

CAHIER D'ÉTUDES WORKING PAPER

N° 143

TOTAL FACTOR PRODUCTIVITY AND THE MEASUREMENT OF NEUTRAL TECHNOLOGY

ALBAN MOURA

MAY 2020



BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTÈME

TOTAL FACTOR PRODUCTIVITY AND THE MEASUREMENT OF NEUTRAL TECHNOLOGY

ALBAN MOURA

ABSTRACT. TFP measures constructed from chain-aggregated output, such as those published by the Bureau of Labor Statistics or Fernald (2014), confound contributions from neutral and sector-specific technology. Therefore, they should not be used to infer the path of neutral technology in presence of investment-specific technical change. Two theory-consistent, utilization-adjusted measures of neutral technology at the quarterly frequency are proposed for the US business sector. Both indicate that neutral technology progress declined dramatically after the mid-1970s. In particular, its contribution to US growth fell from more than 85% before 1973 to less than 25% afterward. The associated welfare loss is enormous: if neutral technology had continued on its pre-1970s trend, 2017 US output would have been 70% higher.

JEL Codes: E22, E23, E32, O41, O47.

Keywords: total factor productivity, neutral technology, investment-specific technology, sources of growth.

May 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. For helpful comments, I thank Patrick Fève, Pablo Garcia Sanchez, Paolo Guarda, and Olivier Pierrard. This paper should not be reported as representing the views of the Banque centrale du Luxembourg or the Eurosystem. The views expressed are those of the authors and may not be shared by other research staff or policymakers at BCL or the Eurosystem.

RÉSUMÉ NON TECHNIQUE

Suite aux contributions de Solow (1956) et Kydland et Prescott (1982), le changement technologique se trouve au cœur des théories macroéconomiques des cycles et de la croissance. Tandis que ces études initiales considéraient seulement un changement technologique “neutre” dans une fonction de production agrégée¹, désigné sous le terme de Productivité Totale des Facteurs (TFP, en anglais), la mise en évidence par Gordon (1990) de la tendance baissière dans le prix relatif des biens d’investissement a attiré l’intérêt sur une seconde source de changement technologique, spécifique à la production des biens d’investissement. Depuis, la littérature macroéconomique débat du rôle relatif du changement technologique neutre et du changement technologique spécifique au secteur de l’investissement, à la fois quant à leur contribution aux cycles et à la croissance économique.

Une difficulté d’ordre pratique pour conclure le débat provient de ce qu’aucun des deux types de technologie n’est directement observable. Par conséquent, l’étude du rôle des deux types de changement technologique s’est longtemps basée sur des restrictions d’ordre théorique. Ainsi, l’on a pu tenter de déduire l’évolution de la technologie neutre et de celle spécifique au secteur de l’investissement à partir des caractéristiques de variables tirées des comptes nationaux, telles le Produit Intérieur Brut et le déflateur des biens d’investissement.

Récemment, certains auteurs ont tenté de discipliner l’inférence en interprétant des mesures de TFP comme des observations directes de la technologie neutre. Cette nouvelle stratégie est rendue possible par la construction de séries de TFP ajustées de l’utilisation des facteurs à fréquence trimestrielle pour l’économie des Etats-Unis, notamment par Fernald (2014). D’après Benati (2014), la série de Fernald est ainsi “largement reconnue comme la meilleure mesure disponible du progrès technologique neutre”.

Dans ce contexte, cet article met en avant deux points. Le premier est qu’il est incorrect d’interpréter les séries de TFP, construites par Fernald et d’autres, comme des observations directes de la technologie neutre lorsque l’économie est aussi affectée par un changement technologique spécifique au secteur de l’investissement. Au contraire, la section 2 de l’article démontre que, dans ce cas, la TFP mesurée agrège et confond la technologie neutre et celle spécifique au secteur de l’investissement. Cet argument soulève de sérieux doutes quant à la validité des conclusions obtenues par les approches identifiant la technologie neutre aux mesures usuelles de TFP.

Le second point est qu’il est immédiat de construire des séries de TFP qui peuvent être correctement interprétées comme reflétant la technologie neutre, y compris en présence de changement technologique spécifique au secteur de l’investissement. La section 3 discute ainsi la construction et les propriétés de deux telles séries pour l’économie américaine.

¹“Neutre” signifie ici que le progrès technologique améliore de manière symétrique la capacité de l’économie de produire tous types de biens. Par opposition, le progrès technologique spécifique à la production des biens d’investissement affecte uniquement la capacité de l’économie de produire ces derniers.

Bien qu'équivalentes en théorie, ces deux mesures de la technologie neutre sont légèrement différentes en pratique. Néanmoins, elles amènent aux mêmes conclusions quant à la nature et à la vitesse du changement technologique aux Etats-Unis. En particulier, les deux séries indiquent un déclin dramatique dans l'intensité du progrès technologique neutre à partir du milieu des années 1970. Il s'agit d'un changement significatif sur le plan statistique : avant 1973, la technologie neutre progressait d'environ 2% par an en moyenne, pour tomber ensuite à un rythme de progression compris entre 0,05% et 0,2% par an seulement. Le changement est également significatif au plan économique : le progrès technologique neutre expliquait plus de 85% de la croissance américaine avant 1973, et moins de 25% par la suite. De plus, si le progrès technique neutre avait continué sur sa tendance initiale après 1973, le PIB américain serait aujourd'hui plus élevé d'environ 70%.

1. INTRODUCTION

Since Solow (1956) and Kydland and Prescott (1982), technical change has been at the heart of macroeconomic theories of growth and business cycles. While early works typically focused on a single neutral technology shifter in an aggregate production function, referred to as Total Factor Productivity (TFP), Gordon’s (1990) finding of a sustained downward trend in quality-adjusted relative prices for durable goods sparked interest in a second source of technical change related to the production of investment goods. Since then, the literature has been debating the relative roles of neutral and investment-specific technology progress, both as regards the sources of economic growth (Greenwood, Hercowitz, and Krusell, 1997; Whelan, 2003; Greenwood and Krusell, 2007) and business cycles (Greenwood, Hercowitz, and Krusell, 2000; Fisher, 2006; Justiniano, Primiceri, and Tambalotti, 2010, 2011).

One practical difficulty to settle the debate is that neither source of technology is directly observable, so that inference about technical change often relies on theoretical restrictions. For instance, Greenwood, Hercowitz, and Krusell (1997) and Whelan (2003) estimate the respective US growth contributions of neutral and investment-specific technology from the behavior of aggregate national account variables, interpreted through the lens of general-equilibrium models.

More recently, some authors have tried to discipline inference by interpreting empirical TFP measures as direct observations of neutral technology. For instance, Beaudry and Lucke (2010) incorporate measured TFP into a VECM analysis focusing on the sources of US business cycles and assume that it responds to disembodied (neutral) technology shocks only. Similarly, Schmitt-Grohé and Uribe (2011, 2012) map TFP into neutral technology in their empirical analysis of real-business-cycle models. This strategy has been made possible by the construction of quarterly utilization-adjusted TFP series by Beaudry and Lucke (2010) and Fernald (2014). According to Benati (2014), the latter “is widely regarded as the best available measure of neutral technological progress.”

In this context, this paper makes two points. The first is that one should not interpret standard TFP series, such as those published by Beaudry and Lucke or Fernald, but also by the Bureau of Labor Statistics (BLS, 2007), as observations on neutral technology in the presence of investment-specific technical change. Instead, I show in Section 2 that in this case measured TFP is a geometric combination of neutral and investment-specific technology processes. It follows immediately that empirical studies mapping measured TFP into neutral technology in multi-sector setups, such as the ones cited above, face potentially serious identification issues. This finding also rationalizes the long-run comovements between measured TFP and the relative price of investment found in US data (Schmitt-Grohé and Uribe, 2011; Benati, 2014), which could originate from their common dependence on investment-specific technology.

The root of the problem lies in the ambiguous nature of output measurement in multi-sector environments. Since Greenwood, Hercowitz, and Krusell (1997), it has become standard to work with one-sector representations of multi-sector models, in which output is measured in consumption units and associated with an aggregate production function shifted by neutral technology. Accordingly, neutral technology should be measured as the ratio between output in consumption units and the input contributions implied by the production function. However, Beaudry and Lucke (2010), Fernald (2014), and the BLS all construct TFP from chain-aggregated output. This measure of production closely resembles a geometric average of real quantities and it grows faster than output in consumption units in presence of investment-specific technological change. It follows immediately that TFP series constructed from chain-aggregated output cannot measure neutral technology if investment-specific technical progress is present. Rather, they correspond to a weighted average of the two processes for neutral and investment-specific technology.

This is not a new result. In particular, the issues related to output measurement and the interpretation of TFP measures are well known in the large literature on aggregate productivity.² However, the examples from Beaudry and Lucke (2010) and Schmitt-Grohé and Uribe (2011, 2012) make it clear that some ambiguity remains for a wider audience of macroeconomists. Another illustration is the survey by Ramey (2016, p. 136), which presents neutral technology shocks as equivalent to TFP shocks.³ This tendency to misinterpret TFP suggests that there is some value in clarifying the matter, even though some authors did use the appropriate interpretation in applied work (for instance Barsky, Basu, and Lee, 2014; Chen and Wemy, 2015).

The second point is that it is straightforward to construct quarterly utilization-adjusted TFP measures that can be interpreted as neutral technology in the presence of investment-specific technical change. In Section 3, I provide two such series for the US economy, which might be of interest for applied researchers. I construct the first from Fernald’s (2014) database, which reports time series on inputs, output, and factor utilization for the US business sector. To obtain a TFP series that corresponds to neutral technology, I simply replace Fernald’s chain-aggregated output by output in consumption units. The second measure is provided directly by Fernald (2014), who decomposes aggregate TFP into distinct consumption- and investment-sector TFP series. I show that Fernald’s TFP for the consumption sector is another theory-consistent estimate of neutral technology in standard two-sector setups. Fernald does not emphasize this interpretation, which is novel.

Although equivalent in theory, these two measures of neutral technology differ slightly in sample. However, they share common implications regarding the nature and pace of

²See, e.g., Hornstein and Krusell (1996, 2000), Fernald (2015), and Byrne, Fernald, and Reinsdorf (2016).

³Ramey’s (2016) exact sentence is: “Neutral technology shocks, or TFP shocks, are shocks to the process driving A_t ,” where A_t is the productivity shifter in an aggregate production function.

technical change in the US economy. In particular, both series indicate a dramatic decline in neutral technology growth starting in the mid 1970s, a period which coincides with a persistent slowdown in US aggregate productivity. Statistical tests for breaks at unknown dates in the sample identify a single significant trend break in the logarithm of both measures of neutral technology, with the same estimated break date: 1973Q1.⁴ The size of the break is economically significant: TFP measured from output in consumption units grew at a yearly rate of 1.9% before 1973, but only 0.04% afterward, while yearly growth in Fernald’s consumption-sector TFP fell from 2.1% to just 0.2%. When embedded in a standard two-sector general-equilibrium model à la Greenwood, Hercowitz, and Krusell (1997), these figures imply that neutral technical progress explained more than 85% of average US growth before 1973, but 25% at most afterward.

These results attribute the 1970s productivity slowdown to the break in neutral technology growth. This is consistent with Hornstein and Krusell (1996) and Greenwood and Yorukoglu (1997), who found that aggregate productivity declined at that time even though investment-specific technical progress apparently accelerated.⁵ This is also consistent with other measures of neutral technology reported in Greenwood, Hercowitz, and Krusell (1997), Marquis and Trehan (2008), or Ireland (2013), which all exhibit slow growth or even decline during the 1970s. Finally, the results put into perspective the swings in aggregate productivity observed during the 1990s-2000s (Fernald, 2015; Byrne, Fernald, and Reinsdorf, 2016; Cette, Fernald, and Mojon, 2016; Syverson, 2017). The welfare loss associated with the slowdown in neutral technology progress is enormous: if technology progress had continued on its pre-1970s trend, 2017 US output would have been higher by about 70%. Relative to this figure, recent movements are of limited magnitude.

Throughout the paper, I focus on real two-sector economies that can be aggregated into a one-sector representation, following Greenwood, Hercowitz, and Krusell (1997) and others. It is well known that such an aggregation requires strong modeling assumptions, but it has become the standard way to think about neutral and investment-specific technical progress. The quarterly measures of neutral technology I propose are valid only for this class of two-sector models; see Basu and Fernald (2002) and Basu, Fernald, and Kimball (2006) for the measurement and interpretation of aggregate productivity in richer multi-sector environments. I also assume that inputs and outputs are correctly measured, abstracting from factor utilization in the theoretical discussion. In the empirical part, I rely on Fernald’s

⁴Fernald (2007) and Marquis and Trehan (2008) also find significant breaks in productivity growth in the early 1970s. Fernald focuses on aggregate TFP, which depends on both neutral and investment-specific technology processes. My results suggest that the break originated from neutral technology.

⁵Both papers propose a theoretical interpretation based on unmeasured learning costs required to master the new investment technology. See also David (1990), Gordon (2012), and Fernald (2016) for the alternative idea that diminishing returns from past inventions contributed to the 1970s productivity slowdown.

(2014) carefully constructed utilization series to obtain TFP measures representing neutral technology.

2. MEASUREMENT IN A TWO-SECTOR MODEL

This section focuses on a simple neoclassical growth model with neutral and investment-specific technical change. The model sheds light on issues related to the measurement and interpretation of output and TFP in aggregate economies with multiple sources of technological progress.

2.1. A stylized model. Consider a stylized stochastic growth model with neutral and investment-specific technological change, adapted from 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 following planning problem:

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

subject to

$$C_t + \frac{I_t}{V_t} \leq A_t K_{t-1}^\alpha H_t^{1-\alpha}, \quad (2)$$

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

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

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

Here, C_t , H_t , I_t , K_t , A_t , and V_t 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 Frisch elasticity of labor supply, α is the capital share, and δ is the depreciation rate. Following Fisher (2006), there is a single capital good, instead of two in Greenwood, Hercowitz, and Krusell (1997, 2000), and neutral and investment-specific technology have stochastic, rather than deterministic, trends. Parameters μ_A and μ_V denote the average growth rates of neutral and investment-specific technology, and ϵ_t^A and ϵ_t^V are the permanent technology shocks.

The discussion below does not require an explicit characterization of the model's general equilibrium. However, it is useful to remark that the following variables are stationary along any balanced growth path, in spite of the presence of stochastic trends:

$$\frac{C_t}{Q_t}, \frac{I_t}{Q_t V_t}, \frac{K_t}{Q_t V_t}, H_t, \quad (4)$$

with $Q_t = A_t^{1/(1-\alpha)} V_t^{\alpha/(1-\alpha)}$.

Although stripped down, this framework constitutes the backbone of many quantitative DSGE models, including those by Smets and Wouters (2007), Justiniano, Primiceri, and

Tambalotti (2010, 2011), or Schmitt-Grohé and Uribe (2011, 2012). Ríos-Rull, Schorfheide, Fuentes-Albero, Kryshko, and Santaaulàlia-Llopis (2012) use a closely related setup to highlight how parameter identification shapes inference about the aggregate effects of technology shocks, while Ireland (2013) considers a two-country extension to study stochastic growth in the US and the euro area. To keep the focus on the measurement issues at the heart of the paper, I abstract from the usual sources of short-run rigidity (consumption habits, investment adjustment costs, costly factor utilization), with no effect on the results.⁶ Assuming trend-stationary rather than integrated technology would also leave the analysis unchanged.

Other omissions are potentially more important. In particular, the two-sector structure of the economy implicit in equation (2) must generate a linear transformation frontier between consumption and investment. While this is a common assumption in the literature about investment-specific technology, Greenwood, Hercowitz, and Krusell (1997) show that a linear transformation curve corresponds to a detailed two-sector model only under strong restrictions, which may not hold in the data. Alternative setups include Guerrieri, Henderson, and Kim (2014, 2016), who propose richer two-sector models more in line with empirical evidence, but in which the distinction between neutral and investment-specific technology is not clear cut.⁷ Introducing sticky prices would also create some difficulties related to the treatment of markups, equilibrium price dispersion, and slow pass-through of technology shocks to market prices (Moura, 2018).

2.2. Output measurement. As discussed in Hornstein and Krusell (2000), Whelan (2003), and Greenwood and Krusell (2007), there are different ways to define, and thus measure, aggregate output in multi-sector models. These accounting differences matter because alternative output measures inherit different properties and require careful interpretation. Unsurprisingly, these measurement issues spread to TFP series computed as residuals from aggregate production functions.

Traditionally, the DSGE literature measures output in consumption units, which is defined as:

$$Y_t := C_t + P_t I_t, \tag{5}$$

where P_t is the relative price of investment goods, that is the price of investment in consumption units. In the simple model from Section 2.1, the equilibrium relative price of investment

⁶As explained in the Introduction, neglecting variable input utilization does not matter for the conceptual discussion. However, in practice it is important to correct for variations in utilization when interpreting observed TFP residuals as technology measures, see Fernald (2014).

⁷Another example is Fisher (2010), who introduces sectoral adjustment costs à la Kim (2003) in an otherwise standard two-sector model. The adjustment costs create curvature in the transformation frontier for producing investment and consumption goods and affect both measured TFP and the RPI in equilibrium.

is the inverse of investment-specific technology:

$$P_t = \frac{1}{V_t}. \quad (6)$$

This relationship has been exploited by a large literature to infer the role of investment-specific technological change in economic growth and business cycles (see Section 1 for a list of references, and Moura (2018) for a criticism of the interpretation of the RPI as a measure of the inverse of investment-specific technology).

Merging equations (2), (5), and (6) yields

$$Y_t = C_t + \frac{I_t}{V_t} = A_t K_{t-1}^\alpha H_t^{1-\alpha}, \quad (7)$$

which directly links Y_t to the aggregate production function $A_t K_{t-1}^\alpha H_t^{1-\alpha}$. This equality explains why DSGE modelers often measure output in consumption units. Another reason is that the trend in Y_t perfectly captures the trend in household utility along a balanced growth path, so that output in consumption units provides useful information about aggregate welfare in the model.⁸ Equation (4) implies that the stochastic trend in Y_t is $Q_t = A_t^{1/(1-\alpha)} V_t^{\alpha/(1-\alpha)}$.

An alternative output measure is the chain-aggregated index of production used to compute real GDP in the US National Income and Product Accounts (NIPAs) since 1996. This measure has received less attention in the DSGE literature, with notable exceptions being Whelan (2003), who maps a two-sector model into NIPA variables, and Edge, Kiley, and Laforge (2008) and Chung, Kiley, and Laforge (2010), who estimate a DSGE model mapping Divisia output into the NIPA real GDP series.⁹ Technically, the NIPAs construct the chain-aggregate for GDP as a Fisher index, i.e. as a geometric average of Paasche and Laspeyres fixed-weighted indexes. As explained in Whelan (2003), in practice the Fisher index for output is well approximated by a Divisia share-weighted index:

$$X_t := C_t^\gamma I_t^{1-\gamma}, \quad (8)$$

where $\gamma = C^*/(C^* + P^*I^*)$ is the steady-state nominal share of consumption in GDP. The approximation supposes that the geometric weights in the Divisia index reflect current GDP shares, so that γ should be time-varying. Here, I use the steady-state values for simplicity

⁸Greenwood and Krusell (2007) emphasize this correspondence between output in consumption units and social welfare. However, Whelan (2003) remarks that welfare generally grows at a different rate when consumers also derive utility from durable goods.

⁹Following Smets and Wouters (2007), several authors have mapped output in consumption units (as a DSGE model variable) into chain-aggregated real GDP (as an observable series used in estimation). Other important papers featuring this approach include Justiniano, Primiceri, and Tambalotti (2010) and Schmitt-Grohé and Uribe (2012). Clearly, this is *not* the same as properly defining chain-aggregated real GDP in the model and mapping it to the appropriate time series from the data.

without affecting significantly any of the results. Below, I neglect this approximation error and indifferently refer to X_t as chain-aggregated or Divisia real GDP.

Contrasting equations (7) and (8) highlights a major difference between X_t and Y_t : while output in consumption units deflates investment by the level of investment-specific technology V_t , chain-aggregated output does not include such a correction. As a result, X_t and Y_t necessarily exhibit different properties in equilibrium. For instance, equation (4) implies that the balanced-growth stochastic trend in X_t is $Q_t V_t^{1-\gamma}$. Since the stochastic trend in Y_t is Q_t , output in consumption units and chain-aggregated output grow at different rates in the presence of investment-specific technological progress ($\mu_V > 0$), with X_t growing faster than Y_t .

2.3. Technology measurement. Consider now the measurement of technology based on growth-accounting procedures. Assuming that factor inputs K_{t-1} and H_t are perfectly observed and that the production function is known, it is straightforward to see that the TFP series constructed from output in consumption units Y_t is exactly the process for neutral technology A_t :

$$TFP_t^Y := \frac{Y_t}{K_{t-1}^\alpha H_t^{1-\alpha}} = A_t. \quad (9)$$

On the other hand, the TFP series constructed from chain-aggregated output X_t is given by

$$TFP_t^X := \frac{X_t}{K_{t-1}^\alpha H_t^{1-\alpha}}. \quad (10)$$

Since $X_t \neq Y_t$, it is clear that $TFP_t^X \neq A_t$. In addition, because X_t and Y_t grow at different rates, TFP_t^X does not even necessarily grow at the same rate as neutral technology. Indeed, simple computations reveal that the stochastic trend in TFP_t^X is

$$\text{trend}(TFP_t^X) = A_t V_t^{1-\gamma}. \quad (11)$$

It follows that TFP series constructed from chain-aggregated output respond to fluctuations in V_t and thus mix contributions from neutral and investment-specific technology in two-sector setups. This is the main result of the paper, stated in the following proposition.

Proposition 1. *In a standard two-sector model with both neutral and investment-specific technology, TFP series constructed from chain-aggregated output do not measure neutral technology and generally grow at a different rate.*

Proposition 1 shows that correctly interpreting a TFP series requires checking the output measure used for the growth-accounting procedure. As explained in the Introduction, the BLS, Beaudry and Lucke (2010), and Fernald (2014) all construct TFP from chain-aggregates for real output. Hence, the resulting TFP series correspond to TFP_t^X and will deviate from neutral technology if investment-specific technology matters. In this case, inference mapping TFP into neutral technology is bound to yield biased results.

2.4. Implications. The theoretical analysis has important implications for applied work about the macroeconomic role of technology shocks. I discuss two simple examples below and devote the next section to a more substantial one.

Inference about technology shocks. As explained in the Introduction, several authors used TFP measures in empirical work on the role of neutral and investment-specific technology shocks in business cycles. Proposition 1 indicates that this strategy is misleading in an environment with investment-specific shocks if a TFP measure based on chain-aggregated output is wrongly interpreted as reflecting only neutral technology. In this case, some movements in measured TFP originate from investment shocks, but the empirical analysis will wrongly attribute them to neutral technology shocks. Consequently, standard exercises are likely to yield biased results, including variance decompositions assessing the respective contributions of the two technology shocks.

Theory also suggests a simple fix, which is to interpret TFP series built from chain-aggregated output according to equation (10) instead of as direct measures of neutral technology. Incorporated in DSGE models, equation (10) captures the differences between A_t and measured TFP in a theory-consistent fashion. It also indicates the appropriate long-run restrictions to identify the effects of neutral and investment-specific technology shocks on TFP, as in structural VARs. For instance, Barsky, Basu, and Lee (2014, p. 248) recognize that, in a two-sector setting, Fernald aggregate TFP corresponds to a weighted average of consumption- and investment-specific TFP series. Chen and Wemy (2015, equation (3)) also acknowledge this property and correctly deduce that “some of the fluctuations in [investment-specific technology] will be identified as fluctuations in aggregate TFP” as a result.

Covariation between TFP and the relative price of investment. Another implication of the theoretical analysis is that, when investment-specific technology has a stochastic trend, TFP_t^X shares a common unit-root process with the relative price of investment P_t . Indeed, comparing equations (6) and (11) shows that the logarithms of TFP_t^X and P_t both respond to $\ln(V_t)$ in the long run. As a result, the common $I(1)$ component corresponds to the stochastic trend in investment-specific technology, which increases TFP_t^X and decreases P_t in the long run. Therefore, theory implies negative long-run covariation between TFP and the relative price of investment.

This relationship provides a simple interpretation for the empirical finding that the TFP series in Beaudry and Lucke (2010) and Fernald (2014) exhibit noticeable long-run comovements with the relative price of investment (Schmitt-Grohé and Uribe, 2011; Benati, 2014). I develop this argument further and show its empirical plausibility in a companion paper (Moura, 2020). For applied work, the main lesson is that TFP measures built from chain-aggregated output and the relative price of investment should not be treated as driven by independent processes in the specification of empirical models.

3. MEASURING NEUTRAL TECHNOLOGY

This section proposes two quarterly theory-consistent measures of neutral technology that solve the identification issue just discussed and can be readily used in applied work. Both series indicate that neutral technology growth in the US experienced a break in the mid-1970s and slowed significantly afterward.

3.1. Two measures of neutral technology. The analysis in Section 2 suggests that TFP built from output in consumption units, TFP_t^Y , corresponds to neutral technology in standard one- and two-sector environments. I show that Fernald’s (2014) dataset makes it straightforward to construct the series as the quarterly frequency. I also discuss an alternative measure of neutral technology, Fernald’s (2014) TFP for the consumption sector.

3.1.1. Constructing TFP_t^Y . Equation (9) implies that the Solow residual of an aggregate production function for output in consumption units is a theory-consistent measure of neutral technology in standard setups. The equation also identifies the elements required to measure TFP_t^Y in the data: first, we need a series for output in consumption units; second, we need a series for the effective capital services and one for labor input, as well as a series on variable factor utilization; third, we need the value of the capital share.

Conveniently, Fernald (2014) provides a database reporting carefully constructed quarterly series for α , $\Delta \ln K$, and $\Delta \ln H$ for the US business sector. Fernald also provides an adjustment factor for endogenous variations in input utilization, which is crucial to convert raw Solow residuals into plausible technology measures. Thus, in order to compute TFP_t^Y from equation (9), and given Fernald’s data, one simply needs a series for output in consumption units.

To ensure consistency with Fernald’s framework, I construct Y_t as business output in consumption units. In practice, I compute the variable as the ratio between nominal business output (from NIPA table 1.3.5) and a consumption price index (obtained by chain-aggregating the price indexes for consumption of non-durable goods and services, from NIPA table 1.1.4). Then, I construct TFP from output in consumption units by removing Fernald’s series for input services and utilization, as in equation (9). According to the theoretical analysis, the resulting series corresponds to neutral technology in the US business sector. It appears in red in Figure 1.

3.1.2. Fernald’s TFP measure for the consumption sector. Fernald (2014, pp. 9-10) proposes a decomposition of aggregate TFP into components related to the consumption and investment sectors. Here, I show that his consumption-sector TFP also corresponds to neutral technology in standard setups.

Fernald bases the decomposition on two equations. The first (p. 9 in Fernald, 2014) equates aggregate TFP with a geometric average of consumption and investment TFP. In

the notation of Section 2, this reads

$$\ln TFP_t^X = \gamma \ln TFP_t^C + (1 - \gamma) \ln TFP_t^I, \quad (12)$$

where γ is the average consumption share of output, and TFP_t^C and TFP_t^I respectively denote TFP in the consumption and investment sectors. The second equation (p. 10 in Fernald, 2014) relates movements in relative prices to changes in relative technology:

$$\ln TFP_t^I = \ln TFP_t^C - \ln \frac{P_t^I}{P_t^C}, \quad (13)$$

where P_t^C and P_t^I are the nominal prices of consumption and investment.

To show the link between consumption-sector TFP and neutral technology, first merge equations (12) and (13), replace TFP_t^X by its expression from equation (10), and rearrange to obtain

$$\ln TFP_t^C = \ln \frac{X_t}{K_{t-1}^\alpha H_t^{1-\alpha}} + (1 - \gamma) \ln \frac{P_t^I}{P_t^C}. \quad (14)$$

Then, use the price index associated with X_t , which verifies $\ln P_t^X + \ln X_t = \ln P_t^C + \ln Y_t$ by construction (both sides represent the log of nominal output since Y_t is measured in consumption units). In addition, the Divisia price index associated with X_t is the share-weighted geometric average of prices, so $\ln P_t^X = \gamma \ln P_t^C + (1 - \gamma) \ln P_t^I$. It follows that

$$\ln X_t = \ln \frac{P_t^C Y_t}{P_t^X} = \ln Y_t - (1 - \gamma) \ln \frac{P_t^I}{P_t^C}.$$

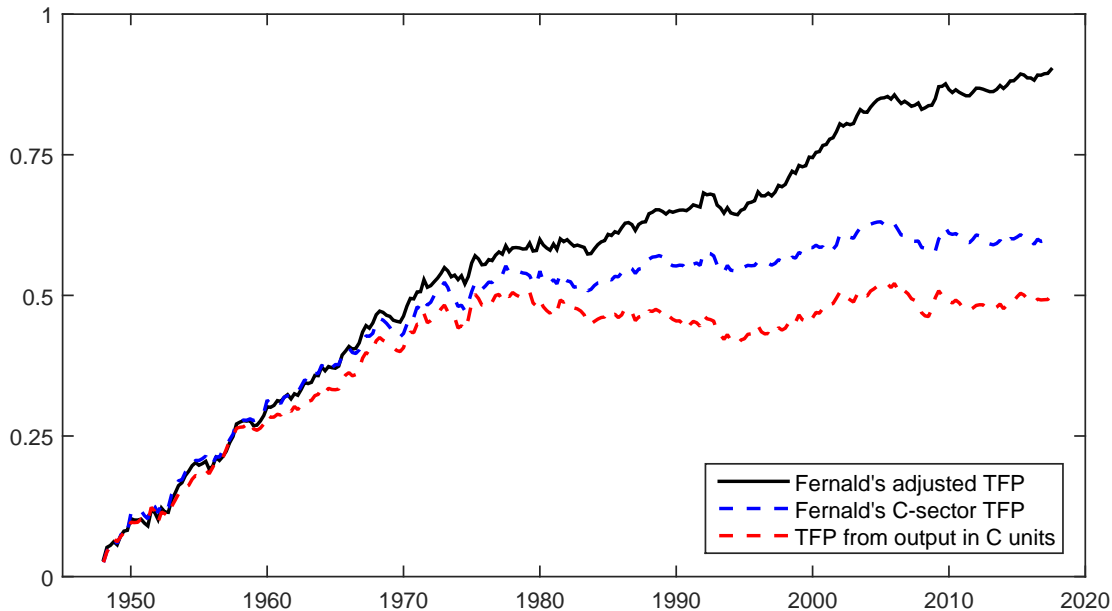
Substituting this expression into equation (14) finally yields

$$\ln TFP_t^C = \ln \frac{Y_t}{K_{t-1}^\alpha H_t^{1-\alpha}} = \ln TFP_t^Y = \ln A_t. \quad (15)$$

This equality proves the equivalence between Fernald's consumption-sector TFP and neutral technology in the model from Section 2.

This result is interesting because it allows us to interpret TFP for the consumption sector as a direct measure of neutral technology. Fernald (2014) does not emphasize this interpretation, which is novel. His original interpretation as consumption-sector technology is obviously equivalent, but the argument formalizes the idea that consumption-sector technology coincides with the standard definition of neutral technology in macroeconomic models. Besides, applied researchers may be interested in knowing they could use TFP_t^C to infer the behavior of A_t . Finally, the result also applies to other measures of consumption-sector TFP discussed in the literature (Ireland and Schuh, 2008; Marquis and Trehan, 2008). In Figure 1, Fernald's consumption-sector TFP is shown in blue.

FIGURE 1. Log TFP Measures



Notes. The black line represents the cumulative sum of Fernald's series for the log-difference of utilization-adjusted TFP for the private business sector; the dashed blue line is the cumulative sum of Fernald's series for the log-difference of utilization-adjusted TFP for the consumption sector; and the dashed red line is the cumulative sum of the log-difference of utilization-adjusted TFP computed from private business output in consumption units (see the text for details).

3.1.3. *Visual representation.* Figure 1 contrasts the two measures of neutral technology with Fernald's (2014) utilization-adjusted TFP series for the US business sector over the 1948-2017 sample. The first column in Table 1 reports the growth rate of each series, as well as further decompositions explained below.

The figure highlights three notable properties of the series. First, the two measures of neutral technology have grown at slightly different rates, especially in the second part of the sample. In particular, Fernald's consumption-sector TFP has grown faster than TFP built from output in consumption units during and after the 1970s. This discrepancy originates from differences in constructing the series, especially as regards the choice of the numéraire. On the one hand, Fernald (2014) defines the price of consumption by deducting an investment price index from the price index for business output. On the other hand, I take the consumption price to be that of consumer spending on non-durable consumption goods and services, in line with the DSGE literature on the relative price of investment (e.g., Fisher, 1999; Justiniano, Primiceri, and Tambalotti, 2011; Moura, 2018). That TFP_t^C exceeds TFP_t^Y during the 1970s and after indicates that my consumption price index has been increasing

faster than Fernald's.¹⁰ Nevertheless, when it comes to the big picture, the two measures of neutral technology display very similar behavior over the postwar period.

Second, the chart shows that since the 1950s neutral technology grew much less rapidly than Fernald's aggregate TFP. Indeed, the statistics in Table 1 indicate that A_t grew between 0.72% and 0.86% a year between 1948 and 2017, depending on whether it is measured from TFP_t^C or TFP_t^Y , whereas TFP_t^X grew nearly 1.30% a year over the same period. Since aggregate TFP combines neutral and investment-specific technology, the difference between Fernald's TFP and the measures of neutral technology originates from investment-specific technical progress. Therefore, Figure 1 conveys visually the significant identification error that arises when aggregate TFP is wrongly interpreted as embodying neutral technology only.

Finally, the figure clearly suggests that the process for neutral technology experienced a break sometime during the 1970s. Indeed, it grew at about the same rate as aggregate TFP up to 1970 but clearly diverged afterward, when growth in neutral technology slowed markedly. Earlier estimates of neutral (Greenwood, Hercowitz, and Krusell, 1997; Ireland, 2013) or consumption-sector (Ireland and Schuh, 2008; Marquis and Trehan, 2008) technology exhibit a similar slowdown after the mid-1970s. Next, I use statistical tests to formally assess the presence of trend breaks and derive implications regarding the sources of US growth in the postwar period.

3.2. The slowdown in neutral technological progress. I test for structural breaks in the trend of the three log TFP series shown in Figure 1 using the approach proposed by Bai and Perron (1998, 2003), which has the double advantage of identifying the number of breaks and providing asymptotic confidence interval(s) for the break date(s). Letting y_t denote log TFP, the estimated model with m breaks is

$$\Delta y_t = \delta_j + u_t, \quad u_t \sim (0, \sigma_j^2), \quad t = T_{j-1} + 1, \dots, T_j,$$

where Δ is the difference operator, $j = 1, \dots, m + 1$ indexes the regimes and T_1, \dots, T_m are the potential break points (with the convention that $T_0 = 0$ and $T_{m+1} = T$). The variance of the disturbance may not be constant across regimes and the setup also allows for auto-correlation in u_t .

The results, detailed in Appendix A, identify a single significant break (at the 1% level) in the trend function of each of the three TFP series. The estimated break date is 1973Q1 for both TFP computed from output in consumption units and Fernald's consumption-sector TFP. This common break date squares well with the idea that these two series correspond to the same theoretical object, namely neutral technology. The estimated break date is 1971Q3 for Fernald's aggregate TFP, but the width of confidence bands makes it impossible to reject

¹⁰See also Beaudry, Moura, and Portier (2015), who discuss how the choice of the numéraire (i.e., the consumption price) shapes the statistical properties of the relative price of investment.

TABLE 1. Average growth rates and contributions to economic growth.

Variable	Sample		
	1948-2017	1948-1972	1973-2017
Fernald's TFP (g_{TFP^X})	1.29%	2.18%	0.79%
<i>Results based on TFP computed from output in consumption units</i>			
Neutral technology (μ_A)	0.72%	1.91%	0.04%
Investment-specific technology (μ_V)	1.77%	0.85%	2.25%
Output in consumption units (g_Y)	1.94%	3.21%	1.20%
Divisia output (g_X)	2.52%	3.48%	1.95%
Contribution of neutral technology to g_Y	55%	88%	5%
Contribution of neutral technology to g_X	42%	81%	3%
<i>Results based on Fernald's consumption-sector TFP</i>			
Neutral technology (μ_A)	0.86%	2.07%	0.17%
Investment-specific technology (μ_V)	1.33%	0.33%	1.85%
Output in consumption units (g_Y)	1.94%	3.21%	1.20%
Divisia output (g_X)	2.37%	3.31%	1.81%
Contribution of neutral technology to g_Y	66%	95%	22%
Contribution of neutral technology to g_X	54%	92%	14%

Notes. The table reports average annual growth rates for Fernald's utilization-adjusted TFP, neutral technology, investment-specific technology, output in consumption units, and chain-aggregated output, over both the 1948-2017 sample and the two sub-samples. The growth rates of Fernald's utilization-adjusted TFP and neutral technology are measured directly from the data, while the others are inferred from balanced-growth relationships. Finally, the table shows the contribution of neutral technological progress to growth in the output variables X_t and Y_t . See the text for details.

the idea that all three TFP measures experienced a break at the same date.¹¹ These findings are well in line with the common wisdom that US productivity growth slowed significantly during the 1970s. Fernald (2007), Marquis and Trehan (2008), and Benati (2014), among others, all identify breaks in trend productivity at about the same period.

The second and third columns in Table 1 report the statistics for the two sub-samples identified by the Bai-Perron tests. They confirm that the US economy experienced a dramatic slowdown in neutral technological progress, which shifted from 1.9%-2.1% a year before 1973 to only 0.05%-0.2% a year afterward.

This structural break has massive implications for the sources of growth in the US economy. To make the point, I rely on a simple accounting exercise designed to measure the respective contributions of neutral and investment-specific technology to aggregate growth, in the spirit of Greenwood, Hercowitz, and Krusell (1997), Whelan (2003), or Greenwood and Krusell (2007). The decomposition works in four steps:

¹¹According to equation (11), any break in the process for neutral technology A_t immediately affects the trend in Fernald's (2014) aggregate TFP.

- First, I compute the growth rate of neutral technology (μ_A) as average growth in either TFP computed from output in consumption units or Fernald's TFP series for the consumption sector.
- Second, I deduce the average growth rate of investment-specific technology (μ_V) from equation (11). Specifically, given average growth in Fernald's aggregate TFP and the nominal investment share in business output $1 - \gamma$, I set $\mu_V = (g_{TFP^X} - \mu_A) / (1 - \gamma)$. I define the investment sector as the sum of fixed investment and consumer durables.
- Third, I use the balanced-growth trends in output in consumption units Y_t and chain-aggregated output X_t to infer their average growth rates. Given estimates of α and γ , I obtain

$$g_Y = \frac{\mu_A}{1 - \alpha} + \frac{\alpha\mu_V}{1 - \alpha}, \quad g_X = g_Y + (1 - \gamma)\mu_V.$$

- Fourth, I compute the contribution of neutral technology to growth in Z as

$$\frac{\mu_A / (1 - \alpha)}{g_Z}, \quad \text{for } Z = X, Y.$$

Because measuring neutral technology using TFP_t^C or TFP_t^Y yields different estimates of μ_A , I perform the accounting exercise twice, once for each of the two measures. In line with the paper's motivation, all the computations rely on interpreting TFP measures as observable combinations of neutral and investment-specific technology and using balanced-growth restrictions to deduce the implied growth rate of output. In contrast, Greenwood, Hercowitz, and Krusell (1997), Whelan (2003), and Greenwood and Krusell (2007) base their computations directly on observed GDP growth. Both approaches are equivalent in theory, but not in practice due to in-sample deviations from balanced growth. Finally, I consider the contributions to growth both for output in consumption units and for chain-aggregated output: the former is interesting because it grows at the same rate as social welfare in the long run in the simple model from Section 2, while the latter grows at the same rate as standard measures of labor productivity computed from X_t .

The successive rows in Table 1 report the results from the decomposition exercise for both the full sample and the two sub-samples identified by the break tests. For each sample, the top row reports average growth in Fernald's aggregate TFP, with the other rows providing average growth in neutral technology, investment-specific technology, output in consumption units, and chain-aggregated output. Finally, the last two rows present the estimated contribution of neutral technology to growth in the output measures X_t and Y_t .

As shown in the first column, the full-sample decomposition attributes an important role to neutral technological progress in driving US growth between 1948 and 2017. Improvement in neutral technology contributed between 55% and 66% of growth in output in consumption units (and welfare) over the sample, and between 42% and 54% of growth in chain-aggregated output. According to the two-sector interpretation of the data, remaining output growth

originated from investment-specific technological change. These contributions differ from those inferred by Greenwood, Hercowitz, and Krusell (1997), who attribute 42% of growth in output in consumption units to neutral technology and 58% from investment-specific technology. However, Whelan (2003, p. 651) argues that Greenwood, Hercowitz, and Krusell's approach is likely to overstate the role of investment-specific technical progress due to their choice of data for the relative price of investment.

The sub-sample analysis reveals substantial changes over time in the sources of US growth, which are masked when looking only at the full sample. The second column in Table 1 shows that aggregate growth originated almost *only* from neutral technological change before 1973, while the third column instead attributes the bulk of growth after 1973 to investment-specific technological change. The raw figures indicate the magnitude of this change: the contribution of neutral technology to growth in output in consumption units fell from more than 85% in the first sub-sample to less than 25% in the second one. For chain-aggregated output, the contribution went from more than 80% to less than 15%.

This slowdown in neutral technology growth also had large welfare consequences for the US economy. To see it, suppose that neutral technology had grown after 1972 at the same rate as during the 1948-1972 sample and take growth in investment-specific technology as given. Based on the statistics in Table 1, output in consumption units in 2017 would have been equal to $(1 + 3.21\%)^{44}Y_{1972}$, instead of the actual $(1 + 1.95\%)^{44}Y_{1972}$ found in the data.¹² This would make output in consumption units in 2017 more than 70% higher than its actual level, which signals a tremendous welfare loss (recall that in the simple model from Section 2.1, welfare has the same trend as Y_t). In comparison, the output and welfare losses due to standard business cycles or even to the Great Recession remain necessarily modest, echoing Lucas's (2003) view that long-run welfare gains or losses exceed by far the costs associated with short-run fluctuations.

Finally, these significant effects of the break in neutral technology growth provide a strong justification for splitting samples in the early 1970s in empirical work on the US economy. Given the arguments in favor of a second structural break in the early 1980s, including a faster decline in the relative price of investment, changes in the conduct of monetary policy, and lower macroeconomic volatility (see Fisher, 2006, and the references therein), the full 1970s decade might well turn out to be a transition period between two different regimes for the US economy, away from the post-World War II economic expansion and into the Great Moderation (Perez-Quiros and McConnell, 2000; Stock and Watson, 2003).

¹²As can be seen from the table, this counterfactual is valid whether the growth rates are computed from TFP in consumption units or from Fernald's consumption-sector TFP.

4. CONCLUSION

Thirty years ago, Prescott (1986) argued that “theory is now ahead of business cycle measurement and [...] should be used to obtain better measures of the key economic time series.” This paper shows that Prescott’s claim still holds true today. Theory implies that standard TFP series do not reflect neutral technology in environments with multiple sources of technical change, contrary to relatively common views. Theory also provides guidance as to how to recover neutral technology in such setups. When appropriately measured, neutral technology growth exhibits a significant slowdown in the early 1970s. In particular, its contribution to US growth fell from up to 85% before 1973 to less than 25% afterward. The welfare loss associated with slower productivity growth is tremendous, as high as 70% of 2017 output levels.

These results call for further research along two lines. First, once one acknowledges that earlier works have mistakenly interpreted TFP as neutral technology in setups with investment-specific technical shocks, the validity of these papers’ conclusions is immediately in question. Reevaluating the empirical findings from, e.g., Beaudry and Lucke (2010) and Schmitt-Grohé and Uribe (2012) might thus be worthwhile. Second, the dramatic slowdown in neutral technology progress after the 1970s has strong and sustained implications for the US economy’s productive capacity and, ultimately, welfare. Understanding the causes and the timing of this break, as well as its magnitude and persistence, is especially important for macroeconomics. Available theories are not fully convincing and more work is warranted to reach persuasive conclusions.¹³

¹³For instance, unmeasured learning costs do not seem able to account for the permanent fall in neutral productivity, as this would require learning costs to represent an ever-growing share of output. The alternative theory based on smaller returns to past inventions can explain the permanent effect but not the brutal shift in neutral technology growth that occurred in the 1970s. Instead, this explanation would imply a gradual downward adjustment that is not apparent in the data.

REFERENCES

- 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.
- BARSKY, R. B., S. BASU, AND K. LEE (2014): “Whither News Shocks?,” in *NBER Macroeconomics Annual 2014, Volume 29*, NBER Chapters, pp. 225–264. National Bureau of Economic Research, Inc.
- BASU, S., AND J. G. FERNALD (2002): “Aggregate Productivity and Aggregate Technology,” *European Economic Review*, 46(6), 963–991.
- BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2006): “Are Technology Improvements Contractionary?,” *American Economic Review*, 96(5), 1418–1448.
- BEAUDRY, P., AND B. LUCKE (2010): “Letting Different Views about Business Cycles Compete,” in *NBER Macroeconomics Annual 2009, Volume 24*, NBER Chapters, pp. 413–455. National Bureau of Economic Research, Inc.
- BEAUDRY, P., A. MOURA, AND F. PORTIER (2015): “Reexamining the Cyclical Behavior of the Relative Price of Investment,” *Economics Letters*, 135(C), 108–111.
- 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.
- BUREAU OF LABOR STATISTICS (2007): “Technical Information About the BLS Multifactor Productivity Measures,” BLS.
- BYRNE, D. M., J. G. FERNALD, AND M. B. REINSORF (2016): “Does the United States Have a Productivity Slowdown or a Measurement Problem?,” *Brookings Papers on Economic Activity*, 47(1 (Spring)), 109–182.
- CETTE, G., J. FERNALD, AND B. MOJON (2016): “The Pre-Great Recession Slowdown in Productivity,” *European Economic Review*, 88(C), 3–20.
- CHEN, K., AND E. WEMY (2015): “Investment-Specific Technological Changes: The Source of Long-Run TFP Fluctuations,” *European Economic Review*, 80(C), 230–252.
- CHUNG, H., M. T. KILEY, AND J.-P. LAFORTE (2010): “Documentation of the Estimated, Dynamic, Optimization-based (EDO) Model of the U.S. Economy: 2010 Version,” Finance and Economics Discussion Series 2010-29, Board of Governors of the Federal Reserve System (U.S.).
- DAVID, P. A. (1990): “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review*, 80(2), 355–361.
- EDGE, R. M., M. T. KILEY, AND J.-P. LAFORTE (2008): “Natural Rate Measures in an Estimated DSGE Model of the U.S. Economy,” *Journal of Economic Dynamics and Control*, 32(8), 2512–2535.

- FERNALD, J. G. (2007): “Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements,” *Journal of Monetary Economics*, 54(8), 2467–2485.
- (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- (2015): “Productivity and Potential Output before, during, and after the Great Recession,” *NBER Macroeconomics Annual*, 29(1), 1–51.
- (2016): “Reassessing Longer-Run U.S. Growth: How Low?,” Working Paper Series 2016-18, Federal Reserve Bank of San Francisco.
- FISHER, J. D. M. (1999): “The New View of Growth and Business Cycles,” *Economic Perspectives, Federal Reserve Bank of Chicago*, 23, 35–56.
- (2006): “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114(3), 413–451.
- (2010): “Comment on Letting Different Views about Business Cycles Compete,” in *NBER Macroeconomics Annual 2009, Volume 24*, NBER Chapters, pp. 457–474. National Bureau of Economic Research, Inc.
- GORDON, R. J. (1990): *The Measurement of Durable Goods Prices*, NBER Books. National Bureau of Economic Research, Inc.
- (2012): “Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds,” NBER Working Papers 18315, 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.
- GUERRIERI, L., D. HENDERSON, AND J. KIM (2014): “Modeling Investment-Sector Efficiency Shocks: When Does Disaggregation Matter?,” *International Economic Review*, 55, 891–917.
- (2016): “Interpreting Shocks to the Relative Price of Investment with a Two-Sector Model,” Finance and Economics Discussion Series 2016-7, Board of Governors of the Federal Reserve System (U.S.).
- 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.

- (2000): “The IT Revolution : Is It Evident in the Productivity Numbers?,” *Economic Quarterly*, Fall, 49–78.
- IRELAND, P., AND S. SCHUH (2008): “Productivity and U.S. Macroeconomic Performance: Interpreting the Past and Predicting the Future with a Two-Sector Real Business Cycle Model,” *Review of Economic Dynamics*, 11(3), 473–492.
- IRELAND, P. N. (2013): “Stochastic Growth In The United States And Euro Area,” *Journal of the European Economic Association*, 11(1), 1–24.
- 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, J. (2003): “Functional Equivalence Between Intertemporal and Multisectoral Investment Adjustment Costs,” *Journal of Economic Dynamics and Control*, 27(4), 533–549.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): “Time to Build and Aggregate Fluctuations,” *Econometrica*, 50(6), 1345–1370.
- LUCAS, R. E. (2003): “Macroeconomic Priorities,” *American Economic Review*, 93(1), 1–14.
- MARQUIS, M. H., AND B. TREHAN (2008): “On Using Relative Prices to Measure Capital-Specific Technological Progress,” *Journal of Macroeconomics*, 30(4), 1390–1406.
- MOURA, A. (2018): “Investment Shocks, Sticky Prices, and the Endogenous Relative Price of Investment,” *Review of Economic Dynamics*, 27, 48–63.
- (2020): “Understanding the Long-Run Relationship Between Total Factor Productivity and the Relative Price of Investment,” Mimeo.
- PEREZ-QUIROS, G., AND M. M. MCCONNELL (2000): “Output Fluctuations in the United States: What Has Changed since the Early 1980’s?,” *American Economic Review*, 90(5), 1464–1476.
- PRESCOTT, E. C. (1986): “Theory Ahead of Business-Cycle Measurement,” *Carnegie-Rochester Conference Series on Public Policy*, 25(1), 11–44.
- RAMEY, V. A. (2016): “Macroeconomic Shocks and Their Propagation,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and H. Uhlig, vol. 2, chap. 2, pp. 71–162. Elsevier.
- RÍOS-RULL, J.-V., F. SCHORFHEIDE, C. FUENTES-ALBERO, M. KRYSHKO, AND R. SANTAEULÀLIA-LLOPIS (2012): “Methods versus Substance: Measuring the Effects of Technology Shocks,” *Journal of Monetary Economics*, 59(8), 826–846.
- 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.
- (2012): “What’s News in Business Cycles,” *Econometrica*, 80(6), 2733–2764.

- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- SOLOW, R. M. (1956): “A Contribution to the Theory of Economic Growth,” *The Quarterly Journal of Economics*, 70(1), 65–94.
- STOCK, J. H., AND M. W. WATSON (2003): “Has the Business Cycle Changed and Why?,” in *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pp. 159–230. National Bureau of Economic Research, Inc.
- SYVERSON, C. (2017): “Challenges to Mismeasurement Explanations for the US Productivity Slowdown,” *Journal of Economic Perspectives*, 31(2), 165–186.
- 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. TESTS FOR STRUCTURAL BREAKS

This appendix reports results from Bai and Perron’s (1998) tests for multiple breaks at unknown dates in the sample in the mean for the log-differences of the three TFP series considered in Section 3.2 (Fernald’s utilization-adjusted TFP, TFP computed from output in consumption units, Fernald’s consumption-sector TFP). The objective is to establish formally the presence of trend breaks in TFP and to estimate the break dates.

Letting y_t denote log TFP, the estimated model with m breaks is

$$\Delta y_t = \delta_j + u_t, \quad u_t \sim (0, \sigma_j^2), \quad t = T_{j-1} + 1, \dots, T_j, \quad (16)$$

where $j = 1, \dots, m + 1$ indexes the regimes and T_1, \dots, 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 and the setup also allows for autocorrelation in u_t .

I implement the tests following Bai and Perron’s (2003) recommendations. In particular, I set the trimming parameter at 15%. The results are reported in Table 2, along with asymptotic critical values for the test statistics. 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 with AR(1) pre-whitening. For all series, the double-maximum test statistics are equal to the $\sup F(1|0)$ statistics, so that I report only the later in Table 2.

As shown in the table, the $\sup F(1|0)$ tests (and the double-maximum tests) find significant breaks in the means of the log-differences of all three TFP series. On the other hand, the $\sup F(2|1)$ statistics are never significantly different from zero, indicating that the data support the presence of a single break in each series. The estimated break dates are late 1971

TABLE 2. Bai-Perron tests of multiple breaks at unknown dates in the trend of TFP series.

Variable	$\text{sup}F(1 0)$	$\text{sup}F(2 1)$	Break date [CI]
Fernald's TFP	23.2	2.53	1971Q3 [63Q3-75Q2]
TFP from output in consumption units	37.4	5.03	1973Q1 [68Q4-76Q1]
Fernald's consumption-sector TFP	44.4	2.97	1973Q1 [70Q1-76Q1]

Notes. The model allows for multiple breaks at unknown dates in the mean of the series and is estimated with 15% trimming. The $\text{sup}F(1|0)$ and $\text{sup}F(2|1)$ tests respectively test the presence of one vs. zero break and two vs. one breaks. Asymptotic critical values at the 10%, 5%, and 1% levels are 8.02, 9.63, and 13.58 for the $\text{sup}F(1|0)$ test, and 9.56, 11.14, and 15.03 for the $\text{sup}F(2|1)$ test. Confidence intervals have 90% asymptotic coverage.

for Fernald's (2014) aggregate TFP,¹⁴ and early 1973 for the two other TFP series. This common break date for TFP computed from output in consumption units and Fernald's TFP series for the consumption sector supports the notion that both variables represent the same underlying object, i.e. neutral technology. In addition, the confidence bands for all break dates overlap significantly, so that it is not possible to reject the idea that all three TFP series experienced a common break.

¹⁴This is slightly different from the estimated break date of 1968Q2 reported by Benati (2014), who used the same series and the same test procedure. However, Benati's sample ends in 2008, while the one considered here runs up to 2017.



BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTEME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu