which Imbs, Mumtaz, Ravn and Rey show to hold theoretically. Instead, I wish to make two points. First, I show that the method the authors use to solve the aggregation bias — the random coefficient model (RCM) — does not match the implicit assumption that they are making about the data generating process underlying the error terms. Second, a simple extension of their data set changes the results considerably using the very same techniques that they apply.

The first point that I make is a statistical one. By controlling for potential aggregation bias, the authors actually force the estimated coefficient of the autocorrelation

*E-mail: jdigio@econ.berkeley.edu. I would like to thank Jean Imbs for kindly providing me with the price and real exchange rate data. I appreciate helpful discussions with Michael Jansson, Maury Obstfeld, Tiago Ribeiro and Andy Rose. I also appreciate feedback from Morten Ravn and Haroon Mumtaz. All remaining errors are my own. Please do not quote.

function (ACF) to be smaller than what one might expect in a standard autoregressive model. In particular, by using the RCM in a dynamic setting, the authors are implicitly assuming that the shocks at the sectoral level (and hence in the aggregated estimation) are correlated across time, even as $T \to \infty$. This is not necessarily a bad assumption, but there was no reference to this fact in the paper. The question to ask then is: What does the noise of the process look like? Does the autocorrelation function asymptote to zero or not (as is assumed in the standard time series models used to examine real exchange rate persistence)? This question can be answered by plotting the autocorrelation function at the sector level. If this function actually asymptotes to zero, as is assumed in the literature (and seems to be assumed in the paper), then the use of a random coefficients model (RCM) by Imbs, Mumtaz, Ravn and Rey is incorrect in its present form. By estimating a RCM, the underlying ACF of the errors does not go to zero by assumption. I shall show why this is so analytically using a very simple example. Specifically, I shall compare a simple AR(1) process to an AR(1) process with an additional (time-invariant) random error. This model can be thought as characterizing the price dynamics at the sector level, where the additional random term would capture the error term from the assumption of a random coefficient. I then present some sample autocorrelation functions that do indeed show that the implicit assumption made by the authors is incorrect. Finally, I suggest a simple estimation that should help control for this potential bias and show that this does not in fact affect the estimated slope coefficients.

The second point I make questions the robustness of the results. The data used by Imbs, Mumtaz, Ravn and Rey were missing real exchange rate data for six countries for pre-1991 or pre-1993 because of missing nominal exchange rate data. Using the original data, the authors find that the RCM yields a half-life point estimate of 14 months with a confidence interval ranging from 5 to 24 months. I simply update the nominal exchange rate data and create the additional real exchange rate data for earlier years. I then run the same RCM model and find a half-life point estimate of 28 months with a 95% confidence interval ranging from 25 to 32 months. This small and simple extension of the data doubles the point estimate and produces a lower confidence interval larger than the original upper confidence interval estimated Imbs, Mumtaz, Ravn and Rey. The RCM model still yields a half-life that is smaller than convention, but this lack of robustness to the additon of easily available data is a concern that the authors should address.

Section 2 presents a simple autoregressive model. Section 3 then extends this simple model to include a random effect error term, and I show the bias that this may create. Section 4 presents some sample autocorrelation functions and discusses their implications. Section 5 presents estimations based on the extended data set, and presents various robustness checks performed by cutting the data at different periods and sectors. Section 6 concludes.

2 Simple Model

First, begin by noting the panel system that Imbs, Mumtaz, Ravn and Rey estimate:

$$q_{i,j,t} = \alpha_{i,j} + \sum_{p=1}^{P} \gamma_{i,j,p} q_{i,j,t-p} + \varepsilon_{i,j,t}, \tag{1}$$

where i indexes the sector, j indexes the country (the US being the base country), p is a given lag, $\gamma_{i,j,p} = \gamma_p + \nu^1_{i,j}$, $\alpha_{i,j} = \alpha + \nu^2_{i,j}$, and $\varepsilon_{i,j,t}$ is white noise. This is simply a RCM, but also characterizes the price dynamics at the sector level independently; i.e., one must believe that this structure exists at the sector level regardless of whether the data are pooled. If not, the pooled regression is not correct.

Now, consider a simple model of the real exchange rate for any given sector and country pair where it follows the simple AR(1) process:

$$q_{i,j,t} = a + b q_{i,j,t-1} + e_{i,j,t}, \tag{2}$$

where $\mathrm{E}(e_{i,j,t}) = 0$, $\mathrm{E}\big(e_{i,j,t}^2\big) = \sigma_e^2$, and $\mathrm{E}(e_{i,j,t}e_{i,j,s}) = 0$ if $t \neq s$.

The autocorrelation function is defined as

$$\rho(\tau) = \frac{E\left[(q_{i,j,t} - \mu)\left(q_{i,j,t-\tau} - \mu\right)\right]}{E\left[\left(q_{i,j,t} - \mu\right)^2\right]}, \tag{3}$$

where the expectation is not conditional on any information, $\mu = \mathrm{E}\left(q_{i,j,t}\right) = \frac{1}{1-b}$, and $\rho(0) = 1$. It is crucial to remember that the expectation is not conditional on any information (besides knowledge of the distribution underlying the shocks) when turning to the random effects model in Section 3.

One can then easily show that the autocorrelation function $\rho(\cdot)$ is:

$$\rho(\tau) = b^\tau, \quad \text{for } \tau = 0, 1, 2, \ldots \tag{4}$$

Therefore, $\hat{b} = \hat{\rho}(1)$, where $\hat{}$ denotes an estimated parameter. There are various methods that can estimate $\hat{b}$, and $\hat{\rho}(1)$ may be estimated from the sample ACF. However, what is crucial to remember is that the true $\rho(1)$ will be the same as the one in the RCM model described below because, by definition, it is independent of the two models. Therefore, the estimated coefficients of the sample ACF should be the same, regardless of what regression model is run on the real exchange rate panel.

3 Random Effects Model

Now consider adding a time-invariant random effect error to (2). This is similar to assuming a RCM of equation (1), but ignores the interaction of the lagged dependent variable and the coefficient's random error term. This specification drastically simplifies the algebra, and will still give the general intuition behind what is happening in the RCM specification.The AR(1) process can therefore be written as:

$$q_{i,j,t} = a + bq_{i,j,t-1} + e_{i,j,t} + u_{i,j}, \tag{5}$$

where $\mathrm{E}(e_{i,j,t}) = 0$, $\mathrm{E}\big(e_{i,j,t}^2\big) = \sigma_e^2$, $\mathrm{E}(e_{i,j,t}e_{i,j,s}) = 0$ if $t \neq s$, $\mathrm{E}(u_{i,j}) = 0$, $\mathrm{E}\big(u_{i,j}^2\big) = \sigma_u^2$, and $\mathrm{E}(e_{i,j,t}u_{i,j}) = 0 \ \forall\ t$. One can then easily show that the autocorrelation function

$\rho(\cdot)$ is:

$$\rho(\tau) = s + (1 - s)b^{\tau}, \tag{6}$$

where $\rho(0) = 1$ and

$$s = \frac{\frac{\sigma_u^2}{(1-b)^2}}{\frac{\sigma_u^2}{(1-b)^2} + \frac{\sigma_e^2}{1-b^2}}.$$

Therefore, $\tilde{b} = \frac{\tilde{\rho}(1)-\tilde{s}}{1-\tilde{s}}$, where $\tilde{}$ denotes an estimated parameter. Please see Appendix A for a derivation of this result as well as an explanation for why σ_u^2 does not drop out of the derivation.

Next, consider the true population parameters. Since s is positive and ρ does not vary across the two models, as discussed above, one arrives at:

$$\tilde{b} = \frac{\tilde{\rho}(1) - \tilde{s}}{1 - \tilde{s}} = \frac{\hat{\rho}(1) - \tilde{s}}{1 - \tilde{s}} < \hat{\rho}(1) = \hat{b}, \tag{7}$$

since the estimated ρ will be less than one as long as the process is stationary.

Equation (7) formalizes the proposition stated in the introduction of this note. In particular, one should expect that the estimated half-life is smaller in the RCM model than a standard autoregressive model because of the extra term (s) that exists in the ACF underlying the pricing dynamics specified in (5). The prices will still revert to the mean once shocked, but will do so faster because of this additional terms. This is turn leads the estimated b to be smaller. In essence, the inclusion of a time-invariant error term effectively "soaks up" some of the autoregressive process that would otherwise be captures by the lagged $q_{i,j,t}$ terms. There is nothing wrong with this if there is in fact some constant autocorrelation in the underlying error structure, but is this really the case? The next section attempts to answer this question graphically, as well as arguing that a "simple fix" may help address the bias and measure to what extent this problem exists in the data.

4 Autocorrelation Functions

Figures 1-4 present sample autocorrelation and partial-autocorrelation functions (ACFs and PACFs) for two sectors — bread (a tradeable sector) and rents (a non-tradeable sector) — for thirteen countries. In examining, the ACF and PACF for bread in Figures 1 and 2 one can see that the (partial)autocorrelations tend towards zero as T gets large. The one outlier appear to be Ireland, but this can explained by its relatively short time-series. Standard tests cannot reject that the autocorrelation tends to zero given a sufficient amount of time for the autocorrelation to dampen out.[1] A similar picture is seen when examining Figures 3 and 4 for rent. These results do not vary across other sectors. Therefore, it would seem that one should assume zero correlation of the underlying error structure. The RCM model estimated in Imbs, Mumtaz, Ravn and Rey implicitly assumes a correlation between the coefficient errors greater than zero over time, which given the graphical evidence just presented in this paper, seems to point towards the estimated RCM coefficients being biased downwards.

How much this bias matters is a question that must still be answered, and I shall not attempt to do so in great rigour in this this note. However, one strategy to do so would be estimate a RCM model where one allows the coefficients to also correlate across time (e.g., see Nicholls and Quinn 1982): $b_{i,j,t} = b + \xi_{i,j,t}$, where $\xi_{i,j,t}$ can be modelled as correlated over time. This will allow more flexibility in the estimation procedure, as well as allowing for the calibration of the model to see how much correlation across errors over time must exist to yield the results of Imbs, Mumtaz, Ravn and Rey; i.e., to match the bias created by the correlation of $u_{i,j}$ over time.

I shall forego this exercise, but there is another potentially easier way to address the issue of bias. In particular, the time-invariant error term would not be a problem

[1]In particular $\Theta \equiv \sqrt{\frac{n(n+2)}{n-k}}\rho(k) \longrightarrow N(0,1)$, where $\rho(k)$ is the k^{th} autocorrelation, and n is the total number of observations. See Ljung and Box (1978) and references within for more details on this derivation. By examining fairly large k's for the different sample autocorrelation functions, I cannot reject that the statistic Θ is equal to zero.

in the RCM estimation in practice if one allows a separate fixed intercept term, $a_{i,j}$, for each country-sector. These terms would soak up most of the fixed correlation from the $u_{i,j}$ terms. However, the RCM estimation procedure used in the paper also specifies that the intercept terms are random; hence, this does not help alleviate the bias. Therefore, I also demean the data and estimate the RCM (without an intercept) as well as examining the data at the country-sector level in Section 5.3 below to get a sense of whether the bias is important or not. Before doing so, however, I turn to analyzing the robustness of the RCM result using an augmented data set.

5 Additional Data

5.1 New Estimates

The Imbs, Mumtaz, Ravn and Rey sample only has nominal exchange rate data dating back to 1991 or 1993 for six countries: Denmark, Finland, Greece, Ireland, Portugal and the Netherlands. These missing observations cut more than ten years of real exchange rate data from the sample per sector (for 10 years and 19 sectors, this would amount to a loss of a potential of $6 \times 10 \times 12 \times 19 = 13680$ observations). I have updated these data and replicated three of the regressions.[2] As can be seen in Table 1, this makes quite a difference for the RCM estimation.

The point estimate for the RCM doubles and now looks more like the fixed effect results, and is closer to matching the Rogoff "consensus view" of three to five years. Furthermore, the lower confidence interval is now larger than the upper confidence interval estimated using the original sample. This result is puzzling given that the OLS and fixed effects estimates do not change greatly given the additional data. I investigate this further in the next section by estimating the the models over different

[2] I also re-estimated the three models using the original data set. I could not exactly replicate the results in the paper, but I was told this was because the results in the paper actually include an additional sector from Norway in the final estimation. However, Haroon Mumtaz told me that my results did match theirs when Norway was excluded, and I ignore this sector in what follows.

Table 1
Half-Life Estimations using updated data

Model	P	$\sum_{j=1}^{P} \gamma_j$	Half-Life	95% C.I.
RCM	5	0.9746	28	(25, 32)
OLS	12	0.9997	2809	(1642, 9766)
Fixed Effects	12	0.9762	33	(29, 36)

Total observations: RCM: 36513, OLS and FE: 34959. Optimal P chosen by same criteria as in Imbs et al. (2002). Confidence intervals are generated from half-life distributions generated from bootstrapping methods for 1500 replications.

sample periods.

5.2 Robustness Check I: Sample Periods

The authors place a good deal of emphasis on how using monthly data is advantageous because the time series will be longer. This is not completely true if one is interested in examining whether the half-life is actually three to five years long. The frequency of the data (monthly, quarterly or annual) should not affect the estimated half-life (in a common metric). Therefore, one would ideally like to have many years of data, regardless of the frequency of these data.[3]

I also estimate the models for quarterly and annual data. The data series seems to be a bit short to say much meaningful about the annual data (the impulse response function fluctuated quite a bit). However, I found the following half-life point estimates for quarterly averaged data: RCM: 10 quarters (or 30 months), OLS: 580 quarters (or 1740 months), and FE: 12 quarters (or 36 months); see Table 2 in Appendix B for more details. These point estimates match the monthly ones in Table 1 quite well which is a nice robustness check, and points to the potential problem of

[3]Of course, the additional information gained from a monthly frequency should yield more precise estimates.

missing ten years of data in the original sample used by Imbs, Mumtaz, Ravn and Rey.

I also estimate the three models at various time intervals. The tables of the estimates are presented in Appendix B, so I shall only briefly discuss the results here. First, I restrict the beginning of the sample for all countries to both 1991 and 1993 (see Table 3). Reducing the sample to 1991-96 lowers the half-lives considerably for both the RCM and FE model, while the OLS estimate implies a non-stationary relative price process.[4] The estimate half-lives decrease even further when the sample is restricted to 1993-96. These results are worrisome given that half the countries used in the original sample have exchange rate data beginning in either 1991 or 1993. Therefore, it would seem that the actual number of years used in the sample matters quite a bit.

To further check this point, I estimate the three models at five and ten year intervals (see Table 4). I begin these sample spans in 1980 as to maximize the number of observations in the five and ten year periods. Again, the OLS estimates do not change greatly over time periods. However, both the RCM and FE estimates have some similar and interesting patterns. First, the 1980-84 five year period has a relatively large half-life compared to the other five year periods for both models. Next, the half-lives estimated for ten years are on average larger than those estimated for five years for the RCM and FE model. This again points to the number of years used being important in these panel regressions. Further research examining the differing behaviour of the real exchange rate in the 1980s and the 1990s might be a fruitful avenue to follow.

[4] I did not compute confidence intervals for these and the rest of the results in this paper given that (a) examining the point estimates suffices for preliminary robustness checks, and (b) the estimates will most probably be more imprecise (i.e., larger confidence intervals) given smaller sample sizes.

5.3 Robustness Check II: Sector Level Estimates

Given my concerns about using the RCM model, I thought it would be interesting to also estimate the individual sector half-lives by estimating an OLS individually for each sector.[5] Doing so (see Tables 5 and 6) yields an arithmetic average for the half-life of 47 months. This number is quite a bit larger than the one estimated by the RCM model in Table 1. However, if one runs a five-lag regression for each sector and then take the average of each coefficient, the estimated half-life is 21 months.[6] This average is similar to running a panel regression where one assumes that the differences in sector coefficients difference are "fixed". The difference in the half-lives estimated from the two approaches is due to Jensen's Inequality. Further work should be done in thinking about this type of estimation (e.g., calculating reliable confidence intervals, which will not be trivial given the small sample sizes for each country-sector time series).[7] Furthermore, if estimation is done at the individual country-sector level, one might also consider using other estimation techniques, such as threshold autoregressive (TAR) models.

Another strategy to check and see how important the bias may be is to run the RCM on demeaned data without an intercept term. In other words, one can define the variable $z_{i,j,t} = q_{i,j,t} - (1/T)\sum_{\tau=1}^{T} q_{i,j,\tau}$, and then run the regression:

$$z_{i,j,t} = \sum_{p=1}^{P} \lambda_{i,j,p} z_{i,j,t-p} + \omega_{i,j,t}, \tag{8}$$

where the $\lambda_{i,j,p}$'s are the random coefficients and the country-sector intercepts have been incorporated as fixed effects by demeaning the real exchange rate. Estimating (8) for five lags actually yields a half-life of 25 months using the expanded data set. This estimate is actually smaller than that estimating using the "plain vanilla" RCM

[5] Imbs, Mumtaz, Ravn and Rey (2003) calculate these using a threshold autoregressive model.

[6] Results do not vary greatly for one- or two-lag regressions.

[7] Using the original data (less the one sector for Norway), the arithmetic average is 30 months, while running separate five-lag regressions and the averaging the coefficients yields an estimated half-life of 9 months.

technique. This evidence, combined with the average half-life estimate of 21 months, points to the bias not being important in the extended sample.

5.4 Tradability

I run AR(1) regressions for traded and non-traded sectors using the extended data. As in Imbs, Mumtaz, Ravn and Rey, I compare the fixed effects and RCM estimations. As can be seen in Table 7 in Appendix B, using the additional data actually reverse the findings of Imbs, Mumtaz, Ravn and Rey. Now, Engel's measure of variability is actually smaller for the fixed effect estimation then for the RCM model (though barely so). One may interpret this as aggregation bias being more severe for non-traded goods (an opposite result to the paper). However, the concern of aggregation bias may not be that important.

5.5 Robustness Check III: Extended Lag Structure

One can recover the shorter RCM half-life by extending the lag length to 36 months. This yield a half-life of 17 months. However, as can be seen in Table 8, the upper-confidence interval now rises to 69 months for the RCM estimation, implying quite imprecise estimation. It is interesting to note, however, that the fixed effects regression yields a similar half-life as the regression in Table 1, and that the confidence intervals remain quite tight. The fixed-effect/RCM model estimated by (8) yields a half-life of 26 months for 36 lags, while the half-life resulting from estimating each sector-country OLS separately and then averaging the coefficients (the group average) is 15 months. I did not calculate the confidence intervals for these estimated half-lives, but given the results from the RCM estimation in Table 8, one should expect that they will be quite wide given the loss of information created by including three years of lags. It is also interesting to note that the fixed-effect/RCM estimated half-life of 26 months is larger than the "plain vanilla" RCM estimate of 17 months, while the group average is 15 months. On the one hand, the result of 26 months points to a positive bias existing in the RCM estimation procedure, while the 15 month estimate points to the opposite holding.

I also ran these three models for 12 and 24 month lags as additional robustness checks.[8] The RCM, fixed-effect/RCM, and group average estimated half-lives for 12 and 24 months are [27,25,21] and [26,26,18], respectively. Again, the RCM estimates of 27 and 26 months are higher than those found in the paper. However, there does not seem to be any bias created by the RCM estimation, and the group average estimates of 21 and 18 months are smaller than the RCM estimated half-lives. Further work on determining the best model (as well as calculating appropriate confidence intervals) is needed before making any definitive claims, but it would appear that the downward bias resulting from the RCM model is not that important.

6 Conclusion

This note examines the interesting recent contribution on the estimation of real exchange rate persistence by Imbs, Mumtaz, Ravn and Rey. The authors points to the existence of an upward bias in half-lives estimates when ignoring aggregation of sector prices. The authors correct for this bias by estimating a random coefficient model across sectors and find a much smaller average half-life.

This note makes two points. First, the RCM estimation procedure imposes certain assumptions about the structure of the error processes underlying the data that do not seem to conform to the actual structure, as seen in the sample autocorrelation functions. This effect does not appear to be quantitatively important when examining alternative estimation procedure (i.e., fixed-effect/RCM or group-average estimates). However, more work might be done to refine these estimates and the potential problem

[8]The AIC points towards an optimal lag structure of about 21 months for the RCM model, but the estimated half-life did not vary greatly compared to the 24 month specification that is being presented.

of the bias should at least be addressed.

Second, a simple extension of the data set appears to show that the RCM result in the original paper is not robust. The estimated half-life using the RCM is quite larger when including the extra data, and though one can rescue the result by including 36 lags, the upper confidence interval becomes much larger. However, group-average estimates of the half-life appear to be similar to the main results in Imbs, Mumtaz, Ravn and Rey. Further work should be done to address why the two estimation techniques yield such differing results, and a deeper examination of the data appears to be in order.[9] Finally, it is interesting to note that the behaviour of the half-life varies between the 1980s and the 1990s. More work addressing this issue might also be interesting.

References

Imbs, Jean, Haroon Mumtaz, Morten O. Ravn, and Hélène Rey, "PPP Strikes Back: Aggregation and the Real Exchange Rate," December 2002. NBER Working Paper No. 9372.

——, ——, ——, and ——, "Non-Linearities and Real Exchange Rate Dynamics," 2003. Forthcoming in *Journal of the European Economic Association.*

Ljung, G. M. and G. E. P. Box, "On a Measure of Lack of Fit in Time Series Models," *Biometrika*, August 1978, *65* (2), 297–303.

Nicholls, Des F. and Barry G. Quinn, *Random Coefficients Autogregressive Models: An Introduction*, Springer-Verlag, 1982.

[9] Plots of the price data also reveal some peculiar patterns.

A Derivation of Random Effects ACF

The derivation of the random effects autocorrelation function, equation (6) in Section 3 is a follows. First, solve for the expectation of $q_{i,j,t}$, μ. Begin by re-writing (5) as:

$$(1 - bL)q_{i,j,t} = a + e_{i,j,t} + u_{i,j}, \tag{9}$$

where L is the lag operator. Dividing both sides of (9) by $(1 - bL)$, one can then solve for $q_{i,j,t}$ as

$$q_{i,j,t} = \frac{a + u_{i,j}}{1 - b} + \sum_{s=0}^{\infty} b^s e_{i,j,t-s}. \tag{10}$$

Next, taking the expectation of (10), one can solve for μ

$$\begin{aligned}
\mu &= \mathrm{E}\left(q_{i,j,t}\right) \\
&= \mathrm{E}\left(\frac{a + u_{i,j}}{1 - b} + \sum_{s=0}^{\infty} b^s e_{i,j,t-s}\right) \\
&= \frac{a}{1 - b},
\end{aligned} \tag{11}$$

since $u_{i,j}$ and $e_{i,j,t}$ have mean zero and are uncorrelated for all t.

Next, subtract (11) from (10), square and take expectations, to begin solving solve for the autocorrelation function:

$$\begin{aligned}
\mathrm{E}\left[(q_{i,j,t} - \mu)^2\right] &= \mathrm{E}\left[\left(\frac{u_{i,j}}{1 - b} + \sum_{s=0}^{\infty} b^s e_{i,j,t-s}\right)^2\right] \\
&= \frac{1}{(1 - b)^2}\mathrm{E}\left(u_{i,j}^2\right) + \frac{1}{(1 - b^2)}\mathrm{E}\left(e_{i,j,t}^2\right) \\
&= \frac{\sigma_u^2}{(1 - b)^2} + \frac{\sigma_e^2}{1 - b^2},
\end{aligned} \tag{12}$$

where the second line follows from the assumptions made that $u_{i,j}$ and $e_{i,j,t}$ are uncorrelated for all t, and that $\mathrm{E}\left(e_{i,j,t}e_{i,j,s}\right] = 0$ for all $t \neq s$. Equation (12) gives the denominator of the autocorrelation functions. Next, solve for the numerator:

$$
\begin{aligned}
\mathrm{E}\left[\left(q_{i,j,t}-\mu\right)\left(q_{i,j,t-\tau}-\mu\right)\right] &= \mathrm{E}\left[\left(\frac{u_{i,j}}{1-b}+\sum_{s=0}^{\infty} b^{s} e_{i,j,t-s}\right)\left(\frac{u_{i,j}}{1-b}+\sum_{s=0}^{\infty} b^{s} e_{i,j,t-\tau-s}\right)\right] \\
&= \frac{1}{(1-b)^{2}} \mathrm{E}\left(u_{i,j}^{2}\right)+\frac{b^{\tau}}{1-b^{2}} \mathrm{E}\left(e_{i,j,t}^{2}\right) \\
&= \frac{\sigma_{u}^{2}}{(1-b)^{2}}+\frac{b^{\tau}}{1-b^{2}} \sigma_{e}^{2},
\end{aligned}
\tag{13}
$$

where the second line again follows from using the properties of the error terms. Finally divide (13) by (12) to yield the autocorrelation function (6). Q.E.D.

It is not obvious why the country-sector specific error term, $u_{i,j}$, is not considered fix in calculating the autocorrelation function. In particular, one might argue that in picking a country-sector to analyze that the country-sector shock has been realized, and the $u_{i,j}$ term should therefore be considered as fixed in calculating the ACF. However, to show why this is not the case, consider the logic of the following experiment that is run in calculating any given country-sector ACF.

Suppose that there are only two country-sector groups. Each country-sector group is hit by a shock $u_{i,j}$ that has the same distribution for each group. Now, flip a coin to decide which country-sector we will calculate the ACF for. Upon realization of the flip of the coin, we know the country-sector of analysis. However, the crucial point to remember is that we only know information about the *distribution* about the shock $u_{i,j}$ that is hitting the country-sector, not what the shock actually is. Therefore, the shock must still be treated as random when calculating the ACF.

B Additional Tables

Table 2

Half-Life Estimations using updated quarterly data

Model	P	$\sum_{j=1}^{P} \gamma_j$	Half-Life	95% C.I.
RCM	2	0.9279	10	(9,11)
			30	(27,33)
OLS	12	0.9988	580	(404,1010)
			1740	(1212,3030)
Fixed Effects	6	0.9370	12	(11,13)
			36	(33,39)

For each model the first line gives the half-life in quarters and the second line in months. Optimal P chosen by same criteria as in Imbs et al. (2002). Confidence intervals are generated from half-life distributions generated from bootstrapping methods for 1500 replications.

Table 3

Estimation for 1991-1996
and 1993-1996 year periods

Model	Years	P	$\sum_{j=1}^{P} \gamma_j$
RCM	1991-96	6	0.9053
	1993-96	11	0.8073
OLS	1991-96	12	1.0001
	1993-96	3	1.0003
Fixed Effects	1991-96	8	0.9058
	1993-96	12	0.7993

Total observations: 1991-96: 12232, 1993-96: 6940. Optimal P chosen by same criteria as in Imbs et al. (2002).

Table 4

Estimation for 5 and 10 year periods

		RCM	
Years	P	$\sum_{j=1}^{P} \gamma_j$	Observations
1980-84	12	0.9754	7407
1985-89	3	0.9381	12315
1990-94	7	0.8927	13182
1980-89	4	0.9514	20871
1985-94	5	0.9518	25411
		OLS	
Years	P	$\sum_{j=1}^{P} \gamma_j$	Observations
1980-84	12	0.9992	7407
1985-89	3	0.9994	12315
1990-94	6	0.9999	13182
1980-89	2	0.9996	21254
1985-94	3	0.9995	25497
		Fixed Effects	
Years	P	$\sum_{j=1}^{P} \gamma_j$	Observations
1980-84	12	0.9709	7407
1985-89	4	0.9401	12272
1990-94	2	0.9046	13182
1980-89	2	0.9674	21254
1985-94	8	0.9479	25282

Optimal P chosen by same criteria as in Imbs et al. (2002).

Table 5
Sector level half-life point estimates

Sector	Belgium	Denmark	Finland	France	Germany	Greece
Bread	22	59	21	31	43	27
	(2)	(2)	(2)	(2)	(2)	(2)
Meat	24	32	15	37	31	22
	(2)	(2)	(2)	(2)	(2)	(1)
Dairy	30	49	20	36	35	22
	(2)	(2)	(2)	(2)	(2)	(2)
Fruits	6	10	15	12	7	3
	(2)	(2)	(2)	(3)	(2)	(2)
Tobbaco	43	41	14	53	40	31
	(1)	(2)	(2)	(2)	(4)	(2)
Drinks	46	36	18	46	614	71
	(1)	(2)	(2)	(2)	(2)	(1)
Clothing	146	84	13	70	85	26
	(3)	(1)	(2)	(2)	(3)	(1)
Footwear	77	164	13	84	111	24
	(2)	(2)	(2)	(2)	(2)	(1)
Rents	38	77	13	61	69	278
	(2)	(2)	(2)	(2)	(2)	(2)
Fuel	15	49		52	44	21
	(1)	(2)		(1)	(1)	(1)
Furniture	98	192		224	141	53
	(2)	(2)		(2)	(2)	(1)
Dom. App	56	69		58	52	17
	(2)	(2)		(2)	(2)	(1)
Vehicles	17	59	6	25	35	6
	(2)	(2)	(2)	(2)	(2)	(2)
Pub. Trans	22	12	15	27	62	12
	(2)	(1)	(2)	(2)	(2)	(1)
Communic	31	88	19	19	41	36
	(2)	(2)	(2)	(2)	(2)	(1)
Sound	41	62	17	52	37	37
	(3)	(2)	(2)	(2)	(2)	(1)
Leisure	33	54	15	35	31	53
	(2)	(2)	(2)	(2)	(2)	(2)
Books	41	71	16	31	39	108
	(2)	(2)	(2)	(2)	(2)	(2)
Hotel	20	13	17	16	16	41
	(2)	(2)	(2)	(5)	(3)	(1)

OLS estimation, with optimal P in (parentheses) chosen by same criteria as in Imbs et al. (2002).

Table 6
Sector level half-life point estimates cont'd

Sector	Ireland	Italy	Netherlands	Portugal	Spain	UK
Bread	11	35	24	54	96	15
	(2)	(2)	(2)	(1)	(2)	(4)
Meat	8	29	29	18	43	10
	(2)	(2)	(2)	(2)	(2)	(2)
Dairy	6	47	34	3	45	21
	(2)	(2)	(2)	(3)	(1)	(2)
Fruits	5	18	9	9	20	9
	(3)	(3)	(3)	(2)	(1)	(2)
Tobbaco	6	17	32	18	24	19
	(2)	(1)	(2)	(1)	(1)	(2)
Drinks	10	53	34	150	124	30
	(2)	(2)	(2)	(2)	(1)	(2)
Clothing	5	60	12	224	117	12
	(3)	(3)	(5)	(3)	(3)	(2)
Footwear	7	65	5	113	104	15
	(2)	(2)	(1)	(2)	(2)	(2)
Rents	7	43	72		31	30
	(2)	(2)	(2)		(2)	(2)
Fuel	8	40	28	52	35	35
	(2)	(2)	(1)	(1)	(1)	(1)
Furniture		113	113	511	183	34
		(2)	(2)	(2)	(2)	(2)
Dom. App	6	57	53	139	82	19
	(2)	(2)	(2)	(2)	(2)	(2)
Vehicles	1	16	29	59	31	10
	(2)	(2)	(2)	(2)	(2)	(2)
Pub. Trans	9	39	29	22	58	16
	(2)	(2)	(2)	(1)	(1)	(2)
Communic	6	30	50	16	42	18
	(2)	(2)	(2)	(2)	(2)	(2)
Sound	13	59	35		57	24
	(2)	(2)	(2)		(3)	(4)
Leisure	9	48	30	21	62	18
	(2)	(2)	(2)	(1)	(2)	(4)
Books	5	51	28	73	54	20
	(2)	(4)	(2)	(2)	(2)	(4)
Hotel	7	16	14	37	23	15
	(2)	(3)	(2)	(2)	(3)	(5)

OLS estimation, with optimal P in (parentheses) chosen by same criteria as in Imbs et al. (2002).

Table 7
Engel's R

Fixed Effects		RCM	
p	R	p	R
6	0.5201	6	0.5279
12	0.5200	12	0.5280
24	0.5200	24	0.5280
60	0.5200	60	0.5280
100	0.5200	100	0.5280
180	0.5199	180	0.5280
240	0.5199	240	0.5280
$\hat{\rho}$	0.9787	$\hat{\rho}$	0.9768
$\hat{\mu}$	0.9768	$\hat{\mu}$	0.9778
$\hat{\sigma}^2$	0.0039	$\hat{\sigma}^2$	0.0043
$\hat{\tau}^2$	0.0039	$\hat{\tau}^2$	0.0037

The following AR(1) regressions are estimated. Traded goods: $x_{j,t} = \rho x_{j,t-1} + \varepsilon_t$, where $\mathrm{Var}(\varepsilon_t) = \sigma^2$ and all $\{\varepsilon_t\}$ are uncorrelated. Non-traded goods: $n_{j,t} = \mu n_{j,t-1} + e_t$, where $\mathrm{Var}(e_t) = \tau^2$ and all $\{e_t\}$ are uncorrelated. Finally, $\{\varepsilon_t\}$ and $\{e_t\}$ are orthogonal.

Table 8
Estimation for 36 Lags

Model	Half-Life	95% C.I.
RCM	17	(15, 69)
OLS	921	(1306, 2401)
FE	28	(21, 30)

Figure 1

Sample autocorrelation functions

Bread

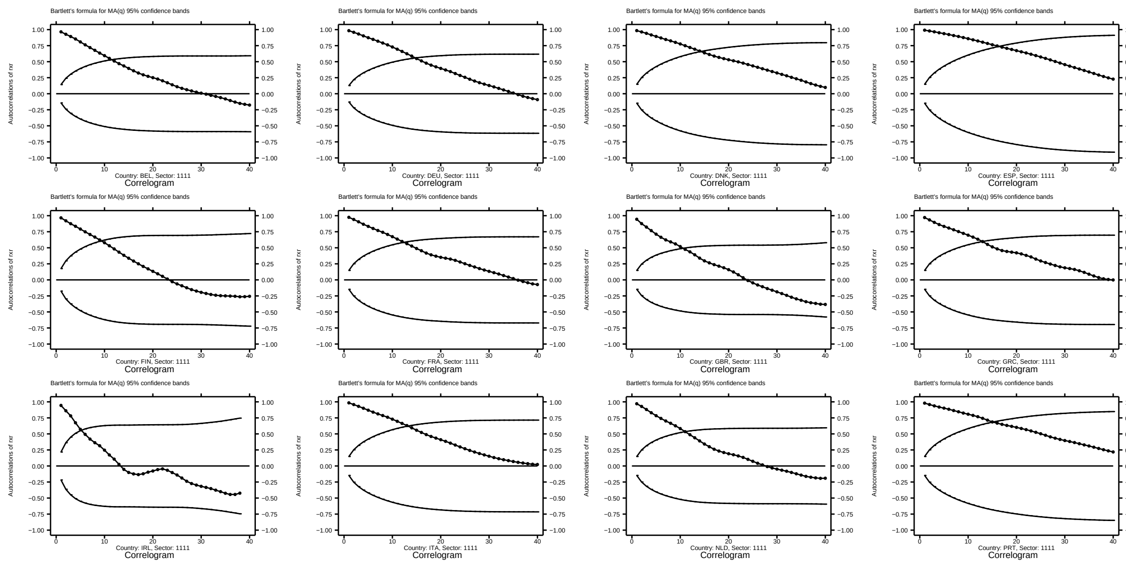

Figure 2

Sample partial-autocorrelation functions

Bread

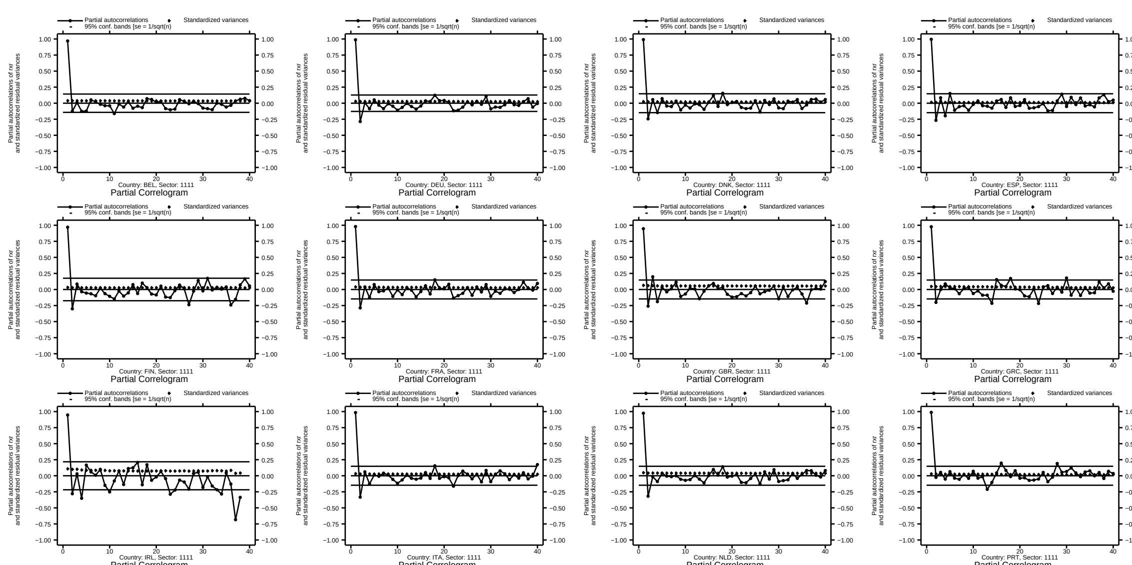

Figure 3

Sample autocorrelation functions

Rents

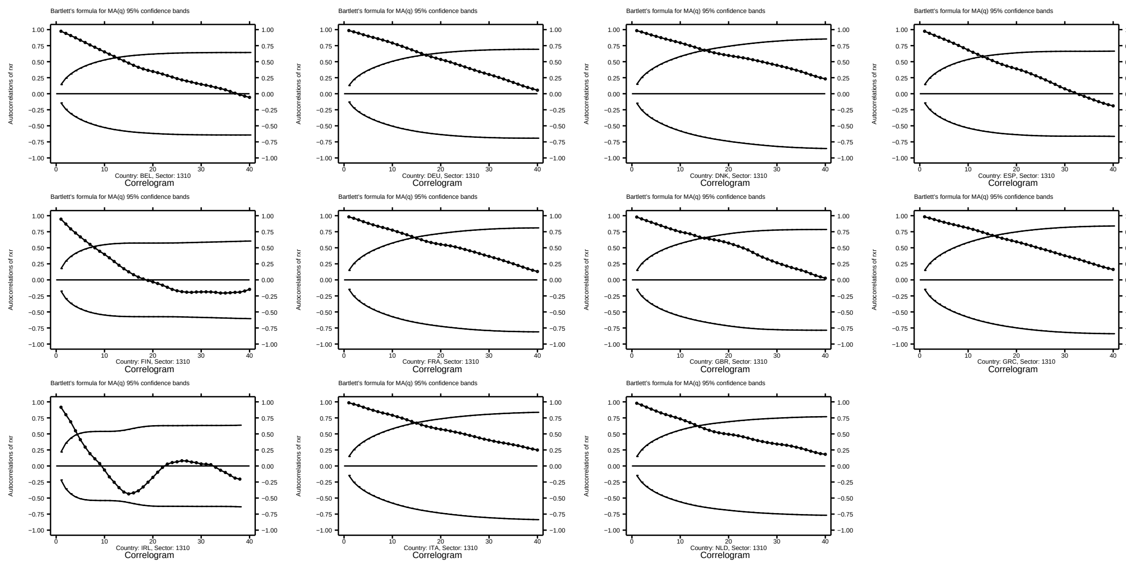

Figure 4

Sample partial-autocorrelation functions

Rents