Appendix to 'Technology shocks around the world'[☆]

Martial Dupaigne[a,b], Patrick Fève[a,c,d,e]

[a] Toulouse School of Economics (GREMAQ), 21 allée de Brienne, F-31000 Toulouse, France
[b] University of Western Brittany, Brest, France
[c] Institut d'Économie Industrielle, Toulouse, France
[d] Institut Universitaire de France
[e] Banque de France, Research Division

Abstract

This appendix presents several robustness experiments, carried on actual and simulated data.

Key words: Technology shocks, Labor input, Vector Autoregressions, Multi–country economy.

A. Robustness of empirical results

We now assess the robustness of the empirical findings presented in Section 1.3 of the main text. We consider larger SVAR models and a more exhaustive measure of labor input.

A.1. Larger SVAR models

In what follows, we systematically compare the results of larger VARs to the responses obtained in Section 1.3 with aggregate–level labor productivity.

First, we include in $\tilde{Z}_t$ an additional variable, the difference between the country–level labor productivity and the G7 aggregate labor productivity. Since a single stochastic trend hits permanently the seven country–level labor productivities, the labor productivity differentials $x_{i,t} - x_{\mathrm{G7},t}$ ($i = 1, ..., 7$) help to capture persistent country–specific components in labor productivities.

The results of this experiment are reported in Figure 8. For each country, the left panel displays the IRF of employment to a permanent productivity shock (thick grey line) in a VAR ($\Delta x_{\mathrm{G7}}, \Delta n_i, x_i - x_{\mathrm{G7}}$), as well as the estimated IRF of Figure 3 (from the main text) with its 95% confidence interval. The right panel displays two smoothed distributions of impact responses: the light grey one refers to VAR ($\Delta x_{\mathrm{G7}}, \Delta n_i$), as in Figure 3; the medium grey one refers to VAR ($\Delta x_{\mathrm{G7}}, \Delta n_i, x_i - x_{\mathrm{G7}}$).

The response of employment in US, Canada Japan and France in the three-variables VAR are very similar to those obtained with the two-variables VAR. In two continental European countries (Germany and Italy), the addition of productivity differentials in the VAR slightly shifts up the responses of labor input. In UK, it shifts them down. The discrepancy between the results of the two VAR models suggest that these three countries have a persistent country–specific component in labor productivity.

Second, we include in $\tilde{Z}_t$ the seven country–level labor inputs: $\tilde{Z}_t = (\Delta x_{\text{G7},t}, \Delta n_{1,t}, \ldots, \Delta n_{7,t})'$. This specification captures in a single SVAR model the joint dynamics of employment in the major seven countries. The estimated responses are displayed in Figure 9. For all countries, the results established in the two–variables VAR are preserved in this eight–variables VAR.

A.2. Hours worked data

As emphasized by Galí (2005), the choice between the different measures of labor input available may have some consequences on the identification of permanent technology shocks. We now use hours worked rather than employment to ensure that (potential) permanent shocks to hours per worker do not have long–run effects on our measure of labor productivity.

In the OECD countries, hours per worker are only available at annual frequency for relatively long homogenous time series. Hours worked are defined as the product of civilian employment and hours per employee, subsequently converted per capita. Labor productivity is here defined as real output per hour worked. We evaluate the response of hours worked at annual frequency over the sample 1972-2004.

We perform two estimations with this annual dataset. First, we compute the IRFs of hours worked using country–level data. Second, we replace country–level labor productivity by G7 aggregate labor productivity. Figure 10 reports the results of the two experiments.[1] The left panel displays the IRFs of hours worked in each country while the right panel reports empirical distributions.

We find the same pattern with hours worked data and with employment data: when we use G7 labor productivity instead of country–level ones, IRFs shift up and the empirical distribution of impact responses moves on the right for most countries. A sharp distinction between US, Canada and UK, on one side, and continental European countries plus Japan, on the other side, emerges. The first group exhibits large responsiveness of hours worked to technology shocks, whereas the second group features small adjustment.

When we use country–level labor productivity, an aggregation anomaly occurs with hours worked data as it does with employment data. The weighted sum of country–level point estimates is negative and lies outside the confidence interval around the G7 aggregate response. For instance, the country–level based impact response equals -0.368, while the impact response obtained at the aggregate level is 0.254. When G7 labor productivity is used instead of country–level labor productivities, the weighted sum of impact responses becomes 0.251.

---

[1] We report the 90% confidence interval, as compared to 95% interval for quarterly data, to take into account the lack of precision due to the short annual sample.

Figure 8: IRFs of employment to an aggregate–level permanent productivity shock in larger VARs – Productivity differentials
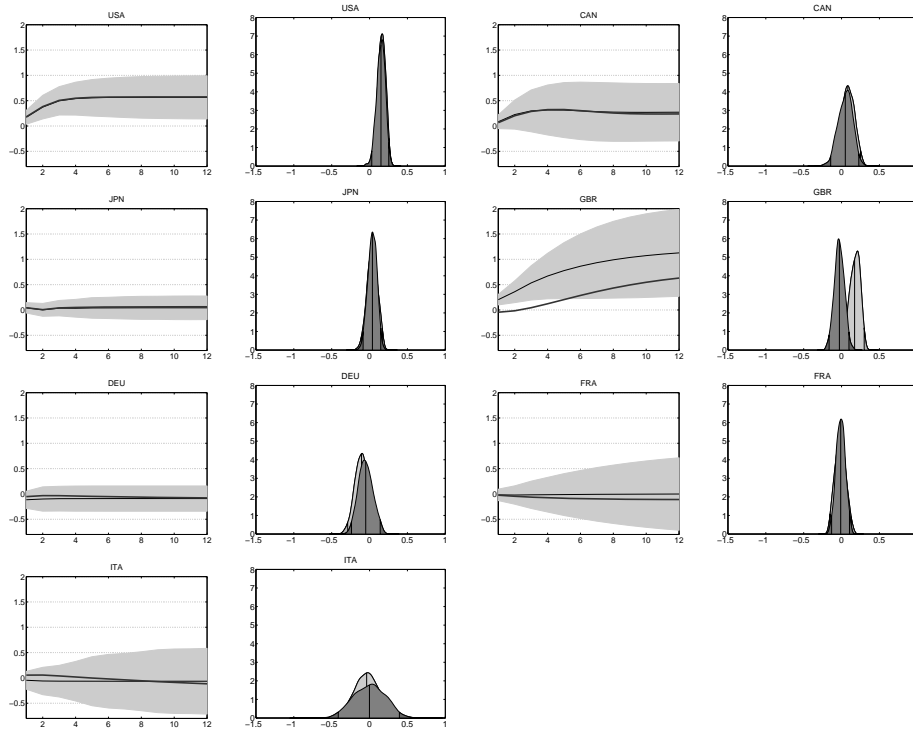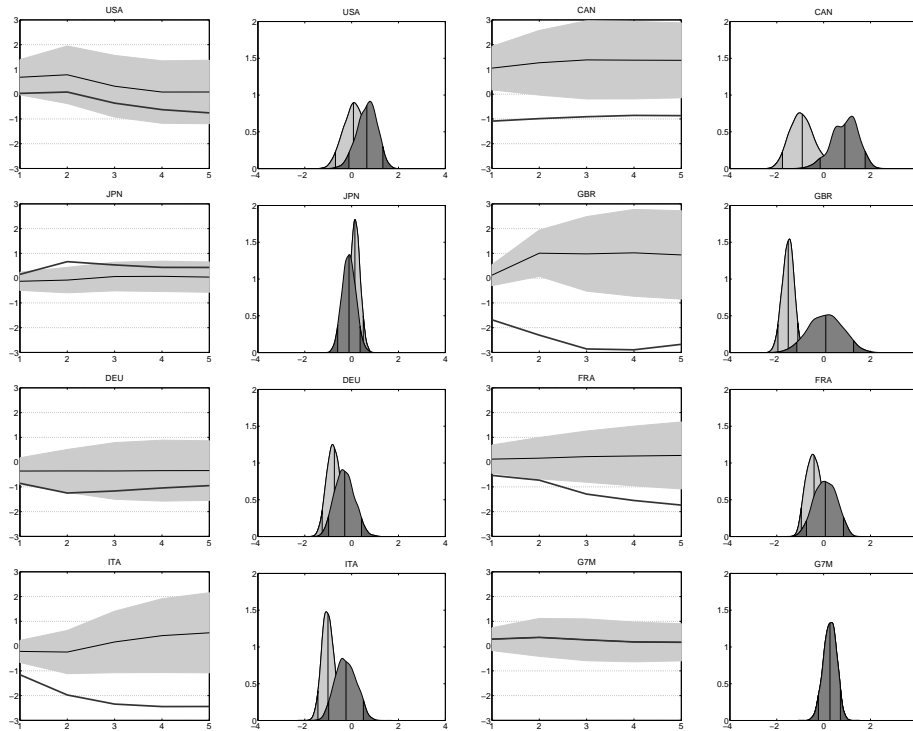

Note: For each country, the left panel displays the impulse response function of employment to a permanent productivity shock in a VAR $(\Delta x_{\mathrm{G7}}, \Delta n_i, x_i - x_{\mathrm{G7}})$ (black thick line), as well as the point estimate impulse response of Figure 3 and its 95% confidence interval. The right panel displays two smoothed distributions of impact responses obtained by standard bootstrap techniques, using 1000 draws from the sample residuals: the light grey one refers to VAR $(\Delta x_{\mathrm{G7}}, \Delta n_i)$, as in Figure 3; the medium grey one refers to VAR $(\Delta x_{\mathrm{G7}}, \Delta n_i, x_i - x_{\mathrm{G7}})$.

3

Figure 9: IRFs of employment to an aggregate–level permanent productivity shock in larger VARs – Seven countries' labor inputs
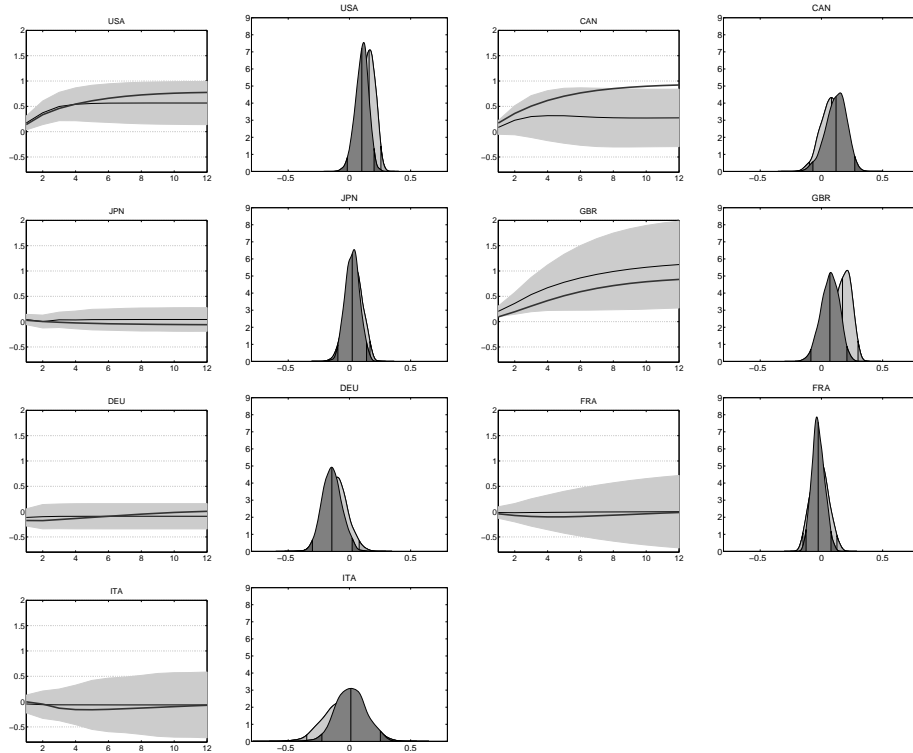
Note: For each country, the left panel displays the impulse response function of employment to a permanent productivity shock in a VAR $\left(\Delta x_{\mathrm{G7}}, \Delta n_i, \Delta n^{-i}\right)$ (black thick line), as well as the point estimate impulse response of Figure 3 and its 95% confidence interval. The right panel displays two smoothed distributions of impact responses obtained by standard bootstrap techniques, using 1000 draws from the sample residuals: the light grey one refers to VAR $\left(\Delta x_{\mathrm{G7}}, \Delta n_i\right)$, as in Figure 3; the medium grey one refers to VAR $\left(\Delta x_{\mathrm{G7}}, \Delta n_i, \Delta n^{-i}\right)$.

Figure 10: IRFs of hours worked to an aggregate–level permanent productivity shock on annual data

Note: For each country, the left panel displays the impulse response function of hours worked to a permanent productivity shock in a two-lags VAR $(\Delta x_{\mathrm{G7}}, \Delta n_i)$ and its 90% confidence interval, as well as the point estimate impulse response of the VAR $(\Delta x_i, \Delta n_i)$. The right panel displays two smoothed distributions of impact responses obtained by standard bootstrap techniques, using 1000 draws from the sample residuals: the light grey one refers to VAR $(\Delta x_i, \Delta n_i)$, as in Figure 1; the medium grey one refers to VAR $(\Delta x_{\mathrm{G7}}, \Delta n_i)$.

B. Estimation results with other datasets

We report in Table 7 ML estimations of the DSGE model with other datasets. We first replace the growth rate of US per capita output with the growth rate of US labor productivity for the same sample (1959:1-2003:4) and for a shorter one (1978:1-2003:4). Second we replace the growth rate of US labor productivity by the growth rate of G7 labor productivity (the latter data being available from 1978).

Table 7: Robustness of estimation results

| | US productivity growth | | G7 productivity growth |
Estimated parameter	1959:1-2003:4	1978:1-2003:4	1978:1-2003:4
φ	0.8511	1.0102	1.5781
	(0.2417)	(0.3432)	(0.1861)
σ_a	0.0176	0.0152	0.0105
	(0.0012)	(0.0015)	(0.0007)
ρ	0.9983	0.9975	0.9994
	(0.0122)	(0.0187)	(0.0022)
σ_ω	0.0031	0.0032	0.0049
	(0.0010)	(0.0014)	(0.0007)
ν	0.9229	0.9295	0.9668
	(0.0405)	(0.0446)	(0.0126)
σ_ϵ	0.0082	0.0074	0.0081
	(0.0006)	(0.0007)	(0.0006)

Note: The vector of observed data includes the growth rates of labor productivity, the growth rates of US hours worked and investment to output ratio. Standard–errors in parentheses.

C. Counterfactual experiments

Another way to understand how country–specific shocks contaminate the response to a permanent TFP improvement is to shut down the international dimensions of the model. We successively alter three margins: the relative volatilities of shocks, their relative persistence and the size of the domestic economy.

In the first counterfactual experiment, we decrease the standard error of both stationary shocks leaving unchanged the volatility of the common permanent shock. More precisely, we make the relative productivity shock and the preference shock half as volatile, compute another set of S equilibrium paths and re-estimate the three structural VARs. The implied response of labor input to the permanent productivity shock are displayed in Figure 12.

The left panel of Figure 12 shows that the bias in the country–level experiment is greatly reduced. An applied researcher would no longer reject the hypothesis that these data have been generated by the theoretical model; nor would she reject that the employment response is positive. The comparison of the left panels of Figure 6 and Figure 12 exemplifies the contamination (or the lack thereof) of the domestic labor productivity when SVARs with long–run restrictions are used to reveal permanent productivity shocks.

The improved performance of the country–level SVAR model is a direct consequence of the greater contribution of the world technology shock to domestic fluctuations. It now accounts for 91% of the variance of labor productivity after one quarter, and 58% of the variance of labor input.

The middle and right panels of Figure 12 do confirm the findings of Figure 6: the use of aggregate labor productivity, instead of country–level labor productivity, shifts the estimated response up towards the true response of employment. The experiment with zone employment delivers virtually the same response as the experiment with domestic employment, but with increased precision.

In a second experiment, we make stationary shocks half as persistent as in the benchmark. The responses of employment to the permanent shock estimated with the three SVAR models are displayed in Figure 13.

The left panel, which corresponds to country–level artificial data, shows that the contamination of permanent productivity shocks is less acute when the other shocks are less persistent. The middle and right panel exhibit the same pattern as before: the use of aggregate–level labor productivity reduces the bias, while the best precision is achieved with aggregate–level employment. Once again, this reduction in bias is a direct consequence of a greater contribution of the world technology shock.

Another dimension in which the contamination of labor productivity is apparent is its persistence. Figure 15 plots the theoretical autocorrelation function of the log–deviations of labor productivity from trend, $\log(Y_{i,t}) - \log(N_{i,t}) - \log(A_t)$. The left panel shows that, in the benchmark model, deviations of the domestic labor productivity are much more persistent than deviations of the zone labor productivity. The discrepancy is much lower when the volatility of country–specific shocks is reduced (middle panel), and it disappears completely when country–specific shocks have a low autocorrelation coefficient (right panel). Figure 15 hence confirms the insights of Figures 6 to 13: in our model of the US economy, labor productivity persistently deviates from its stochastic trend and these deviations perturb identification based on long–run restrictions. Aggregate–level productivity is less distorted by persistent, non–permanent shocks.

Our third, and last, counterfactual experiment concerns the relative size of the domestic economy, π_1. The benchmark case portrays the US as home country, implying $\pi_1 = 0.40$. The US has the largest population in the seven countries we consider in the empirical analysis, while Canada has the lowest population share, 0.04. In the model, the size of the domestic economy controls for its exposure to foreign shocks, as well as for the distinction between the domestic economy and the zone economy.
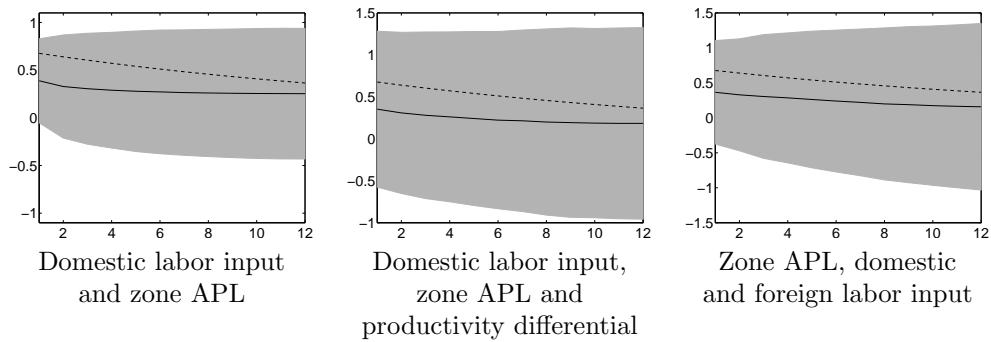
Figure 14 plots the estimated response of employment for values of π_1 between 0.05 and 0.8. The left panel shows that the response estimated in a VAR with country–level data monotonously increases with the size of the domestic economy, reducing the bias. As π_1 goes from low values to 0.5, the contamination of domestic labor productivity by the idiosyncratic productivity choice is reduced and eventually disappears. When $\pi_1 = 0.5$, the only remaining contamination is due to the preference shock. For values of π_1 above 0.5, relative productivity shock do perturb again the identification of the common permanent shock to TFP.

The remaining two experiments, displayed in the middle and right panel of Figure 14, respectively, concern the SVAR on zone data and the SVAR combining aggregate–level labor productivity with country–level employment. Those experiments illustrate that the difference between country–level and aggregate–level labor productivity is high for

low values of π_1 and disappears as π_1 tends to unity. Hence, the estimated responses of employment in the three VAR models are close when the domestic economy represents a large share of world economy, as can be seen when $\pi_1 = 0.8$.
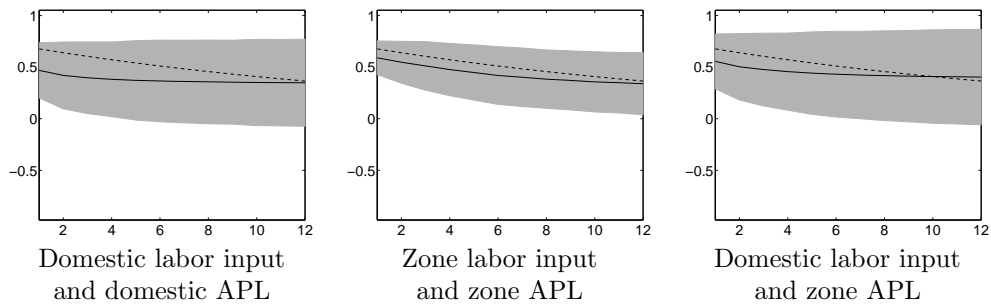
Overall, our set of experiments yields a clear conclusion. When the contribution of country–specific shocks is of the size observed on US data or larger, the response of labor input to permanent technology improvements estimated using country–level data is severely biased. The use of aggregate–level labor productivity instead of country–level labor productivity shifts the estimated response towards their true value. These quantitative results give strong support to the empirical analysis using G7–aggregates put forward in this paper.

Figure 11: Estimated response of employment in the robustness experiments

Domestic labor input and zone APL

Domestic labor input, zone APL and productivity differential

Zone APL, domestic and foreign labor input

Note: VAR estimated on data generated by the two–country model with parameter values described in Table 1 ($S = 500$ simulations of 104 observations), yielding the VMA representation (6).
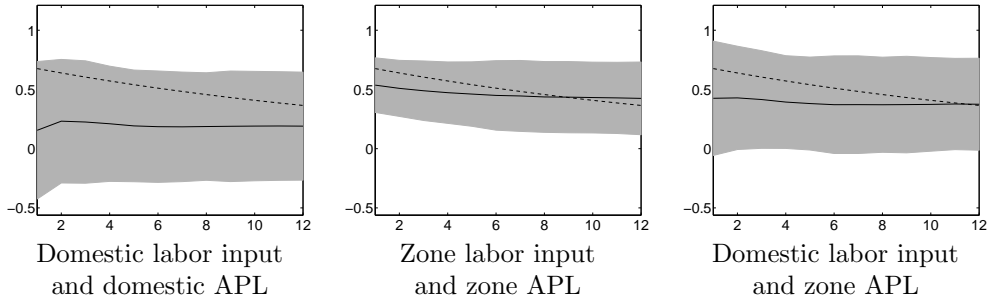
Figure 12: Estimated response of employment with less volatile shocks

Domestic labor input and domestic APL

Zone labor input and zone APL

Domestic labor input and zone APL

Note: VAR estimated on data generated by the two–country model with parameter values described in Table 1 ($S = 500$ simulations of 104 observations), yielding the VMA representation (6).
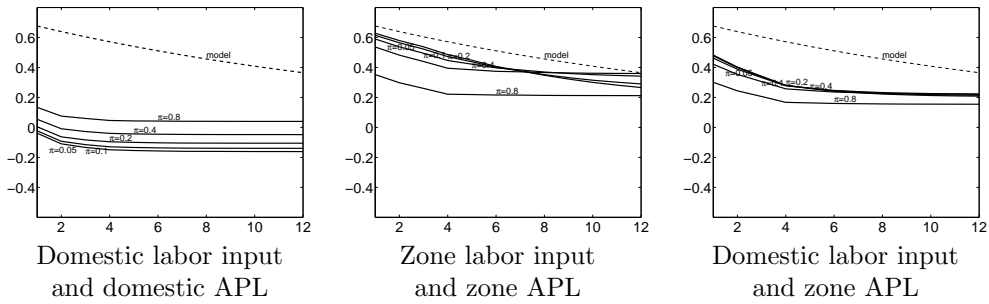
Figure 13: Estimated response of employment with less persistent shocks

| Domestic labor input
and domestic APL | Zone labor input
and zone APL | Domestic labor input
and zone APL |

Note: VAR estimated on data generated by the two–country model with parameter values described in Table 1 ($S = 500$ simulations of 104 observations), yielding the VMA representation (6).

Figure 14: Estimated response of employment according to the size of the 'home' country

| Domestic labor input
and domestic APL | Zone labor input
and zone APL | Domestic labor input
and zone APL |

Note: VAR estimated on data generated by the two–country model with parameter values described in Table 1 ($S = 500$ simulations of 104 observations), yielding the VMA representation (6).

Figure 15: Autocorrelation functions of labor productivity

| Benchmark | Half as volatile | Half as persistent |
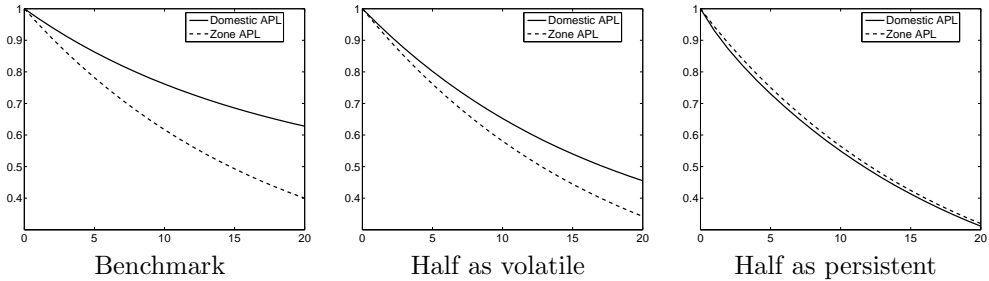
Note: Autocorrelation functions of labor productivity in deviation from the stochastic trend, as implied by the two–country model with parameter values described in Table 1.