CAHIER D'ÉTUDES
WORKING PAPER

N° 152

ARE NEUTRAL AND INVESTMENT-SPECIFIC TECHNOLOGY SHOCKS CORRELATED?

ALBAN MOURA

JANUARY 2021

BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTÈME

ARE NEUTRAL AND INVESTMENT-SPECIFIC TECHNOLOGY SHOCKS CORRELATED?

ALBAN MOURA

ABSTRACT. The joint behavior of Total Factor Productivity (TFP) and the Relative Price of Investment (RPI) in the data lead several authors to conclude that neutral technology shocks are positively correlated with investment-specific technology shocks, challenging the specification of standard macroeconomic models. This paper rejects the correlated-shocks hypothesis using both parametric and non-parametric methods and controlling for structural breaks. The data suggests moderately negative long-run covariation between the RPI and TFP constructed from chain-linked output, but the RPI is orthogonal to TFP in consumption units. These results are consistent with a simple two-sector model in which neutral technology shocks and investment-specific technology shocks are uncorrelated, while models with correlated shocks cannot account for the second result. I conclude that it is not necessary to adapt macro models to allow for correlated technology processes

JEL Codes: E30, E32, O41.

Keywords: total factor productivity, relative price of investment, neutral technology, investment-specific technology, long-run covariability, structural VARs.

December 2020. Banque centrale du Luxembourg, Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg (e-mail: alban.moura@bcl.lu). I am grateful to Luca Benati for sharing his codes and to Mark Watson for useful advice on long-run covariation. For helpful comments and suggestions, I thank Patrick Fève, William Fox, Pablo Garcia Sanchez, Paolo Guarda, and Olivier Pierrard.

Résumé Non Technique

Depuis une quinzaine d'années, nombre de banques centrales ont adopté des modèles d'équilibre général dynamiques et stochastiques (modèles DSGE, en abrégé) pour analyser les effets de la politique monétaire. Une implication robuste de ce type de modèles est que la politique optimale dépend de la nature des chocs affectant l'économie au cours du temps. Par exemple, il est bien connu que la politique monétaire répond de manière différente aux chocs d'offre et de demande. De manière plus fine, des chocs d'offre affectant des secteurs différents de l'économie appellent également des réponses distinctes. Globalement, la compréhension de la nature et des propriétés des chocs est cruciale au choix de la politique la plus appropriée.

Cet article porte sur les propriétés des chocs technologiques. Traditionnellement, les modèles du cycle économique font la différence entre le changement technologique neutre et le changement technologique spécifique au secteur des biens d'investissement. Le changement neutre a un impact symétrique sur la production de tous les biens, tandis que le changement spécifique au secteur de l'investissement affecte uniquement la production des biens d'investissement.[1] Dans ce cadre, il est usuel de supposer que les deux types de technologie évoluent de manière indépendante, ce qui permet d'identifier leurs propriétés statistiques dans les données.

Récemment, plusieurs auteurs ont questionné cette hypothèse d'indépendance en étudiant la relation empirique entre la productivité totale des facteurs (PTF) et le prix relatif de l'investissement (PRI). Leur démarche exploite une propriété fondamentale des modèles macroéconomiques, selon laquelle la PTF reflète l'évolution de la technologie neutre et le PRI celle de la technologie spécifique au secteur des biens d'investissement. Dans les données US, ces deux variables présentent une importante corrélation et semblent répondre aux mêmes chocs. Ces propriétés paraissent contredire l'hypothèse d'indépendance habituellement imposée et ont été attribuées au phénomène de diffusion progressif des innovations technologiques entre les secteurs.

Ces résultats empiriques ont d'importantes conséquences pour la modélisation macroéconomique. Si les chocs affectant les différents types de technologie sont corrélés, tous les travaux antérieurs basés sur l'hypothèse d'indépendance sont mal spécifiés et leurs conclusions peuvent être remises en question. De plus, la présence de chocs corrélés pourrait compliquer la conduite de la politique économique : puisque la politique monétaire optimale réagit de manière différente à un choc technologique neutre et à un choc technologique

[1] Le changement technologique spécifique au secteur des biens d'investissement trouve son origine dans le progrès rapide des technologies de l'information et de la communication. Ainsi, il est bien connu que le prix d'un ordinateur ou d'un téléphone a considérablement chuté au cours du temps, reflétant l'efficacité accrue du processus de production. La montée en puissance de la digitalisation de l'économie est tributaire de la poursuite de ce type de progrès technologique.

spécifique au secteur de l'investissement, il devient compliqué de stabiliser l'économie si les chocs sont corrélés.

Au contraire, cet article démontre que les données ne sont pas compatibles avec la présence de chocs technologiques corrélés. Cette conclusion découle d'une réévaluation critique des résultats avancés dans la littérature pour justifier l'hypothèse alternative de chocs corrélés. Deux points principaux sont à noter. Premièrement, la présence de co-movements entre la PTF et le PRI dans les données ne signale pas forcément la présence de chocs technologiques corrélés : elle peut s'expliquer, même en présence de chocs indépendants, par les propriétés des agrégats économiques mesurant la production réelle dans les comptes nationaux. Deuxièmement, une analyse empirique conduite sur un long échantillon (1950-2019) confirme l'absence de corrélation entre des mesures directes de la technologie neutre et de la technologie spécifique au secteur de l'investissement.

Ainsi, ces résultats valident les modèles macroéconomiques basés sur l'hypothèse usuelle de chocs technologiques orthogonaux. Il s'agit d'une conclusion importante pour bien comprendre le rôle macroéconomique des chocs technologiques et pour analyser et guider la politique économique, par exemple dans les banques centrales.

1. Introduction

Following Greenwood, Hercowitz, and Krusell (2000), the literature studying the role of technology shocks in business cycles focused on the distinction between neutral and investment-specific technology: neutral shocks affect efficiency in production of all goods in a symmetric fashion, while investment-specific shocks make the production of investment goods more efficient than that of consumption goods. A standard assumption in this literature is that the two kind of shocks evolve independently, a restriction that motivates identification strategies for structural vector auto-regressions (Fisher, 2006) and estimated DSGE models (Smets and Wouters, 2007; Justiniano, Primiceri, and Tambalotti, 2010, 2011).

Recently, several papers questioned the notion that neutral and investment-specific technology follow independent processes. Using different empirical methods, Schmitt-Grohé and Uribe (2011), Benati (2014), Chen and Wemy (2015), and Guerrieri, Henderson, and Kim (2020) all argue that *neutral and investment-specific technology comove positively in the long run.* To make the point, Schmitt-Grohé and Uribe and Benati study the long-run relationship between total factor productivity (TFP) and the relative price of investment (RPI) using both standard cointegration tests and more advanced statistical methods, while Chen and Wemy and Guerrieri, Henderson, and Kim exploit structural vector autoregressions (SVARs). All find that permanent shocks driving neutral technology are positively correlated with permanent shocks driving investment-specific technology, calling into question the standard distinction between the two.[2] In addition, Chen and Wemy and Guerrieri, Henderson, and Kim provide a theoretical explanation, according to which correlated technology shocks arise from inter-sectoral linkages and spillovers.

As discussed in Benati (2014), these empirical results have far-reaching consequences for business-cycle theory. If the stochastic processes for neutral and investment-specific technology are indeed correlated, earlier work imposing independence is misspecified and may provide biased estimates of the role of technology shocks in aggregate fluctuations. Furthermore, standard macroeconomic models should be amended to account for the correlation structure found in the data. Finally, this change would modify the welfare implications of DSGE models: for instance, it is well known that a benevolent central bank should react differently to neutral and investment-specific technology shocks (Basu and De Leo, 2016), so that stabilizing the economy might prove more difficult if the shocks are correlated.

In this paper, I take the opposite view and argue that there is in fact little empirical support for correlated shocks to neutral and investment-specific technology. I reach this conclusion by conducting a critical re-evaluation of the empirical results supporting the correlated-shocks hypothesis. Overall, my findings suggest that the usual assumption of independent

[2]Benati's (2014) empirical conclusions are mixed because of the uncertainty related to structural trend breaks in the logarithms of TFP and the RPI. However, his preferred specification implies that neutral and investment-specific shocks are positively correlated.

technology shocks is well in line with the behavior of TFP and the RPI observed in US data. This outcome is reassuring for the specification of standard macroeconomic models, as well as for the large stream of papers that evaluated the effects of technology shocks and their contribution to business cycles under the orthogonality assumption.

In Section 2, I start by clarifying a measurement issue that complicates the interpretation of the relationship between TFP and the RPI. The premise of both Schmitt-Grohé and Uribe (2011) and Benati (2014) is the notion that TFP reflects neutral technology and that the RPI corresponds to the inverse of investment-specific technology. I stress that this mapping is violated by the way TFP is measured in the data: standard TFP series are constructed as Solow residuals from chain-aggregated output and do *not* correspond to neutral technology in multi-sector models.[3] Instead, they combine neutral and investment-specific technology, implying that negative long-run comovements between TFP and the RPI are expected even if technology shocks are orthogonal. Both Schmitt-Grohé and Uribe and Benati overlooked this issue, weakening their argument that a non-zero correlation between TFP and the RPI implies correlated technology shocks.

This criticism calls for a re-evaluation of the joint behavior of TFP and the RPI in the data. This is the topic of Section 3, which applies econometric techniques recently developed by Müller and Watson (2018, 2019) to study the long-run properties of economic time series. These methods focus on the second moments of low-frequency transformations of the original series, allowing proper inference without relying on parametric models.[4] To deal with the measurement issue highlighted above, I consider three TFP series: first, Fernald's (2014) utilization-adjusted TFP for the US business sector, constructed from chain-aggregated output (quantity TFP hereafter); second, Fernald's (2014) consumption-sector TFP; third, an alternative series for consumption TFP proposed by Moura (2020). A standard two-sector model with orthogonal technology shocks implies that quantity TFP responds to both neutral and investment-specific technology in the long run, while consumption TFP responds only to neutral technology. Thus, comparing the comovements of each TFP series with the RPI provides valuable information about the long-run relationship between neutral and investment-specific technology.

The empirical results reveal three properties of the data. First, when one does not control for structural breaks in the series, the RPI exhibits significant positive long-run comovements with both quantity and consumption TFP series. Taken at face value, this pattern eliminates

[3]This property is not new: the issues related to output measurement and TFP interpretation are well known in the literature on aggregate productivity (see, e.g., Oulton, 2007, and Greenwood and Krusell, 2007). In a companion paper (Moura, 2020), I provide a simple way to correct standard TFP measures to recover neutral technology.

[4]The alternative is to fit a parametric time series model, typically a VAR, to estimate the autocovariances of the data over both short and long horizons. I consider this approach in Section 4.

the possibility of positive correlation between the stochastic trends in neutral and investment-specific technology, since this would imply a negative relationship between TFP and the RPI. Second, the RPI exhibits slightly negative long-run comovement with quantity TFP when controlling for structural breaks. Third, the RPI exhibits essentially zero long-run correlation with consumption TFP when controlling for breaks.

The second result echoes the findings from Schmitt-Grohé and Uribe (2011) and Benati (2014), but my interpretation is different: I show that the estimated correlation between quantity TFP and the RPI is well in line with an economy with independent technology shocks, once one acknowledges the properties of chain-aggregated output. This interpretation is confirmed by the third property: if neutral and investment-specific technology were correlated, then the RPI would *also* comove with consumption TFP, which is not the case. I conclude that the long-run characteristics of the data do not support the view that technology shocks are correlated.

Finally, Section 4 proposes a robustness exercise based on a structural VAR. Following Chen and Wemy (2015), I identify the shocks that contribute most to the long-run forecast error variance of TFP and the RPI. The results are in line with the outcome of the Müller-Watson procedure: neutral technology (as measured by consumption TFP) and investment-specific technology (as measured by the inverse of the RPI) are negatively correlated when ignoring structural breaks, and not at all correlated over a sample running from 1983 to 2019, i.e. after the last estimated break date. I conclude that neutral and investment-specific technology are well characterized as orthogonal processes in the long run, so that introducing correlated technology shocks in general-equilibrium models is not warranted.[5]

2. Long-Run Implications of a Two-Sector Model

To motivate the empirical analyses conducted in Sections 3 and 4, this section re-examines the long-run implications of a standard two-sector growth model with independent technology shocks for the comovements between TFP and the RPI.

2.1. **Model.** The model is based on Greenwood, Hercowitz, and Krusell (1997, 2000) and Fisher (2006). The economy is closed and all agents behave competitively. The general equilibrium corresponds to the solution of the planning problem:

$$\max \quad E_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\ln C_t - \theta \frac{H_t^{1+1/\kappa}}{1+1/\kappa} \right) \right] \tag{1}$$

[5]I also investigate the discrepancy between my results and Chen and Wemy's (2015) estimates of a near perfect correlation between the long-run shocks to TFP and the inverse of the RPI. I show that updates in Fernald's (2014) TFP series, together with a longer estimation sample, explain the difference.

subject to

$$C_t + \frac{I_t}{V_t} = A_t K_{t-1}^{\alpha} H_t^{1-\alpha}, \tag{2}$$

$$K_t = (1-\delta)K_{t-1} + I_t, \tag{3}$$

$$\ln \frac{A_t}{A_{t-1}} = \mu_A + \epsilon_t^A, \tag{4}$$

$$\ln \frac{V_t}{V_{t-1}} = \mu_V + \epsilon_t^V. \tag{5}$$

Here, C, H, I, K, A, and V denote consumption, hours worked, investment, the capital stock, neutral technology, and investment-specific technology. β is the household discount factor, θ is a preference weight, κ is the elasticity of labor supply, α is the capital share, and δ is the depreciation rate.

This economy is driven by the stochastic trends in neutral and investment-specific technology. Unit-root processes are required for the notion of long-run comovements to make sense; one could allow for additional transitory components in A_t and V_t with no effect on the conclusions.[6] Parameters $\mu_A > 0$ and $\mu_V > 0$ denote the average growth rates of neutral and investment-specific technology, while ϵ_t^A and ϵ_t^V are independent technology shocks with standard deviations σ_A, $\sigma_V > 0$. Thus, the processes for neutral and investment-specific technology defined by equations (4)-(5) are orthogonal.

2.2. **Measurement.** The model counterpart of the RPI measured in the data is the inverse of investment-specific technology:

$$P_t = \frac{1}{V_t}. \tag{6}$$

This is because equation (2) implies that one unit of consumption trades against V_t units of investment.

For TFP, things are complicated by the ambiguous nature of output measurement in multi-sector economies: different concepts of aggregate production coexist, with distinct properties, and these differences spread to TFP series computed as residuals from production functions. This issue is well known in the productivity literature, but has been sometimes overlooked in the business-cycle literature.[7]

Consumption TFP is constructed from output expressed in consumption units, defined as $Y_t := C_t + P_t I_t$. Using equations (2) and (6), it is straightforward to see that consumption

---

[6]Fisher (2006) emphasizes that the long-run implications of the simple model are quite general, in the sense that they follow only from the assumptions on preferences and technology necessary for balanced growth. Including additional short-term frictions or temporary shocks would leave the key implications unchanged.

[7]Whelan (2003) contains a clear exposition of the modeling challenges raised by the chain-aggregated variables featured in the US National Income and Product Accounts. Moura (2020) provides a non-exhaustive list of papers mistakenly mapping quantity TFP into neutral technology in the context of multi-sector models.

TFP exactly recovers neutral technology:

$$TFP_t^Y := \frac{Y_t}{K_{t-1}^{\alpha} H_t^{1-\alpha}} = A_t. \tag{7}$$

Quantity TFP, on the other hand, is constructed from chain-aggregated real output (or quantity output), which has been the standard measure of real GDP in the US National Income and Product Accounts (NIPAs) since 1996. Using Whelan's (2003) result that chain-aggregates are well approximated by share-weighted Divisia indexes, the model counterpart of chain-aggregated output is $D_t := C_t^{\gamma} I_t^{1-\gamma}$, where $\gamma = C^{\star}/(C^{\star} + P^{\star} I^{\star}) \in (0,1)$ is the steady-state nominal share of consumption in GDP.[8] Quantity TFP follows:

$$TFP_t^D := \frac{D_t}{K_{t-1}^{\alpha} H_t^{1-\alpha}},$$

and since $Y_t \neq D_t$, it is clear that $TFP_t^D \neq A_t$. Straightforward computations confirm that the stochastic trend in TFP_t^D is

$$\text{trend}(TFP_t^D) = A_t V_t^{1-\gamma}, \tag{8}$$

so that quantity TFP *rises* with investment-specific technology in the long run (recall that $\gamma < 1$).

2.3. **Implications.** Clarifying the difference between TFP_t^Y and TFP_t^D is key to inferring the statistical properties of neutral and investment-specific technology from observations on TFP and the RPI.

For instance, both Schmitt-Grohé and Uribe (2011) and Benati (2014) assumed that TFP reflects neutral technology and interpreted the presence of long-run comovements between TFP and the RPI as a proof that neutral and investment-specific technology shocks correlate. Theory implies that this view would be valid if the empirical analysis was based on consumption TFP. Unfortunately, both Schmitt-Grohé and Uribe and Benati used quantity TFP. This is an issue because quantity TFP increases with investment-specific technology: since the RPI decreases with V_t, negative long-run comovements between TFP_t^D and the RPI are expected even when neutral and investment-specific technology are orthogonal. This property largely weakens the empirical evidence presented by Schmitt-Grohé and Uribe and Benati in favor of correlated technology shocks.

Theory also suggests that studying the long-run relationship between *consumption TFP* and the RPI provides a fruitful avenue to evaluate the likelihood of correlated shocks to neutral and investment-specific technology. The next two sections are devoted to this task.

[8] The true Divisia index has geometric weights reflecting current GDP shares. Here, I use the steady-state values to avoid minor complications that do not affect the results.

3. Long-Run Comovements Between TFP and the RPI

This section characterizes the long-run comovements between TFP and the RPI in the data using recent econometric techniques. The result provide little support for the view that neutral and investment-specific technology are correlated in the long run. Instead, they appear consistent with the implications of orthogonal technology processes.

3.1. **Data.** The empirical analysis is based on time series ranging from 1950Q1 to 2019Q4. This sample adds more than 10 years of additional observations with respect to Schmitt-Grohé and Uribe (2011) and Benati (2014), which provides useful supplementary information about the "long run."

As usual, I measure the RPI as the ratio of quarterly seasonally-adjusted chain-aggregated deflators for investment and consumption goods, both derived from the National Income and Product Accounts (NIPA).[9] Following Justiniano, Primiceri, and Tambalotti (2011), Chen and Wemy (2015), and Moura (2018), investment is the sum of expenditures on durable consumption goods and fixed investment, while consumption is the sum of expenditures on nondurable consumption goods and services.

As in Benati (2014) and Chen and Wemy (2015), I construct quantity TFP by cumulating the quarterly growth rate of Fernald's (2014) seasonally-adjusted TFP adjusted for factor utilization. A look at Fernald's appendix confirms that this series is derived from chain-aggregated real output, so that it corresponds to TFP^D.

Turning to TFP^Y, I work with two measures of consumption TFP. The first comes from Fernald (2014), who decomposes quantity TFP into components related to the consumption and investment sectors. Simple computations confirm that this consumption-sector TFP corresponds to TFP^Y in standard two-sector models (see Moura, 2020). However, one could criticize this variable on the ground that Fernald's decomposition exploits the relative price of *equipment*, whereas the RPI series has broader coverage and also includes durable household goods, structures, and residential investment. To address this consistency issue, I consider a second measure of consumption TFP, constructed as in Moura (2020): building upon from Fernald's database, I keep his series for inputs and utilization and simply replace the original chain-aggregated output measure by output in consumption units, computed as the ratio between nominal output and the consumption price used to deflate the RPI. The resulting TFP series verifies equation (7) by construction.

---

[9] Some authors, including Gordon (1990) and Cummins and Violante (2002), have argued that NIPA series do not correctly incorporate quality improvements in equipment goods and yield a biased measure of the RPI. They have also proposed alternative investment deflators to improve quality adjustment. Because these alternative price indexes are only available up to the mid-2000s, this paper uses NIPA deflators to obtain a longer sample. Besides, Benati (2014) already provides a thorough empirical analysis based on the shorter quality-adjusted RPI series constructed by Liu, Waggoner, and Zha (2011) and Schmitt-Grohé and Uribe (2011).

Figure 1. TFP and the RPI in US data

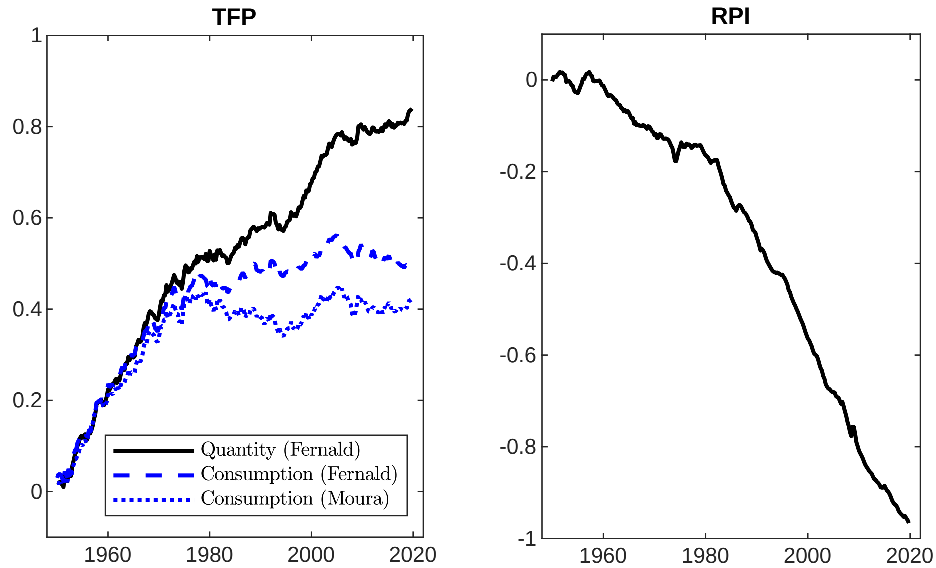

Notes. All series are in log and TFP series are adjusted for utilization. See the text for the sources.

Figure 1 presents the series. The left panel highlights that quantity TFP and consumption TFP grew at about the same rate up to 1970 but diverged afterward. In particular, both series for TFP^Y exhibit a clear downward trend break during the 1970s, which reflects a slowdown in average neutral technology growth. In addition, it is clear that the two measures of consumption TFP display very similar behavior over the postwar period. Technology slowdown is also apparent in quantity TFP, but faster investment-specific technology growth illustrated by the downward trend break in the RPI in the early 1980s (right panel) partly offset it. Overall, visual inspection suggests different long-run properties for quantity and consumption TFP and identifies trend breaks as a potential nuisance for the empirical analysis.

3.2. **Preliminary analysis.** I start by formally assessing the presence of trend breaks and unit roots in TFP and the RPI, as well as their cointegration properties. This is a required step because inference about the long run is sensitive to the presence of breaks in deterministic trends (Perron, 1989). Furthermore, the discussion of long-run comovements would be moot with stationary series.

I reach the same conclusions as Benati (2014), so I delegate the material related to this preliminary analysis to Appendix A.1. The results are as follows:

- There is substantial empirical evidence that both TFP series and the RPI experienced a single trend break in postwar data. The estimated break dates are 1973Q1 for TFP and 1982Q2 for the RPI. Because the confidence intervals are wide, the empirical evidence does not rule out the possibility of common break in 1982Q2.

- There is no evidence in the data against a unit root in both TFP and the RPI, irrespective of the presence of breaks.
- The data do not support the idea that TFP and the RPI are cointegrated when controlling for trend breaks. However, this does not rule out the possibility that TFP and the RPI share a common $I(1)$ component.

Uncertainty about trend breaks was problematic for Benati (2014) because he found long-run inference to be very sensitive to their timing. Instead, the empirical approach considered next appears remarkably robust to this issue.

3.3. **Long-run correlation.** To characterize the long-run relationship between TFP and the RPI, I apply a procedure recently developed by Müller and Watson (2018, 2019). Their approach is intuitive and only exploits the low-frequency characteristics of the data, providing reliable inference about long-run comovements between two time series for a wide range of persistence patterns.

Müller and Watson's procedure involves two stages. In the first, the original series are transformed using linear projections onto deterministic low-frequency cosine waves. The resulting transforms resemble standard low-pass filters that eliminate the effects of short-run disturbances, but their statistical properties are easier to derive. In the second stage, inference about the long run exploits the second moments of the low-frequency transforms, expressed in the form of a long-run correlation coefficient ρ_T and a long-run linear regression coefficient β_T.[10] These parameters have the usual interpretation: positive values for β_T and ρ_T signal that the original series tend to rise and fall together at low frequencies, implying positive long-run comovements, and the converse is true for negative values. When the original series are $I(0)$ as postulated here (the series are the log-differences of TFP and the RPI), standard finite-sample normal linear regression formulas apply, allowing for straightforward inference.

Table 1 presents estimates and confidence sets for the coefficients ρ_T and β_T characterizing the long-run relationship between TFP and the RPI. As in Müller and Watson (2018), I focus on periods above 11 years, which are longer than the typical business cycle and should be enough for spillovers to operate.[11] The choice of this threshold determines the number of cosine regressors in the first stage, namely $q = 12$. For robustness, I have also estimated the

[10]Following Müller and Watson's notation, these moments are indexed by the sample size T. This is because the number of cosine waves used as regressors in the first stage increases with T, changing the population properties of the low-frequency transforms.

[11]For instance, the estimates reported by Chen and Wemy (2015) indicate that the responses of TFP and the price of investment to technology shocks stabilize after 5 to 10 years.

TABLE 1. Long-run comovements between TFP and the RPI.

Specification of break(s)		TFP series					
		Quantity		Consumption			
				Fernald		Moura	
No break	$\widehat{\rho}_T$	0.35	[0.04, 0.60]	0.63	[0.39, 0.79]	0.60	[0.36, 0.77]
	$\widehat{\beta}_T$	0.39	[0.07, 0.70]	0.60	[0.37, 0.82]	0.55	[0.33, 0.78]
Idiosyncratic (TFP)	$\widehat{\rho}_T$	−0.22	[-0.50, 0.10]	0.01	[-0.30, 0.32]	0.04	[-0.27, 0.35]
	$\widehat{\beta}_T$	−0.35	[-0.82, 0.13]	0.02	[-0.57, 0.62]	0.08	[-0.46, 0.61]
Idiosyncratic (TFP, RPI)	$\widehat{\rho}_T$	−0.15	[-0.45, 0.17]	0.11	[-0.21, 0.41]	0.18	[-0.14, 0.47]
	$\widehat{\beta}_T$	−0.11	[-0.34, 0.11]	0.11	[-0.17, 0.38]	0.15	[-0.10, 0.40]
Common	$\widehat{\rho}_T$	−0.17	[-0.46, 0.15]	0.02	[-0.30, 0.33]	0.07	[-0.25, 0.38]
	$\widehat{\beta}_T$	−0.10	[-0.28, 0.08]	0.01	[-0.18, 0.20]	0.04	[-0.13, 0.21]

Notes. The sample is 1950Q1-2019Q4 and the procedure exploits comovements at periods longer than 11 years. Entries in the 'Quantity' column describe the long-run relationship between Fernald's quantity TFP and the RPI, while entries in the subsequent columns describe the long-run relationship between the two measures of consumption TFP and the RPI. $\widehat{\rho}_T$ is the estimated long-run correlation coefficient, $\widehat{\beta}_T$ is the estimated coefficient in the linear long-run regression of TFP on the RPI, and brackets report 67% confidence sets. See Appendix A for details on trend breaks.

coefficients focusing on periods longer than 15 years to limit the influence of medium-term cycles: the results are very similar, both qualitatively and quantitatively (see Appendix B.1).

The main rows in the table correspond to different treatments of the structural breaks in the series. I consider four cases: ignoring the breaks, controlling only for the break in TFP, controlling for idiosyncratic breaks in both TFP and the RPI, and controlling for a common break. In his empirical study, Benati (2014) considered the first, third, and fourth cases. The rationale for the second case, which controls only for the idiosyncratic break in TFP, comes from equation (8). A break in the average growth rate of investment-specific technology affects both quantity TFP and the RPI, so that it creates meaningful comovements and should not be eliminated. On the other hand, it is appropriate to control for the break in TFP if the latter originated from neutral technology and did not spillover to the RPI.

Three results stand out in Table 1. The first is that both TFP series exhibit *positive* long-run comovements with the RPI when structural breaks are ignored. The estimated long-run correlation coefficients and long-run linear regression coefficients are statistically above zero, even though confidence intervals are wide. This uncertainty reflects the limited long-run information in the sample: inference effectively relies on the $q = 12$ pairs of projection coefficients.

Interpreted through equations (6) and (7), the positive relationship between the RPI and consumption TFP suggests a negative long-run correlation between neutral and investment-specific technology. This pattern is consistent with two well-known features of the data: the 1970s productivity slowdown and the 1980s increase in the rate of investment-specific technology growth identified by the tests for structural breaks.[12] Taken at face value, this property nullifies the possibility that neutral and investment-specific technology shocks are positively correlated. However, failing to account for trend breaks is likely to bias this conclusion.

The second result is that quantity TFP shows moderate *negative* long-run comovements with the RPI when controlling for structural breaks. This negative relationship holds irrespective of the handling of breaks: using the Müller-Watson approach, correcting for an idiosyncratic break in TFP only, for idiosyncratic breaks in both TFP and the RPI, or for a common break in the series yields a very similar picture of the long-run relationship between quantity TFP and the RPI. The estimated long-run correlations range from -0.15 to -0.22, while estimated long-run regression coefficients range from -0.10 to -0.35. The estimates are not statistically significant, but the coherence of the results across methods suggests that a weakly negative long-run relationship between quantity TFP and the RPI is a robust feature of the data.

It would be wrong to interpret these negative comovements as reflecting correlated technology shocks. Quantity TFP is growing with investment-specific technology in the long run while the RPI is falling with it, so that negative comovements do not necessarily reflect correlated shocks. In fact, the model with independent technology shocks rationalizes well the estimates reported in Table 1. Using small font to denote logarithms, equations (6) and (8) imply

$$\operatorname{cov}\left(\Delta tfp_t^D,\, \Delta p_t\right) = \operatorname{cov}\left(\Delta a_t + [1-\gamma]\Delta v_t,\, -\Delta v_t\right) = -(1-\gamma)\operatorname{var}(\Delta v_t),$$

where the last equality follows from the assumption that neutral and investment-specific technology are orthogonal. By focusing on stochastic trends, this expression defines a theoretical counterpart to the long-run covariance estimated by the Müller-Watson procedure. It is possible to deduce a theoretical long-run correlation coefficient, given by

$$\operatorname{corr}^{LR}\left(\Delta tfp_t^D,\, \Delta p_t\right) = -(1-\gamma)\frac{\sigma^{LR}(\Delta p_t)}{\sigma^{LR}(\Delta tfp_t^D)},$$

where $\sigma(.)$ denotes the standard deviation and the LR superscript indicates long-run moments. The average investment-to-output ratio $1-\gamma$ is equal to 0.22 in Fernald's dataset and the ratio of estimated long-run standard deviations $\sigma^{LR}(\Delta p_t)/\sigma^{LR}(\Delta tfp_t^D)$ ranges from 0.6 to 1.6 depending on the handling of breaks. Thus, theory implies that estimated long-run correlations between Fernald's TFP and the RPI should lie between -0.13 and -0.35 if technology shocks are orthogonal. A simple look at Table 1 confirms that this implication holds

[12] See Hornstein and Krusell (1996) and Greenwood and Yorukoglu (1997) for interpretations relating both the productivity slowdown and the acceleration in the rate of decline of the RPI to the rise in new information technologies.

true when controlling for structural breaks. This is a case in which simple theory accounts well for the properties of the data, both qualitatively and quantitatively.

Finally, the third result is that consumption TFP, as measured by either of Fernald's or Moura's series, exhibits *zero* or *weakly positive* comovement with the RPI in the long run when controlling for trend breaks. The estimated correlation coefficients range from 0.01 to 0.18 and the estimated regression coefficients lay between 0.01 and 0.15, while all confidence sets are roughly centered around zero. Hence, in this case too the outcome of the Müller-Watson approach is robust to the nature and timing of the structural breaks.

Recall from Section 2 that the stochastic trend in TFP_t^Y is exactly neutral technology and the stochastic trend in P_t is exactly the inverse of investment-specific technology. Therefore, the estimated long-run correlation and linear regression coefficients imply that neutral and investment-specific technology are largely orthogonal at low frequencies. This property directly contradicts the correlated-shocks hypothesis, which would implies counterfactual comovements between consumption TFP and the RPI. I conclude that, viewed through the lenses of the Müller-Watson procedure, the data provide no support for the notion that shocks to neutral and investment-specific technology correlate in the US.

4. A Structural VAR Check

This last section evaluates the long-run relationship between TFP and the RPI using a structural VAR. The results confirm those from the Müller-Watson approach: in the long run, the shocks driving consumption TFP and the RPI are decorrelated in recent data, indicating that shocks to neutral and investment-specific technology are different.

4.1. **Empirical approach.** Consider the reduced-form moving-average representation of a VAR system:

$$\mathbf{Y}_t = \mathbf{C}(L)\mathbf{u}_t,$$

where $\mathbf{Y}_t$ is the vector gathering the m observable variables, $\mathbf{C}(L) = \mathbf{I} + \sum_{j=0}^{\infty} \mathbf{C}_j L^j$ is a matrix polynomial in the lag operator L, and $\mathbf{u}_t$ is the vector of m residuals with variance matrix $\mathbf{\Sigma}$. In the application, $\mathbf{Y}_t$ includes the log-levels of TFP and the RPI, as well as other macro variables.

Identification is based on the notion that permanent technology shocks dominate fluctuations in TFP and the RPI at low frequencies. Following Chen and Wemy (2015), I identify two shocks using the Maximum Forecast Error Variance (MFEV) approach (Uhlig, 2004; Francis, Owyang, Roush, and DiCecio, 2014): the first identified shock is the largest contributor to the forecast error variance of the RPI at the horizon of 80 quarters; the second identified shock is the largest contributor to the forecast error variance of TFP at the same horizon. In addition, the RPI shock is normalized to generate a negative long-run response of the RPI, while the TFP shock is normalized to generate a positive long-run response of productivity. In the standard model of Section 2, the first identified shock corresponds to the innovation ϵ^V to investment-specific technology. The second shock corresponds to the innovation ϵ^A to neutral technology if consumption TFP is included in the VAR, and to a combination of ϵ^A and ϵ^V if quantity TFP is used instead.

This SVAR strategy usefully complements the empirical analysis proposed in the previous section: while the Müller-Watson approach is non-parametric and does not base inference on a specific data generating process, the SVAR approach is parametric and infers long-run properties from short-run dynamics based on the structure of the model. There are obvious advantages to non-parametric inference, for instance robustness, but the efficiency loss can be large relative to an appropriate parametric model. Therefore, contrasting the results of the two approaches is interesting: if both agree, the findings are strengthened; if they disagree, there is a puzzle to investigate.

As in Chen and Wemy, the estimated VAR includes TFP and the RPI, together with GDP, consumption, investment, and hours worked, all in log-levels. TFP and the RPI are the series described in Section 3.1.[13] GDP is chain-aggregated real output from the NIPA, consumption is the chain-weighted real aggregate of personal consumption of nondurable goods and services, and investment is the chain-weighted real aggregate of fixed private investment and personal consumption expenditures on durable goods. Hours worked correspond to the nonfarm business sector. GDP, consumption, investment, and hours worked are expressed in per-capita terms using the civilian non-institutional population aged 16 and above.[14] The estimation sample is 1950Q1 to 2019Q4, longer than the period 1961-2008 considered by Chen and Wemy.

4.2. **Results.** Figure 2 reports the estimated IRFs to the identified long-run shocks to the RPI (solid line) and TFP (dashed line) when the estimated VAR includes quantity TFP. The RPI shock resembles a standard investment-specific technology innovation (Fisher, 2006). The RPI falls on impact and keeps trending down until it settles at a new level, about 0.25% below its original value. This movement is associated with an economy-wide expansion, as GDP, consumption, investment, and hours worked all rise together. Finally, the RPI shock triggers a permanent increase in quantity TFP.

The long-run shock to TFP has the flavor of an aggregate supply disturbance. In particular, it is associated with permanent positive responses of both TFP and investment-specific technology, the latter represented by a fall in the RPI. In the short run, this positive supply shock generates standard business-cycle comovements across macroeconomic aggregates.

[13]To stay as close as possible to Chen and Wemy (2015), the empirical analysis in this section uses Fernald's (2014) series for quantity and consumption TFP. I have verified that using the alternative measure of consumption TFP yields identical results.

[14]The consumption and investment series are constructed from the NIPA. Hours worked come from Valerie Ramey's webpage and the population series is published by the Bureau of Labor Statistics.

FIGURE 2. Impulse-responses to the long-run shocks to TFP and the RPI — VAR with quantity TFP.

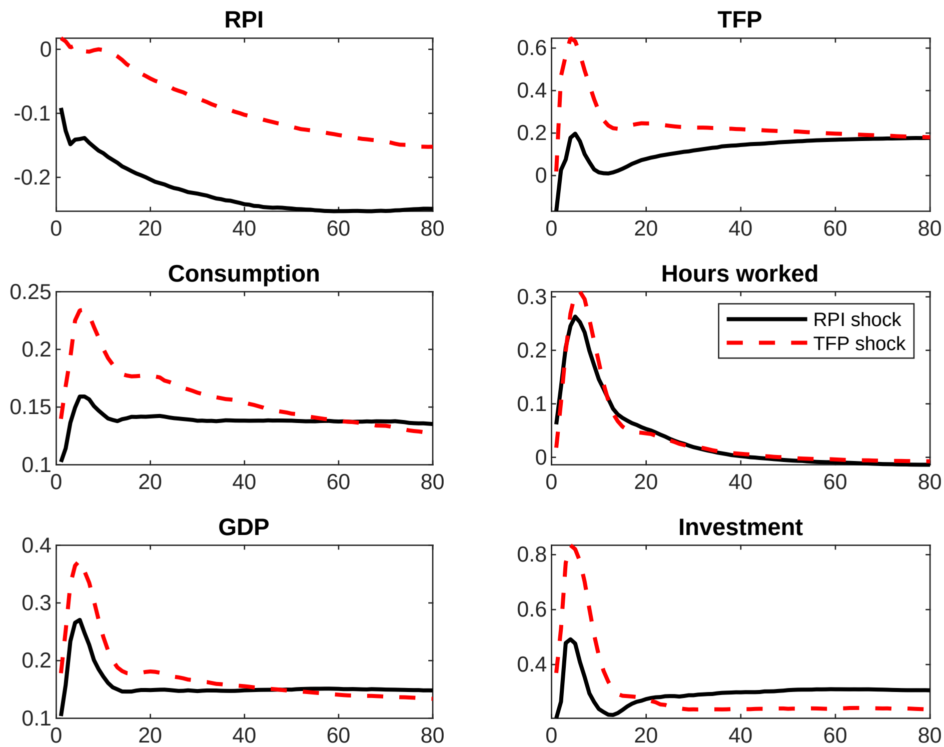

Notes. IRFs to the shocks that contribute most to the FEV of the RPI (solid black line) and TFP (dashed red line) at the forecast horizon of 80 quarters. The reported responses are the median from the bootstrapped distributions with 1,000 replications. The VAR includes quantity TFP as observable.

It is difficult to read these estimates as evidence that technology shocks are correlated. Since quantity TFP increases with both neutral and investment-specific technology, it is not surprising that a positive investment shock generates both a permanent fall in the RPI and a permanent rise in TFP. In addition, the identification strategy backing out a single TFP disturbance is bound to recover a combination of neutral and investment-specific shocks, which rationalizes the estimated permanent effects on TFP and the RPI.

Another important result is that, for all variables in the VAR, the responses to the TFP and RPI shocks are quite different from each other. For instance, the RPI does not react to the TFP shock during the first 10 quarters, while it responds immediately to the RPI shock. Likewise, TFP, consumption, and investment display hump-shaped movements after the TFP shock, while they tend to jump and stabilize quickly at new levels after the RPI shock. These different dynamics contrast starkly with Chen and Wemy's (2015) finding of a near perfect correlation between the responses to the two shocks. The discrepancy is also clear when looking directly at the identified shocks. For instance, the first entry in Table 2 indicates a positive median correlation of 0.42 between the long-run shocks to quantity TFP

Table 2. Correlation coefficients between the long-run shocks to TFP and the RPI.

TFP series	Sample	Correlation coefficient
Quantity	1950-2019	0.42 [0.05, 0.77]
Consumption	1950-2019	−0.25 [-0.55, 0.19]
Consumption	1950-1972	0.56 [0.08, 0.88]
Consumption	1983-2019	−0.01 [-0.49, 0.49]

Notes. Correlation coefficients between the shocks that contribute most to the FEV of the RPI (solid black line) and TFP (dashed red line) at the forecast horizon of 80 quarters. The reported statistic is the median from the bootstrapped distribution with 1,000 replications.

and the RPI: while positive, this statistics is well below the 0.97 correlation reported by Chen and Wemy, suggesting that the two shocks are far from identical.

Figure 3 shows the estimates from the VAR including consumption TFP. As in Section 3, relying on consumption TFP allows for easier interpretation because it exploits a direct observation on neutral technology. The results indicate a clearer distinction: shocks to the RPI have no long-run effect on consumption TFP, and shocks to consumption TFP have no long-run effect on the RPI. Table 2 indicates that the two shocks are correlated, so that long-run movements in neutral and investment-specific technology remain related. However, the correlation is negative (-0.25), so that a RPI shock increasing investment technology in the long run is associated, on average, with a TFP shock lowering neutral technology. It is interesting that the Müller-Watson approach finds a similar negative relationship between neutral and investment-specific technology when ignoring trend breaks. (This is the correct benchmark since the VAR is fitted over the whole sample, without controlling for structural breaks.)

To assess the role of trend breaks, I repeat the exercise with consumption TFP based on a split sample. The first subsample ends in 1972, just before the estimated break date for TFP, and the second starts in 1983, immediately after the estimated break date for the RPI. To save on space, I only report the correlations between the identified long-run shocks to TFP and the RPI in Table 2, and relegate the estimated IRFs to Appendix B.2.

The results indicate a strong positive correlation of 0.56 between the TFP and RPI shocks over the 1950-1972 sample, which partly supports Chen and Wemy's (2015) view. However, the correlation remains much smaller than those estimated by Chen and Wemy, which are close to 0.95. Furthermore, it seems hazardous to estimate comovements at the 20-year horizon from a sample of 22 years only. Over the longer and more recent 1983-2019 subsample, the estimated correlation is zero: the shock that drives consumption TFP in the long run appears totally orthogonal to the long-run shock to the RPI. In addition, Figure 6 in Appendix B.2 shows that the two shocks trigger completely different dynamics in the

FIGURE 3. Impulse-responses to the long-run shocks to TFP and the RPI — VAR with consumption TFP.

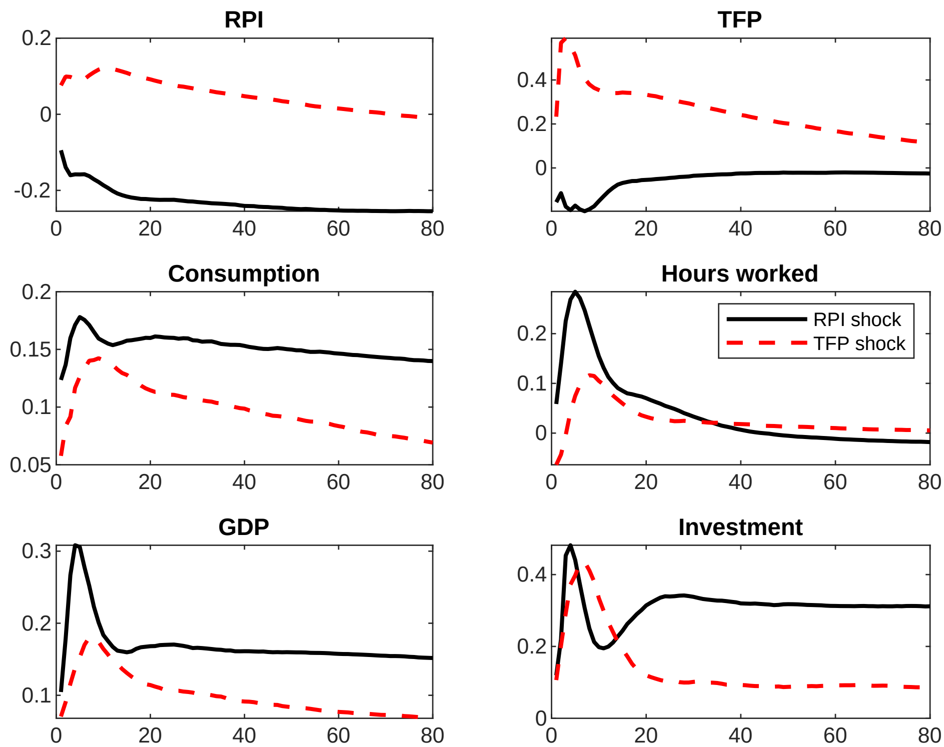

Notes. IRFs to the shocks that contribute most to the FEV of the RPI (solid black line) and TFP (dashed red line) at the forecast horizon of 80 quarters. The reported responses are the median from the bootstrapped distributions with 1,000 replications. The VAR includes Fernald's (2014) consumption TFP as observable.

RPI, TFP, GDP, consumption, investment, and hours worked. It follows that neutral and investment-specific technology have been driven by distinct shocks in recent years.

Overall, I conclude that the SVAR approach delivers the same insights as the more flexible method proposed by Müller and Watson (2018, 2019). In particular, the two strategies concur that, when cleansing the data from structural breaks, neutral technology (as measured by consumption TFP) and investment-specific technology (as measured by the RPI) are largely decorrelated in the long run in recent data. This provides strong empirical evidence against the hypothesis that technology shocks are correlated in the US.

4.3. **Explaining the discrepancy with Chen and Wemy (2015).** Finally, I investigate the discrepancy between my estimates and the results reported by Chen and Wemy (2015), who followed from the same empirical approach, applied to datasets constructed from identical sources, and found instead a strong relationship between the long-run shocks to TFP and the RPI.

Figure 4. Different vintages of Fernald's (2014) consumption TFP

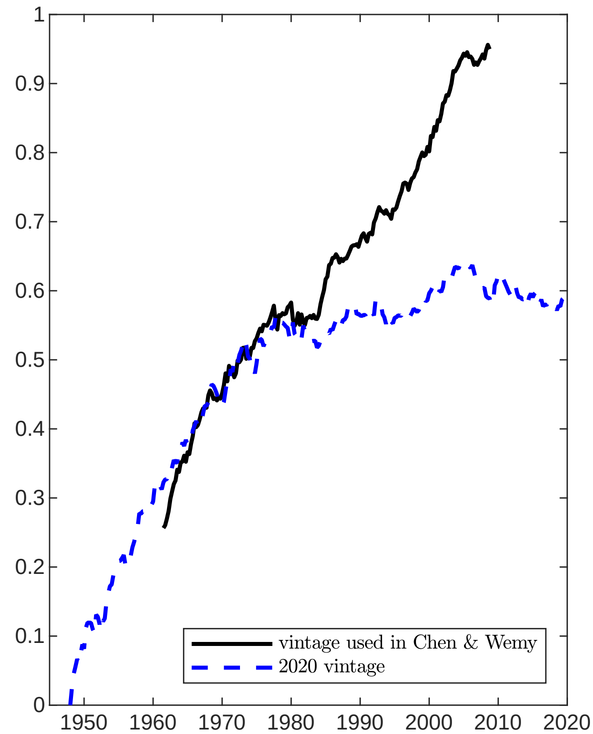

Notes. All series are in log and adjusted for utilization. See the text for the sources.

A first explanation is that adding more observations, both at the start and the end of the estimation sample, changes in an important way the properties of the data. When I estimate the VAR over the 1961-2008 sample, as did Chen and Wemy, the median correlation between the identified long-run shocks to consumption TFP and the RPI is 0.31, much higher than the -0.25 estimated over the longer 1950-2019 sample. This suggests that Chen and Wemy's conclusions are partly driven by their estimation sample and do not hold for longer datasets.

Another, and more fundamental, cause originates from significant revisions in Fernald's (2014) TFP dataset over time. This can be seen from Figure 4, which compares two vintages of Fernald's consumption TFP, the one used by Chen and Wemy and the most recent one used in this paper. It is clear that the two series behave differently in the long run, with the earlier vintage failing to capture the 1970s slowdown in neutral technology growth apparent from the recent estimate.[15]

To evaluate the impact of this revision in consumption TFP on Chen and Wemy's results, I perform a simple experiment: I reproduce their empirical exercise based on *their* original

[15]Kurmann and Sims (2017) also document large revisions in Fernald's (2014) TFP. However, their analysis focuses on changes in high-frequency properties of quantity TFP caused by a revision in the estimation of factor utilization. Instead, I emphasize a shift in the long-run behavior of consumption TFP, while that of quantity TFP does not change much across vintages.

dataset, in which I simply replace TFP by the recent vintage. The results indicate that this change alone goes a long way toward accounting for the different conclusions: when updating the series, the median correlation between the identified long-run shocks to consumption TFP and the RPI drops from 0.99, as reported in Chen and Wemy, to 0.34. To the extent that more recent vintages are improved, this estimate is more reliable and signals once more that long-run shocks to neutral and investment-specific technology are much less correlated than suggested by Chen and Wemy (2015).

5. Conclusion

Results from both parametric and non-parametric methods indicate that, when controlling for structural breaks, the relative price of investment exhibits negative long-run covariation with total factor productivity constructed from chain-aggregated output, but appears orthogonal to consumption TFP. These empirical facts are consistent with a simple two-sector model in which neutral and investment-specific technology shocks are orthogonal. On the other hand, they cannot be explained by models emphasizing correlated technology shocks.

This conclusion is important for two reasons. First, it shows that the long-run comovements between TFP and the RPI, which have sparked interest in the recent years, can be explained without questioning the independence of neutral and investment-specific technology. This is reassuring for a host of studies assuming orthogonal technology shocks, whose conclusions would have been called into question otherwise. Second, the results make it clear that interpreting Fernald's quantity TFP in the context of multi-sector models requires caution, as TFP constructed from chain-aggregated output responds to both neutral and investment-specific technology. Clarifying this point for a wide audience of applied macroeconomists is an important contribution of this paper.

References

Bai, J., R. L. Lumsdaine, and J. H. Stock (1998): "Testing For and Dating Common Breaks in Multivariate Time Series," *Review of Economic Studies*, 65(3), 395–432.

Bai, J., and P. Perron (1998): "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 66(1), 47–78.

——— (2003): "Computation and Analysis of Multiple Structural Change Models," *Journal of Applied Econometrics*, 18(1), 1–22.

Basu, S., and P. De Leo (2016): "Should Central Banks Target Investment Prices?," Boston College Working Papers in Economics 910, Boston College Department of Economics.

Benati, L. (2014): "Do TFP and the Relative Price of Investment Share a Common I(1) Component?," *Journal of Economic Dynamics and Control*, 45(C), 239–261.

Chen, K., and E. Wemy (2015): "Investment-Specific Technological Changes: The Source of Long-Run TFP Fluctuations," *European Economic Review*, 80(C), 230–252.

Cummins, J. G., and G. L. Violante (2002): "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5(2), 243–284.

Fernald, J. G. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fisher, J. D. M. (2006): "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, 114(3), 413–451.

Francis, N., M. T. Owyang, J. E. Roush, and R. DiCecio (2014): "A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks," *The Review of Economics and Statistics*, 96(4), 638–647.

Gordon, R. J. (1990): *The Measurement of Durable Goods Prices*, NBER Books. National Bureau of Economic Research, Inc.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87(3), 342–62.

——— (2000): "The Role of Investment-Specific Technological Change in the Business Cycle," *European Economic Review*, 44(1), 91–115.

Greenwood, J., and P. Krusell (2007): "Growth Accounting with Investment-Specific Technological Progress: A Discussion of Two Approaches," *Journal of Monetary Economics*, 54(4), 1300–1310.

Greenwood, J., and M. Yorukoglu (1997): "1974," *Carnegie-Rochester Conference Series on Public Policy*, 46(1), 49–95.

Gregory, A. W., and B. E. Hansen (1996a): "Residual-Based Tests for Cointegration in Models with Regime Shifts," *Journal of Econometrics*, 70(1), 99–126.

——— (1996b): “Tests for Cointegration in Models with Regime and Trend Shifts,” *Oxford Bulletin of Economics and Statistics*, 58(3), 555–560.

Guerrieri, L., D. Henderson, and J. Kim (2020): “Interpreting Shocks to the Relative Price of Investment with a Two-Sector Model,” *Journal of Applied Econometrics*, 35(1), 82–98.

Hornstein, A., and P. Krusell (1996): “Can Technology Improvements Cause Productivity Slowdowns?,” in *NBER Macroeconomics Annual 1996, Volume 11*, NBER Chapters, pp. 209–276. National Bureau of Economic Research, Inc.

Justiniano, A., G. Primiceri, and A. Tambalotti (2010): “Investment Shocks and Business Cycles,” *Journal of Monetary Economics*, 57(2), 132–145.

——— (2011): “Investment Shocks and the Relative Price of Investment,” *Review of Economic Dynamics*, 14(1), 101–121.

Kim, D., and P. Perron (2009): “Unit Root Tests Allowing for a Break in the Trend Function at an Unknown Time under both the Null and Alternative Hypotheses,” *Journal of Econometrics*, 148(1), 1–13.

Kurmann, A., and E. Sims (2017): “Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks,” NBER Working Papers 23142, National Bureau of Economic Research, Inc.

Liu, Z., D. F. Waggoner, and T. Zha (2011): “Sources of Macroeconomic Fluctuations: A Regime-Switching DSGE Approach,” *Quantitative Economics*, 2(2), 251–301.

Moura, A. (2018): “Investment Shocks, Sticky Prices, and the Endogenous Relative Price of Investment,” *Review of Economic Dynamics*, 27, 48–63.

——— (2020): “Total Factor Productivity and the Measurement of Neutral Technology,” BCL working papers 143, Banque centrale du Luxembourg.

Müller, U. K., and M. W. Watson (2018): “Long-Run Covariability,” *Econometrica*, 86(3), 775–804.

——— (2019): “Low-Frequency Analysis of Economic Time Series,” Draft chapter for the *Handbook of Econometrics*, Vol. 7.

Oulton, N. (2007): “Investment-Specific Technological Change and Growth Accounting,” *Journal of Monetary Economics*, 54(4), 1290–1299.

Perron, P. (1989): “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis,” *Econometrica*, 57(6), 1361–1401.

Perron, P., and T. Yabu (2009): “Testing for Shifts in Trend With an Integrated or Stationary Noise Component,” *Journal of Business & Economic Statistics*, 27(3), 369–396.

Qu, Z., and P. Perron (2007): “Estimating and Testing Structural Changes in Multivariate Regressions,” *Econometrica*, 75(2), 459–502.

Schmitt-Grohé, S., and M. Uribe (2011): "Business Cycles With A Common Trend in Neutral and Investment-Specific Productivity," *Review of Economic Dynamics*, 14(1), 122–135.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

Uhlig, H. (2004): "Do Technology Shocks Lead to a Fall in Total Hours Worked?," *Journal of the European Economic Association*, 2(2-3), 361–371.

Whelan, K. (2003): "A Two-Sector Approach to Modeling U.S. NIPA Data," *Journal of Money, Credit and Banking*, 35(4), 627–656.

Appendix A. Statistical Appendix

This appendix updates Benati's (2014) study of trend breaks in TFP and the RPI. The results overwhelmingly support the presence of breaks, but evidence is mixed as to whether these are common or idiosyncratic. Then, the appendix applies robust unit-root and cointegration tests. The results suggest that TFP and the RPI are $I(1)$ and do not cointegrate.

A.1. **Testing for trend breaks: Bai-Perron and Bai-Lumsdaine-Stock tests.** The analysis starts by documenting the presence of trend breaks in the logarithms of TFP and the RPI. I consider two tests: Bai and Perron's (1998) test for multiple breaks at unknown dates in the mean of a univariate stationary series and Qu and Perron's (2007) test for a common break in the mean of a multivariate stationary process.[16]

A.1.1. *Bai-Perron tests.* Table 3 reports results from Bai and Perron's (1998) tests for multiple breaks at unknown dates in the mean for the log-differences of TFP and the RPI. Denoting the log-difference of interest by y_t, the estimated model with m breaks is

$$y_t = \delta_j + u_t, \qquad u_t \sim iid(0, \sigma_j^2), \qquad t = T_{j-1} + 1, \ldots, T_j, \tag{9}$$

where $j = 1, \ldots, m+1$ indexes the regimes and $T_1, \ldots, T_m$ are the break points (with the convention that $T_0 = 0$ and $T_{m+1} = T$). The variance of the disturbance needs not be constant across regimes.

I implement the tests following Bai and Perron's recommendations; in particular, I set the trimming parameter at 15%. The only differences with their setup are that I use bootstrap procedures based on 1,000 replications to compute p-values and to construct confidence intervals for the estimated break dates. Specifically, I apply the test to 1,000 replications of equation (9) fitted for the actual estimated break date(s) and use the bootstrapped quantiles of the estimated break date(s) to form confidence intervals. For completeness, the table also reports asymptotic critical values. For each series, I start by performing double-maximum tests checking the presence of at least one break, allowing for a maximum of $m = 2$ breaks. To determine the number of breaks, I then compute the $\sup F(1|0)$ and $\sup F(2|1)$ statistics testing the presence of one vs. no break and two vs. one breaks. All test statistics are constructed using Newey-West standard errors. For all series, the double-maximum test statistics are equal to the $\sup F(1|0)$ statistics, so that I report only the latter.

As shown in Table 3, the $\sup F(1|0)$ tests (and the double-maximum tests) find significant breaks in the means of the log-differences of both TFP constructed from output in consumption units and the RPI. On the other hand, the $\sup F(2|1)$ test statistics are never

[16] Following Benati (2014), I also considered Perron and Yabu's (2009) test for a single break at an unknown date in the trend function for the logarithms of TFP and the RPI and Bai, Lumsdaine, and Stock's (1998) test for a common break in the intercept of a bivariate VAR model for the log-differences of TFP and the RPI. These tests also strongly support the presence of trend breaks.

Table 3. Bai-Perron tests of multiple breaks at unknown dates in the mean of the log-difference.

Variable	sup$F(1\|0)$	sup$F(2\|1)$	Break date [90% CI]
Log-difference of TFP			
Quantity (Fernald)	17.9 (0.00)	2.17 (0.64)	1973Q1 [61Q4-90Q2]
Consumption (Fernald)	39.9 (0.00)	3.56 (0.42)	1973Q1 [65Q3-81Q2]
Consumption (Moura)	33.7 (0.00)	2.82 (0.64)	1973Q1 [66Q2-83Q1]
Log-difference of the RPI			
NIPA	31.8 (0.00)	1.62 (0.86)	1982Q2 [74Q3-87Q1]

Notes. The sample is 1950Q1-2019Q4. The model allows for multiple breaks at unknown dates in the mean of the series and is estimated with 15% trimming. The sup$F(1|0)$ and sup$F(2|1)$ tests respectively test the presence of one vs. zero break and two vs. one breaks. Parentheses report bootstrapped p-values based on 1,000 replications. Asymptotic critical values at the 10%, 5%, and 1% levels are 8.02, 9.63, and 13.58 for the sup$F(1|0)$ test, and 9.56, 11.14, and 15.03 for the sup$F(2|1)$ test. Confidence intervals for the break dates are bootstrapped estimates based on 1,000 replications.

significant, signaling that a single break in each series is enough. The estimated break dates are 1973Q1 for all TFP series and 1982Q2 for the RPI, but the confidence intervals are wide. In particular, the confidence bands for the break dates in TFP and the RPI overlap significantly, making it possible that the two series experienced a common break.

A.1.2. *Qu-Perron tests.* Table 4 reports results from Qu and Perron's (2007) tests for a single common break at an unknown date in the mean for the bivariate vector containing the log-differences of TFP and the RPI. Denoting the vector of interest by $\mathbf{y}_t$, the estimated model is

$$\mathbf{y}_t = \delta_j + \mathbf{u}_t, \qquad \mathbf{u}_t \sim \mathcal{N}(0, \Sigma_j), \qquad t = T_{j-1}+1, \ldots, T_j, \tag{10}$$

where $j = 1, 2$ indexes the regimes and T_1 is the break point (with the convention that $T_0 = 0$ and $T_2 = T$). The variance of the disturbance needs not be constant across regimes.

I implement the tests following Qu and Perron's recommendations; in particular, I set the trimming parameter at 15%. As before, I use bootstrap procedures based on 1,000 replications to compute p-values and to construct confidence intervals for the estimated break dates, but I also report asymptotic critical values. The sup-Wald test statistics are constructed using Newey-West standard errors.

As shown in Table 4, the tests find strong evidence for the notion that the log-differences of TFP and the RPI experienced a common break. The estimated break data, 1982Q2, is the same irrespective of the TFP series considered. It is also the same as that obtained when considering the RPI in isolation (see the previous section). However, interpretation is complicated by the fact that the test might reject the null when a single series presents a break large enough to justify splitting the sample for the bivariate process.

Table 4. Qu-Perron tests of a common break in the mean.

Variables			
Log-diff. of TFP	Log-diff. of the RPI	exp-Wald	Break date [90% CI]
Quantity (Fernald)	NIPA	54.5 (0.00)	1982Q2 [80Q1-93Q1]
Consumption (Fernald)	NIPA	58.8 (0.00)	1982Q2 [80Q1-90Q2]
Consumption (Moura)	NIPA	49.0 (0.00)	1982Q2 [78Q4-91Q3]

Notes. The sample is 1950Q1-2019Q4. The model allows for a single break at an unknown date in the mean of the vector and is estimated with 15% trimming. The null hypothesis for the sup-Wald statistic is that of no break. Parentheses report bootstrapped p-values based on 1,000 replications. Asymptotic critical values at the 10%, 5%, 1% levels are 7.04, 8.58, and 12.29. The confidence interval for the break date is a bootstrapped estimate based on 1,000 replications.

A.1.3. *Summing up.* There is strong empirical support for the idea that the logarithms of TFP and the RPI experienced trend breaks in the postwar period. According to Bai-Perron tests considering each series in isolation, the estimated break dates are 1973 for TFP and 1982 for the RPI. According to Qu-Perron tests considering a bivariate process for TFP and the RPI, the estimated break date is 1982. An important takeaway is that trend breaks, if ignored, are likely to bias inference about the long-run properties of TFP and the RPI.

A.2. **Testing for non-stationarity: Kim-Perron tests.** Table 5 reports results from Kim and Perron's (2009) unit-root tests allowing for a trend break at an unknown date under both the null and alternative hypothesis. Denoting the variable of interest by y_t, the estimated model is

$$y_t = \mathbf{x}_t'\boldsymbol{\beta} + u_t, \qquad A(L)u_t = B(L)\epsilon_t, \qquad \epsilon_t \sim iid(0, \sigma^2), \tag{11}$$

where $\mathbf{x}_t' = [1, t, d_t]$ and $d_t = I(t > \tau) \cdot (t - \tau)$, with τ denoting the break date. $A(L)$ and $B(L)$ are two polynomials in the lag operator L of orders $p+1$ and q. $A(L)$ can be factored as $A(L) = (1 - \alpha L)A^\star(L)$, with both $A^\star(L)$ and $B(L)$ stationary. The null hypothesis is: $H_0 : \alpha = 1$.

I implement the tests following Kim and Perron's recommendations. Taking the above results as indicative of trend breaks in the logarithms of TFP and the RPI, I proceed to estimate the break dates using the 'dynamic' regression

$$y_t = \gamma_0 y_{t-1} + \gamma_1 I(t = T_b + 1) + \gamma_2 I(t \geq T_b) + \mathbf{x}_t'\boldsymbol{\beta} + \widetilde{u}_t,$$

estimated by OLS for each possible break date T_b within the interval $[0.15T] - [0.85T]$. The estimated break date is the quarter associated with the minimal sum of squared residuals. As explained in Kim and Perron, this estimate of the break date converges faster than that obtained from the original regression (11). Conditional on the estimated break date $\widehat{\tau}$, I construct the associated regressors $\mathbf{x}_t$ and detrend y_t to obtain the residual $\widehat{u}_t = y_t - \mathbf{x}_t'\widehat{\boldsymbol{\beta}}$.

Table 5. Kim-Perron unit root tests allowing for breaks.

Variable	Lag order $p = 1$	$p = 2$	$p = 4$
Logarithm of TFP			
Quantity (Fernald)	-2.71 (0.52)	-2.73 (0.49)	-2.73 (0.48)
Consumption (Fernald)	-3.23 (0.36)	-3.42 (0.33)	-3.18 (0.41)
Consumption (Moura)	-2.97 (0.37)	-3.00 (0.34)	-2.81 (0.44)
Logarithm of the RPI			
NIPA	-1.83 (0.85)	-1.56 (0.94)	-1.68 (0.90)

Notes. The sample is 1950Q1-2019Q4. The model has an intercept and a time trend and allows for a break at an unknown date in the trend function under both the null hypothesis (unit root) and the alternative (stationarity). Parentheses report bootstrapped p-values based on 1,000 replications. Asymptotic critical values depend on the estimated break date; see Kim and Perron (2009).

Finally, I implement ADF-like tests on $\widehat{u}_t$, considering three possible lag orders for the test: $p = 1$, 2, and 4. Just as for the other tests, I provide both asymptotic critical values and bootstrapped p-values based on 1,000 replications under H_0.

The results reported in Table 5 indicate that it is not possible to reject the null hypothesis of a unit root in either TFP or the RPI when controlling for structural breaks. Standard ADF tests also fail to reject the presence of unit roots in the series, confirming the notion that both TFP and the RPI should be considered as $I(1)$.

A.3. **Testing for cointegration: Gregory-Hansen tests.** Table 6 reports results from the cointegration test proposed by Gregory and Hansen (1996a,b), which allows for a break at an unknown date in the intercept, the trend function, and the cointegrating vector. Denoting the two $I(1)$ variables of interest by $y_{1,t}$ and $y_{2,t}$, the estimated model is

$$y_{1,t} = \mu_1 + \beta_1 t + \alpha_1 y_{2,t} + I(t > \tau)\left(\mu_2 + \beta_2 t + \alpha_2 y_{2,t}\right) + e_t, \tag{12}$$

where τ is the break date and e_t the residual. In the empirical application, $y_{1,t}$ is the logarithm of TFP and $y_{2,t}$ is the logarithm of the RPI. The null hypothesis is that of no cointegration, that is: $H_0 : e_t \sim I(1)$.

I implement the tests following Gregory and Hansen's recommendations. For each possible break date within the interval $[0.15T] - [0.85T]$, I estimate equation (12) by OLS to obtain the residual $\widehat{e}_t$. Then, I compute the augmented Dickey-Fuller statistic $ADF(\tau)$, as well as the two Phillips statistics $Z_t(\tau)$ and $Z_\alpha(\tau)$. As for the earlier ADF tests, I consider three different values for the lag order, $p = 1$, 2, and 4. Finally, I obtain the three test statistics, ADF, Z_t, and Z_α, defined as the minimum of each test's sequence. As usual, I report both asymptotic critical values and bootstrapped p-values based on 1,000 replications under H_0.

Table 6. Gregory-Hansen cointegration tests allowing for breaks.

Variables		Statistics		
Log-diff. of TFP	Log-diff. of the RPI	ADF	Z_t	Z_α
Quantity (Fernald)	NIPA	-3.99 (0.71)	-4.05 (0.72)	-30.9 (0.73)
Consumption (Fernald)	NIPA	-3.76 (0.82)	-3.75 (0.80)	-26.5 (0.82)
Consumption (Moura)	NIPA	-4.62 (0.35)	-4.60 (0.38)	-40.4 (0.35)

Notes. The sample is 1950Q1-2019Q4. The model allows for a break at an unknown date in the intercept, the trend function, and the cointegrating vector. The null hypothesis is that of no cointegration between the logarithms of the RPI and TFP. ADF is the maximal statistic (in absolute value) among those obtained for lag orders equal to 1, 2, 4, and 6. Parentheses report bootstrapped p-values based on 1,000 replications. Asymptotic critical values at the 10%, 5%, and 1% levels are -5.24, -5.50, and -6.02 for the ADF and Z_t statistics, and -53.3, -58.6, and -69.4 for the Z_α statistics.

The results in Table 6 provide no evidence against the notion that TFP and the RPI are not cointegrated when controlling for trend breaks: all statistics are well above their asymptotic critical levels and the bootstrapped p-values confirm the lack of rejection.

A.4. **Conclusion.** Three conclusions emerge from the empirical analysis:

- There is substantial empirical evidence that both TFP and the RPI experienced a trend break in postwar data. Evidence is mixed as to whether the breaks were idiosyncratic or common.
- There is no evidence against the notion that TFP and the RPI contain a unit root, irrespective of the presence of breaks.
- There is no support for the idea that TFP and the RPI are cointegrated when controlling for trend breaks. This does not rule out the possibility that TFP and the RPI share a common $I(1)$ component.

Appendix B. Robustness Analysis

This appendix gathers robustness checks for the empirical analyses conducted in the paper.

B.1. **Long-run correlation: Periods above 15 years.** The results discussed in Section 3 characterized the long-run relationship between TFP and the RPI based on projections capturing periods longer than 11 years. To demonstrate the robustness of the findings, this appendix presents estimates based on projections capturing periods above 15 years. This cutoff date cleanses part of what is sometimes called the medium run, at the cost of limiting further the amount of information that can be extracted from the data.

Table 7 presents the results. The three properties highlighted in Section 3 remain valid when focusing on longer periods: (i) all TFP series exhibit *positive* long-run comovements with the RPI when structural breaks are ignored; (ii) quantity TFP shows moderate *negative*

Table 7. Long-run comovements between TFP and the RPI — Periods longer than 15 years.

Specification of break(s)		TFP series		
		Quantity	Consumption	
			Fernald	Moura
No break	$\widehat{\rho}_T$	0.41 [0.03, 0.68]	0.70 [0.43, 0.85]	0.65 [0.36, 0.82]
	$\widehat{\beta}_T$	0.46 [0.08, 0.85]	0.67 [0.42, 0.92]	0.60 [0.34, 0.86]
Idiosyncratic (TFP)	$\widehat{\rho}_T$	−0.22 [-0.55, 0.17]	0.08 [-0.31, 0.44]	0.05 [-0.33, 0.42]
	$\widehat{\beta}_T$	−0.39 [-1.02, 0.24]	0.19 [-0.69, 1.07]	0.10 [-0.61, 0.81]
Idiosyncratic (TFP, RPI)	$\widehat{\rho}_T$	−0.18 [-0.52, 0.22]	0.20 [-0.20, 0.54]	0.17 [-0.22, 0.52]
	$\widehat{\beta}_T$	−0.14 [-0.44, 0.15]	0.22 [-0.18, 0.62]	0.16 [-0.17, 0.48]
Common	$\widehat{\rho}_T$	−0.17 [-0.51, 0.22]	0.08 [-0.31, 0.44]	0.08 [-0.30, 0.45]
	$\widehat{\beta}_T$	−0.11 [-0.33, 0.12]	0.05 [-0.20, 0.30]	0.05 [-0.17, 0.27]

Notes. The sample is 1950Q1-2019Q4 and the procedure exploits comovements at periods longer than 15 years. Entries in the 'Quantity' column describe the long-run relationship between Fernald's quantity TFP and the RPI, while entries in the subsequent columns describe the long-run relationship between the two measures of consumption TFP and the RPI. $\widehat{\rho}_T$ is the estimated long-run correlation coefficient, $\widehat{\beta}_T$ is the estimated coefficient in the linear long-run regression of TFP on the RPI, and brackets report 67% confidence sets. See Appendix A for details on trend breaks.

long-run comovements with the RPI when controlling for structural breaks; (iii) consumption TFP exhibits *weakly positive* comovement with the RPI when controlling for trend breaks. Therefore, the main findings about the long-run relationship between TFP and the RPI are robust to the definition of the low-frequency band used to construct the projections.

B.2. **SVAR results: Subsample analysis.** Figures 5 and 6 report the subsample estimates from the VAR including consumption TFP.

Over the first subsample, the identified long-run shocks to TFP and the RPI are positively correlated (the correlation is 0.56, see Table 2). This strong link between the two shocks can be seen from Figure 5, since the shocks trigger similar responses from most variables in the VAR. However, there remains some important differences: for instance, in the short run, both TFP and the RPI are more responsive to their own shock than to the other one.

Over the second, more recent subsample, the identified shocks are completely decorrelated (the correlation is -0.01). Figure 6 shows that the estimated responses to the shocks are very different for all variables, making it clear that consumption TFP and the RPI are driven by totally different shocks in recent data.

FIGURE 5. Impulse-responses to the long-run shocks to TFP and the RPI — VAR with consumption TFP, 1950-1972 subsample.

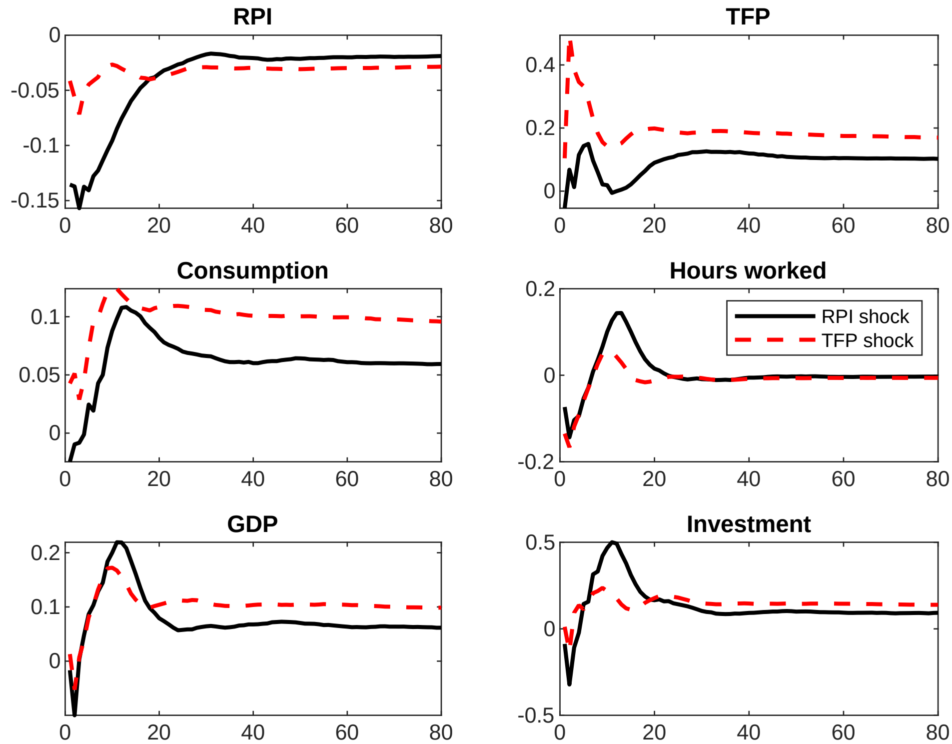

Notes. IRFs to the shocks that contribute most to the FEV of the RPI (solid black line) and TFP (dashed red line) at the forecast horizon of 80 quarters. The reported responses are the median from the bootstrapped distributions with 1,000 replications. The VAR includes Fernald's (2014) consumption TFP as observable and is estimated on the subsample ranging from 1950 to 1972.

FIGURE 6. Impulse-responses to the long-run shocks to TFP and the RPI — VAR with consumption TFP, 1983-2019 subsample.

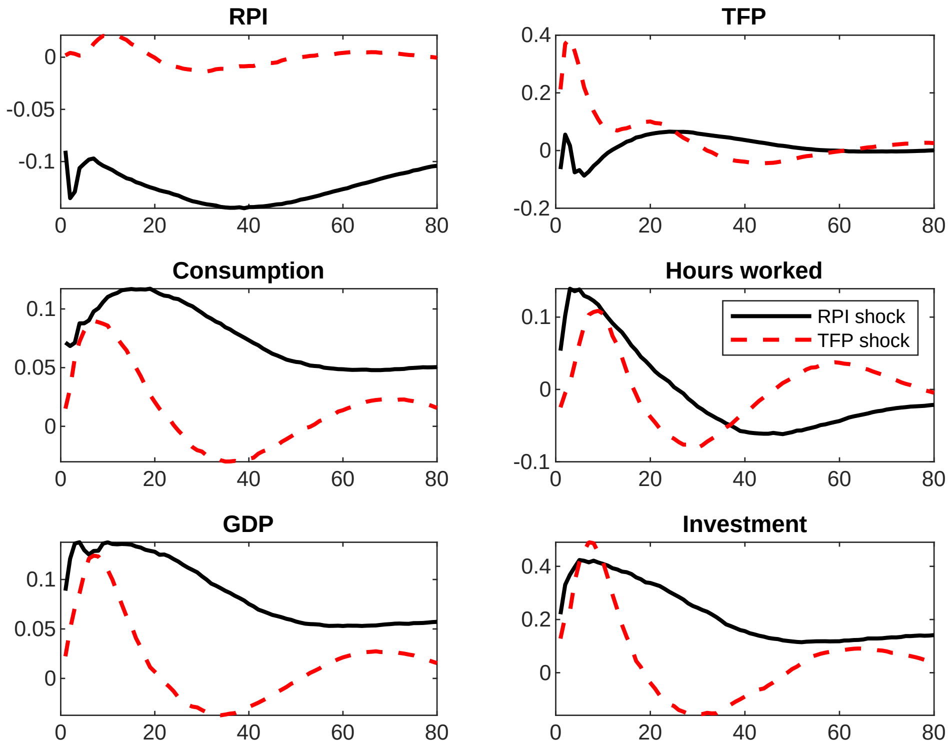

Notes. IRFs to the shocks that contribute most to the FEV of the RPI (solid black line) and TFP (dashed red line) at the forecast horizon of 80 quarters. The reported responses are the median from the bootstrapped distributions with 1,000 replications. The VAR includes Fernald's (2014) consumption TFP as observable and is estimated on the subsample ranging from 1983 to 2019.

BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTÈME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu