UNIVERSITY OF ALBERTA
FACULTY OF ARTS
Department of Economics

Working Paper No. 2021-11

Information and Communication Technologies and Medium-Run Fluctuations

Marco Brianti
University of Alberta

Laura Gáti
European Central Bank

December 2021

Information and Communication Technologies and Medium-Run Fluctuations[*]

Marco Brianti[†] Laura Gáti[‡]

December 8, 2021

Abstract

This paper explores the possibility that productivity improvements in information and communication technologies (ICT) are a source of medium-run fluctuations in total factor productivity (TFP). We document in a structural VAR setting that innovations in ICT investment are followed by hump-shaped increases in TFP. Following the ICT literature, we use a two-sector model to suggest a mechanism behind the hump-shaped TFP response: that ICT is a general-purpose technology (GPT). Using impulse-response matching, we show that a model with a spillover from ICT capital is able to match the hump-shaped TFP response, hinting at the importance of the diffusion of ICT.

JEL classification: E3
Keywords: information and communication technologies, general-purpose technologies, two-sector models, total factor productivity

[*]We are deeply grateful to Susanto Basu, Ryan Chahrour, Peter Ireland, and Fabio Schiantarelli for their generous guidance and support. We also thank Vito Cormun, Pierre de Leo, Mario Forni, Pablo Guerron and Jean-Paul L'Huillier.

[†]Department of Economics, University of Alberta, 8-14 Tory Building, Canada T6G 2H4; mail: brianti@ualberta.ca; phone: +1 780 717 8095; website: sites.google.com/site/marcobriantieconomics

[‡]Directorate General Research, European Central Bank, Sonnemannstrasse 20 60314, Frankfurt am Main, Germany; mail: laura_veronika.gati@ecb.europa.eu; website: laura-gati.net

1 Introduction

Are information and communication technologies (ICT) a driver of medium-run economic fluctuations? That there is reason to suspect so can be gleaned from Figure 1, which plots the estimated coefficient of a naive regression of total factor productivity (TFP) on ICT investment. In particular, as Equation 1 highlights, we regress TFP at various horizons h on contemporaneous ICT investment (ICTI), four lags of TFP and ICTI, and an error term u for each horizon. The figure plots $\hat{\theta}_h$, which can be interpreted as the effect of a contemporaneous increase in ICTI on TFP at various horizons h.

$$TFP_{t+h} = \theta_h ICTI_t + \sum_{j=1}^{4}(ICT_{t-j}, \quad TFP_{t-j})\Gamma_j + u_{t,t+h} \tag{1}$$

What the estimated coefficient $\hat{\theta}_h$ in Figure 1 is telling us is that increases in ICT investment are associated with future TFP growth. In particular, while contemporaneous TFP does not significantly respond to ICT investment increases, TFP about four years ahead increases significantly, and the positive effect builds up with the horizon, reaching its peak about seven years out. In other words, a contemporaneous increase in ICT investment triggers a slow build-up of positive effects on TFP in the medium-run.

This paper investigates the positive relation between ICT investment and future TFP suggested in Figure 1. We explore this question through a structural vector autoregression (SVAR) approach, by identifying a structural shock that we call an “ICT shock” - a shock that increases ICT investment. Our baseline identifying assumption is that the ICT shock is the only shock that is permitted to affect ICT investment on impact. In other words, we use a Cholesky decomposition in which the ICT shock is the first shock in the Cholesky ordering, and ICT investment is the first variable in the VAR.

Figure 1: Dynamic relation between ICT Investment and Total Factor Productivity

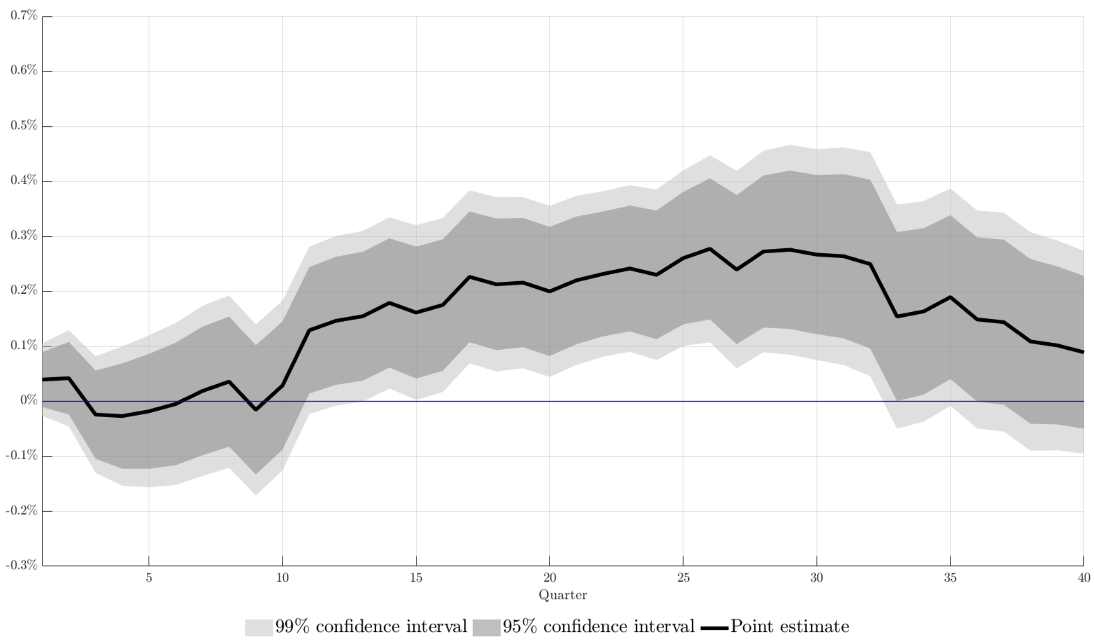

Notes. Dynamic regression (local projection) of utilization-adjusted total factor productivity (TFP) on ICT investment, controlling for four lags of TFP and ICT investment. Point estimate is the OLS estimation of θ_h from Equation 1, for 1989 to 2020. Confidence bands are estimated using standard errors.

On the one hand, this ensures that the shock we back out is relevant for ICT investment. On the other, it leaves the TFP impact response unrestricted, just like the local projection from Figure 1. Unsurprisingly, we find that the TFP impact response to the ICT shock is not significantly different from zero. Over time, however, TFP increases gradually, peaking approximately six years after the ICT shock hit. In other words, an identified ICT shock leads to increases in TFP that are nil on impact, but slowly build up over time, exhibiting a hump shape and peaking about six years after the shock.

We perform a range of tests to ascertain ourselves that we are not confounding our shock with other structural forces in the economy. First, we consider a Cholesky decomposition in which the ICT shock is ordered second, and the in vector of observables in the VAR, TFP is the first variable, and ICT investment the second. This amounts to the assumption that the ICT shock is the only shock to affect ICT investment *without* affecting TFP on impact. Not only are impulse responses very similar to those of our baseline identification, but we also show that the correlation between the backed-out shock series and those of the baseline is 96%.

To investigate the nature of our identified shock, we augment our VAR to include relative prices, defined as the consumption goods price index divided by the price index of ICT goods. Implementing the otherwise unchanged identification strategy, we find that after a positive ICT shock, relative prices drop. Since prices and quantitites are moving in opposite directions, this implies that our ICT shock can be interpreted as a supply shock in the ICT sector.

We then check whether other possible interpretations of the shock are reasonable. In other words, we ask whether our shock is driven by other structural forces in the economy and thus suffers from an endogeneity problem. To this end, we examine the correlation of our backed-out structural shock series with a range of structural shocks identified in the SVAR literature (Ben Zeev and Pappa, 2017, Ramey, 2011, Leeper et al., 2013, Mertens and Ravn, 2011, Wieland and Yang, 2020). We see the fact that our shock exhibits near zero correlation with any of the alternative structural shock series as reassuring that the interpretation of an ICT supply shock is an adequate one, and we are not confounding it with other structural disturbances.

We then turn to a possible mechanism behind the observed hump-shaped response of TFP to the ICT shock. A robust finding in studies of ICT at the sectoral and firm level (such as Oliner and Sichel, 2000 or Stiroh, 2002) is that ICT has a general-purpose technology (GPT) character. Bresnahan and Trajtenberg (1995) define general-purpose technologies as technologies that on top of the direct productivity gains stemming from their use also lead to indirect productivity improvements in sectors that use them. Such indirect improvements could come from many sources, such as complementary investment (Basu and Fernald, 2007), reorganization of production (Brynjolfsson et al.,

1994) or time-to-build phenomena, including the time it takes for the human workforce to learn to use novel technologies efficiently (Atkeson and Kehoe, 2007).

What all these stories of GPTs have in common is that productivity gains in the general-purpose technology spill over to its users. We therefore introduce a spillover from ICT capital in an anotherwise standard two-sector model in the spirit of Greenwood et al. (1997) and Oulton (2007), where one sector produces consumption goods, and the other produces ICT goods. Estimating the spillover elasticity using impulse-response matching, we show that the spillover from ICT capital is what allows the model to match the hump-shaped response of TFP to innovations in ICT productivity. What we take away from this exercise is that since the aggregate data suggest that ICT acts as a GPT in the economy, what matters for the medium-run effects of ICT innovations is the extent and speed of the diffusion of new ICT technologies.

The intuition behind the role of ICT for medium-run fluctuations can be summarized using a simple example. Think of two researchers in the late 1980s or early 1990s who are considering to buy a computer, an ICT good that has recently become available on the market. When the first one purchases a computer, it is with the idea of using the new ICT capital good for her own production process. But when the second researcher also buys a computer with the same motivation, they can now exchange emails and work together on research projects. It is this indirect effect that renders the aggregate economy more productive. Thus, as more and more researchers purchase computers, the effect of the new ICT good on TFP builds up. In this way, the effect plays out at the medium horizon, in parallel with the diffusion of the ICT good.

1.1 Related Literature

This paper builds on two distinct literatures. The first is the literature on the drivers of growth in aggregate total factor productivity. Research on endogenous growth models points to the role of innovation in expanding the productive capacity of the economy. This has lead researchers to focus on research & development (R&D) as well as the adoption of new technologies (Romer, 1990, Aghion et al., 2005, Comin and Gertler, 2006, Bianchi et al., 2014, Anzoategui et al., 2016, Moran and Queralto, 2017, Cao and L'Huillier, 2018, Jinnai, 2014, among many others). We contribute to this line of research by bringing in a particular kind of innovation, namely innovation in ICT, and investigating whether innovations in this specific kind of technology exhibit different behavior than R&D and adoption in general. Unlike Hall et al. (2012), who examine whether R&D or ICT matter more for growth using Italian firm-level data, we ask whether there is something different about ICT than technological innovations in general.

Of course, we are not the first to suggest that ICT may have a distinct role to play in the growth process. Many researchers examine this question in depth in the context of firm-level or industry-level production, or in a growth-accounting framework. Following Greenwood et al. (1997)'s seminal work that indicates that the relative price of investment goods reflects the sector's productivity, Cummins and Violante (2002) construct quality-adjusted price indices of equipment & software (E&S) goods to measure technological progress in the E&S sector. They find that growth in this sector is a significant component of the TFP resurgence in the 1990s. Oliner and Sichel (2000) use a growth-accounting framework to come to the same conclusion. In a similar vein, Basu et al. (2004) compare the UK and US growth experience in the 1990s in a growth-accounting framework to shed light on the role of ICT for the diverging growth paths of the two economies. Fisher (2006) explores the role of investment-specific technological change at the business cycle frequency. Stiroh (2002) compares sectoral productivities of industries that were users of ICT with those that were not, finding significant productivity differentials. Finally, Oulton (2007) proposes a truly two-sector generalization of Greenwood et al. (1997) that allows for a national-income-accounts (NIPA) consistent construction of GDP and TFP.

A subset of this literature takes up the view that ICT is a GPT, and probes the possible mechanisms through which the GPT character of ICT manifests itself. Basu and Fernald (2007) focus on unobserved intangible investments triggered by increasing ICT productivity in sectors that use ICT goods. As in other related work (Basu et al. (2010)), they find that in the short run, sectoral improvements in ICT are contractionary because resources are soaked up in doing the necessary complementary investments and retraining of the workforce. Bresnahan et al. (2002) use firm-level data to uncover the workplace reorganization effects that new ICT technologies often bring with them, highlighting the positive effect on the demand of high-skilled labor. Indeed, Black and Lynch (2004) find that the reorganization and retraining of the labor force following ICT improvements have more sizable impact on TFP than the pure ICT shock itself. Brynjolfsson et al. (1994)'s analysis of industry-level data suggests that such workplace reorganization tends to lead to smaller firms, and thus implies a shift in the firm size distribution toward smaller sizes. Similar to Moran and Queralto (2017)'s formulation of a spillover in R&D adoption, we do not take a stance on where the GPT nature of ICT may be coming from, but consider a spillover formulation sufficient to capture the "general-purpose" nature of a GPT. Another paper to consider a model with spillovers from capital is Chen and Wemy (2015). These authors also interpret the spillover as a reduced-form way of capturing the GPT nature of capital. However, their focus is not strictly on ICT, but on investment-specific technological change in general.

Unlike the vast majority of the ICT literature, we examine the role of ICT at the aggregate level, and instead of relying on growth accounting, employ structural VAR methods. This has the advantage that we can study the frequencies where Comin and Gertler (2006) have prominently argued most of the action is: the medium run. To the best of our knowledge, only three other papers investigate ICT or investment-specific shocks in a SVAR setting. For one, Jafari Samimi and Roshan (2012) estimate the effect of ICT shocks at business cycle frequences and find that ICT increases TFP but reduces hours worked.

For the other, Ben Zeev and Khan (2015) show that news shocks to investment-specific technology generate comovement in output, consumption, investment and hours. We instead focus on improvements in the contemporaneous productivity of ICT, and highlight that the dynamics of the TFP response to such improvements helps us tease out the mechanism between ICT and medium-run fluctuations.

The finding that conteporaneous ICT shocks lead to TFP increases with a delay aligns with the conlusion of the literature on general-purpose technologies that it takes time for the beneficial productivity effects of GPTs to unfold (Jovanovic and Rousseau, 2005, David, 1990). This is widely viewed as a possible resolution to the so-called "productivity paradox," which reflects the idea that improvements in the productivity of the ICT sector should be visible in measured TFP contemporaneously.

The importance of the diffusion of ICT also has implications for the Gordon-Mokyr debate on the growth outlook for the 21st century. The pessimist view of Robert Gordon holds that since all major inventions have already been made, one should not expect high growth in the future (Gordon, 1996, Gordon, 2012). The optimist view à la Joel Mokyr is that GPTs will continue to lead to new waves of innovation (Mokyr, 2014a, Mokyr, 2014b). Our work emphasizes that one should indeed expect GPTs like ICT to lead to growth in the medium run. Whether one subscribes to the pessimist or to the optimist view then depends on the ability of the economy to churn out new general-purpose technologies - an innovation process that is beyond the scope of our paper. However, seen from the lens of our work, the fact that the Covid-19 crisis forced innovation and investment in ICT technologies provides reasons to be optimistic as it suggests that as these technologies diffuse in the economy, we should expect TFP to rise over the medium run.

The rest of the paper is structured as follows. Section 2 depicts our methodology to identify ICT shocks in a structural VAR context, and presents the main result of the paper: the hump-shaped response of TFP. Section 3 shows that in a two-sector model in the spirit of Greenwood et al. (1997), one needs a spillover effect from ICT capital to rationalize the hump-shaped TFP response. Section 4 concludes.

2 Empirics

This section builds on the message of Figure 1: that innovations in ICT productivity lead to hump-shaped, persistent increases in TFP. We first impose more structure by estimating a structural VAR and backing out the ICT shock, a shock that leads ICT investment to increase. We then rely on a VAR augmented with the relative price of ICT goods to interpret this shock as an ICT supply shock. We back up this interpretation by investigating a series of alternatives and showing that alternative structural shock series are uncorrelated with ours. Lastly, we use our VAR to show how a positive ICT supply shock affects the economy through its effects on real GDP.

2.1 Baseline

We start by the simplest possible strategy to identify our shock of interest. We consider a three-variable VAR with (the logs of) real ICT investment, utilization-adjusted TFP and real GDP. Nominal ICT investment is defined as the expenditures of firms in US dollars on ICT goods that are purchased for use in production. This number is then deflated using a price index for ICT goods to obtain the real measure. The dataset comprises quarterly US data and ranges from 1989-Q1 to 2020-Q1.[1] For a detailed description of the data, see Appendix A.

The order of the variables in the observation vector is: ICT investment, TFP and GDP. Our baseline identification is a Cholesky factorization of the variance-covariance matrix of the shocks, where the ICT shock is the first shock. On other words, our identifying assumption is that the ICT shock is the only shock to have a contemporaneous effect on ICT investment.

The intuition behind this assumption is as follows. We are interested in backing out a shock that is specific to the ICT sector. Since this is a sectoral shock, it should trigger movements in ICT investment. By contrast, aggregate structural shocks should not affect ICT investment. In other words, it should not only be

[1] We have also considered a specification with nominal ICT investment. While this increases the sample size (starting in 1982-Q1), results are virtually unchanged.

the case that the ICT shock affects ICT investment, but also that it should be the only shock to do so. Therefore our identification comes from the inclusion of the sector-specific variable, ICT investment.

Figure 2 shows what we obtain using our baseline specification. Panel (a) depicts impulse responses following an ICT shock, while Panel (b) illustrates the forecast error variance of each of the variables explained by the identified shock. As expected, ICT investment increases on impact and exhibits a slow decay, becoming insignificant after about ten years. The continued increase for about three years after impact hints at the presence of adjustment costs in ICT. In line with an increasing ICT capital stock, GDP responds on impact somewhat, but displays a more substantial increase over time as the ICT capital stock is accumulated.[2]

The main result of the paper is the response of TFP. Despite impact effects of the ICT shock on GDP, the effect on TFP on impact is not significantly different from zero. On the one hand, this is not surprising in light of Figure 1. On the other hand, if one suspects that ICT investment may be increasing because the ICT sector became more productive, and one recalls that aggregate TFP is a share-weighted average of sectoral TFPs, one would expect the ICT shock to show up in contemporaneous TFP. According to the Bureau of Economic Analysis (BEA), however, the real value added share of ICT goods in GDP was a little below 5% between 1997-2018, which, while not insignificant, is a relatively small contribution.[3] This helps rationalize why TFP is not moving on impact.

But the dynamic response of TFP is far from zero. Instead, it mirrors the pattern seen in Figure 1. Over time, TFP begins increasing. Approximately ten quarters after the ICT shock, the response becomes significant, and continues increasing until about 30 quarters out, at which point it reverses and falls,

[2] Appendix B contains an extension with the additional variables real consumption, real (total) investment and total hours worked in levels.

[3] This number comes from the June 2021 release of the BEA's "Measuring the Digital Economy" report, which computes the share of the "digital economy" in GDP. Approximately half of the BEA's concept of "digital economy" corresponds to ICT goods and services. (See https://www.bea.gov/data/special-topics/digital-economy.)

Figure 2: Effect of an ICT investment shock on the U.S. economy - baseline

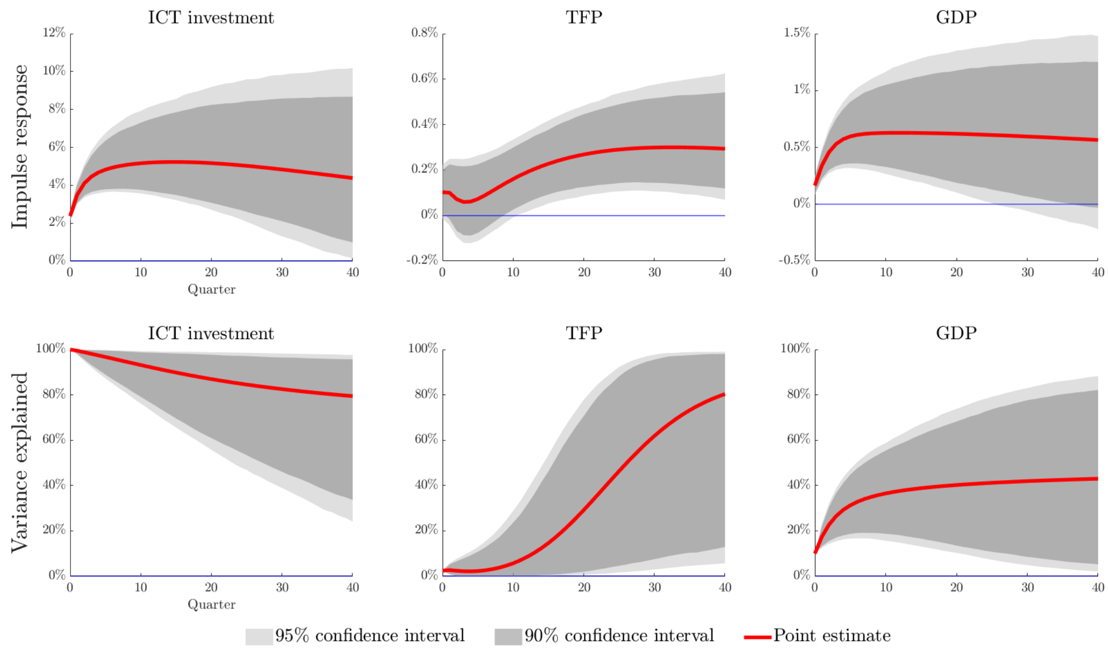

Notes. Impulse responses to a one percent ICT investment shock and forecast error variance explained by the shock series. Data range: 1989-2020. VAR system: log-transformation of real ICT investment, level of the utilization-adjusted TFP, and log-transformation of real GDP. VAR lags: 2 according to the AIC, BIC, and HQ criteria. Identification: first shock of the Cholesky decomposition of the variance-covariance matrix of the residuals of the reduced-form VAR; ICT investment is placed on top. Inference: confidence bands are obtained using Bayesian techniques à la Sims and Zha (1999) (see also Basu and Bundick (2017)'s Online Appendix for more details). See Appendix A for variable descriptions.

vanishing after about 30 quarters. Given the small share of the ICT sector in the economy and the transitory nature of the ICT shock, it is difficult to argue that total factor productivity increases one-to-one with ICT investment. Comparing the response of ICT investment with that of TFP, it is clear that the bulk of the ICT investment response is already over when the TFP response peaks at about 30 quarters. Similarly, the ICT shock explains the bulk of the variance of ICT investment at short horizons, while for TFP the picture is the opposite. Initially, the ICT shock explains nothing of the dynamics of TFP. Over time, however, 40 quarters out, the variance explained is almost 80%. The variance explained of

GDP mirrors that of TFP. The ICT shocks starts out by only explaining about 10% of the GDP response, which steadily increases to 40% after 40 quarters.

All of this dismisses the idea of a strong direct link between the ICT shock and TFP, as the ICT shock is all but gone by the time the effects on TFP kick in, both in terms of the size of the response and the variance explained. Instead, there seems to be some propagation mechanism in the background that is still active once the shock itself has already died out. In Section 3, we will use a two-sector model to think deeper about such a propagation mechanism.

But before doing so, we first explore how to interpret our identified shock. To this end, we first include the relative price of ICT goods, defined as the ICT price index divided by the consumption goods price index, in the VAR. We then identify the ICT shock using our baseline specification and check whether the response of the relative price series makes sense. Figure 3 shows the response of relative prices from this exercise, along with the responses of the other variables.

The main takeaway from this figure is that after a positive ICT shock, relative prices drop. The opposite movement of prices and quantities continues as more ICT capital is accumulated (ICT investment is positive). This suggests that the shock we are picking up is a supply shock in the ICT sector. We therefore conclude that we can interpret our ICT shock as an ICT supply shock. This motivates our choice in Section 3 to model the ICT shock as a shock to the productivity of ICT goods.

2.2 Robustness

In this section, we experiment with alternative identification strategies and robustness checks. The first check is to impose the zero-impact restriction on TFP that we found in the responses to the baseline specification. We thus now identify the ICT shock as the second shock of the Cholesky decomposition for a VAR in which we place TFP first, and ICT investment second.

Figure 3: Effect on relative prices

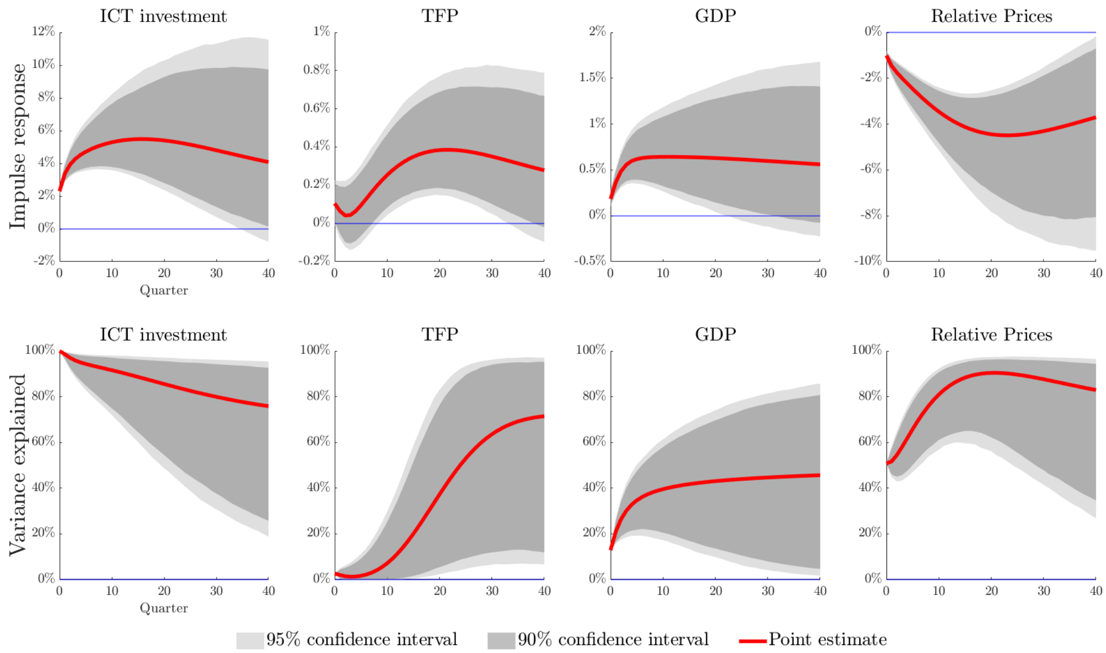

Notes. Impulse responses to a one percent ICT investment shock. Data range: 1989-2020. VAR system: log-transformation of real ICT investment, level of the utilization-adjusted TFP, log-transformation of real GDP, and the ratio of ICT prices over CPI. VAR lags: 2. Identification: first shock of the Cholesky decomposition of the variance-covariance matrix of the residuals of the reduced-form VAR; ICT investment is placed on top. Inference: confidence bands are obtained using Bayesian techniques à la Sims and Zha (1999) (see also Basu and Bundick (2017)'s Online Appendix for more details). See Appendix A for variable descriptions.

Figure 4 shows the obtained impulse responses and variance decompositions. As we can see, the responses are practically identical to those of the baseline specification. The only difference is that now the TFP impact response is by assumption a hard zero. But since in the unrestricted baseline, an insignificant TFP response on impact was a feature of the impulse responses, imposing this as a restriction does not introduce additional structure to the VAR. Instead, this specification is useful to bring the point home that the positive link between ICT productivity and TFP is *not* mechanically coming from the identity that TFP is just the weighted sum of sectoral productivities. Another way to say this is that even if the share of the ICT sector in GDP would be infinitesimal, TFP

Figure 4: Control for contemporaneous TFP

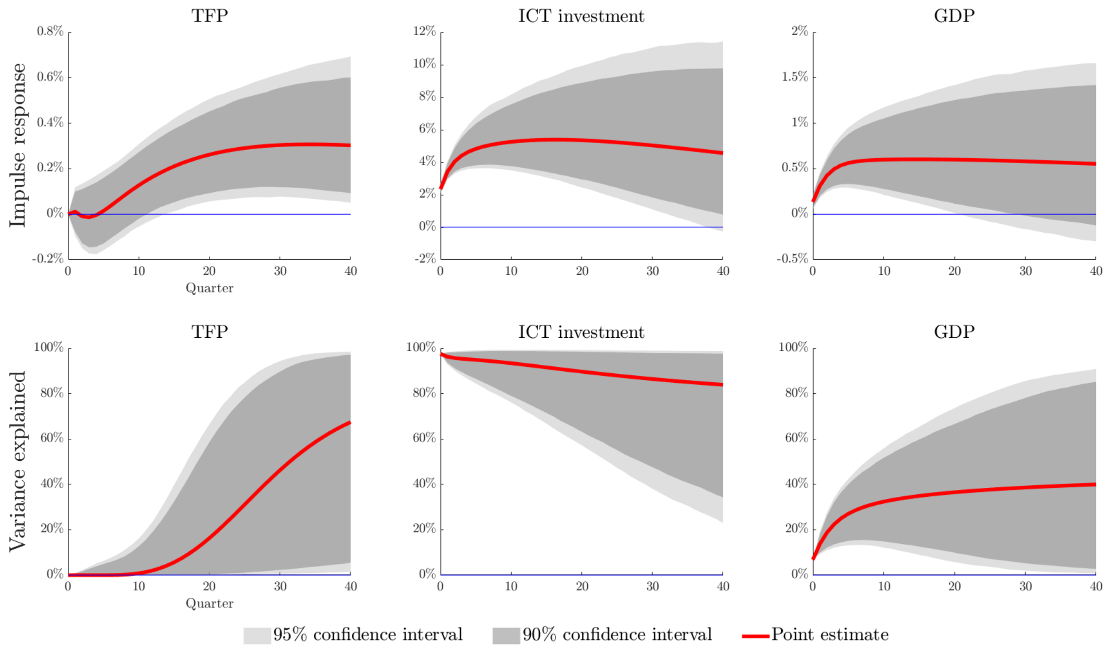

Notes. Impulse responses to a one percent ICT investment shock and forecast error variance explained by the shock series. Data range: 1989-2020. VAR system: log-transformation of real ICT investment, level of the utilization-adjusted TFP, and log-transformation of real GDP. VAR lags: 2. Identification: second shock of the Cholesky decomposition of the variance-covariance matrix of the residuals of the reduced-form VAR; ICT investment is placed right after TFP. Inference: confidence bands are obtained using Bayesian techniques à la Sims and Zha (1999) (see also Basu and Bundick (2017)'s Online Appendix for more details). See Appendix A for variable descriptions.

would still increase in a hump-shaped manner following an ICT shock. This specification, like the baseline, is thus pointing to some additional, indirect relationship between ICT and TFP.

Figure 5 explores the robustness of the baseline specification to altering the number of lags used in the VAR. Instead of the two lags of the baseline, we here plot impulse responses to the identified ICT shock for a specification with four lags (top panel) and eight lags (bottom panel). Clearly, increasing the number of lags results in more bumpy responses as more parameters need to be estimated using the same data. But the key features of the impulse responses

Figure 5: Different number of lags

(a) VAR(4)

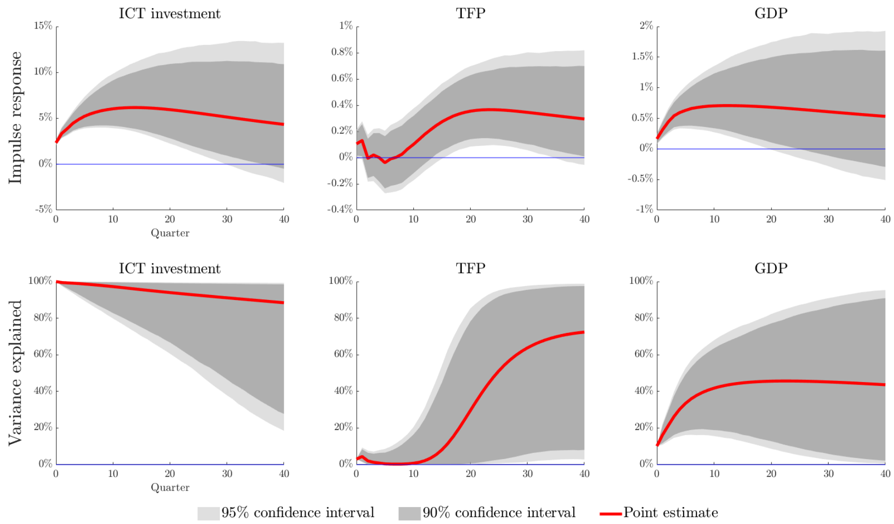

(b) VAR(8)

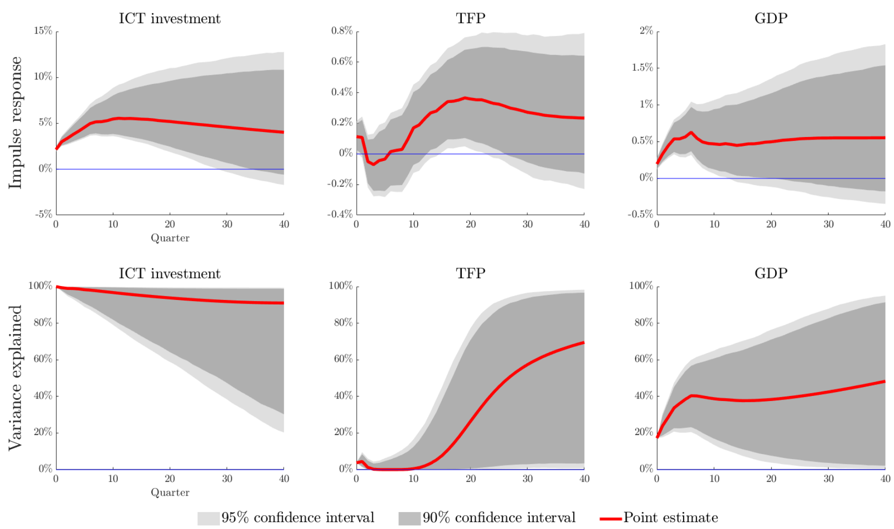

Notes. Impulse responses to a one percent ICT investment shock. Top panel: 4 lags. Bottom panel: 8 lags. Data range: 1989-2020. VAR system: level of the utilization-adjusted TFP, and the log-transformation of ICT investment, and real GDP. Identification: first shock of the Cholesky decomposition of the variance-covariance matrix of the residuals of the reduced-form VAR; ICT investment is placed on top. Inference: confidence bands are obtained using Bayesian techniques à la Sims and Zha (1999) (see also Basu and Bundick (2017)'s Online Appendix for more details). See Appendix A for variable descriptions.

Table 1: Correlation with other structural shocks

Structural shock	Correlation	P-value	Source
Military News	0.11382	0.24765	Ben Zeev and Pappa (2017)
Military News	0.14508	0.13977	Ramey (2011)
Expected Tax	-0.078419	0.5314	Leeper et al. (2013)
Unanticipated Tax	-0.15574	0.18516	Mertens and Ravn (2011)
Anticipated Tax	0.084568	0.47376	Mertens and Ravn (2011)
Romer&Romer Monetary Policy	0.14637	0.21659	Wieland and Yang (2020)

Notes. Correlation of the ICT investment shock from the baseline specification with other structural shocks previously identified by the literature. p-values indicate the probability of obtaining the estimated correlation under the null hypothesis of no correlation. We fail to reject the null hypothesis at conventional level in all cases.

stay the same. In particular, the TFP response displays a hump shape in each case, with a small, mainly insignicant response on impact.

Could it be however that we are mistakenly interpreting the shock we are picking up as an ICT productivity shock, while under the true data-generating process, it has a different structural interpretation? To investigate this question, we gather a set of structural series identified previously in the SVAR literature, and examine the correlation with our ICT shock series. Table 1 plots the correlations we obtain, as well as p-values of a two-sided test with a null of zero correlation.

For none of the considered shock series do we find evidence of a relation with our ICT shock. It is not surprising that our shock does not resemble a monetary policy shock from Wieland and Yang (2020) or an unanticipated tax shock from Mertens and Ravn (2011). After all, the latter two are demand shocks, while our shock clearly has supply-side features. The distinction between our shock and other shocks that pertain to information about the future is less clear-cut. What the tiny correlations in Table 1 allow us to state with confidence is that our ICT shock does not carry information about aspects of the future that are unrelated to productivity, such as taxes and other demand-side disturbances.

However, if the impulse responses of TFP examined so far adequately capture the true data-generating process, then that carries the implication that following a positive ICT innovation, an agent with rational expectations should understand the shape of the TFP impulse response, and thus come to *expect* future increases in TFP. In other words, one way to think of our ICT shock is as a possible source of news shocks about future TFP.

Whereas in the classical news shock literature (Beaudry and Portier, 2006, Barsky and Sims, 2011), news about future TFP are thought of as exogenous, our shock provides and endogenous source of news. Jinnai (2014) makes this very point in the context of sectoral improvements in the R&D sector, suggesting that they contain information about future TFP increases. Görtz and Tsoukalas (2018) also highlight that sector-specific innovations serve as predictors of future TFP growth. Both papers point out that the predictive power of sectoral technology improvements works through the diffusion of the new technologies. In the next section, we use a structural model to suggest how important the diffusion of ICT is for its effect on TFP in the medium run.

3 Theory

Our empirical investigation in Section 2 uncovered a surprising link between aggregate TFP and innovations in ICT productivity. While ICT shocks are not associated with contemporaneous increases in TFP, they do lead to a slow build-up of TFP over the medium run. What is the mechanism behind this?

As summarized in the introduction, a number of papers have argued that ICT is a general-purpose technology. Authors have arrived at this conclusion from many methodological approaches, ranging from industry- or firm-level estimation strategies to growth accounting at various levels of aggregation. Similarly, researchers have postulated numerous alternative formulations of the way the productivity enhancements of a GPT spill over to its users. Whether a particular productivity improvement in a GPT triggers complementary investments, a retraining of the workforce or a reorganization of the production process, in each case it leads to some sort of diffusion process. This implies that without taking a stance on through what exact channel a GPT diffuses in the economy, a convenient way to model the GPT character of a technology is to postulate some spillover from the GPT, just as Moran and Queralto (2017) do in the case of R&D.

We explore whether we can think of ICT as a GPT in the aggregate economy relying on this idea. We build a model in the vein of Greenwood et al. (1997) and Oulton (2007) with two sectors: a consumption-goods-producing sector and an ICT-goods-producing one. On top of sector-specific technologies, both production functions feature a spillover from the aggregate stock of ICT capital. We bring this model to the data to estimate the spillover elasticity, and show that only with a positive spillover can the model generate the hump-shaped TFP impulse response we saw in Section 2.

3.1 Model

Our model is a version of the Greenwood et al. (1997) growth model, augmented with a spillover from ICT capital. In terms of exposition, however, we follow Oulton (2007), who presents a two-sector formulation of the Greenwood et al. (1997) model. This has the advantage that the model offers theoretical counterparts to both GDP and TFP, constructed exactly as in the national income accounts.

3.1.1 The Supply Side: Two Sectors

The key element of the model is that there are two kinds of goods: a consumption good, y_c, and an ICT good, y_i. The consumption good can be used for consumption, c, or for investment in hard capital, i_c, while the ICT good can only be used for investment in ICT capital, i_i. Accordingly, the resource constraints of the two kinds of goods are

$$y_{c,t} = c_t + i_{c,t}\left(1 + \Phi_c\left(\frac{i_{c,t}}{i_{c,t-1}}\right)\right), \tag{2}$$

$$y_{i,t} = i_{i,t}\left(1 + \Phi_i\left(\frac{i_{i,t}}{i_{i,t-1}}\right)\right), \tag{3}$$

where sector x, with $x = \{c, i\}$, is subject to investment adjustment costs according to the quadratic cost function $\Phi_x(\frac{i_{x,t}}{i_{x,t-1}}) = \frac{\phi_x}{2}\left(\frac{i_{x,t}}{i_{x,t-1}} - g^x\right)^2$. Here g^x is the balanced growth path (BGP) growth rate of sector x. This is included in the adjustment cost function to make sure that adjustment costs are zero along the BGP.

The two goods are produced using hard capital, $k_{c,t}$, ICT capital, $k_{i,t}$, and labor hours, h_t, as well as a neutral productivity parameter, η_t, and sector-specific productivities, $\theta_{c,t}$ and $\theta_{i,t}$. Denoting capital of type x used in sector y at time t by $k_{x,y,t}$, the production functions of the two sectors are

$$y_{c,t} = \eta_t\theta_{c,t}(k_{i,t})^{\gamma}(k_{c,c,t})^{a}(k_{i,c,t})^{b}(h_{c,t})^{(1-a-b)}, \tag{4}$$

$$y_{i,t} = p_t \eta_t \theta_{i,t} (k_{i,t})^{\gamma} (k_{c,i,t})^{a} (k_{i,i,t})^{b} (h_{i,t})^{(1-a-b)}. \quad (5)$$

There are several things to note about the production functions in the two sectors. First, the production function of the ICT good involves the relative price of ICT goods, $p_t \equiv p_{i,t}/p_{c,t}$, where $p_{c,t}$ is normalized to 1. This expresses ICT goods in units of consumption goods. Second, following Oulton (2012), we assume that apart from the sectoral technologies, θ_c and θ_i, the two production functions are equal. This simplifying assumption has the advantage that along the BGP, the growth rate of relative prices is simply the ratio of sector-specific productivites. This permits a neat distinction between the two sectors, while keeping the model tractable and preserving the one-to-one mapping to the Greenwood et al. (1997) framework. Third and most importantly, we assume that the aggregate stock of ICT capital, $k_{i,t}$, enters both production functions with a spillover elasticity γ.

The assumption that there is a spillover from the ICT capital stock to the productivities of both sectors is our reduced-form way of modeling the GPT nature of ICT, and follows Moran and Queralto (2017)'s assumption in the case of R&D. The spillover captures the idea that as ICT capital is accumulated, it triggers either complementary investments, a retraining of the workforce, or a reorganization of production in ways that render the sectors using ICT capital more productive. As a simple example, think of an entrepreneur who purchases a mobile phone for her business. There is a direct effect on her productivity, because she is now using a higher stock of ICT capital in production. But if her trading partner, another entrepreneur, also purchases a mobile phone, then the two entrepreneurs can more effectively communicate with one another, yielding a second, indirect boost to the first entrepreneur's productivity. It is this second, additional effect that we capture with our spillover assumption. The objective of this section is to assess whether this effect, that we think of as the GPT character of ICT, is present in the data and, importantly, whether it is necessary to explain the hump-shaped TFP response to an ICT productivity shock.

We assume the sectoral productivity in sector x follows a deterministic trend, denoted by Γ_x. Moreover, we allow transitory shocks to ICT productivity:

$$\theta_{c,t} = (\Gamma_c)^t, \tag{6}$$

$$\theta_{i,t} = (\Gamma_i)^t e^{\zeta_t}, \tag{7}$$

$$\zeta_t = \rho_\varepsilon \zeta_{t-1} + \sigma_\varepsilon \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0,1). \tag{8}$$

Here, ε_t is the transitory shock to ICT productivity, with standard deviation σ_ε and persistence ρ_ε.

The market clearing conditions of the model are

$$k_{c,t} = k_{c,c,t} + k_{c,i,t}, \tag{9}$$

$$k_{i,t} = k_{i,c,t} + k_{i,i,t}, \tag{10}$$

$$h_t = h_{c,t} + h_{i,t}, \tag{11}$$

and the laws of motion of the two capital types are

$$k_{c,t+1} = (1-\delta_c)k_{c,t} + i_{c,t}, \tag{12}$$

$$k_{i,t+1} = (1-\delta_i)k_{i,t} + i_{i,t}. \tag{13}$$

We assume that firms are identical *within* the two sectors. Then the problem of the firms in the consumption good sector and the ICT good sector can be stated as

$$\max_{k_{c,c,t},k_{i,c,t},h_{c,t}} \eta_t \theta_{c,t}(k_{i,t})^\gamma (k_{c,c,t})^a (k_{i,c,t})^b (h_{c,t})^{(1-a-b)} - w_t h_{c,t} - r_{c,t}k_{c,c,t} - r_{i,t}k_{i,c,t}, \tag{14}$$

and

$$\max_{k_{c,i,t},k_{i,i,t},h_{i,t}} p_t\eta_t\theta_{i,t}(k_{i,t})^\gamma(k_{c,i,t})^a(k_{i,i,t})^b(h_{i,t})^{(1-a-b)} - w_th_{i,t} - r_{c,t}k_{c,i,t} - r_{i,t}k_{i,i,t}, \tag{15}$$

where r_x is the rental rate of capital of type x, and w is the wage.

3.1.2 Households

There is a single representative household that uses income from working in both sectors and renting out capital to consume and to invest. The household problem reads

$$\begin{aligned}
&\max_{c_t,h_t,k_{c,t+1},k_{i,t+1},i_{c,t},i_{i,t}} \mathbb{E}_0\sum_{t=0}^{\infty}\beta^t u(c_t,h_t),\\
&\text{s.t}\quad c_t + i_{c,t} + \Phi_c\left(\frac{i_{c,t}}{i_{c,t-1}}\right)i_{c,t} + p_ti_{i,t} + p_t\Phi_i\left(\frac{i_{i,t}}{i_{i,t-1}}\right)i_{i,t} = w_th_t + r_t^ck_t^c + r_t^ik_t^i,\\
&\text{and}\quad k_{c,t+1} = (1-\delta_c)k_{c,t} + i_{c,t},\\
&\text{and}\quad k_{i,t+1} = (1-\delta_i)k_{i,t} + i_{i,t}.
\end{aligned} \tag{16}$$

Here we assume that $u(x_t,h_t) = \log x_t - \nu\frac{1}{1+\frac{1}{\chi}}h_t^{1+\frac{1}{\chi}}$, so that χ is the Frisch elasticity of labor hours to the wage and ν is a scaler of the disutility of labor.

3.1.3 Competitive Equilibrium

The competitive equilibrium of the model is the sequence of capital allocations, investments, labor and interest rates in both sectors, as well as the wage, relative prices and consumption that satisfy the first order conditions of the firms in both sectors and of households, market clearing conditions, production functions, laws of motion of capital and resource constraints. In other words, the equilibrium consists of the sequence $\{k_{c,t}, k_{i,t}, k_{c,c,t}, k_{c,i,t}, k_{i,c,t}, k_{i,i,t}, i_{c,t}, i_{i,t}, h_t, h_{c,t}, h_{i,t}, r_{c,t}, r_{i,t}, w_t, p_t, c_t\}_{t=0}^{\infty}$ that solves the problems (14), (15) and (16), the market clearing conditions (9), (10) and (11), the production functions (4) and (5), the laws of motion (12) and (13), and the resource constraints (2) and (3).

3.2 Estimating the GPT Nature of ICT

With our model in hand, we now turn to our investigation of the role of the spillover from ICT to TFP. Following Christiano et al. (2005), we split the model's parameters into two groups. We calibrate the parameters in the first group using standard values from the literature, or to match steady-state relationships. We estimate those in the second group by minimizing the distance between model-implied impulse responses following an innovation in ICT productivity (ε_t) and the empirical impulse responses from Section 2. In particular, we use the impulse responses of TFP, real ICT investment, and real GDP to estimate three parameters: the standard deviation and persistence of the ICT productivity shock, σ_ε and ρ_ε, as well as the key parameter capturing the GPT character of ICT, the spillover elasiticity γ.

3.2.1 Calibrated Parameters

Table 2: Calibrated parameters

	Value	Interpretation	Source
β	0.99	Discount factor	Woodford (2003)
a	0.6	Share of hard capital	Match ICT value added share ($\approx$ 5%)
b	0.03	Share of ICT capital	Oulton (2012)
Γ_c	1.00328	Growth rate, hard capital tech.	Match growth rates of c and p
Γ_i	1.032	Growth rate, ICT tech.	Match growth rates of c and p
δ_c	0.056	Depreciation, hard capital	BEA, Karabarbounis and Neiman (2014)
δ_i	0.124	Depreciation, ICT capital	BEA, Oliner (1992)
χ	1	Frisch elasticity of labor supply	Chetty et al. (2011)
ν	7.534	Scaler of labor disutility	Match steady state hours of 0.3
ϕ_c	5.9	Invest. adj. cost, hard capital	Smets and Wouters (2007)
ϕ_i	5.9	Invest. adj. cost, ICT capital	Smets and Wouters (2007)

Table 2 shows the calibrated parameters. Where possible, we rely on values from the literature, such as for the discount factor β and the investment adjustment cost parameters ϕ_c, ϕ_i. The output share of ICT capital, b, comes from Oulton (2012). When selecting the hard capital share, a, we face a tension. On the one hand, the overall capital share is around $1/3$. On the other hand, in our simple model, this number would render the ICT production a too large component of GDP. Since our focus here is on the role of ICT for TFP (and thus for GDP) *despite* the small size of the ICT sector, we opt to increase a so as to get the size of the ICT sector relative to GDP right.

We select the trend growth rates of the consumption and the ICT sector to match the growth rates of consumption and relative prices along the balanced growth path. The quarterly depreciation come from the BEA, Karabarbounis and Neiman (2014) and Oliner (1992). Here the two capital categories of the model (hard capital and ICT capital) are somewhat too rough to be mapped one-to-one into the measurements of the BEA or the estimates of Oliner (1992). We therefore strike a compromise between the measurements from the data and Karabarbounis and Neiman (2014)'s calibration of 0.02 and 0.2 for low- and high-depreciating capital respectively. For the Frisch elasticity χ, there is a well-known tension between micro and macro estimates, as documented in Chetty et al. (2011). While micro estimates tend to be around 0.25, macro estimates are an order of magnitude larger, between 1 and 2. Since we are matching aggregate data, we opt for the lowest number compatible with macro estimates.

3.2.2 Estimated Parameters

Let Ω denote the vector of parameters to be estimated, $\Omega = (\gamma, \sigma_\varepsilon, \rho_\varepsilon)$, and define $\Psi(\Omega)$ as the theoretical impulse responses of the model given this set of parameters.[4] Finally, with $\hat{\Psi}$ denoting the empirical impulse responses, we estimate Ω as

$$\hat{\Omega} = \arg\min \, (\hat{\Psi} - \Psi(\Omega))'\Lambda(\hat{\Psi} - \Psi(\Omega)), \tag{17}$$

where Λ is a diagonal weighting matrix, which is the inverse of the variance matrix of the empirical impulse responses.

Table 3 shows the estimation results. The obtained values for σ_ε, 0.05, and for ρ_ε, 0.99, suggest that the ICT technology shock needs to be sizable and

[4] For ease of notation, we suppress the fact that $\Psi(\Omega)$ also depends on the calibrated parameters.

persistent for the model to adequately match the empirical responses. Our object of interest, however, is γ. The estimated value, 0.03, is greater than zero, indicating that the data favor the presence of a spillover from ICT capital to TFP.

Table 3: Estimated parameters

	Estimated value	Interpretation
γ	0.0271	Spillover elasiticity
σ_ε	0.0494	ICT productivity shock standard deviation
ρ_ε	0.9929	ICT productivity persistence

To understand why the data are selecting a strictly nonzero spillover parameter, and the role that γ plays for model dynamics, consider the impulse responses of the model to an ICT shock, presented in Figure 6. The figure shows the theoretical responses of ICT investment, TFP and GDP to a one percent shock, alongside the empirical responses (in red). The shaded areas are 95% and 90% confidence intervals for the VAR. In the top row, the blue line represents the model's impulse responses conditional on the estimated $(\hat{\gamma}, \hat{\sigma}_\varepsilon, \hat{\rho}_\varepsilon)$. In the bottom row, the black dashed line instead shows the model impulse responses for the same estimated $(\hat{\sigma}_\varepsilon, \hat{\rho}_\varepsilon)$, but setting $\gamma = 0$.

The first observation is that, despite its simplicity, the model does a good job at matching the targeted moments. The impulse responses of ICT investment, TFP and GDP have reasonable magnitudes and dynamics, relative to the empirical responses. Comparing the top row with the bottom row, it becomes apparent what is the key to the model's success: a positive value of the spillover parameter γ. While selecting $(\hat{\sigma}_\varepsilon, \hat{\rho}_\varepsilon)$ appropriately is sufficient to capture the dynamics of ICT investment, as one moves from the top row to the bottom, the TFP and GDP impulse responses both deteriorate in fit.

The most crucial point the $\gamma = 0$ case cannot capture is the hump-shaped response of TFP. Of course, this is not surprising given that absent a spillover effect, the only relationship between ICT and TFP is the productivity of ICT.

Figure 6: Model-implied impulse responses

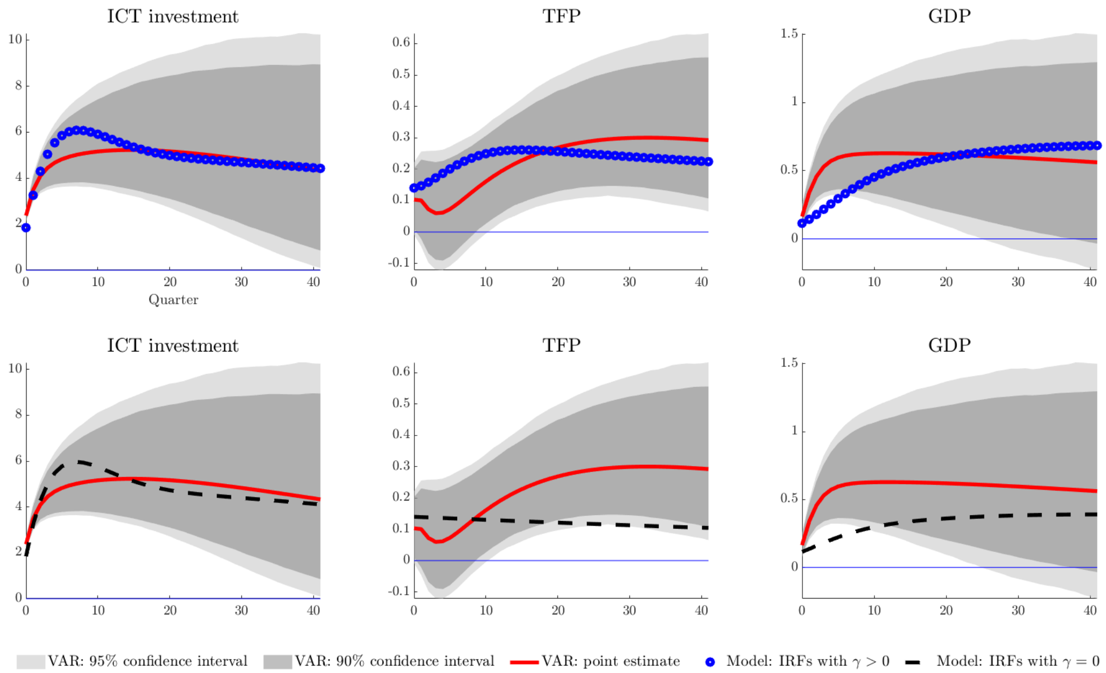

Notes. Theoretical impulse responses together with empirical impulse responses to a one percent ICT investment shock. Model-implied impulse responses are estimated via impulse response matching estimation on parameters: γ, σ_ε, ρ_ε. Data-implied impulse responses are from the baseline presented on Figure 2.

When $\gamma = 0$, TFP increases mechanically and instanteneously in response to an improvement in ICT productivity, and as the shock dies out, so does the effect on TFP. What a positive spillover does, then, is it creates an additional link between the ICT sector and TFP. In this case, as the economy accumulates more ICT capital, all the sectors that use ICT capital become more productive, leading to a second-round positive effect on TFP. What this allows us to conclude is that it is the diffusion of the new ICT good in the economy that is the key to generating the hump-shaped TFP response. Whatever the mechanism behind the GPT-nature of ICT is, it leads to a diffusion process which explains why TFP builds up slowly over time. It is thus this diffusion process which rationalizes the link between ICT and medium-run fluctuations in TFP.

Interestingly, with $\gamma = 0$, the model also gets a quite different GDP response than if γ equals its estimated value. Comparing the black dashed line ($\gamma = 0$) with the blue one ($\gamma = 0.03$) in the right column of Figure 6, one observes that without the spillover effect, the model predicts about 15 basis points lower GDP 10 quarters after the shock, and about 30 basis points lower 40 quarters out. These are sizable differences, implying that if one is interested in the economic growth implications of ICT from the lens of a structural model, it is crucial to get the link between ICT and TFP right so as to then correctly capture the GDP response.

3.3 Discussion

The empirical analysis in Section 2 showed that following an innovation in ICT, TFP exhibits a hump-shaped response. We have seen from the impulse-response matching exercise that the feature of a two-sector growth model that can rationalize this TFP response is a spillover from ICT capital to TFP. Above, we interpreted this as saying that what matters for the size and timing of the aggregate productivity gains following the invention of new information technology goods is how extensively and how fast these technologies diffuse in the economy.

This is related to the so-called "productivity paradox," the notion that despite breakthroughs in the productivity of ICT goods, we do not see contemporaneous increases in measured TFP. The fact that our results emphasize the importance of the diffusion of ICT technologies offers a solution to this paradox, namely that the productivity gains take a long time to materialize precisely because diffusion is a slow process. In this sense, our work provides corroborating evidence to the idea that general-purpose technologies take time to affect economy-wide productivity, as emphasized already in the literature on GPTs (Jovanovic and Rousseau, 2005, David, 1990).

The idea that GPTs offer productivity benefits that only show up over time also has implications for the growth outlook of the 21st century. There are essentially two views on this question. The pessimistic view, prominently held by Robert Gordon, emphasizes that since the technological frontier has already been reached, no great inventions remain to be discovered, leading to a slowdown in producitivity growth (Gordon, 1996, Gordon, 2012). The optimists, with Joel Mokyr in their lead, instead emphasize that on the contrary, the nature of modern inventions is that they lead to waves of subsequent inventions, boding well for future productivity growth (Mokyr, 2014a, Mokyr, 2014b).

Our result that ICT diffusion is key for the productivity enhancement in TFP sees value in both views. On the one hand, it suggests that the diffusion process that characterizes a general-purpose technology may indeed trigger waves of subsequent innovation. On the other hand, the prerequisite for such second-round effects is that new GPTs are invented in the future. Investigating how probable this is is beyond the scope of our paper. However, the slow diffusion of new ICT technologies also means that observed slow TFP growth at a particular point in time is not necessarily a cause for concern. The reason is that it could be that simultaneously to slow current TFP growth, new general-purpose technologies are invented whose positive effects are yet to show up in TFP in the future. One can thus end on the optimistic note that slow TFP growth today need not be taken as indicative of slow TFP growth in the future.

Lastly, our results also have implications for recent episodes such as the Covid-19 crisis. Work-from-home regimes and the necessity of virtual communication have forced both ICT innovation and investment, and arguably also led to a faster diffusion of ICT technologies. Our VAR predicts that if the speed of diffusion were the same as previously, this should show up in higher TFP in about five years. With a faster diffusion process, the time span could shorten.[5]

[5] Our choice to end our sample before the Covid-19 crisis reflects that such a large shock would violate stability assumptions underlying the VAR approach. See Lenza and Primiceri (2020).

4 Conclusions

Are information and communication technologies a source of medium-run fluctuations? A simple VAR analysis with an identified shock to the producitivity of ICT goods suggests yes. Following a positive ICT shock, identified as the only shock to affect ICT investment on impact, TFP does not initially react, but slowly increases over the medium run. The response exhibits a hump shape, peaking after about six years.

We also show that in a simple two-sector model in the spirit of Greenwood et al. (1997), the feature that can rationalize the hump-shaped response of TFP to an ICT shock is a spillover from ICT capital. The spillover is a reduced-form way of capturing the idea adovated by the ICT literature that ICT is a general-purpose technology in the sense of Bresnahan and Trajtenberg (1995). This invites an interpretation of the hump shape of the TFP dynamics as a diffusion process through which new ICT goods are put to use throughout the economy.

This diffusion process resolves the productivity paradox, as ICT innovations should not be expected to show up in TFP measures contemporaneously, but with a lag of about six years. It also provides arguments in favor of both the optimistic and the pessimistic view of the Gordon-Mokyr debate concerning the future of producitivity growth in the 21st century. Lastly, an interesting implication of the role of the ICT diffusion process is that the necessitated investments in ICT technologies during the Covid-19 crisis may imply significant TFP improvements down the road.

What we draw for conclusions for policy from our work is the need to subsidize ICT investment. Since our estimated model from Section 3 points to the importance of the spillover from ICT capital, it is clear that a positive externality is present. In such a case, policymakers can improve on the market allocation by subsidizing the good in question; in this case, ICT goods. A prudent medium-run policy, then, is a subsidy on ICT.

References

AGHION, P., N. BLOOM, R. BLUNDELL, R. GRIFFITH, AND P. HOWITT (2005): "Competition and Innovation: An Inverted-U Relationship," *The Quarterly Journal of Economics*, 120, 701–728.

ANZOATEGUI, D., D. COMIN, M. GERTLER, AND J. MARTINEZ (2016): "Endogenous technology adoption and R&D as sources of business cycle persistence," Tech. rep., National Bureau of Economic Research.

ATKESON, A. AND P. J. KEHOE (2007): "Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions," *American Economic Review*, 97, 64–88.

BARSKY, R. AND E. SIMS (2011): "News shocks and business cycles," *Journal of Monetary Economics*, 58, 273–289.

BASU, S. AND B. BUNDICK (2017): "Uncertainty shocks in a model of effective demand," *Econometrica*, 85, 937–958.

BASU, S. AND J. FERNALD (2007): "Information and communications technology as a general-purpose technology: Evidence from US industry data," *German Economic Review*, 8, 146–173.

BASU, S., J. FERNALD, J. FISHER, AND M. KIMBALL (2010): "Sector-specific technical change," .

BASU, S., J. G. FERNALD, N. OULTON, AND S. SRINIVASAN (2004): *The Case of the Missing Productivity Growth, or Does Information Technology Explain Why Productivity Accelerated in the United States But Not in the United Kingdom?*, The MIT Press, vol. Volume 18, 9–82.

BEAUDRY, P. AND F. PORTIER (2006): "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96, 1293–1307.

BEN ZEEV, N. AND H. KHAN (2015): "Investment-specific news shocks and US business cycles," *Journal of Money, Credit and Banking*, 47, 1443–1464.

BEN ZEEV, N. AND E. PAPPA (2017): "Chronicle of a war foretold: The macroeconomic effects of anticipated defence spending shocks," *The Economic Journal*, 127, 1568–1597.

BIANCHI, F., H. KUNG, AND G. MORALES (2014): "Growth, slowdowns, and recoveries," Tech. rep., National Bureau of Economic Research.

BLACK, S. E. AND L. M. LYNCH (2004): “What’s driving the new economy?: the benefits of workplace innovation*,” *The Economic Journal*, 114, F97–F116.

BRESNAHAN, T. F., E. BRYNJOLFSSON, AND L. M. HITT (2002): “Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence,” *The Quarterly Journal of Economics*, 117, 339–376.

BRESNAHAN, T. F. AND M. TRAJTENBERG (1995): “General purpose technologies ‘Engines of growth’?” *Journal of econometrics*, 65, 83–108.

BRYNJOLFSSON, E., T. W. MALONE, V. GURBAXANI, AND A. KAMBIL (1994): “Does Information Technology Lead to Smaller Firms?” *Management Science*, 40, 1628–1644.

CAO, D. AND J.-P. L’HUILLIER (2018): “Technological revolutions and the three great slumps: A medium-run analysis,” *Journal of Monetary Economics*, 96, 93–108.

CHEN, K. AND E. WEMY (2015): “Investment-specific technological changes: The source of long-run TFP fluctuations,” *European Economic Review*, 80, 230–252.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): “Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins,” *American Economic Review*, 101, 471–75.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy,” *Journal of political Economy*, 113, 1–45.

COMIN, D. AND M. GERTLER (2006): “Medium-term business cycles,” *American Economic Review*, 96, 523–551.

CUMMINS, J. G. AND G. L. VIOLANTE (2002): “Investment-Specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, 5, 243 – 284.

DAVID, P. A. (1990): “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *The American Economic Review*, 80, 355–361.

FISHER, J. D. (2006): “The dynamic effects of neutral and investment-specific technology shocks,” *Journal of Political Economy*, 114, 413–451.

GORDON, R. J. (1996): “Problems in the measurement and performance of service-sector productivity in the United States,” .

——— (2012): "Is US economic growth over? Faltering innovation confronts the six headwinds," Tech. rep., National Bureau of Economic Research.

GÖRTZ, C. AND J. D. TSOUKALAS (2018): "Sectoral TFP news shocks," *Economics Letters*, 168, 31–36.

GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): "Long-Run Implications of Investment-Specific Technological Change," *The American Economic Review*, 87, 342–362.

HALL, B. H., F. LOTTI, AND J. MAIRESSE (2012): "Evidence on the Impact of R&D and ICT Investment on Innovation and Productivity in Italian Firms," Working Paper 18053, National Bureau of Economic Research.

JAFARI SAMIMI, A. AND Y. E. ROSHAN (2012): "The Impact of ICT Shocks on Business Cycle Some Evidence from Iran," *Iranian Economic Review*, 16, 123–145.

JINNAI, R. (2014): "R&D Shocks and News Shocks," *Journal of Money, Credit and Banking*, 46, 1457–1478.

JOVANOVIC, B. AND P. L. ROUSSEAU (2005): "General Purpose Technologies," Working Paper 11093, National Bureau of Economic Research.

KARABARBOUNIS, L. AND B. NEIMAN (2014): "Capital depreciation and labor shares around the world: measurement and implications," Tech. rep., National Bureau of Economic Research.

LEEPER, E. M., T. B. WALKER, AND S.-C. S. YANG (2013): "Fiscal foresight and information flows," *Econometrica*, 81, 1115–1145.

LENZA, M. AND G. E. PRIMICERI (2020): "How to Estimate a VAR after March 2020," Tech. rep., National Bureau of Economic Research.

MERTENS, K. AND M. O. RAVN (2011): "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks," *Review of Economic dynamics*, 14, 27–54.

MOKYR, J. (2014a): "The Next Age of Invention: Technology's future is brighter than pessimists allow." *City journal*, 24, 12–21.

——— (2014b): "Secular stagnation? Not in your life," *Secular Stagnation: Facts, Causes and Cures*, 83.

MORAN, P. AND A. QUERALTO (2017): "Innovation, productivity, and monetary policy," *Journal of Monetary Economics*.

OLINER, S. D. (1992): "Estimates of depreciation and retirement for computer peripheral equipment," in *Workshop on the Measurement of Depreciation and Capital Stock at the Conference on Research in Income and Wealth*, vol. 5.

OLINER, S. D. AND D. E. SICHEL (2000): "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives*, 14, 3–22.

OULTON, N. (2007): "Investment-specific technological change and growth accounting," *Journal of Monetary Economics*, 54, 1290–1299.

——— (2012): "Long term implications of the ICT revolution: Applying the lessons of growth theory and growth accounting," *Economic Modelling*, 29, 1722–1736.

RAMEY, V. A. (2011): "Identifying government spending shocks: It's all in the timing," *The Quarterly Journal of Economics*, 126, 1–50.

ROMER, P. M. (1990): "Endogenous technological change," *Journal of political Economy*, 98, S71–S102.

SIMS, C. A. AND T. ZHA (1999): "Error bands for impulse responses," *Econometrica*, 67, 1113–1155.

SMETS, F. AND R. WOUTERS (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 97, 586–606.

STIROH, K. J. (2002): "Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?" *American Economic Review*, 92, 1559–1576.

WIELAND, J. F. AND M.-J. YANG (2020): "Financial dampening," *Journal of Money, Credit and Banking*, 52, 79–113.

WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

A Data

Table 4: Details on aggregate US data

Variable	Description	Source
TFP	Utilization-adjusted total factor productivity (dtfp_util)	San Francisco Fed
GDP	Real Gross Domestic Product, Billions of Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate (GDPC1)	FRED
Investment	Real Gross Private Domestic Investment, Billions of Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate (GPDIC1)	FRED
Consumption: durables	Real Personal Consumption Expenditures: Durable Goods (RCOND)	Philadelphia Fed
Consumption: non-durables	Real Personal Consumption Expenditures: Non-durable Goods (RCONND)	Philadelphia Fed
Consumption: services	Real Personal Consumption Expenditures: Services (RCONS)	Philadelphia Fed
Hours	Nonfarm Business Sector: Hours Worked for All Employed Persons, Index 2012=100, Quarterly, Seasonally Adjusted (HOANBS)	FRED
ICT Investment	Private fixed investment in information processing equipment and software, Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate (A679RC1Q027SBEA)	FRED
Real ICT Investment	Real gross private domestic investment: Fixed investment: Nonresidential: Information processing equipment and software, Billions of Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate (A679RX1Q027SBEA)	FRED
ICT Prices	Consumer Price Index for All Urban Consumers: Information Technology, Hardware and Services in U.S. City Average, Index Dec 1988=100, Quarterly, Not Seasonally Adjusted (CUUR0000SEEE)	FRED
CPI	Consumer Price Index: All Items for the United States, Index 2015=100, Quarterly, Not Seasonally Adjusted (USACPIALLMINMEI)	FRED
Population	Total Population: All Ages including Armed Forces Overseas, Thousands, Quarterly, Not Seasonally Adjusted (POP)	FRED
S&P 500	Standard & Poor 500 Index, last observation of the quarter	Yahoo Finance
Recessions	NBER based Recession Indicators for the United States from the Period following the Peak through the Trough, +1 or 0, Quarterly, Not Seasonally Adjusted (USRECD)	FRED

B Other empirical results

Figure 7: ICT shock on the US economy

(a) Impulse responses

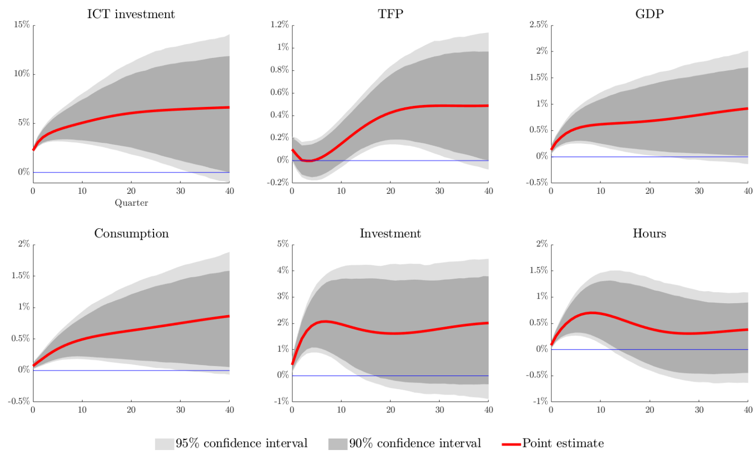

(b) Variance decomposition

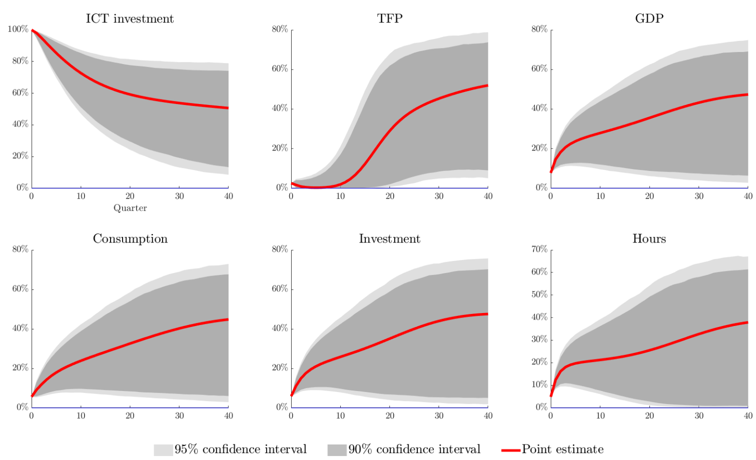

Notes. Impulse responses to a one percent ICT investment shock. Data range: 1982-2020. VAR system: level of the utilization-adjusted TFP, and the log-transformation of ICT investment, real GDP, real consumption, real investment, and total hours. Identification: first shock of the Cholesky decomposition of the variance-covariance matrix of the residuals of the reduced-form VAR; ICT investment is placed on top. Inference: confidence bands are obtained using Bayesian techniques à la Sims and Zha (1999) (see also Basu and Bundick (2017)'s Online Appendix for more details). See Appendix A for variable descriptions.

Department of Economics, University of Alberta
Working Paper Series

2021-10: Are Bank Bailouts Welfare Improving? **– Shukayev, M.,** Ueberfeldt, A.

2021-09: Information, Belief, and Health Behavior: Evidence from China **–** Lei, X., Song, G., **Su, X.**

2021-08: Stay at Home if You Can: COVID-19 Stay-at-Home Guidelines and Local Crime – Díaz, C., **Fossati, S.,** Trajtenberg, N.

2021-07: Financial Frictions, Equity Constraints, and Average Firm Size Across Countries – Bento, P., **Ranasinghe, A.**

2021-06: Should Canada's Capital Gains Taxes be Increased or Reformed? – **McMillan, M.**

2021-05: Financial Shocks, Uncertainty Shocks, and Monetary Policy Trade-Offs – **Brianti, M.**

2021-04: Valuing Elementary Schools: Evidence from Public School Acquisitions in Beijing – **Su, X.,** Yu, H.

2021-03: Colonel Blotto's Tug of War – **Klumpp, T.**

2021-02: Stockpiling and Shortages (the "Toilet Paper Paper") – **Klumpp, T.**

2021-01: When Social Assistance Meets Market Power: A Mixed Duopoly View of Health Insurance in the United States – **Ranasinghe, A., Su, X.**

2020-15: First to $15: Alberta's Minimum Wage Policy on Employment by Wages, Ages, and Places – **Fossati, S., Marchand, J.**

2020-14: The Impacts of Load-Following Forward Contracts – **Brown, D.,** Sappington, D.

2020-13: Loss of Life and Labour Productivity: The Canadian Opioid Crisis – **Cheung, A., Marchand, J.,** Mark, P.

2020-12: Equilibrium Indeterminacy and Extreme Outcomes: A Fat Sunspot Ta(i)l(e) - **Dave, C.,** Sorge, M.

2020-11: Marginal Entrants and Trade-Liberalization Effects Across Models of Imperfect Competition - **Alfaro, M.,** Lander, D.

2020-10: Export Conditions in Small Countries and their Effects On Domestic Markets - **Alfaro, M.,** Warzynski, F.

2020-09: Trade Liberalizations with Granular Firms – **Alfaro, M.,** Warzynski, F.

2020-08: Routine Tasks were Demanded from Workers during an Energy Boom – **Marchand, J.**

2020-07: Financial Frictions, Borrowing Costs, and Firm Size Across Sectors – Bento, P., **Ranasinghe, A.**

2020-06: Bank Lending, Monetary Policy Transmission, and Interest on Excess Reserves: A FAVAR Analysis – **Dave, C.,** Dressler, S., Zhang, L.

2020-05: Vertical Integration and Capacity Investment in the Electricity Sector – **Brown, D.,** Sappington, D.

2020-04: Residential Electricity Pricing in Texas's Competitive Retail Market – **Brown, D.,** Tsai, C., Woo, C., Zarnikau, J., Zhu, S.

2020-03: Competition in Higher Education – Kaganovich, M., Sarpca, S., **Su, X.**

2020-02: Misallocation across Establishment Gender – **Ranasinghe, A.**

2020-01: Does Locational Marginal Pricing Impact Generation Investment Location Decisions? An Analysis of Texas's Wholesale Electricity Market – **Brown, D.,** Zarnikau, J., Woo, C.