European University Institute
DEPARTMENT OF ECONOMICS

Three Essays in Macroeconomics

Robert Goodhead

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Florence, 25 May 2018

European University Institute
Department of Economics

Three Essays in Macroeconomics

Robert Goodhead

Thesis submitted for assessment with a view to obtaining the degree of
Doctor of Economics of the European University Institute

Examining Board

Prof. Peter Reinhard Hansen, University of North Carolina, Chapel Hill, Supervisor
Prof. Juan J. Dolado, EUI
Prof. Alessia Campolmi, University of Verona
Prof. Jean-Paul L'Huillier, EIEF

European University Institute

DEPARTMENT OF ECONOMICS

Researcher declaration to accompany the submission of written work

I Robert Goodhead certify that I am the author of the work "Three Essays in Macroeconomics" I have presented for examination for the PhD thesis at the European University Institute. I also certify that this is solely my own original work, other than where I have clearly indicated, in this declaration and in the thesis, that it is the work of others.

I warrant that I have obtained all the permissions required for using any material from other copyrighted publications.

I certify that this work complies with the *Code of Ethics in Academic Research* issued by the European University Institute (IUE 332/2/10 (CA 297).

Statement of inclusion of previous work (if applicable):

I confirm that chapter 1 was jointly co-authored with Dr. Benedikt Kolb and I contributed 50% of the work.

Chapters 2 and 3 are entirely my work.

Signature and Date:

Robert Goodhead

07/05/2018

Introduction

The first chapter of this thesis (which represents joint work with Benedikt Kolb) quantifies the potentially different effects of monetary policy communication on macroeconomic variables, relative to monetary policy actions, using U.S. data. We employ a decomposition of high-frequency jumps in financial futures around meeting days of the FOMC, in the vein of Gürkaynak et al. (2005), but with an additional weighted averaging procedure to reduce noise and to allow the series to be entered into a simple hybrid VAR of the Romer and Romer (2004) type. We find that only the communication shocks deliver responses of Industrial Production consistent with theory. The second chapter of the thesis re-examines the argument of Asker et al. (2014), who suggest that recently developed measures of resource misallocation largely document the dynamic adjustment of firms in the presence of capital adjustment costs and idiosyncratic shocks to demand. The study extends the analysis of Asker et al. (2014) to the case of labour adjustment costs, and finds that such arguments are well able to deliver the levels of labour misallocation recorded using Italian firm-level data. The model struggles with the correlation between misallocation measures and firm-size. The model is then used to investigate a regulatory threshold, representing the first structural investigation of factor adjustment costs that explicitly incorporates features of the regulatory environment into estimation. The effects of regulatory reform for TFP are shown to be potentially quite large. The third chapter of the thesis evaluates whether the financial crisis of 2008 precipitated positive or negative changes in levels of misallocation, as would follow either from the "cleansing" or the "sullying" views of recessions respectively. The paper uses European data from 13 economies. Regression analysis is able to detect some evidence for cleansing redistributions of value-added and employment, for the manufacturing sample.[1]

Acknowledgments

Thank you to my supervisors, Peter Hansen and Juan Dolado, for all their valuable advice and support during my time as a researcher at the EUI. Thank you also to Evi Pappa and to Árpád Ábrahám for their help throughout, and particularly during the last year of my Ph.D. Thank you to Petia Topalova at the IMF, and Marko Melolinna at the Bank of England

[1]Note that the opinions expressed in this thesis are those of the author and do not necessarily reflect the views of the Central Bank of Ireland.

for allowing me to undertake work on my thesis as an intern at their respective organizations. An especial thank you to my co-author, Benedikt Kolb. Thanks to the army of fellow researchers and friends, office-mates and flatmates who helped me along my way, who are too numerous to mention and hopefully know who they are. Lastly, thank you to my family: Tim, Laurie, and Harriet.

References

Asker, J., A. Collard-Wexler, and J. D. Loecker (2014). Dynamic Inputs and Resource (Mis) Allocation. *Journal of Political Economy 122*(5).

Gürkaynak, R. S., B. Sack, and E. Swanson (2005). Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking 1*(1), 55–93.

Romer, C. D. and D. H. Romer (2004, September). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review 94*(4), 1055–1084.

Chapter 1:

Monetary Policy Communication Shocks and the Macroeconomy

Represents joint work between Robert Goodhead and Benedikt Kolb.[1]

1 Introduction

On December 16, 2015, the Board of Governors of the Federal Reserve decided to increase the federal funds target rate range for the first time since June 2006. The move came as little surprise to financial markets, however. While in the previous Federal Open Market Committee (FOMC) meeting of October the target rate had been held constant, policymakers had indicated that a rate rise was likely, subject to a continuation of recent positive macroeconomic developments.[2]

Although the December rate hike officially marked the end of the zero-lower bound (ZLB) period of monetary policy for the United States, it was the October FOMC meeting that saw a revival in trading of near maturity federal funds futures contracts, which are used by market participants to bet on future Fed target rates. The market for federal funds futures, which began operating in 1988 and quickly became deep and liquid, has been used extensively to identify surprises in U.S. monetary policy and analyse their effect on financial markets, and more recently the real economy. The idea is simple: given that the prices of such contracts incorporate all information available to markets, they ought to embody the market expectation of future policy rates. Changes in the futures rate during the course of FOMC meeting days can thus be credibly interpreted as policy surprises, or mon-

[1]We thank Jason Allen, Geert Bekaert, Gabriel Bruneau, Fabio Canova, Wouter den Haan, Juan Dolado, Uros Djuric, Zeno Enders, Andrea Gazzani, Peter Hansen, Peter Karadi, Leonardo Melosi, Athanasios Orphanides, Evi Pappa, Esteban Prieto, Alejandro Vicondoa and Shengxing Zhang for helpful comments. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of Deutsche Bundesbank or Central Bank of Ireland.

[2]The press release of the October meeting states that "[t]he Committee anticipates that it will be appropriate to raise the target range for the federal funds rate when it has seen some further improvement in the labor market and is reasonably confident that inflation will move back to its 2 percent objective over the medium term". See http://www.federalreserve.gov/newsevents/press/monetary/20151028a.htm.

etary policy shocks. Data from futures markets therefore provide the econometrician with a means to separate the effects of changes in monetary policy from the underlying changes in macroeconomic conditions to which policy makers respond.

This so-called "high-frequency identification" approach has been used to examine the effects of changes in monetary policy rates on financial and macroeconomic variables, and also to highlight the role of information dissemination during FOMC meetings.[3] Instead, we focus on the informational content of the futures data to discern the effects of unanticipated actions and communication on macro variables. Given that the available maturity spectrum of futures rates spans the known dates of several future policy meetings, we can use differences in futures price reactions across the maturity spectrum to discern market expectations about future monetary policy moves.[4] We show that these changes in expectations in response to communication by the FOMC are powerful drivers of economic activity. Barakchian and Crowe (2013) rightly point to the increasingly forward-looking nature of monetary policy, insofar as the Federal Reserve (and other central banks) rely more heavily on forecasting when designing policy. However, we argue that there is also an important anticipation effect in the other direction: If financial markets are similarly forward-looking in their judgment of FOMC communication, and given that Federal Reserve communication has become more detailed about its future policy course, markets should react to announcements in a way that is reflected systematically over the spectrum of federal fund future maturities.

We propose a method to differentiate monetary-policy action and communication shocks from futures rates of different maturities. We employ a linear decomposition of futures price movements on FOMC meeting days in combination with an institutional arrangement: Since 1994, the FOMC has published its meeting days well in advance, so that market participants know the earliest possible date when future policy actions can be taken. We can therefore transform the movements in the various maturities of monthly federal funds funds rates (reflecting anticipated average target rates in future months) into movements in anticipated rates before and after the first future FOMC meeting.

Since changes in the target rate tend to persist, surprise rate changes today affect rates across the whole spectrum of available maturities into the same direction. However, additional information about potential policy changes in future meetings ("surprise communica-

[3]E.g. Barakchian and Crowe (2013), Gertler and Karadi (2015), Nakamura and Steinsson (2018). We expect interest by economic researchers in the federal funds futures market to increase in the near future, as the series available after the end of the zero-lower bound episode will soon be long enough for analysis again.

[4]For our sample, there are six monthly contracts liquid enough for analysis, from those betting on the average federal funds rate valid in the concurrent month to five months into the future.

tion") should affect only futures maturities after that date, and not before. We are therefore able to employ a simple yet credible recursive scheme to orthogonalise monetary policy "action shocks" (surprises about the actual target rate decision announced at an FOMC meeting) from monetary policy "communication shocks" (anything that changes market expectations about potential target rate decisions taken at future FOMC meetings during the current meeting).[5] So communication shocks are the linear component of observed variation that does not affect preceding futures maturities. Finally, we orthogonalise our shocks measure with respect to internal information revelation by regressing them on the Fed Greenbook data.

In this sense our approach is in the spirit of Gürkaynak et al. (2005), who offer a two-factor interpretation of monetary policy surprises and convincingly argue that a "target factor" (an effect similar across all maturities) and a "path factor" (increasing over maturities) are sufficient to explain futures rate movements on announcement days. Indeed, Barakchian and Crowe (2013) choose to only use the first one (a "levels effect" similar to the "target factor" in Gürkaynak et al., 2005), reasoning that "[s]ince the transmission of monetary policy is generally thought to occur via the impact of short rate changes on longer term (real) rates, it is this portion of the new information on rates that corresponds most closely to the relevant policy shock." (p. 959). We argue that this interpretation may not be justified, since some maturities react more strongly and consistently during FOMC meetings; these are the ones at the upper end of the six-month spectrum. This is difficult to align with the "levels effect" interpretation of Barakchian and Crowe (2013), which seems to leave out important information about how monetary policy shocks affect the economy. Instead, when explicitly distinguishing between action and communication shocks, we find the latter to have larger and more significant effects on macro aggregates. While this is in line with the interpretation of Gürkaynak et al. (2005), our method of obtaining the shocks allows for identification of more precisely defined monetary policy communication shocks that pertain to given dates in the future, since we do not apply a factor structure to the data.

When we purge the shocks of the component that relates to the revelation of hidden information by Fed policymakers during their announcements, we show that monetary policy communication shocks create contractions of industrial production, which are larger and more significant than for action shocks. Communication shocks also explain a larger share of variation in both production at business cycle frequencies and have effects that are more in line with narrative accounts of changes in the monetary policy stance during the

[5]Note that these shocks represent changes in expectations that may or may not be accurate *ex post* (i.e., news and noise shocks).

period. We show that our main findings are robust to a variety of specifications.

Finally, we extend our analysis to cover the zero-lower bound (ZLB) period using Eurodollar futures, which are available at longer maturities than the federal funds futures contracts; importantly, the contracts remain sufficiently liquid for analysis during the period. We use a decomposition similar to the one used for the federal funds futures and study the effects of the derived shocks on macro variables in a sample from 1994 to 2016. We find the effects of longer-term Eurodollar-derived communication shocks (three years into the future) to be stronger for inflation than for industrial production, with a price puzzle observed only for the more short-term Eurodollar-derived communication shock (one year into the future). This finding underlines our key message that central bank communication has significant macroeconomic effects.

2 Related Literature

Our analysis relates to the growing field of high-frequency identification literature, as well as to the related topics of central-bank information revelation and forward guidance.

General high-frequency identification literature. The literature on identification of monetary policy shocks employing high-frequency data goes back to Rudebusch (1998) and Kuttner (2001). Söderström (2001) argues that movements in the federal futures rates around an FOMC meeting are in fact a good predictor of target rate changes implemented in the following meeting. Faust et al. (2004) were the first to incorporate a structural shock series identified via changes in federal funds future rates into a VAR together with financial and macro variables.

Related to our analysis, Gürkaynak et al. (2005, GSS in the following) aggregate the informational content of the futures using factor analysis, and argue that two factors are sufficient to capture the correlation over the maturity spectrum. Analogous to the yield curve literature, they refer to these as the “current federal funds rate target factor” and “future path of policy factor”. GSS argue that the “path factor” reflects soft information on future policy actions during FOMC meetings and is important for the analysis of monetary policy statements. They perform an orthogonalisation of their shocks similar to ours, but do not use their shocks in a study of the macroeconomic system. Another decomposition of the movements in futures prices can be found in Gürkaynak (2005), who identifies “timing”, “level” and “slope” surprises. The decomposition omits a factor structure, and assumes observed variation to be a linear function of three structural shocks. In contrast, our restrictions identify two shocks of a different nature: each has the same interpretation, they merely apply to different future horizons. More recently, Swanson (2017) uses a factor analysis similar to GSS to distinguish between surprises in federal funds rate changes, forward guidance and LSAP effects. Three factors sufficiently describe the dynamics of underlying high-frequency changes in various returns on meeting days. The factors are then identified by rotating them such that the forward guidance and LSAP factor have no influence on yields of short-term assets, and by minimising the variance explained by the LSAP factor before the ZLB episode.[6] Swanson's analysis clearly shows the importance of monetary-policy communication for asset prices during the ZLB episode.

Our paper is also related to Barakchian and Crowe (2013, BC in the following), who show that for samples starting in 1988, monetary policy shocks identified via widely used recursive schemes lead to significant *increases* in output following supposedly contractionary monetary policy shocks. In contrast, a VAR with cumulated high-frequency shocks, computed as a single factor of the maturity spectrum, yields contractionary effects on industrial production in response to contractionary policy. BC suggest this might be due to a more forward-looking monetary policy after the 1980s, under which policy rates react contemporaneously to, or even before, changes in economic activity.[7] Our results, however, suggest that we cannot rule out that the monetary-policy shock in BC is in fact primarily driven by central-bank communication. Gertler and Karadi (2015) use an external-instruments approach to safeguard against simultaneity in a VAR including both a monetary policy shock measure and credit costs. However, the authors note that in the case that no further financial variables are considered, a recursive VAR such as the one we employ is appropriate for an analysis of monetary policy shocks. Moreover, the authors reject the GSS shocks as weak instruments, while we are able to show the importance of communication using a different methodology.

Fed information revelation. Romer and Romer (2000) are the first to show that the Fed transmit important information during FOMC meetings. Romer and Romer (2004) control for such information using the Fed-internal Greenbook forecasts, while Thapar (2008) introduced the approach to the high-frequency literature. More recently, Miranda-Agrippino

[6]While we find this method intuitive, we are not sure whether this identification would not also allow an alternative interpretation of the factors as action surprise, a pre-crisis communication component and a post-crisis communication component.

[7]For a similar argument, see Cochrane and Piazzesi (2002).

and Ricco (2017) adjust the instrument in Gertler and Karadi (2015) to account for autocorrelation and central-bank information revelation (again via Greenbook forecasts). They find unequivocally negative effects of a contractionary monetary policy shock, however, they do not identify the effect of communication surprises. Nakamura and Steinsson (2018) use a small structural model to distinguish two counteracting effects of FOMC announcements (say a surprise tightening of rates): They find that the (expansionary) Fed information effect on the natural interest rate is in fact larger than the (tightening) increase in real rates over the natural rate. Similarly, Jarocinski and Karadi (2018) distinguish monetary-policy action and information shocks by their different high-frequency effects on interest rates and stock prices, finding a strong expansionary role for information transmission in the Euro Area and U.S. While the authors mostly talk about central-bank information revelation and communication interchangeably, we identify communication shocks orthogonal to information revelation. We argue that over and above information revelation, communication about the near-future course of monetary policy matters, as markets update their expectations about the central bank's course of action.[8]

Forward guidance. Since our baseline sample runs from 1994 to 2008, we analyse communication shocks during times of conventional monetary policy. However, in our extension using Eurodollar futures we also find an important role of communication shocks, in particular for inflation.[9] Our study thus relates to a growing literature on forward guidance, i.e. the deliberate steering of the public's expectations by central banks sharing internal forecasts or committing to longer-term policies. Campbell et al. (2012) distinguish between "Delphic" forward guidance, or transmission of private central-bank information, and "Odyssean" forward guidance, which represents explicit commitments to a future policy course.[10] Andrade and Ferroni (2016) differentiate Delphic and Odyssean forward guidance for the euro area as follows: For a high-frequency increase in the term structure (measured by overnight-index swaps), inflation expectations (measured by inflation-linked swaps) will increase for Delphic, but fall for Odyssean forward guidance. Lakdawala (2016) uses a proxy-SVAR in conjunction with the two GSS shocks to distinguish between federal funds rate and forward guidance shocks. He finds an expansionary effect for contractionary

[8]Ben Zeev et al. (2017) identify "monetary policy news shocks" as in the TFP literature, but do not distinguish between information revelation and communication.

[9]This strong effect, which increases in the horizon of communication, is in line with the predictions of DSGE models, as studied by Del Negro et al. (2016) and McKay et al. (2016).

[10]See also Hansen and McMahon (2016) for a complementary approach using computational linguistics analysis to distinguish FOMC communication regarding current economic conditions and forward guidance.

forward guidance communication, which is rendered insignificant when controlling for the information set of the Fed using Greenbook and Bluechip forecast data. This contrast with our results might be explained by the shorter horizon of our communication shocks (within six months instead of one year), or by the way his external instruments approach includes data from earlier periods (back to 1979). Bundick and Smith (2016) use movements in 12-month ahead federal funds futures rates as a measure of forward guidance during the ZLB period in a VAR, and find the effects in line with predictions by a structural model. While the use of just one future interest rate facilitates comparison to their model, it rules out a distinction between monetary policy actions and communication as in our study.

3 Methodology

This section introduces our data, and outlines how we obtain changes in anticipated policy rates from changes in the price of futures contracts defined over calendar months. We then present a Cholesky decomposition that delivers identification of both the monetary-policy action and communication shock. Finally, we explain how we incorporate our shocks into a structural VAR model in order to examine their effect on macroeconomic variables.

3.1 Federal Funds Futures

Federal funds futures contracts were introduced on October 3, 1988, by the Chicago Board of Trade, and are the most widely used futures contract tied to the federal funds rate.[11] The use of these futures limits our sample to the period before the ZLB, since trading in the shorter maturity contracts effectively ceased at the onset of this period.[12]

Federal funds futures contracts allow market participants to place a bet in month t on the average effective federal funds rate during the concurrent or future months, denoted by $\bar{r}_{t+m}$, with $m \geq 0$. A buyer of the contract on day d in month t can commit to borrow federal funds at a fixed rate at the end of the month $t+m$, and we denote this futures rate by $f_{d,t}^{(m)}$.[13] Under no arbitrage, we have that the futures rate $f_{d,t}^{(m)}$ reflects the market

[11] See Moore and Austin (2002).

[12] We examine the ZLB period using Eurodollar futures in Section 5. However, our short baseline sample makes the likelihood of structural breaks in the transmission of monetary policy less likely (see e.g. Boivin and Giannoni, 2006, and D'Amico and Farka, 2011).

[13] Throughout we let t refer to the month, which is the frequency of our VAR, and we let d denote the particular day in given month t.

expectations of the average effective federal funds rate $\bar{r}_{t+m}$:

$$f_{d,t}^{(m)} = \mathbb{E}_{d,t}[\bar{r}_{t+m}] + \delta_{d,t}^{(m)}, \quad \forall m \geq 0,$$

where $\delta_{d,t}^{(m)}$ is a risk-premium term. Since Kuttner (2001), many authors have argued that the movements in the federal funds futures market observed on FOMC meeting days (labelled "jumps" in the following) capture a surprise component of monetary policy. Let us assume no change in the risk-premium $\delta_{d,t}^{(m)}$ for that short time window.[14] Then, a policy surprise can be computed as the difference in the futures rate at the end of the FOMC meeting day from that at the end of the previous day:[15]

$$\Delta f_{d,t}^{(m)} \equiv f_{d,t}^{(m)} - f_{d-1,t}^{(m)} = \Delta\mathbb{E}_{d,t}[\bar{r}_{t+m}], \quad m > 0$$

Although federal funds funds contracts are now available for maturities as far as three years into the future, only the first six maturities of futures are considered liquid enough to be treated as efficient financial markets over our time-period (see BC, p. 959). We use daily changes in futures rates around FOMC dates for the maturities $m \in \{0, 5\}$. GSS find that using intraday or daily data makes virtually no difference for the post-1994 sample.[16]

3.2 From Futures Rate Changes to Expected Policy Rate Changes

The federal funds futures prices give us changes in market expectations about policy rates on FOMC meeting days. To analyse surprises regarding current monetary-policy actions

[14] Piazzesi and Swanson (2008) have shown that the risk-premium in the federal funds futures market is sizeable and time-varying, but only at business-cycle frequencies.

[15] Note that for contracts on the current month, agents will already have observed a component of the realization of $\mathbb{E}_{d,t}[\bar{r}_t]$, because $d-1$ days of that month have already elapsed – we take into account that the overnight rate refers to the night *after* day d, (Hamilton, 2008, p. 378). We follow Kuttner (2001) in scaling the futures rate for the concurrent month, $\Delta f_{d,t}^{(0)}$, by the ratio of number of days in the month, M, over the number of days remaining after the meeting, $M-(d-1)$. Thus we obtain a corrected measure $\Delta f_{d,t}^{*(0)}$:

$$\Delta f_{d,t}^{*(0)} = \frac{M}{M-(d-1)} \cdot \Delta f_{d,t}^{(0)}.$$

This scaling factor becomes very large at the end of the month (up to 31 for $M = d = 31$). We therefore again follow Kuttner (2001, pp. 529f.) and use the change in the futures rate of next month ($\Delta f_{d,t}^{(1)}$ in place of $\Delta f_{d,t}^{*(0)}$) for meetings within the last three days of a month, provided there is no meeting next month.

[16] "[F]or samples that exclude employment report dates, or samples that begin in 1994, the surprise component of monetary policy announcements can be measured very well using just daily data." (GSS, p. 66). Nakamura and Steinsson (2018) argue for the use of shorter time windows, but they employ longer-term interest rates and a sample that covers more long-term forward guidance.

and communication about future rate decisions, we first need to extract measures of the market expectation of average rates within two intervals: between the current and the next FOMC meeting, and after the next meeting.

However, our six usable futures maturities are defined over calendar months, while meeting days are unevenly spread out across the months in the maturity spectrum.[17] Although we are able to use six rate jumps that span the next five months into the future (and therefore at least three future meetings), the futures contracts cannot represent six individual policy surprises, since monetary policy can change at most another three times. To obtain average rates expected by the markets between meeting dates, we follow a linear extraction method. Similar methods are used in GSS and Gürkaynak (2005), however we add an iterative weighted averaging procedure to reduce noise and use all available information.

Let $\Delta\rho^0_{d,t}$ denote the change in the expected rate between the current and next FOMC meeting, and $\Delta\rho^1_{d,t}$ the change in the expected rate after the next meeting.[18] Figure 1 illustrates the timing with an example: the FOMC meeting taking place on May 17, 1994, and the three meetings that followed (those of July 6, August 16 and September 27). The figure displays the five calendar months into the future from the month of the meeting, and the jumps in the futures rate for the contract associated with that month.

So our starting point is to obtain the change in market expectations on policy rates valid between FOMC meetings, $\Delta\rho^j_{d,t}$, from changes in futures prices $\Delta f^{(i)}_{d,t}$. To do so, we work iteratively forward, starting with $\Delta\tilde{\rho}^0_t$, which is set equal to the corrected jump in the futures rate for the concurrent contract[19],

$$\Delta\tilde{\rho}^0_{d,t} = \Delta f^{*(0)}_{d,t}.$$

Since contracts are defined over average interest rates for calendar months, we know the price change of the futures contract for the month of the next meeting (July in our example), $f^{(I)}_{d-1,t}$, must be a weighted average of the expected interest rate carried forward from the previous meeting, and that expected to be set in the next (indexed 0 and I, respectively),

[17]FOMC meetings take place roughly every six weeks, usually in late January, April, July and October and mid March, June, September and December. The meetings for late July and October often take place in early August and November instead.

[18]We could potentially also create measures of expectation changes for policy rates between the next meeting and the one after that. Experimentation showed that the results for the two communication shocks looked very similar, so we focus on the case of one communication shock here.

[19]$\Delta\tilde{\rho}^0_{d,t}$ and the uncorrected $\Delta\tilde{\rho}^1_{d,t}$ would correspond to $mp0$ and $mp1$ in Gürkaynak (2005).

Figure 1: Illustration of the Conversion of $\Delta f_{d,t}^{(m)}$ to $\Delta\rho_{d,t}^{j}$

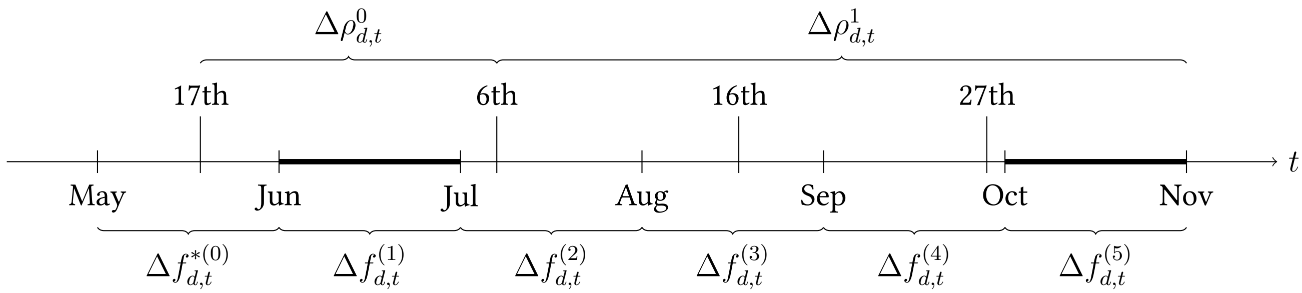

Notes: The timeline shows the months May to November 1994, as labelled below the axis. Above the axis are the days of FOMC meetings. The jumps in the monthly futures rates, $\Delta f_{d,t}^{(m)}$, are indicated below the axis, above it are the jumps in expected federal funds rates between meetings, $\Delta\rho_{d,t}^{j}$. Months without FOMC meetings are marked by a thick line.

$$\Delta f_{d,t}^{(I)} = \frac{d_I - 1}{M_I} \cdot \Delta\tilde{\rho}_{d,t}^{0} + \frac{M_I - (d_I - 1)}{M_I} \cdot \Delta\tilde{\rho}_{d,t}^{1},$$

where d_I refers to the day of the next meeting, and M_I to the number of days in the month of the next meeting. Therefore:

$$\Delta\tilde{\rho}_{d,t}^{1} = \frac{M_I}{M_I - (d_I - 1)} \left(\Delta f_{d,t}^{(I)} - \frac{d_I - 1}{M_I} \cdot \Delta\tilde{\rho}_{d,t}^{0} \right). \tag{1}$$

Because the futures rate jumps are likely to be noisy, and since such noise could be weighted up by the scaling terms, we utilise the extra information represented by changes in futures rates for calendar months without meetings. Thus, if there is no meeting in the month following the meeting, we create a final version of $\Delta\rho_{d,t}^{0}$ by taking a weighted average of this measure with the jump in the next month's futures rate, as follows:

$$\Delta\rho_{d,t}^{0} = \frac{M_0 - (d_0 - 1)}{M_0 - (d_0 - 1) + M_1} \cdot \Delta\tilde{\rho}_{d,t}^{0} + \frac{M_1}{M_0 - (d_0 - 1) + M_1} \cdot \Delta f_{d,t}^{(1)}.$$

We are therefore using the fact that the jump in the price of next month's contract (June in our example) is an equally valid measure of the surprise in the cases that there is no meeting next month (since a single target rate will hold over the whole period). We employ the same strategy to create $\Delta\rho_{d,t}^{1}$ whenever there is no meeting in the month following a given meeting.[20] This approach ensures that the futures rate changes that occur towards the

[20] In the case that there is a meeting next month we do not perform the weighting. Further, we perform

end of the month (with higher d) will get a smaller weighting in the convex combination. Thus, the procedure will reduce potential idiosyncratic noise in the futures changes.

As mentioned by Gürkaynak (2005), a potential limitation of this method is the possibility of rate changes during unscheduled meetings. The FOMC can deviate from its published meeting schedule if circumstances require it and has done so several times in our sample.[21] We only rely on scheduled meetings here. However, if markets were to incorporate an endogenous probability of emergency meetings into their pricing, this could be problematic for our identification scheme. However, given that we take differences of futures prices on meeting days, the occurrence of unscheduled meetings will only bias our shock measures when the market expectations about the likelihood of an unscheduled meeting are changed *during the day of the previous* (scheduled) *FOMC meeting*. From inspection of the minutes, the committee has never mentioned unscheduled meetings during the meetings that preceded them. Therefore, we do not believe the effect of unscheduled meetings presents a serious concern.

3.3 From Expected Policy Rate Changes to Structural Shocks

Given the surprises in the policy rates, $\Delta\rho^j_{d,t}$, we want to obtain the structural shocks that generate these changes in expectations. Target rate changes by the Fed are highly persistent (as shown in the paper of Coibion and Gorodnichenko, 2012, for example), and therefore any rate decision communicated during the FOMC meeting will shift market expectations across the spectrum of maturities. This is what GSS and BC refer to as their "target factor" and "level factor", respectively. Thus an unexpected policy rate change by the FOMC will lead to an updating of expectations about the current as well as about future rates, as the policy rate is likely to persist: Without any additional information about the future course of policy action, markets can take the policy rate to be the new *status quo*. We refer to these surprise announcements of immediate policies as action shocks.

We also aim to quantify the effects of an important second component to FOMC meetings, namely communication about potential future actions. We therefore posit the existence of orthogonal information about future policy changes contained within the announcement. The central banker may reveal surprise information about a rate change, and simultaneously deliver independent surprise information relating to future policy. Ad-

this operation during the iterative extraction, in the sense that where appropriate the weighted version of the previous surprise is used to extract the next, which then may be weighted, etc.

[21] The dates were 04/18/1994, 10/15/1998, 01/03/2001, 04/18/2001, 09/17/2001, 08 and 17/10/2007, 01 and 22/09/2008, 03/11/2008 and 10/08/2008, see Appendix F for details.

ditional surprise communication about potential policy actions in these future meetings ought to affect all futures rates *after* the future meeting (and associated expected policy move), i.e. $\Delta\rho^1_{d,t}$, but not rates *before* them, $\Delta\rho^0_{d,t}$. In contrast, action shocks affect both expectations $\Delta\rho^0_{d,t}$ and $\Delta\rho^1_{d,t}$ similarly. This recursive system motivates the use of a Cholesky decomposition of the expectations jump vector.

Formally, the changes in expectations about the future monetary policy rate, $\Delta\rho^0_{d,t}$ and $\Delta\rho^1_{d,t}$, are decomposed into two orthogonal shocks: surprises about monetary decisions today (the action shock, $\tilde{\varepsilon}^A_{d,t}$) and surprise communication about potential futures actions (the communication shock, $\tilde{\varepsilon}^C_{d,t}$) as follows:

$$\mathbf{\Delta R}_{d,t} \equiv \begin{bmatrix} \Delta\rho^0_{d,t} \\ \Delta\rho^1_{d,t} \end{bmatrix} = \begin{bmatrix} m_{11} & 0 \\ m_{21} & m_{22} \end{bmatrix} \cdot \begin{bmatrix} \tilde{\varepsilon}^A_{d,t} \\ \tilde{\varepsilon}^C_{d,t} \end{bmatrix} = \mathbf{M} \cdot \mathbf{E}_{d,t}. \tag{2}$$

Rearranging, we obtain the expression for the vector of structural shocks:

$$\mathbf{E}_{d,t} = \mathbf{M}^{-1} \cdot \mathbf{\Delta R}_{d,t},$$

where $\mathbf{M} = \text{chol}\left(\text{var}(\mathbf{\Delta R}_{d,t})\right)$. Note that these operations are conducted at the frequency of the meetings, in the sense that we extract structural shocks from a jump vector with observations only on meeting days. Because we restrict our analysis to days with scheduled meetings, there is never more than one meeting per month, meaning that we can drop the d subscript from our shock series. We enter a zero value to the shock series for the months without meetings, as in BC.

Figure 2 shows the action and communication shock series, as well as the BC shock series (the first principal component over all maturities) to serve as a basis for comparison.[22] First, we see that the size of both action and communication surprises during FOMC meetings are relatively small in size (with standard deviations of 3.91 and 4.42 basis points). This implies that markets generally anticipate decisions with a high precision. The shock series show increased volatility around 2001, after the bursting of the dot-com bubble and September 11, and during the immediate run-up to the financial crisis. Generally, the fac-

[22]Our action shock is significantly positively correlated with both factors in BC and the first GSS factor, while our communication shock is positively correlated with only the first BC factor and the second GSS factor. Both shocks are positively correlated with the shock of Nakamura and Steinsson (2018) and with an updated Romer and Romer (2004) shock series. We conclude that our shocks capture information from these existing shock series, but are not reducible to any of them. Furthermore, the "level factor" interpretation of BC regarding their shock may be questioned, given its significant positive correlation with our communication shocks. For details, see Appendix A.

tor approach amalgamates the more idiosyncratic movements our two shock series. The larger movements in the communication shock series in the early part of the sample are not present in the BC shock, however.

Figure 2: Shock Series

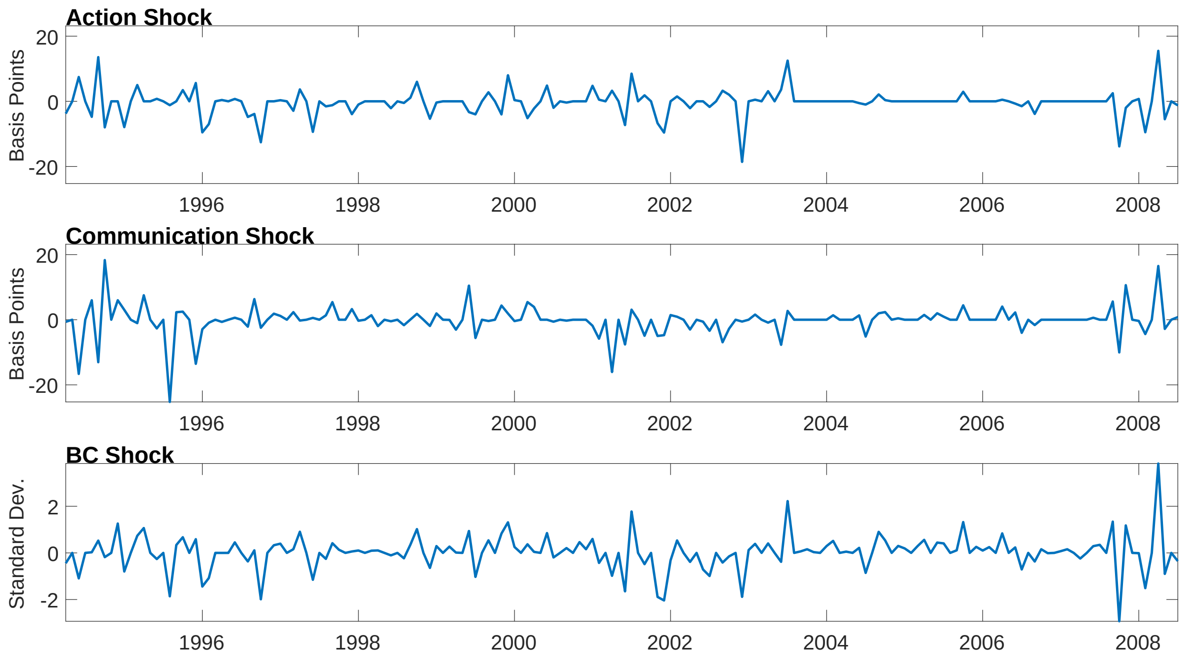

Notes: Our action and communication shock series, $\tilde{\epsilon}_t^A$ and $\tilde{\epsilon}_t^C$. We also display the shock series of Barakchian and Crowe (2013, "BC"), formed of the first principal component of the six federal funds rate maturities, for reference. The R^2 from regressions of the BC shock on the action and communication shock are 0.54 and 0.35.

In our baseline analysis below we will orthogonalise our shocks $\tilde{\varepsilon}_t^j$ with respect to Fed-internal Greenbook forecasts.[23] This will cleanse the potentially superior central-bank information which could be transmitted to the public during FOMC meetings.[24] Greenbook forecasts are made public after with a lag of five years and therefore are not known by markets at the time of central bank announcements.[25] The cleansed shocks ε_t^j are the residual

[23]Again, such orthogonalisations have been used since Romer and Romer (2004).

[24]If the FOMC had a tendency to reveal new positive forecasts regarding output and inflation at the same time as it increased interest rates, then this would likely bias our estimation of the contractionary effects of policy towards zero, making our results under-estimate the true responses.

[25]We include similar Greenbook variables as BC, although like Ramey (2016) we use only the Greenbook forecasts, while BC employ the difference between Blue Chip and Greenbook indicators. The variables used are: (1) contemporaneous unemployment, (2) contemporaneous output growth and its lag and first two leads; (3) the GDP deflator and its lag and first two leads; (4) the previous values of the output growth forecasts; (5) the previous values of the GDP deflator forecasts.

from an OLS regression of $\tilde{\varepsilon}_t^j$ on a vector of Greenbook forecasts GB_t:[26]

$$\varepsilon_t^j \equiv \tilde{\varepsilon}_t^j - \hat{\boldsymbol{\beta}}^{OLS} \cdot \text{GB}_t. \tag{3}$$

Finally, we cumulate the shocks over time to form a monthly time series of policy surprises in levels, as in BC and Romer and Romer (2004)[27], labelled S_t^A and S_t^C, where

$$S_t^j = \sum_{i=0}^{t} \varepsilon_t^j, \quad j \in \{A, C\}.$$

3.4 VAR Setup

We want to gauge the effect of our two measures of policy surprises on (seasonally adjusted) monthly industrial production (IP) and consumer price inflation (CPI) with this structural VAR:

$$\mathbf{Y}_t \equiv \begin{bmatrix} \log(\text{IP}_t) \\ \log(\text{CPI}_t) \\ S_t^A \\ S_t^C \end{bmatrix} = \mathbf{C}_c + \mathbf{C}_d \cdot t + \sum_{l=1}^{\text{lags}} \mathbf{C}_l \mathbf{Y}_{t-l} + \mathbf{D} \cdot \epsilon_t \tag{4}$$

We estimate the model with a constant $\mathbf{C}_c$ and a deterministic trend $\mathbf{C}_d$, using twelve lags in our baseline model.[28] Given the small size of our VAR following the literature, one could be concerned that our four-variable VAR system might be too small to be informationally sufficient, we run the "fundamentalness" test suggested by Forni and Gambetti (2014) on our system. The test fails to reject the null of informational sufficiency of the system (and thereby also the shocks) in various specifications (see Appendix B).

As in Romer and Romer (2004) and BC, the VAR is recursive, so that monetary policy surprises cannot affect IP and CPI in the same period (but are allowed to react to them).

[26]Note that in our VAR analysis, we will deal with the issue of generated regressors via a system bootstrap.

[27]Note that these series are I(1) by construction, and will be entered directly into the VAR in this form (as in Romer and Romer, 2004, and BC). However, the argument of Sims et al. (1990) should hold, insofar that "the OLS estimator is consistent whether or not the VAR contains integrated components, as long as the innovations in the VAR have enough moments and a zero mean, conditional on past values of [the vector of endogenous variables]" (p. 113).

[28]The Bayesian information criterion proposes one lag, and the likelihood ratio test 14. We settle for twelve lags as in Faust et al. (2004). We show that our results are robust to different numbers of lags in Section 4.3. We gratefully acknowledge the use of code from the VAR toolbox by Ambrogio Cesa-Bianchi, made available on his personal website: `https://sites.google.com/site/ambropo/MatlabCodes`.

We need to make the assumption that markets do not observe the monthly observations on industrial production and inflation in real time (as e.g. also in Bundick and Smith, 2016), which we find plausible.

Given the recursiveness of the VAR, it becomes more efficient to simply enter in purged versions of $\Delta\rho^{j}_{d,t}$ (cumulated) into the SVAR system of Equation (4), since the orthogonalization represented by Equation (2) is then conducted in a single step, without the need to adjust the VAR standard errors for the fact that the decomposed shocks result from estimation. Thus, to effect the analysis we replace S^{A}_{t} and S^{C}_{t} with purged (and cumulated) versions of $\Delta\rho^{A}_{d,t}$ and $\Delta\rho^{C}_{d,t}$ respectively. Confidence bands for the VAR do need to be adjusted for the regression to purge the shocks of Fed internal information, as in Equation (3), and this is done via a system bootstrapping procedure (for all specifications involving purged shocks).

4 Results

Here, we present evidence that short-term surprise communication by the Fed matters substantially for the business cycle. In fact, based on impulse responses, forecast-error variance and historical decompositions, we argue that the effects on macro variables of what are generally termed "monetary-policy shocks" are much more driven by communication about the short-term path of policy than surprise actions. This result holds also when controlling for potential transmission of Fed-internal information during FOMC press conferences (the information-revelation channel).

4.1 A Study of Shocks without Purging Hidden Information

First, we compare our results to BC using our cumulated shocks not orthogonalised with respect to the Greenbook foreceasts, i.e. $\tilde{S}^{j}_{t} = \sum_{i=0}^{t} \tilde{\varepsilon}^{j}_{t}$, where $j \in \{A, C\}$. This allows us to compare action and communication shocks with the single factor in the BC baseline results, as well as to contrast the responses to our orthogonalised baseline shocks below.

Figure 3 shows the impulse responses of (log) industrial production and consumer prices to an action shock $\tilde{S}^{A}_{t}$ and a communication shock $\tilde{S}^{C}_{t}$. Throughout the structural shocks are 10 basis points rate increases.[29]

[29] 10 basis points correspond more than two standard deviations of our shocks, which are 3.91 ($\tilde{S}^{A}_{t}$) and 4.42 basis points ($\tilde{S}^{C}_{t}$). Moreover, a 10 basis point increase can be straightforwardly translated into a 25 or 100 basis point increase given the linearity of the model.

Figure 3: Responses of $\log(\text{IP}_t)$ and $\log(\text{CPI}_t)$ to Our Shocks

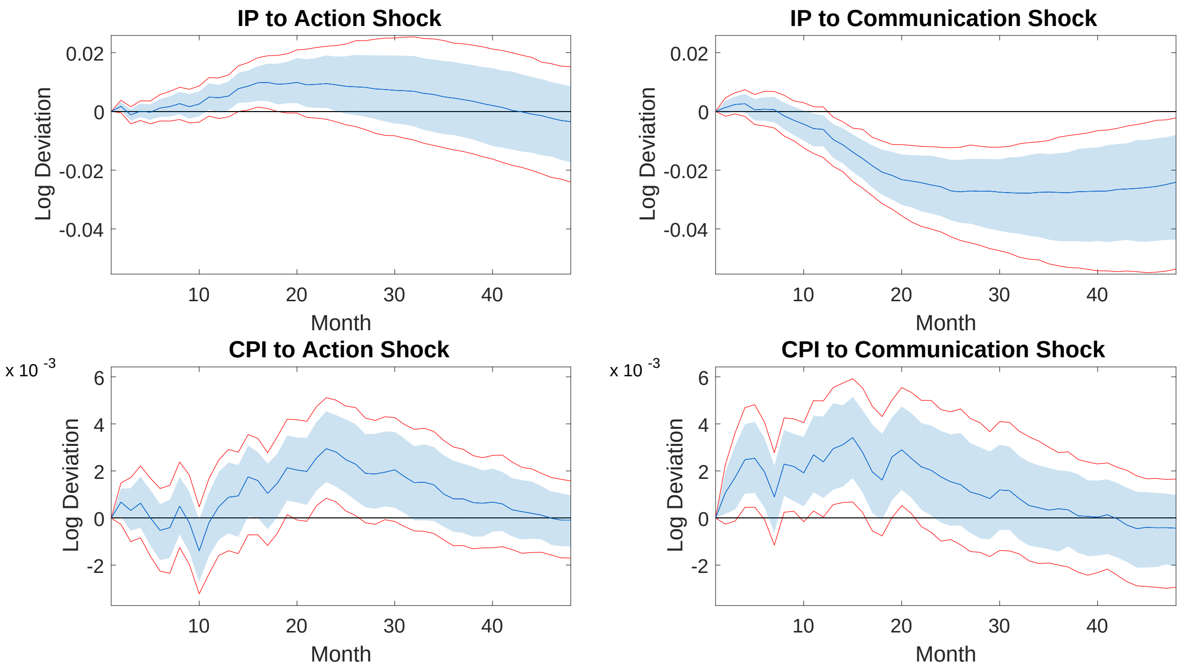

Notes: Impulse responses from our four-variable VAR, with $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$ and the cumulated action and communication shock series, $\tilde{S}_t^A$ and $\tilde{S}_t^C$ (not orthogonalised to Greenbook forecasts). The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels. Responses are shown to a 10 basis point positive shock to the interest rate.

We see that the reaction of production (IP) to a surprise increase in the expected policy rate is negative at the 90% confidence level only for the communication shock. Here IP falls by about three percent after a ten basis points surprise increase in the funds rate, which is a sizable effect (we discuss the magnitudes of the responses below for our baseline model). The IP reaction to the action shock, on the contrary, displays a counter-intuitive, significant positive reaction after around 18 months. Inflation shows an equally counter-intuitive positive reaction to both shocks. This price puzzle is a widespread finding in the high-frequency literature (see e.g. Thapar, 2008, and BC). The increase in production after a supposedly contractionary surprise monetary-policy action is more problematic for our interpretation of the shock, but can be explained by information transmission by the Fed during FOMC meetings once we control for Greenbook data. However, here we already see that the communication shock does not display this counter-intuitive effect on production.

To contrast our findings with those from the existing high-frequency identification literature, we compare our shock responses to the ones from the single factor over the six federal funds futures as used in BC. Figure 4 superimposes the results from such a three-variable VAR (repeated over columns) on those from Figure 3 above. We see that the responses to the BC factor shock, in particular its significant negative effect on industrial production, are in line with our communication shock, but not the action shock.

Figure 4: Comparison to Barakchian and Crowe (2013)

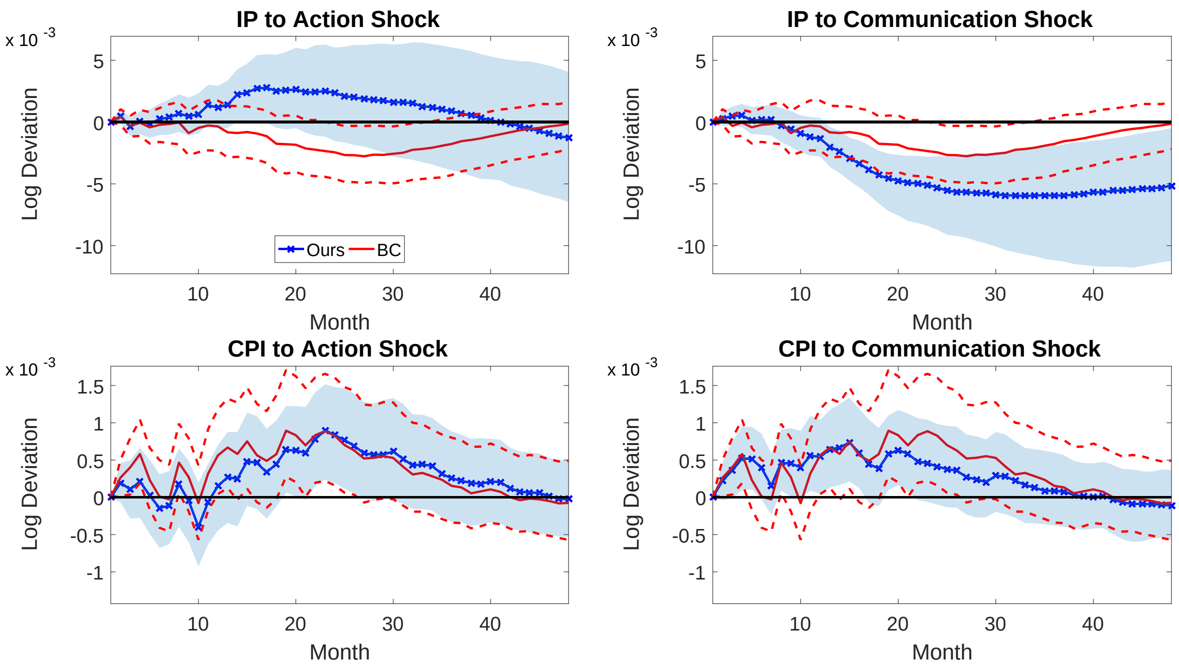

Notes: Impulse responses from our four-variable VAR, with $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$ and the cumulated action and communication shock series, $\tilde{S}_t^A$ and $\tilde{S}_t^C$ (not orthogonalised to Greenbook forecasts), together with the responses to the factor ("level") shock from Barakchian and Crowe (2013) in red, estimated in a 3-variable system (thus identical responses are repeated across each row). Responses are to a 10 basis point positive shock to the interest rate. The median and confidence intervals (at 90%, blue for our VAR and dashed red for BC) were obtained from bootstrapping each VAR model 500 times. Responses are to a one standard deviation positive shock to interest rates.

BC make a very convincing case that monetary policy in the U.S. has become more forward-looking after 1994.[30] However, we believe that also Fed communication, and its reception by the markets, has become more forward-looking during this time. This is reflected in our finding that, in post-1994 data, it is not monetary-policy surprise rate changes themselves, but rather surprise central bank communication about its future course of ac-

[30]They show that therefore wide-spread recursive identification schemes fail for samples starting in the late 1980s. We confirm this finding in Appendix C: A recursive identification as in Christiano et al. (1996) within a three-variable SVAR including log IP, log CPI, and the federal funds rate, delivers a significantly positive response to both production and inflation after a surprise policy-rate increase.

tion that affects economic activity in the way expected from a "standard" monetary policy shock. Therefore, while an interpretation of a first factor as "level" effect might hold for the GSS sample, it should not be generalised. Instead, we argue that it is rather surprise communication about the short-term path of monetary policy that leads to effects usually associated with "monetary policy shocks".

4.2 Baseline Results: Controlling for Fed-Internal Forecasts

The shocks above likely contain two elements: Surprise action or communication by the Fed, but also transmission of Fed-internal information about the likely future course of the macroeconomy. This information-revelation channel[31] might counteract the pure action or communication effect: A surprise hike in rates will likely dampen the economy, but could also signal a Fed outlook on the economy which is more expansionary than markets anticipated, and thus have stimulative effects.[32] To isolate our action and communication shocks from such information transmission, we show results from such shocks orthogonalised on Fed-internal Greenbook data, i.e. the ε_t^j shocks in (3) above.

[31] Again, see e.g. Romer and Romer (2004), Campbell et al. (2012), or Nakamura and Steinsson (2018).

[32] See Jarocinski and Karadi (2018) for examples of such "information shocks".

Figure 5: The Effects of Shocks Orthogonal to Greenbook Forecasts

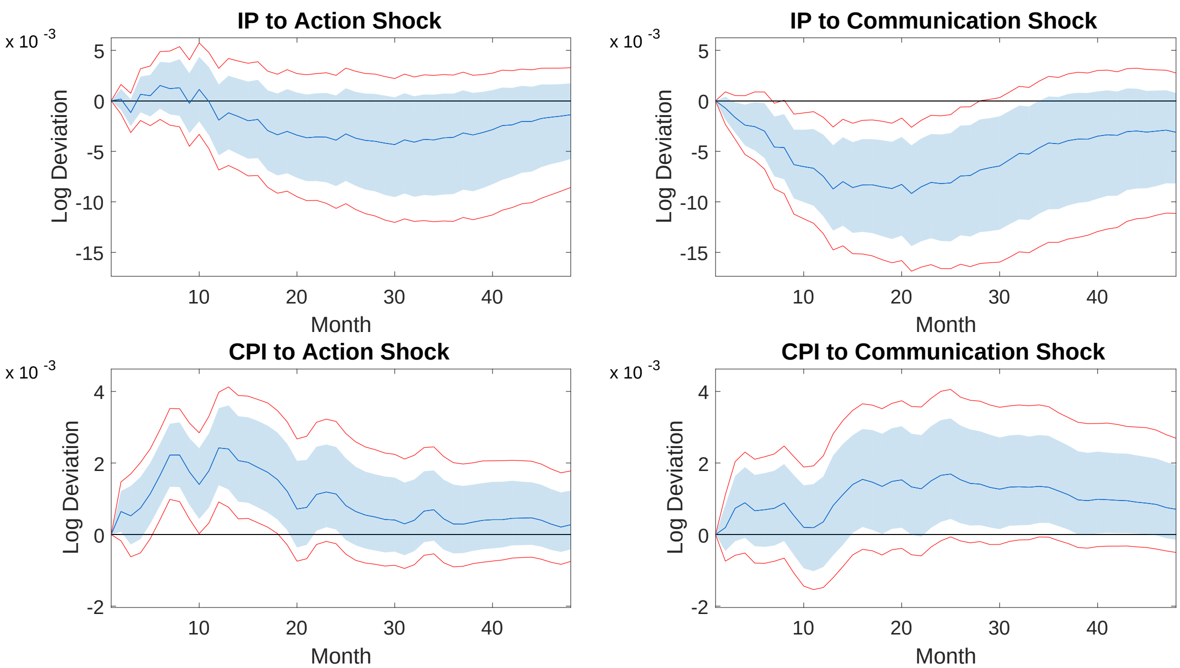

Notes: Impulse responses from our four-variable VAR, including $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$ and the cumulated action and communication shock series, S_t^A and S_t^C, orthogonalised to the Fed's Greenbook forecasts. The median response and confidence intervals were obtained from bootstrapping the VAR 5000 times through both the purging regression and the VAR, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels. Responses are to a 10 basis point positive shock to the interest rate.

Figure 5 shows that results are affected by this orthogonalisation. In fact, for the action shock, information revelation seems to explain most of the counter-intuitive expansionary effect on IP: After the purging, we have a positive effect on IP which is much smaller in size and significant only in the short run. In contrast, the contractionary effect of communication on IP is slightly smaller when controlling for information. Nevertheless, the effect of communication on IP remains significant and sizable with the expected sign. The price puzzle is not solved by controlling for information transmission, but is reduced to only borderline significance for the communication shock. Thus we find that our key result, i.e. the importance of FOMC communication about its near-future policy for the business cycle, is preserved even in the case that we control for potential contemporaneous information revelation on the part of the FOMC.

In magnitude the effects of our shocks are large relative to the literature. Ramey (2016) summarises existing estimates of the effects on industrial production of 100 basis points rises in the federal funds rate, and the maximum reported decrease is typically less than 5% (from BC), and usually closer to 1% (Christiano et al., 1996, find 0.7% after 24 months). Our communication shock would deliver a negative 9.17% trough 21 months after the shock hits. However, a 100 basis point surprise would exceed 25 standard deviations of our shocks (3.27 and 3.94 basis points for the purged action and communication shock respectively). Since our shock series are measures of purely unanticipated changes in the federal funds rate on the FOMC meeting day, they are small relative to the shock series employed in existing research that does not use high-frequency identification. The stronger effect of our shock series relative to that of BC is interesting, and is partially explained by the fact that our communication shock has stronger, negative effects than the action shock. To the extent the BC shock amalgamates both our shock series, it would follow that their estimated effect should be smaller than ours.

Forecast-Error Variance Decomposition. Table 1 depicts the shares of our purged shocks in a forecast-error variance decomposition of both macro variables at different horizons. The share of the communication shock is larger than that of the action shock for industrial production at longer horizons. Central bank surprise communication seems to have larger effects on production than surprise actions, though the action shocks have large effects on the variability of inflation.

Table 1: Forecast-Error Variance Decomposition at Business Cycle Frequency

	Horizon (months)	IP_t	CPI_t	S_t^A	S_t^C
IP_t	12	0.7412	0.0618	0.0138	0.1832
	18	0.4913	0.1591	0.0469	0.3028
	24	0.3166	0.2437	0.0861	0.3537
	36	0.1908	0.3408	0.1410	0.3275
CPI_t	12	0.1128	0.6271	0.2308	0.0292
	18	0.1378	0.5013	0.3082	0.0528
	24	0.1338	0.4721	0.3146	0.0795
	36	0.1121	0.4451	0.3090	0.1338

Notes: Contribution of our shocks to a forecast-error variance decomposition of IP and CPI at the 12, 18, 24 and 36 month horizons from our baseline four-variable purged VAR. The identified two shocks are $S^j_{d,t}$, $j \in \{A, C\}$. "IP" and "CPI" shocks are not identified.

Historical Decomposition. We present the historical decomposition of industrial production with respect to the purged shocks in Figure 6.[33] Prior to 2002 the share of production swings associated with action shocks is larger than that of the communication shock. The effects of the shocks seem to move in four larger cycles over the sample: two hawkish ones, 1995 to mid-1997 and 2001 to 2006, and two dovish ones, around 1998 to 2001 and 2006 to 2008.

Figure 6: Historical Decomposition of $\log(\mathrm{IP}_t)$

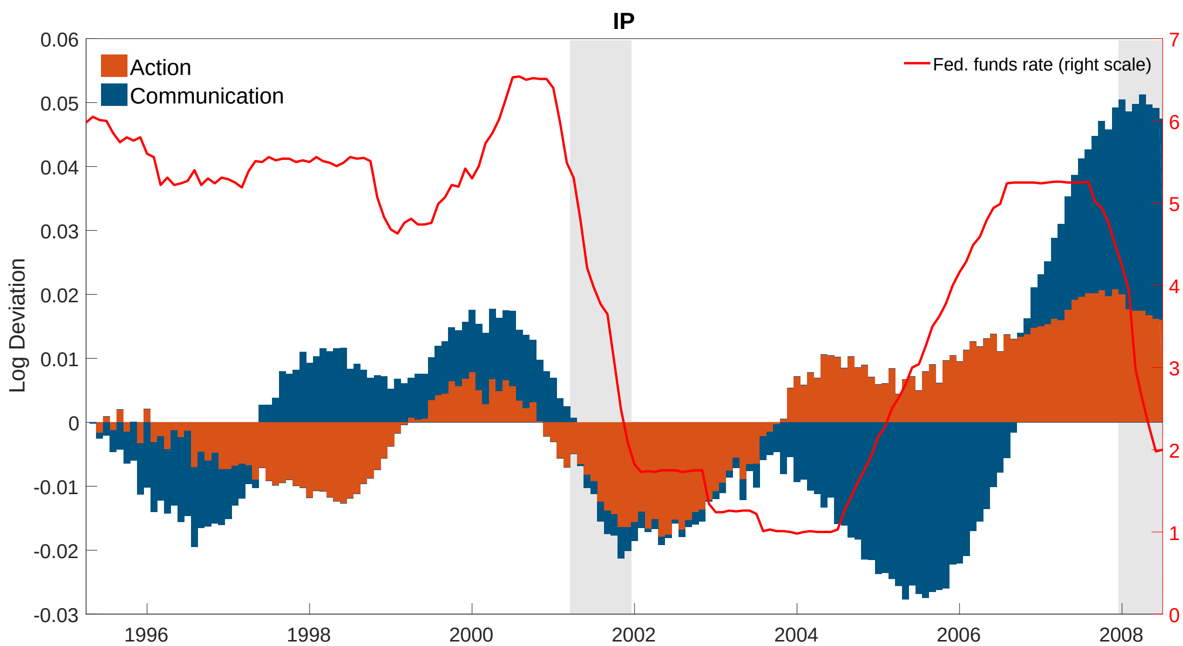

Notes: Historical decomposition of $\log(\mathrm{IP}_t)$ in our four-variable VAR, including the variables $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ and the cumulated purged action and communication shock series, S_t^A and S_t^C . The bar plots are stacked, so their height above (below) the zero-axis represents the cumulative historical contribution of our monetary shocks to industrial production above (below) its unconditional mean. We also display the federal funds rate (right-hand scale) for reference. Grey areas denote NBER recession periods.

The first expansionary episode (1998 to 2001) coincides with the last phase of the so-called "Greenspan put", i.e. the conjecture that the Fed systematically eased policy in reaction to deteriorating stock market conditions during the period. The second contractionary episode (2001 to 2006) was one of unstable growth and several corporate scandals involving American enterprises. Even though the Fed kept policy rates at low levels, markets seem

[33]Given the puzzling responses of prices, we choose to examine only variation in industrial production. However, we discuss the decomposition of CPI for the Eurodollar case below.

to have expected more policy easing before 2004.

Generally we find mixed evidence for the “monetary excesses” view of John Taylor, who argues that monetary policy had remained too lax for too long and contributed to an unsustainable housing boom in the U.S. during the period preceding the financial crisis (Taylor, 2009). Between 2002 and 2004, the effects of monetary surprises are predominantly contractionary in their contribution to fluctuations in industrial production. During the gradual increase of the federal funds rate from June 2004 onwards, the rate increases seem to have been of little surprise to markets (the action component is small), whereas the communication of the continued increases seems to have lowered market sentiment (as visible in the negative contribution of this shock over the period). Only after 2006 there is some expansionary effect as markets seem to have expected more contractionary actions. Finally, after the beginning financial-market turmoil in August 2007, both actions and communication successfully helped to stabilise production until futures markets became stuck at their zero-lower bound in mid-2008.

To summarise, distinguishing between central-bank action and communication shocks adds detail to our understanding of the recent monetary policy history of the U.S., and our communication shock seems broadly in line with common narratives.

4.3 Robustness

Figure 7 contrasts the impulse responses of our standard four-variable VAR with twelve lags to those of versions with 8, 10 and 14 lags.[34] The responses of the latter mostly fall into the 90% confidence bands of our baseline model.

[34] 14 is the number suggested by the likelihood test. Note that we cannot use more lags (BC use 36), since we would have insufficient degrees of freedom for our four-variable VAR.

Figure 7: Robustness Check – Different lags

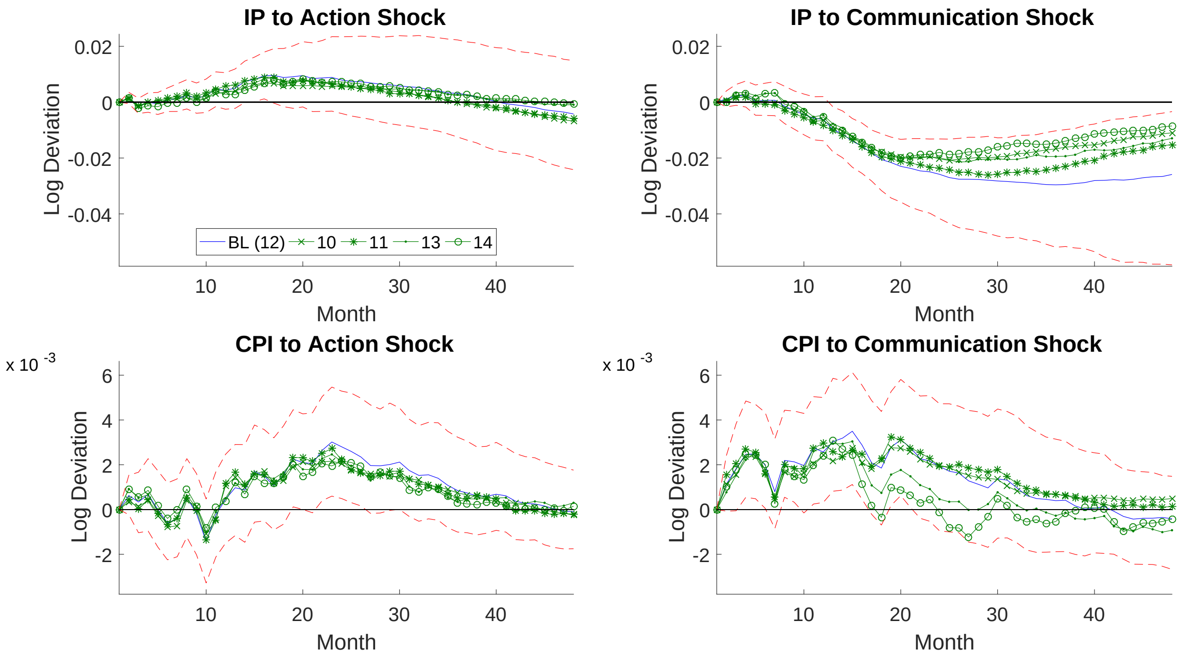

Notes: Impulse responses of $\log(\mathrm{IP}_t)$ in our $\log(\mathrm{CPI}_t)$ in our four-variable VAR with twelve lags (median blue, 90% confidence band in red), as well as VARs with 8, 10 and 14 lags (median response). The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at the 90% significance level. Responses are to a 10 basis point positive shock to the interest rate.

We also estimate separate three-variable VAR systems, loading in one shock at a time. We do this to respond to any concerns regarding the efficiency of our baseline VAR: In our four-variable system, the shocks are allowed to respond endogenously to each other, when in fact these interaction effects should be limited, given the shocks are externally identified and orthogonal to each other by construction. Figure 8 show that our main findings are robust to this setup: The action shock has no more counter-intuitive positive effect on IP, while the communication shock continues to have a significant effect.

Figure 8: Robustness Check – Separate 3-Variable VAR Systems

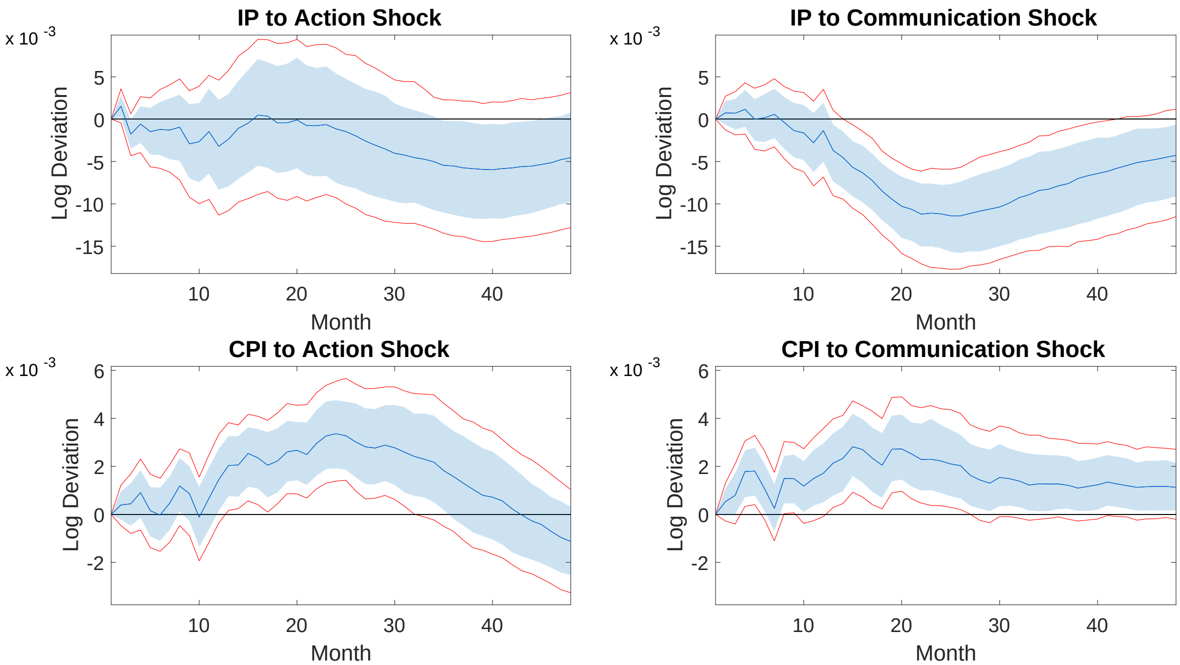

Notes: Impulse responses from two three-variable VAR, including $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$ and one of the cumulated action and communication shock series each, $\tilde{S}_t^A$ and $\tilde{S}_t^C$. The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels. Responses are to a 10 basis point positive shock to the interest rate.

Furthermore, since our shocks are identified outside the VAR system, it is not necessary to estimate all interactions between variables as part of a VAR. In fact, our analysis lends itself to the local projection approach of Jordà (2005). We run separate forecasting regressions of our macroeconomic variables using the shock as a predictor, while controlling for lags of macroeconomic variables. This approach avoids compounding any potential errors, as can happen during the iterative procedure used to compute VAR responses.

Following Ramey (2016), we project $\mathbf{Y}_{t+q}$ onto $(\boldsymbol{Y}_{t-1}, \boldsymbol{Y}_{t-2}, ..., \boldsymbol{Y}_{t-l})$:

$$\mathbf{Y}_{t+q} = D_c + \sum_{l=1}^{\text{lags}} D_l^q \mathbf{Y}_{t-l} + D_0^q \varepsilon_t^j + u_{t+q}^q, \quad q = 0, 1, 2, ..., Q \text{ and } j \in \{A, C\},$$

where D_c is a vector of constants, and the D_l^q are coefficient matrices for given lags l, and horizons q up to Q. We include two lags of the endogenous variables (which we load without first cumulating them). Further, we add the contemporaneous value ε_t^j of the shock itself, thus assuming recursiveness as in the VAR specification. The series $\{D_0^q\}_1^Q$ then gives the impulse of the shock on the variables IP and CPI, and standard errors are computed using the heteroskedasticity and autocorrelation robust Newey and West (1987) standard errors.

Figure 9: Robustness Check – Local Projection Approach

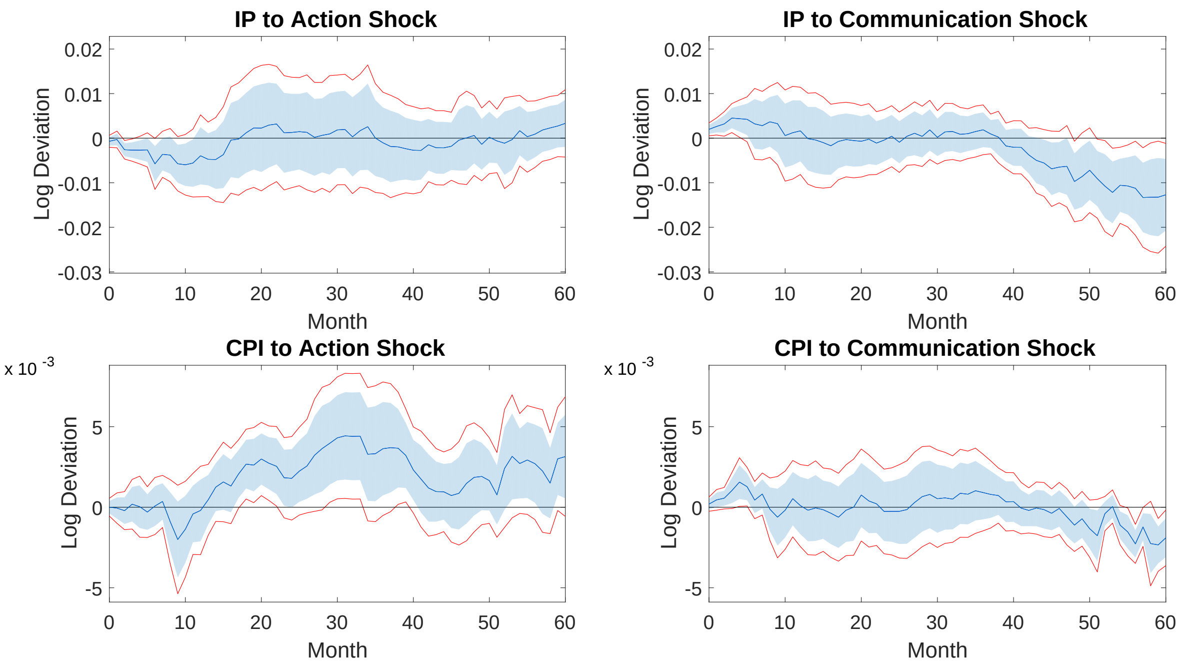

Notes: Impulse responses of $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ to the action and communication shocks under the local projection approach. 90% and 75% confidence intervals were obtained using Newey-West standard errors. Responses are to a 10 basis point positive shock to the interest rate.

Figure 9 shows the counter-intuitive initial increase of the action shock in our VAR disappears under local projections, albeit only with a very long lag. The action shock also displays the price puzzle under this specification.

To conclude, our main results are robust to different specifications, in particular the stronger effect of communication on IP and the smaller associated price puzzle.

5 Covering the Zero-Lower Bound Episode

As discussed previously, following the start of the ZLB period during the financial crisis of 2008, our decomposition of federal funds futures movements into action and communication shocks is no longer possible as these futures prices cease to move. However, certain longer maturity interest-rate futures contracts which remained liquid during this time, namely Eurodollar futures, can help us to analyse communication shocks during the ZLB period.

Eurodollar contracts are defined over quarters and not months, so it is no longer possible to extract expectations regarding particular meetings for these contracts.[35] Neither are we able to clearly identify an "action shock" in this case, since the contemporaneous Eurodollar future reflects expectations regarding both the most recent meeting and all future meetings within one quarter. Furthermore, the underlying for the contracts is the 3-month rate on dollar-denominated assets held abroad, as opposed to the overnight federal funds rate, which means the contracts are less tightly linked to the policy decisions of the FOMC. However, these contracts trade in highly liquid markets and the pricing of these contracts does move systematically on FOMC meeting days, implying that market participants were updating their expectations for future shorter-term interest rates in reaction to central bank communication.

We propose a linear decomposition of the Eurodollar contracts similar to that used in our previous analysis. To capture longer-term communication as used for forward guidance, we employ contracts betting on the Eurodollar rate one, two and three years out (ED4, ED8 and ED12).[36]

Here, we use a recursive scheme directly on the selected futures maturities as follows:

$$\begin{bmatrix} \Delta ED_{d,t}^{(4)} \\ \Delta ED_{d,t}^{(8)} \\ \Delta ED_{d,t}^{(12)} \end{bmatrix} = \begin{bmatrix} k_{11} & 0 & 0 \\ k_{21} & k_{22} & 0 \\ k_{31} & k_{32} & k_{33} \end{bmatrix} \cdot \begin{bmatrix} \varepsilon_{d,t}^{NED} \\ \varepsilon_{d,t}^{MED} \\ \varepsilon_{d,t}^{FED} \end{bmatrix} = \mathbf{K} \cdot \mathbf{E}_{d,t}, \tag{5}$$

where $\Delta ED_{d,t}^{(h)}$ is the daily difference of the Eurodollar contract futures rates at horizon h on the FOMC meeting day indexed by day d and month t. We call these shocks "near ED shock", "medium ED shock", and "far ED shock". It is important to note that the near ED shock is quite different to the action shock discussed previously. Given that the near ED shock represents the combined effects of all FOMC communication regarding interest rates during the next year, it contains (among others) the effects of action and communication

[35]Three month Eurodollar futures take as their underlying the 3-month interest rate on time-deposits of U.S. dollar denominated assets held outside of the U.S. The exact interest rate used comes from the dollar-denominated LIBOR rate. Unlike the federal funds futures contracts, these contracts are defined relative to the interest rate prevailing on the third Wednesday of the expiration month, and are available across quarterly horizons, for the next 10 years.

[36]Our results are robust to other ED futures maturities ([ED4, ED8, ED18] and [ED4, ED12, ED18]). FEVD analysis underlines the stronger effect on prices relative to industrial production of communication shocks captured by ED futures. For these results, see Appendix E.

shock above.

The shock series are displayed in Figure 10, for our sample period covering March 1994 to September 2016. A striking feature is the marked shift in volatility from the near ED shock to the longer-term ED shocks during the ZLB period: This suggests that before the Great Recession, markets were less likely to receive important surprise information about monetary policy more than one year into the future during FOMC meetings. However, with the onset of unconventional monetary policy, surprise information about the potential course of central bank decisions two or three years into the future became increasingly important. We can also see a period of larger volatility of the medium ED shock during the dot-com bust, which does not translate into far ED shock volatility.

Figure 10: Eurodollar Shock Series

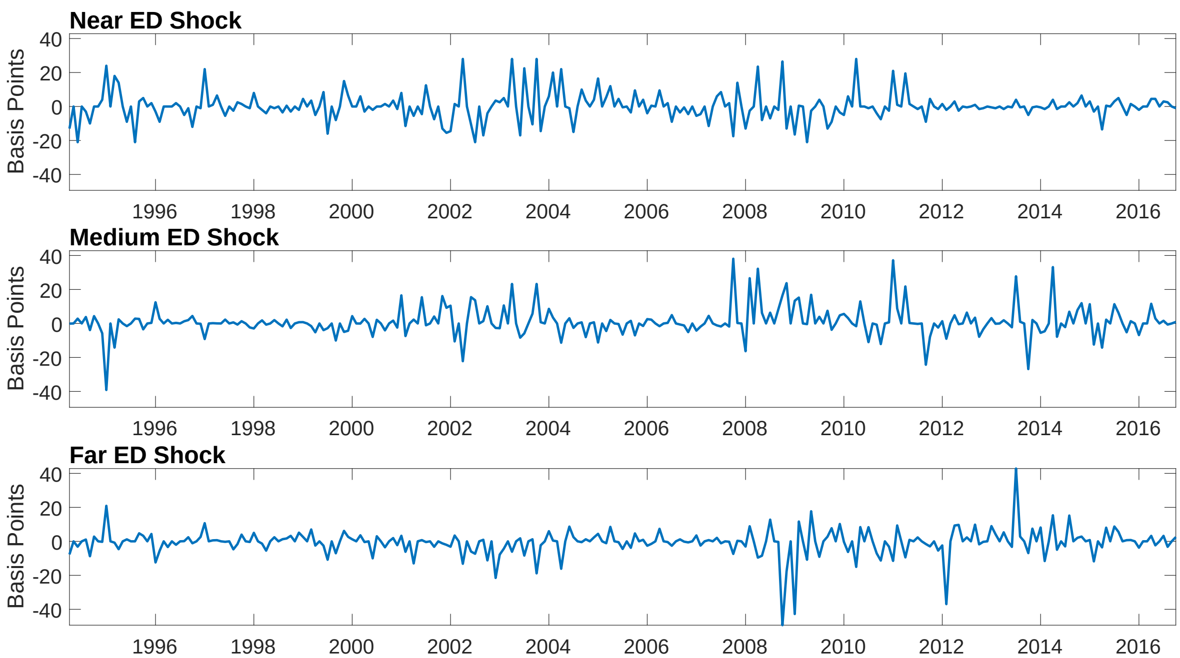

Notes: The figure displays the three Eurodollar shocks S_t^j, $j \in \{NED, MED, FED\}$ – dubbed "Near ED", "Medium ED", and "Far ED" shock, respectively.

We follow the baseline VAR specification above, entering our three cumulated shock series at the end of a vector including IP and CPI, with 12 lags. Again we eschew the need for a two-step estimation by entering the movements $\Delta ED_{d,t}^{(4)}$, $\Delta ED_{d,t}^{(8)}$, and $\Delta ED_{d,t}^{(12)}$ directly into a recursive VAR, meaning the decomposition of Equation (5) is performed within the SVAR. The impulse response functions are displayed in Figure 11. None of the Eurodollar shocks has significant effects on IP – the fact that responses are smaller and less significant than in our previous analysis may be due to Eurodollar futures being a noisier measure of monetary policy stance than federal funds futures. The near ED shock has no significant effect on prices, if anything, there is a borderline insignificant price puzzle.

Figure 11: Responses of $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ to Eurodollar Shocks

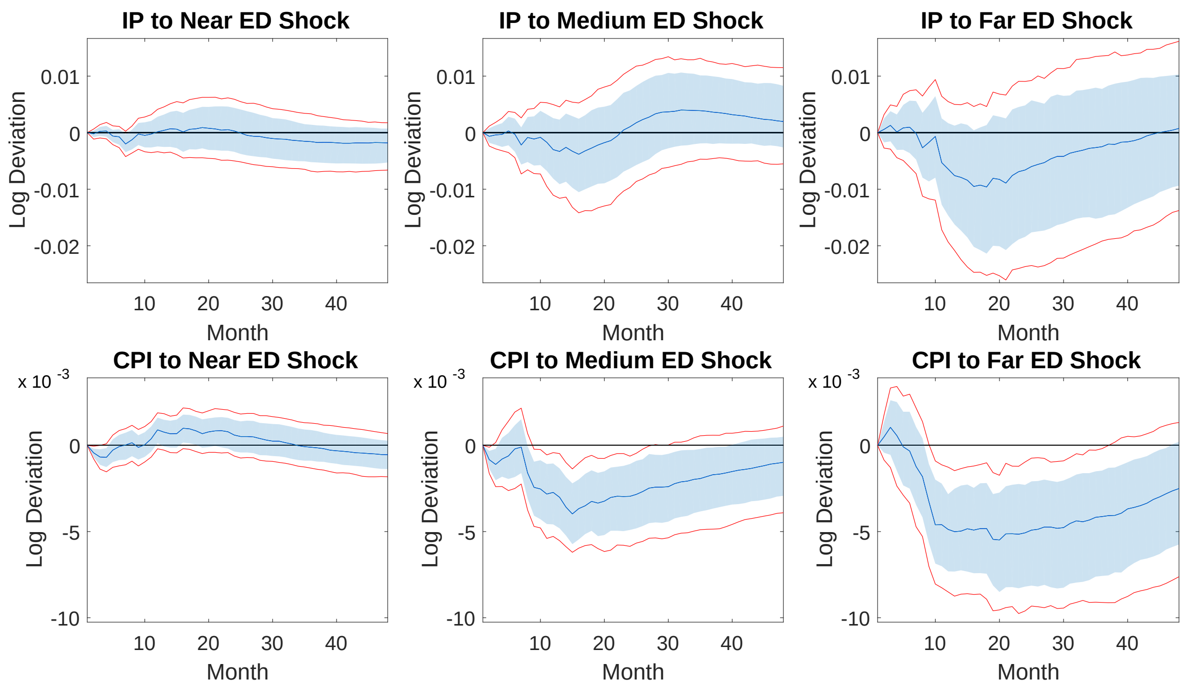

Notes: Impulse responses from our five-variable VAR, including $\log(\mathrm{IP}_t)$, $\log(\mathrm{CPI}_t)$ and the three cumulated shock series S_t^j, $j \in \{NED, MED, FED\}$ – near ED, medium ED and far ED shock, respectively. The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels. Responses are to a 10 basis point positive shock to the interest rate.

Further, the medium and far ED shocks show significant contractions to CPI, of increasing strength. This would match the predictions of the New Keynesian literature regarding the effects of forward guidance at increasing horizons on inflation (Del Negro et al., 2016; McKay et al., 2016). We reach the conclusion that forward guidance surprises at longer horizons seem to have a stronger and more persistent effect on CPI than on industrial production.

Moreover, historical decompositions of IP and CPI (see Figures **??** and 15 in Appendix E) show that from the announcements of asset purchases and forward guidance in September 2012 onwards, all three ED shocks have an expansionary effect on IP and CPI, which is evidence for an important role of the Fed in boosting inflation expectations, despite the fact that its policy rate remained at the ZLB during the period. Indeed, the timing of these later expansionary contributions almost exactly coheres with the timing of the FOMC's explicit forward guidance statements (2012-2015). These decompositions also help explain the relatively muted effect of monetary policy on IP here: In the direct aftermath of the financial crisis (2009-11), the communication shocks had a strong stabilising effect on inflation (with the stimulating medium and far ED shocks more than outweighing the contractionary near ED one), while all three are contractionary for IP. This could perhaps be explained in terms of central bank communication contributing to "anchored expectations", thereby accounting for the "missing disinflation" during this period (Bernanke, 2010). Moreover, the explicit long-term commitments communicated under forward guidance after 2012 have a much more marked effect on CPI than IP. Both episodes partially explain the larger impact on inflation relative to economic activity in response to our ED shocks. Overall, our analysis of Eurodollar futures supports our findings regarding the importance of central bank communication for the macroeconomy, especially at longer horizons.

6 Conclusion

In this paper, we have investigated the effect of communication surprises during FOMC meetings on the macroeconomy, and contrasted them with surprises about actual policy decisions. To distinguish surprise action from surprise communication, we use a simple Cholesky decomposition of changes within certain maturity segments of federal funds futures contracts.

For our sample from 1994M3 to 2008M6, we find that communication surprises play a more important role in macroeconomic fluctuations than action shocks. Communication shocks lead to the expected contractionary reaction of industrial production, explain a larger share of variance in macro variables, and can be easily aligned with the recent history of U.S. monetary policy. These findings are robust to various changes in specification. Both shocks have negative and significant effects on IP when purged of their information content. Overall, our findings emphasise the crucial importance of central bank communication, and of forward-looking information reception by market participants, even for the period before the explicit adoption of forward guidance as a policy tool by the Federal Reserve. In fact, our analysis suggests that researchers ought to think of "monetary policy shocks", of the type extensively studied in the literature, more in terms of central bank communication rather than unanticipated rate changes.

Our baseline analysis based on federal funds futures is only meaningful before conventional monetary policy hit the ZLB, and these futures markets became illiquid. Therefore we use longer-term communication shocks derived from Eurodollar futures to cover the period until late 2016. We find a shift in the volatility of the shock series to the longer horizons, suggesting a stronger focus on long-term communication by the FOMC. Moreover, there are large effects of central bank communication on inflation, with the size of the effect increasing in the horizon of the shocks, implying that forward guidance has indeed had a strong influence on price stability in the U.S.

Our analysis has shown the importance of central bank communication regarding future actions for the macroeconomy. However, it is likely that Fed policy has become *gradually* more forward-looking over the last twenty years. This would imply an increasing role for our communication shocks–and a decreasing role for our action shocks–within our sample. Indeed, this is partially reflected in the larger contributions of the communication shocks in the later part of our historical decompositions (while the action shock yielded large fluctuations in production only in the pre-2001 part of the sample). Moreover, the Federal Reserve switched to explicit long-term commitments under a policy of forward guidance after the financial crisis. Thus, a time-varying parameter specification, or one with different parameter regimes for the periods before and after explicit forward guidance, may represent a promising area of future research for the high-frequency VAR literature.

References

Andrade, P. and F. Ferroni (2016, October). Delphic and Odyssean monetary policy shocks: Evidence from the euro-area. School of Economics Discussion Papers 1216, School of Economics, University of Surrey.

Barakchian, S. M. and C. Crowe (2013). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics 60*(8), 950–966.

Ben Zeev, N., C. M. Gunn, and H. Khan (2017, March). Monetary News Shocks. Carleton Economic Papers 15-02, Carleton University, Department of Economics.

Bernanke, B. S. (2010, August). "The Economic Outlook and Monetary Policy", speech delivered at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, 27 August. Available under-http://www.federalreserve.gov/boarddocs/speeches/2005/20050505/.

Boivin, J. and M. P. Giannoni (2006, August). Has Monetary Policy Become More Effective? *The Review of Economics and Statistics 88*(3), 445–462.

Bundick, B. and A. L. Smith (2016, January). The dynamic effects of forward guidance shocks. Research Working Paper RWP 16-2, Federal Reserve Bank of Kansas City.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012). Macroeconomic Effects of Federal Reserve Forward Guidance. *Brookings Papers on Economic Activity 44*(1), 1–80.

Christiano, L. J., M. Eichenbaum, and C. Evans (1996, February). The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds. *The Review of Economics and Statistics 78*(1), 16–34.

Cochrane, J. H. and M. Piazzesi (2002, May). The Fed and Interest Rates - A High-Frequency Identification. *American Economic Review 92*(2), 90–95.

Coibion, O. and Y. Gorodnichenko (2012, October). Why Are Target Interest Rate Changes So Persistent? *American Economic Journal: Macroeconomics 4*(4), 126–62.

D'Amico, S. and M. Farka (2011). The Fed and the Stock Market: An Identification Based on Intraday Futures Data. *Journal of Business & Economic Statistics 29*(1), 126–137.

Del Negro, M., M. Giannoni, and C. Patterson (2016). The Forward Guidance Puzzle. 2016 Meeting Papers 143, Society for Economic Dynamics.

Faust, J., E. T. Swanson, and J. H. Wright (2004, September). Identifying VARS based on high frequency futures data. *Journal of Monetary Economics 51*(6), 1107–1131.

Forni, M. and L. Gambetti (2014). Sufficient information in structural VARs. *Journal of Monetary Economics 66*(C), 124–136.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

Gürkaynak, B. Sack, and E. T. Swanson (2005, May). Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking 1*(1), 55–93.

Gürkaynak, R. S. (2005). Using federal funds futures contracts for monetary policy analysis. Finance and Economics Discussion Series 2005-29, Board of Governors of the Federal Reserve System.

Hamilton, J. D. (2008). Assessing monetary policy effects using daily federal funds futures contracts. *Federal Reserve Bank of St. Louis Review 1*(Jul), 377–394.

Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics 99*(S1), 114–133.

Jarocinski, M. and P. Karadi (2018, March). Deconstructing Monetary Policy Surprises - The Role of Information Shocks. CEPR Discussion Papers 12765, C.E.P.R. Discussion Papers.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Kuttner, K. N. (2001, June). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics 47*(3), 523–544.

Lakdawala, A. (2016, November). Decomposing the Effects of Monetary Policy Using an External Instruments SVAR. MPRA Paper 78254, University Library of Munich, Germany.

McKay, A., E. Nakamura, and J. Steinsson (2016, October). The Power of Forward Guidance Revisited. *American Economic Review 106*(10), 3133–3158.

Miranda-Agrippino, S. and G. Ricco (2017, April). The transmission of monetary policy shocks. Bank of England working papers 657, Bank of England.

Moore, J. and R. Austin (2002). The Behavior of Federal Funds Futures Prices over the Monetary Policy Cycle. *Federal Reserve Bank of Atlanta Economic Review Second Quarter*, 45–61.

Nakamura, E. and J. Steinsson (2018, forthcoming). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*.

Newey, W. K. and K. D. West (1987, May). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*(3), 703–708.

Piazzesi, M. and E. T. Swanson (2008, May). Futures prices as risk-adjusted forecasts of monetary policy. *Journal of Monetary Economics 55*(4), 677–691.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2 of *Handbook of Macroeconomics*, Chapter 2, pp. 71–162. Elsevier.

Romer, C. D. and D. H. Romer (2004, September). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review 94*(4), 1055–1084.

Romer, D. H. and C. D. Romer (2000, June). Federal Reserve Information and the Behavior of Interest Rates. *American Economic Review 90*(3), 429–457.

Rudebusch, G. D. (1998, November). Do Measures of Monetary Policy in a VAR Make Sense? *International Economic Review 39*(4), 907–31.

Sims, C. A., J. H. Stock, and M. W. Watson (1990, January). Inference in Linear Time Series Models with Some Unit Roots. *Econometrica 58*(1), 113–144.

Söderström, U. (2001). Predicting monetary policy with federal funds futures prices. *Journal of Futures Markets 21*(4), 377–391.

Swanson, E. T. (2017, April). Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets. NBER Working Papers 23311, National Bureau of Economic Research, Inc.

Taylor, J. B. (2009, January). The Financial Crisis and the Policy Responses: An Empirical Analysis of What Went Wrong. NBER Working Papers 14631, National Bureau of Economic Research, Inc.

Thapar, A. (2008, May). Using private forecasts to estimate the effects of monetary policy. *Journal of Monetary Economics 55*(4), 806–824.

Appendix

A Correlation with Other Shock Series

In order to better understand how our shock series relate to other high-frequency monetary shock series available in the literature, this section examines correlations with different measures.

Table 2: Correlation of Our Shocks with Other Monetary-Policy Shocks

	Action Shock	Comm. Shock	# obs.
BC 1st Factor	0.7378***	0.5941***	115
BC 2nd Factor	0.5330***	-0.6118***	115
GSS Target Factor	0.8788***	0.0299	82
GSS Path Factor	-0.1318	0.6025***	82
RR Shock (original)	0.5743***	−0.0978	23
RR Shock (updated)	0.3351***	0.0517	111
NS Shock	0.6564***	0.6164***	68

Notes: BC: Barakchian and Crowe (2013); GSS: Gürkaynak et al. (2005); RR: Romer and Romer (2004); NS: Nakamura and Steinsson (2018). The updated RR series is from Wieland and Yang (2016).

We first examine the relation between our shock series and that of BC. From Table 2 we note that all our shocks are positively correlated with the shock of BC, i.e. the first principal component across maturities of contract. In fact, the action shock is most strongly correlated with the BC shock. This is somewhat surprising, since when we examine the impulse-response functions, we find that the communication shocks yield responses closest to those of the BC shock. With respect to the second principal component extracted by BC from the futures jumps, we can see that only the action shock is positively correlated with the second factor. The communication shock instead is negatively correlated, so it is not true that our communication shock series merely reflect information captured by the second factor of BC. This is heartening, since we know from BC that the second factor explains only a small fraction of the variance of the federal funds futures contracts. We can conclude that our structural decomposition offers different kinds of information relative to the two factors of BC (although they restrict their analysis to the first factor).

When we examine the relation between our shocks and those of GSS we find largely expected results. Our action shock is strongly and significantly correlated with the GSS target shock. Our communication shock shows positive correlation with the GSS path factor, although it is smaller, at 0.60. Therefore our shocks should be understood to be closely related, but not reducible, to the factors of GSS.

We moreover find that the action shock is significantly correlated with the Romer and Romer (2004) shock, but the communication shocks are not. When we examine a longer period, using the series computed by Wieland and Yang (2016), we find similar results. Finally, all our shocks are positively correlated to that of Nakamura and Steinsson (2018). The fact that the correlation structure looks much like those of our shocks with that of the BC shock, with the greatest correlation for the action shock, is unsurprising since the Nakamura and Steinsson (2018) shock is also the first principal component, although the bundle of futures jumps includes longer horizon Eurodollar contracts also.

B Testing our VAR for informational sufficiency

Preliminaries. We use our baseline VAR setting with $\boldsymbol{Y}_t = [\log(\text{IP}_t), \log(\text{CPI}_t), S_t^A, S_t^C]'$ as in (4), with 12 lags and 2 deterministic regressors. For the principal components, we use the monthly "FRED-MD" data set by McCracken and Ng (2016). The data set comprises 128 monthly macro-financial time series and has been proposed specifically for the purpose of factor analysis. We obtain 10 principal components (PCs), as $P = 10$ is the maximum amount of PCs used by Forni and Gambetti (2014, FG in the following).

A first step: F-test whether lagged PCs explain our shock series. In their simulation design, FG propose an F-test as an initial step to check whether principal components of a larger data set help explain their shock series. We follow this procedure, calculating the p-values of 12 F-tests, for both of our shocks S_t^A and S_t^C with their specifications:

$$\varepsilon_t^j = \sum_{i=1}^{P} \sum_{l=\{2,4\}} \phi_{il} f_{it-l} + v_t, \ \forall j = \{A, C\},$$

where we use the non-cumulated shock series ε_t^j, the lags $L = \{2, 4\}$ and number of PCs $P = \{1, 2, 4\}$, and where f_t are the estimated PCs.

The p-values for the test of the lagged PCs *not* explaining the shocks ($H_0 : \phi_{il} = 0 \forall \{i, l\}$) are reported in Table 3. For all of the tests, the null of no joint explanatory power of the lagged PCs cannot be rejected at the usual significance levels.

Table 3: p-values of an F-test whether lagged PCs help explain our shocks

	# lags	1 PC	2 PCs	4 PCs
ε_t^A	2	0.9443	0.9980	0.7197
	4	0.6434	0.9485	0.9258
ε_t^C	2	0.2593	0.5679	0.8837
	4	0.6095	0.8693	0.9502

The FG test for informational sufficiency. The FG test is a multivariate Granger-causality test checking whether PCs from a data set large enough to capture economic agents' expectations help predict variables in our VAR. If they do, then the econometrician's data set in the VAR (here, our four variables in $\boldsymbol{Y}_t$) is not informationally sufficient to capture the dynamics based on the economic agents' decisions. The major advantage of the FG fundamentalness test is that "rejection of PCs Granger-causing model variables" is not only a necessary condition for fundamentalness (as in other testing procedures), but also a sufficient one: If the state-space of the economy is captured by our PCs, then fundamentalness is *implied* by the failure to reject Granger-causality of the PCs (see FG for details and other tests available).

Technically, FG follow the multivariate test for Granger-causality in Gelper and Croux (2007, GC in the following), see FG p. 16.[37] The test proceeds as follows (compare GC pp. 3ff.): Set up an unrestricted or "full" model,

$$\boldsymbol{Y}_t = \phi_c + \sum_{l=1}^{L} \phi_l \boldsymbol{Y}_{t-l} + \sum_{l=1}^{L} \psi_l \mathrm{PC}_{t-l}(P) + \varepsilon_{f,t},$$

where $\mathrm{PC}_{t-l}(P)$ are the lth lags of P PCs (for the actual specification, see below). Also, set up a restricted model assuming the coefficients on the PCs are zero,

$$\boldsymbol{Y}_t = \phi_c + \sum_{l=1}^{L} \phi_l \boldsymbol{Y}_{t-l} + \varepsilon_{r,t}.$$

Now the idea is to compare forecast errors of both models: If the PCs matter in forecasting $\boldsymbol{Y}_t$, the forecasts should be significantly different for both models. So split the sample into $T - H$ pseudo in-sample periods and H pseudo out-of-sample periods. Then obtain

[37] In particular, they use the "Regression" implementation, which GC show to have the largest power. We also use this statistic.

the (recursive one-step-ahead) forecast error of both models for the H periods:

$$u_{f,t} = \boldsymbol{Y}_t - \hat{\boldsymbol{Y}}_{f,t}$$
$$u_{r,t} = \boldsymbol{Y}_t - \hat{\boldsymbol{Y}}_{r,t}$$

GC follow the univariate implementation of the test by Harvey et al. (1998) in defining no Granger causality as zero correlation between $u_{r,t}$ and $u_{r,t} - u_{f,t}$, i.e.

$$u_{r,t} = \lambda(u_{r,t} - u_{f,t}) + e_t \tag{6}$$

has $\lambda = 0$ (see GC p. 4). A likelihood ratio test with the statistic

$$\text{Reg} = H\left(\log(|u_r' u_r|) - \log(|\hat{e}'\hat{e}|)\right),$$

where $\hat{e}$ is the estimated residual from (6), will reject the null of no Granger causality ($H_0 : \lambda = 0$) is very large.[38]

As "the limiting distribution of the multivariate out-of-sample tests under the null hypothesis of no Granger causality is unknown" (GC, p. 16), we have to resort to a bootstrap procedure like GC: "The percentage of bootstrap statistics exceeding the test statistic computed from the observed time series is an approximation of the p-value." (GC, ibid.)

We choose the following setup for our testing procedure: Like FG, we let the Akaike information criterion decide on the lag length of the VAR.[39] We include 2 deterministic regressors (constant and linear trend) as in our main VAR specification. Like FG, we report the p-values of the test for the number of factors included into the system $P = \{4, 6, 8, 10\}$, forecast horizons of $H = \{24, 48, 68\}$ and 500 bootstrap replications for the test statistics.[40] The corresponding p-values are summarised in Table 4 – the null of no Granger causality is never rejected at the 5%, and just once at the 10% significance level. We thus conclude that our four-variable VAR system is informationally sufficient.

[38]The p-value of the test would be given by the cdf of the statistic Reg at the χ^2 distribution with P times the number of lags ($P \cdot L$) degrees of freedom (the dofs are given by the difference in dofs of the full and restricted VAR). However, see the notes on the bootstrap below.

[39]The AIC consistently chooses one lag. If instead we want to use 12 lags for our VAR, we run into degree-of-freedom problems when including more than $P = 4$ factors or choosing a forecast horizon of $H > 24$. However, in a setting of $P = 4$ and $H = 24$ the test again fails to reject fundamentalness (the p-value is 0.5980).

[40]FG use $H = 80$ for their (simulated) VAR of 200 periods (footnote 6 on p. 131), and we scale down the maximal H in line with the smaller number of periods (172 for our analysis).

Table 4: p-values of the Forni-Gambetti (2014) test for informational sufficiency

	$P = 4$	$P = 6$	$P = 8$	$P = 10$
$H = 24$	0.3280	0.1180	0.4220	0.6740
$H = 48$	0.3460	0.9940	0.8200	**0.060**
$H = 68$	0.3760	0.8180	0.5800	0.5980

Notes: H_0 : no Granger causality (i.e. VAR is informationally sufficient); lags chosen by the Akaike Info Criterion.

C Responses to Recursively Identified Shocks

When we use a simple recursive identification scheme with the ordering $Y_t = [\text{IP}_t, \text{CPI}_t, \text{FFR}_t]'$, where FFR_t is the federal funds rate, we obtain the counter-intuitive responses reported in BC, see Figure 12

Figure 12: IRFs from a Recursively Identified Shock

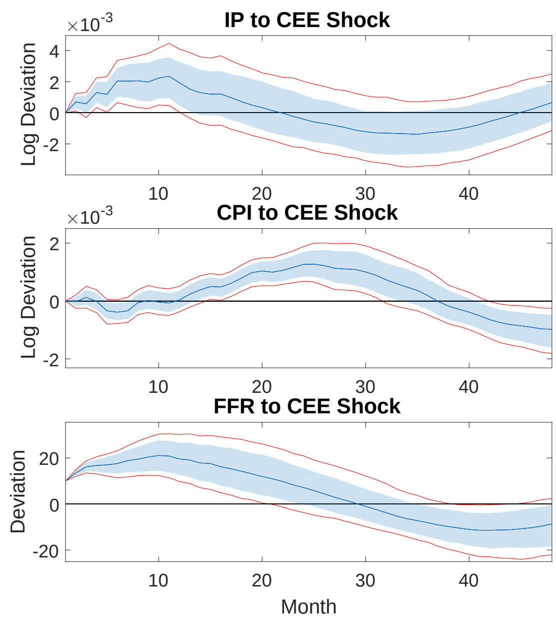

Notes: Responses of $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$, and the Federal Funds Rate to a 10 basis point contractionary shock, identified via the lower triangular restriction of Christiano et al. (1996). The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels.

D Using commodity prices to control for inflation expectations

We examine whether the inclusion of commodity prices is a viable alternative to orthogonalising shocks with respect to Greenbook data. Christiano et al. (1996) suggested to include commodity prices in VARs to correct for forward-looking monetary policy, since commodity prices are strong predictors of future inflation. We can see that the inclusion of commodity prices (ordered first in the VAR) has a notable effect on the impulse responses: the reaction of IP to the action shock is no longer significantly positive and the price puzzle is only borderline significant. However, only our communication shock shows a significant contraction on IP and no significant price puzzle. We thus find our main findings robust to using commodity prices instead of Greenbook forecasts to control for forward-looking monetary policy.[41]

[41]Thapar (2008) makes the argument that if Fed-internal Greenbook forecasts should strictly dominate commodity prices as a means to control for central-bank expectations and thus to resolve a price puzzle. We find that both methods resolve the price puzzle for our communication shock, but fail to do so for the action shock.

Figure 13: Impulse Responses of $\log(\text{IP}_t)$ and $\log(\text{CPI}_t)$ in a VAR with Commodity Prices

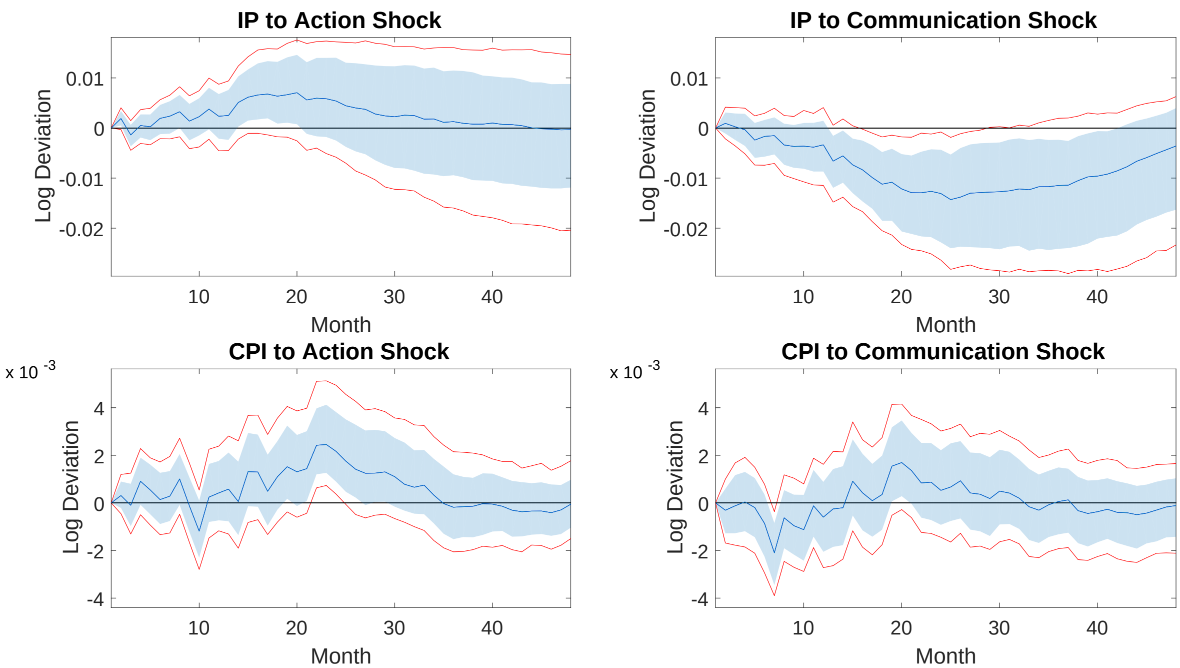

Notes: Impulse responses from our five-variable VAR, including commodity prices $\log(\text{PCOMM}_t)$, $\log(\text{IP}_t)$, $\log(\text{CPI}_t)$ and the cumulated action and communication shock series not orthogonalised to Greenbook forecasts, $\tilde{S}_t^A$ and $\tilde{S}_t^C$. The median response and confidence intervals were obtained from bootstrapping the VAR 500 times, the graph depicts the latter at 90% (red) and 75% (blue shadow) significance levels. Responses are shown to a 10 basis point positive shock to the interest rate.

E Further Results from Analysis Using Eurodollars

E.1 Robustness Checks for the Eurodollar Specification

We assess the robustness of our results to different selections of Eurodollar contracts, namely $[ED4,\ ED8,\ ED18]$ and $[ED4,\ ED12,\ ED18]$. We display the IRFs in Figure 14. Results are qualitatively unaffected by our choices.

Figure 14: Responses of $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ to Eurodollar Shocks

Contracts 4, 8, 18:

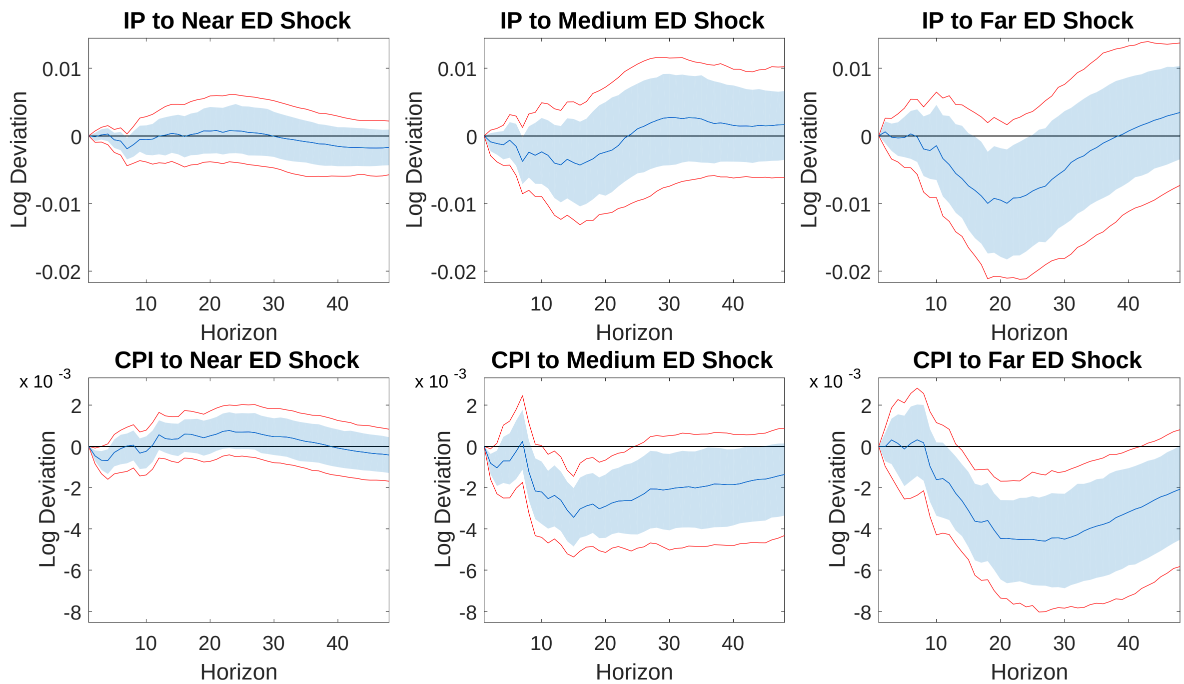

Contracts 4, 12, 18:

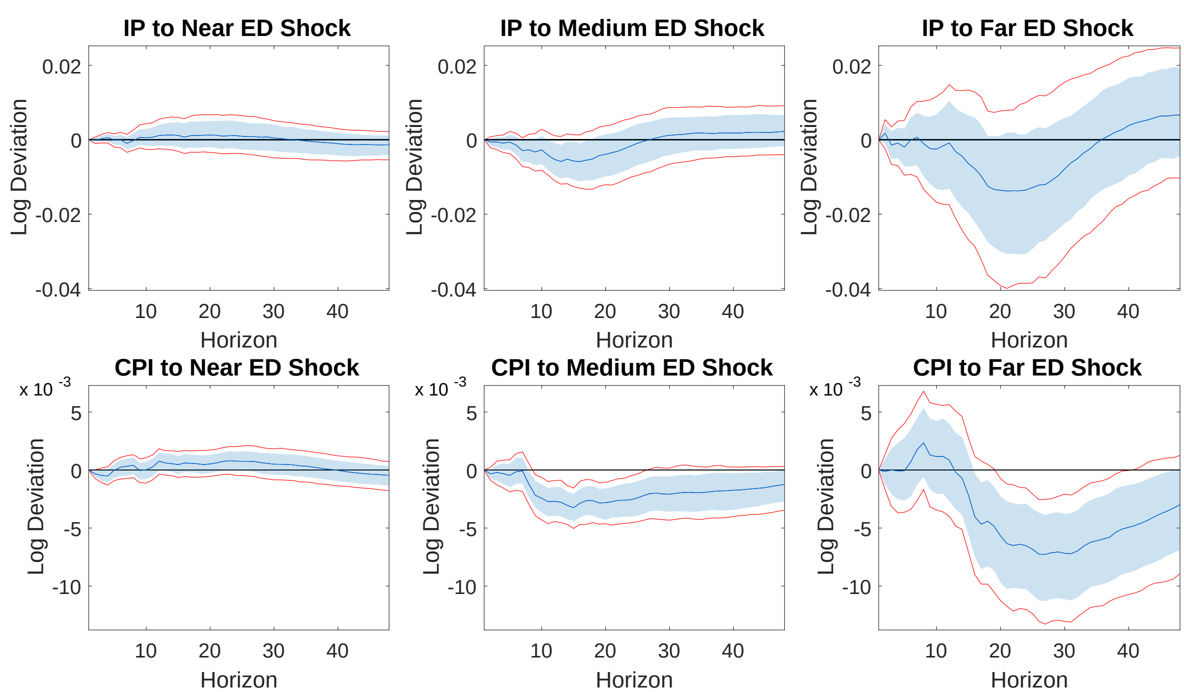

Notes: See Figure 11 in the main text.

E.2 Forecast-Error Variance Decomposition

We also examine forecast-error variance decompositions of the contribution of our Eurodollar-derived shocks to movements in macro variables, which are displayed for the 12, 18, 24 and 36 month horizons in Table 5. We chart economically significant differences between the contributions of shocks according to their horizon, with the further forward Eurodollar shocks typically having a larger contribution. In general, movements in the medium-term ED communication shock have a particularly strong forecasting power relative to the other two communication shocks.

Table 5: Forecast-Error Variance Decomposition at Business Cycle Frequency

	Horizon (months)	IP	CPI	S_t^{NED}	S_t^{MED}	S_t^{FED}
IP_t	12	0.8611	0.1218	0.0058	0.0052	0.0062
	18	0.6935	0.2645	0.0036	0.0117	0.0267
	24	0.6078	0.3362	0.0036	0.0116	0.0406
	36	0.5743	0.3571	0.0071	0.0123	0.0493
CPI_t	12	0.0405	0.8167	0.0172	0.0620	0.0637
	18	0.0388	0.6217	0.0402	0.1528	0.1465
	24	0.0315	0.4752	0.0587	0.2128	0.2219
	36	0.0283	0.3260	0.0582	0.2707	0.3169

Notes: Contribution of our shocks to a forecast-error variance decomposition of IP and CPI at the 12, 18, 24 and 36 month horizons from our baseline five-variable Eurodollar VAR. The identified three shocks are $S_{d,t}^{j}$, $j \in \{NED, MED, FED\}$. "IP" and "CPI" shocks are not identified.

E.3 Historical Decompositions for the Eurodollar Specification

Here we discuss in detail the results of historical decompositions of our ED futures analysis for both industrial production and prices. What is perhaps most interesting for us is the decomposition of industrial production during and after the Great Recession, shown in Figure **??**. Surprisingly, the model suggests that from the onset of the crisis in early 2008 to around 2010, communication at all three horizons (1, 2 and 3 years) had a recessionary impact. This would likely reflect information about the Fed's negative outlook for the economy superseding its communication that these conditions were "likely to warrant exceptionally low levels of the federal funds rate for some time".[42] However, almost exactly from the onset of the asset purchases made in September 2012 onwards, all three ED shocks have an expansionary effect on IP, speaking for an inflation-expectations boosting effect even in the absence of movements in the federal funds rate, which remained close to zero until late 2015. However, at the same time as the Fed's mention of a possible exit of the ZLB in late 2015, the far ED shock shows a contractionary effect on IP again. We therefore conclude that the ED shocks seem well-suited to track the recent history of Fed announcements on market expectations and thus the real economy, just as our federal funds futures shocks do for the pre-2008 period. The historical decomposition of CPI, shown in Figure 15, shows a strong positive contribution of the further forward ED shocks to inflation during the ZLB period. This suggests that despite generally weaker effects on output in the aftermath of the financial crisis, FOMC communication was supportive of inflation and inflation expectations.

[42] See https://www.federalreserve.gov/faqs/what-is-forward-guidance-how-is-it-used-in-the-federal-reserve-monetary-policy.htm.

Figure 15: Historical Decomposition of $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ with Eurodollar Shocks

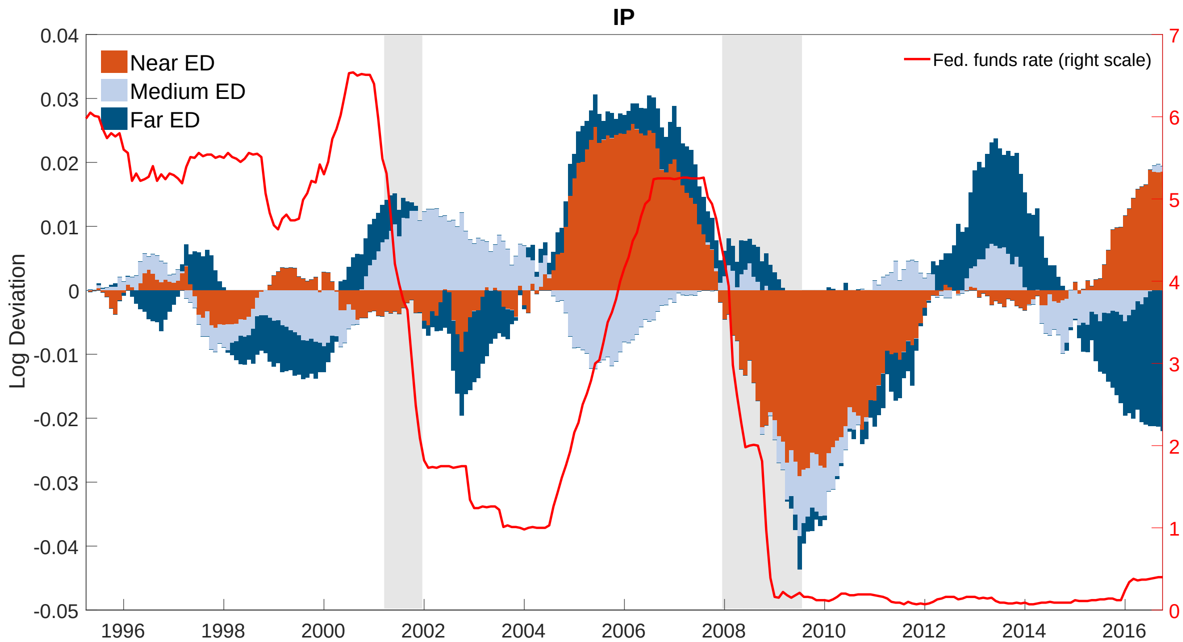

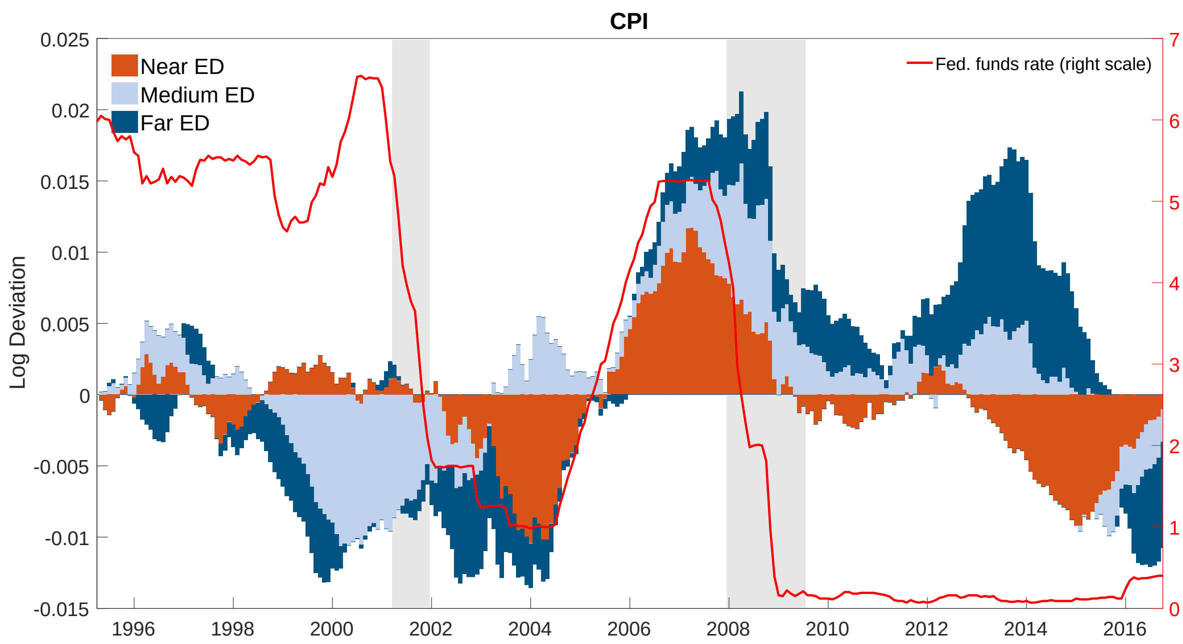

Notes: Historical decomposition of $\log(\mathrm{CPI}_t)$ in our five-variable VAR, including the variables $\log(\mathrm{IP}_t)$ and $\log(\mathrm{CPI}_t)$ and three cumulated shock series S_t^j, $j \in \{NED, MED, FED\}$ – near ED, medium ED and far ED shock respectively. The bar plots are stacked, so their height above the zero-axis represents the cumulative historical contribution of our monetary shocks to industrial production above its unconditional mean. Similarly for their height below the zero-axis. We also display the federal funds rate (using the right-hand scale) for reference. NBER recession periods are shown as grey areas.

F Details on FOMC meeting days

Here we give some details on which FOMC announcements we do not consider to be scheduled (and therefore do not use in our analysis). We also compare these to the Appendix 2 of the working paper version of GSS, Gürkaynak et al. (2004, GSSWP in the following), which contains a detailed summary (up to May 2004).

4/18/1994. Unscheduled conference call; from the minutes from March 22, 1994[43]: "It was agreed that the next meeting of the Committee would be held on Tuesday, May 17. 1994." GSSWP lists this date as an "intermeeting move".

10/15/1998. Unscheduled conference call. From the meeting statement of the previous meeting on Sept. 29th[44], it is not fully clear whether the meeting was scheduled: "In a telephone conference held on October 15, 1998, the Committee members discussed recent economic and financial developments and their implications for monetary policy. (...) At the conclusion of this discussion, the Chairman indicated that he would instruct the Federal Reserve Bank of New York to lower the intended federal funds rate by 25 basis points, consistent with the Committee's directive issued at the meeting on September 29, 1998. It was agreed that the next meeting of the Committee would be held on Tuesday, November 17, 1998." However, we choose not to consider this date as GSSWP have it as an "intermeeting move".

1/3/2001. Unscheduled conference call. From the December 19th (2000) FOMC minutes[45]: "This meeting adjourned at 1:35 p.m. with the understanding that the next regularly scheduled meeting of the Committee would be held on Tuesday-Wednesday, January 30-31, 2001."

4/18/2001. Unscheduled conference call. From the March 20th FOMC minutes[46]: "It was agreed that the next meeting of the Committee would be held on Tuesday, May 15, 2001." GSSWP: "intermeeting move".

9/17/2001. Unscheduled conference call. From the August 21st FOMC minutes[47]: "It was agreed that the next meeting of the Committee would be held on Tuesday, October 2, 2001."

[43] https://www.federalreserve.gov/fomc/MINUTES/1994/19940322min.htm
[44] https://www.federalreserve.gov/monetarypolicy/fomchistorical1998.htm
[45] https://www.federalreserve.gov/fomc/minutes/20001219.htm
[46] https://www.federalreserve.gov/fomc/minutes/20010320.htm
[47] https://www.federalreserve.gov/fomc/minutes/20010821.htm

GSSWP: “intermeeting move”.

8/10/2007 and 8/17/2007. Both dates were unscheduled conference calls. From the August 7th FOMC minutes[48]: “It was agreed that the next meeting of the Committee would be held on Tuesday, September 18, 2007.”

1/9/2008 and 1/22/2008. Unscheduled conference call on 9th and 22nd, but meeting on 30th was scheduled. From the Dec. 11th, 2007 FOMC minutes[49]: “It was agreed that the next meeting of the Committee would be held on Tuesday-Wednesday, January 29-30, 2008.”

3/11/2008. Meeting on 18th, unscheduled conference call on the 11th. From the Jan. 30th FOMC minutes[50]: “It was agreed that the next meeting of the Committee would be held on Tuesday, March 18, 2008.”

10/08/2008. Meeting on 29th, unscheduled conference call on 7th. From the Sept. 16th FOMC minutes[51]: “It was agreed that the next meeting of the Committee would be held on Tuesday-Wednesday, October 28-29, 2008.”

References

Afonso, G., A. Kovner, and A. Schoar (2011). Stressed, not frozen: The federal funds market in the financial crisis. *Journal of Finance* 66(4), 1109–1139.

Christiano, L. J., M. Eichenbaum, and Ch. Evans (1996). The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds. *The Review of Economics and Statistics* 78(1), 16-34.

Forni, M. and L. Gambetti (2014). Sufficient information in structural VARs. *Journal of Monetary Economics* 66(C), 124–136.

Gelper, S. and C. Croux (2007, April). Multivariate out-of-sample tests for Granger causality. *Computational Statistics & Data Analysis* 51(7), 3319–3329.

[48] https://www.federalreserve.gov/fomc/minutes/20070807.htm
[49] https://www.federalreserve.gov/monetarypolicy/fomcminutes20071211.htm
[50] https://www.federalreserve.gov/monetarypolicy/fomcminutes20080130.htm
[51] https://www.federalreserve.gov/monetarypolicy/fomcminutes20080916.htm

Gürkaynak, R., B. P. Sack, and E. T. Swanson (2004). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. Finance and Economics Discussion Series 2004-66, Board of Governors of the Federal Reserve System.

Harvey, D. S., S. J. Leybourne, and P. Newbold (1998). Tests for forecast encompassing. *Journal of Business & Economic Statistics* 16(2), 254–259.

McCracken, M. W. and S. Ng (2016). FRED-MD: A Monthly Database for Macroeconomic Research. *Journal of Business & Economic Statistics* 34(4), 574–589.

Moore, J. and R. Austin (2002). The Behavior of Federal Funds Futures Prices over the Monetary Policy Cycle. *Federal Reserve Bank of Atlanta Economic Review Second Quarter*, 45–61.

Nakamura, E. and J. Steinsson (2018, forthcoming). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*.

Robertson, J. C. and D. L. Thornton (1997). Using federal funds futures rates to predict Federal Reserve actions. *Federal Reserve Bank of St. Louis Economic Review* 1 (Nov), 45–53.

Romer, C. D. and D. H. Romer (2004, September). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review* 94(4), 1055–1084.

Thapar, A. (2008, May). Using private forecasts to estimate the effects of monetary policy. *Journal of Monetary Economics* 55(4), 806–824.

Wieland, J. F. and M.-J. Yang (2016, March). Financial Dampening. NBER Working Papers 22141, National Bureau of Economic Research, Inc.

Chapter 2:

Factor Misallocation and Adjustment Costs: Evidence from Italy

1 Introduction

The literature on misallocation, following the study of Hsieh and Klenow (2009), has used measures of within-industry dispersion in the marginal revenue product to factors as evidence for inefficient allocations of resources across firms. The basic intuition is that, in a static model, and given that firms within an industry are otherwise homogenous, differences in marginal revenue products must imply the existence of firm-level distortions that prevent the market from equilibriating returns to inputs. The model of Hsieh and Klenow (2009) provides researchers with a means to quantify the gains from removing such distortions.

This "static" view of misallocation was recently challenged by the study of Asker et al. (2014), who argue for a "dynamic" understanding of why marginal revenue product dispersion might exist in firm-level data. The authors argue that marginal revenue products may in fact merely be the result of firms adjusting inputs in a "lumpy" manner, in response to idiosyncratic shocks and in the presence of adjustment costs. When firms receive productivity shocks they invest in more capital, but they do so in stages, meaning they do not equilibriate marginal product and marginal cost. The authors show that a dynamic partial-equilibrium model of investment is able to qualitatively deliver the level cross-sectional dispersion in the marginal revenue product of capital observed in multiple firm-level datasets for developed and developing countries.

The broad aim of this paper is to subject the dynamic view of misallocation to closer scrutiny, with a focus on labour demand, using firm-level data from Italy that provide an excellent environment to examine the nature of misallocation. Indeed, Italy's productivity performance in recent years has been poor in comparison to other advanced economies. We can see from the left panel of Figure 5 that Italian labour productivity has been largely stagnant, increasing by only 3.5% over the 1998-2013 period.[1] More worryingly, we can see from the right panel of Figure 5 that Italian TFP (measured according to a standard Solow accounting exercise) has in fact been falling a cumulative 7.5% over the last decade. These trends are not observed in many other developed economies (with the exception of Spain where a similar TFP trend is in evidence). While the analysis in this paper is principally static, I do also quantify trends in measures of misallocation in my dataset that can at least partially account for long-term dynamics.

If it is true that within-industry misallocation is best understood in terms of firm's responses to idiosyncratic shocks, then there are certain qualitative and quantitative features of the data the model must meet. Whilst the model of Asker et al. (2014) is indeed able to explain variance in the marginal revenue product of capital, there is more to misallocation than a single moment. In particular, it is important that the model match dispersion measures in revenue products to other factors, and their relation to firm size, which I shall test in this study. A further goal of the paper is to qualitatively examine the role of firing costs as an explanation for the poor TFP performance of the Italian economy in recent years. The paper also exploits the existence of a specific size-threshold in the application of Italian labour law (Article 18), to aid estimation of the structural model, and examines the potential effect on TFP of policy-reform using model simulations.

To this end, I re-examine the dynamic view of misallocation across several dimensions, using Italian firm-level data:

1. I examine the ability of an estimated model of firm dynamics to explain the misallocation of labour. While the inclusion of a frictional labour decision does not change the nature of the argument of Asker et al. (2014), it would be of interest to estimate the impact of frictional adjustment in the labour market on allocative efficiency.

2. One of the most useful features of the Hsieh and Klenow (2009) framework is that the measures of idiosyncratic distortions can be backed out from the data and potentially related to real world factors. One of the more obvious is whether the distortions are correlated with the size of firms. I investigate the capacity for a model of Asker et al. (2014) type to deliver similar correlations in the case of labour.

3. I use the existence of specific size-dependent policies to estimate a structural model of firm dynamics, in order to assess the ability of such a model to explain misallocation in the Italian case. The use of the threshold should allow for more precise estimation

[1]Here I refer to estimates from Lanau and Topalova (2015).

of the parameters of such a model than in previous papers, which do not employ structural features of the regulatory environment, as in this paper. It also allows for the use of a dynamic model of firm-behaviour, of the Cooper and Haltiwanger (2006) type, to track closely and analyse a real world policy reform, which is a first for the literature.

The principal findings of this paper are that the dynamic model of labour demand is well able to deliver the levels of misallocation of labour seen in the data, suggesting (at first glance) that the Italian experience of labour misallocation is best understood in terms of hiring and firing costs. However, the model does a poor job at replicating the relation between misallocation and firm-size, which suggests a role for size-dependent distortions as a further means to explain sources of misallocation of labour in Italy. Once precise estimates of the effects of a regulatory threshold on firing costs of labour in the data are attained, the effects on labour productivity of policy-reforms of a type close to that enacted by the administration of Matteo Renzi in 2014 are found to be potentially quite large, albeit in a study undertaken in partial equilibrium.

In this paper I proceed as follows. In Section 2 I discuss related literature. In Section 3 I discuss the Italian firm-level data used in the study. In Section 4.1 I outline the misallocation model of Hsieh and Klenow (2009). In Section 4 I analyse misallocation in my sample. Section 5 lays out a structural model and estimation strategy, based on an extension of the model of Asker et al. (2014). Section 6 adapts the model to a consideration of the effects of Article 18 legislation in Italy. Section 7 documents time trends in the measures of misallocation in Italian data. Section 8 concludes.

2 Literature Review

This paper adds to the literature on quantifying the misallocation of resources, that began with Hsieh and Klenow (2009). Using their framework, the authors found that reducing misallocation to U.S. levels would lead to increases in manufacturing TFP of 30-50% for China, and 40-60% for India. Ziebarth (2013) investigates the U.S. case using 19th century data and finds misallocation levels of India and China to be comparable to that of the U.S. at an earlier stage of development. Oberfield (2013) investigates the 1982 Chilean financial crisis with an alternative, but related, framework and finds between-industry misallocation accounts for most of the observed TFP falls.

My analysis also closely relates to several recent applications of these methods to misallocation in Southern Europe. Dias et al. (2015) find that within-industry misallocation almost doubled between 1996 and 2011 using Portuguese data, employing a version of the Hsieh and Klenow (2009) framework with a three-factor production function. García-Santana et al. (2016) also find increases in misallocation in Spanish firm-level data for the period 1995-2007, concluding that in the absence of such deterioration, average TFP growth would have been around 0.8% a year, as opposed to the *falls* in Spanish TFP seen in aggregate data. Gopinath et al. (2017) is closely related to this study in the sense that they study several Northern and Southern economies, including Italy, also employing ORBIS data. Their analysis focuses on time trends and points to large increases in the dispersion of the marginal revenue product of capital in the Italian manufacturing sector between 1999 and 2013, suppressing TFP. They focus on Spanish sub-sample and argue that a fall in the real interest rate for Spain lead to capital flows toward relatively unproductive firms, given size-dependent borrowing constraints. Their model also includes adjustment costs to capital and time-to-build, although it is calibrated, not estimated.

A very closely related study is a preliminary paper by Bayer et al. (2015), who decompose the marginal revenue products into permanent and transitory components. They find that the persistent components of labor and capital productivies are negatively correlated, while their transitory components are positively correlated. They explain their findings they develop a dynamic partial-equilibrium model in which firms are able to make frictional decisions regarding their own production technologies, while in the short-run they operate Leontief production functions.[2]

With respect to the state of research into the Italian case, there is much debate in the literature as to the nature of the productivity issues it has been facing in recent decades. Pellegrino and Zingales (2014) summarize several leading hypotheses. They note that Italian productivity trends are somewhat hard to square with several benign features of the economic conditions in Italy in the early 2000's period: the low and stable interest rates and inflation, the not particularly restrictive fiscal policy (with an average deficit of 3.7%), and political stability in the sense of the longest surviving government of the post-WWII period. As candidate explanations for low and falling productivity they cite Italy's specialization in relatively low-tech sectors, with increased exposure to Chinese import competition. They also refer to Italy's reliance on small firms, which could suppress productivity given evidence of a positive correlation between firm size and productivity (van Ark and Monnikhof, 1996). Pellegrino and Zingales (2014) also posit issues of high regulatory protection of labor,

[2]See also Hawkins et al. (2015) and Catherine et al. (2017), for recent examples of papers employing estimated structural models of adjustment costs.

high corruption, low rule of law, low human capital, and low adult literacy in comparison to OECD averages. With respect to the slowdown in productivity growth over time, the cross-country regression analysis, with complementary firm-level data work, of Pellegrino and Zingales (2014) points to an interaction effect between small firms and the effects of Chinese import competition, and a failure to take full advantage of the ICT revolution.

Recent work has applied structural models to the case of regulatory thresholds. As in this study, the paper of Garicano et al. (2016) also analyses a size-dependent threshold in labour regulations, using an heterogeneous-firm model, with an aim to studying its distortionary impact on TFP. Garicano et al. (2016) study the case of France, where they describe how firms who expand to 50 workers or more must adhere to a series of policies, including creating a works council and a health and safety committee, and appointing a union representative. French firms with more than 50 employees also incur higher liabilities for accidents, and they must also undertake a formal professional assessment for each worker older than 45. Firms affected by the regulation also face higher firing costs for dismissals of 10 workers or more. The authors find the welfare costs of the regulation to be 3.4% of GDP. For the purposes of the study of size-dependent regulation, the Italian case has a specific advantage relative to the French one: only the firing costs of labour differ as firms increase above 15 employees. This allows easier isolation of the effects of employment protection legislation specifically, as opposed to the bundle of policies examined by Garicano et al. (2016). Further, the analysis of Garicano et al. (2016) is principally static, based on a span-of-control framework, as opposed to the dynamic setting studied here. While the authors do study a dynamic extension to their model, their approach differs to that used in this paper insofar as we model the effects of a threshold in an adjustment cost parameter, as opposed to the imposition of a static tax on labour with a fixed cost used in their study. Garicano et al. (2016) do include quadratic adjustment costs to labour in their dynamic model, but this paper studies a model in which there are non-continuous adjustment costs to factor adjustment costs. Such differences in theoretical approach are necessary given the unique aspects of the Italian legislation. This paper also estimates adjustment cost parameters using simulated method of moments, while the paper of Garicano et al. (2016) explicitly leaves this for future work, and calibrates parameters using evidence from the literature.

Recently, the literature has turned its attention to evaluating misallocation for the Italian case specifically. Calligaris (2015) uses a sample of manufacturing firms from CERVED, which provides information on the universe of limited liability firms, and documents an increasing trend in misallocation during the sample period of 1993-2011, as was also reported in Gopinath et al. (2017). Calligaris (2015) also related measures of misallocation in Italy negatively to firm-characteristics such as age and size. The paper of Calligaris et al. (2016) extended the analysis to the non-manufacturing sector and documented similar trends, also using CERVED data. Calligaris et al. (2016) used "*Inquiry into manufacturing and service firms*" of Banca d'Italia (INVIND) as well as "*Centrale dei Bilanci*" balance sheet data to relate misallocation in Italy to additional firm characteristics, via regression analysis. The principal contribution of our paper is to estimate a dynamic model of misallocation on Italian data for the first time, utilizing the special opportunity represented by a regulatory threshold. Neither Calligaris (2015) nor Calligaris et al. (2016) quantify the effects of hiring and firing costs on misallocation of Italy directly.[3] Given the importance of potential trends in misallocation, it is of interest to corroborate results from different data sources and sectors, although for manufacturing the ORBIS and CERVED samples are likely similar.[4] Section 7 therefore presents the time dynamics of misallocation in our sample. From the perspective of robustness the estimates attained for this paper are also similar in magnitude to those of Calligaris (2015), who reports potential reallocation gains of 58% in 1993, 67% in 2006, and 80% in 2011 as baseline estimates – this paper attains an estimate of 48.02% in 1995, 66.12% in 2007, and 81.59% in 2013.

3 Data

The firm-level data are taken from the Italian subset of ORBIS by Bureau van Dijk, which is a large commercial dataset containing information on over 75 million companies worldwide.[5] Importantly, the dataset includes both private and public companies, which means the analysis is not restricted to the case of relatively large, listed firms.[6] The data are a version of publically available company reports, that are then harmonized across countries by Bureau van Dijk. The data contain information on the industries in which firms are active, financial information from the balance sheet and profit and loss account, as well as many other variables not employed in this study (mergers and acquisitions data, patents, owner-

[3]Further, this paper documents the effects of the specific distortion at 15 employees, which is not studied in Calligaris (2015) nor Calligaris et al. (2016). As mentioned, Gopinath et al. (2017) also examines Italian ORBIS data for manufacturing, but not for services, and focuses on time-trends.

[4]See Socio and Michelangeli (2015).

[5]http://www.bvdinfo.com/en-gb/our-products/company-information/international-products/orbis

[6]In this respect ORBIS differs to the similar COMPUSTAT dataset for North American listed companies, for example. These data are used in Bloom (2009).

Table 1: Dataset Coverage (%)

	MANUFACTURING				*SERVICES*			
	Operating Revenue		Number of Employees		Operating Revenue		Number of Employees	
Year	*Raw*	*Clean*	*Raw*	*Clean*	*Raw*	*Clean*	*Raw*	*Clean*
2008	63.40	50.97	48.13	42.35	52.58	18.57	34.83	15.25
2009	66.66	50.93	47.41	41.23	55.26	19.20	34.42	15.46
2010	68.20	51.72	49.32	43.37	55.45	18.05	34.08	14.80
2011	71.18	61.75	65.98	56.54	58.96	22.79	50.83	21.89
2012	70.48	60.95	68.09	58.22	58.92	23.30	52.87	22.08
2013	72.19	62.67	69.69	59.36	60.91	22.64	52.23	21.79

Notes: I compute the ORBIS totals according to the raw data, which has not undergone a cleaning process (other than to remove consolidated firm balance sheet information), and for the dataset post-cleaning. Both data use NACE r2. industry definitions of Manufacturing and Services aggregates.

ship information, for example). The data include measures of firm capital as an observation, which is an advantage.[7] The analysis in this paper will principally focus on the Italian manufacturing sector, for data which cover the period 2000-2013 inclusive, however there will be some discussion of the nature of misallocation in the services sector also.[8]

I do choose to restrict to observations with employment information at the baseline, in order to consider firms with more thorough accounting standards, given large heterogeneity in the ORBIS sample. The data undergo a cleaning process, as is standard in the literature, and I largely follow the procedures outlined in **?**.[9] It is important to note that I drop firms with one employee at the baseline, as I am concerned about spurious filings for tax reasons, and about the inclusion of professionals in the service sector (e.g. dentists) who may be registered individually, while in fact operating in a larger practice, which is (to all intents and purpose) the real "firm".

The coverage of the dataset is summarized in Table 1, which details the size of the aggregates computed from the ORBIS dataset relative to those found in the Structural Business Statistics from Eurostat, computed across the relevant sectors. The sample is restricted to

[7]For many firm-level datasets, only investment is observable, so the value of the capital stock must be inferred from investment series using the perpetual inventory method. Such usually methods require an initial value of capital to be assigned to firms, which is typically imputed from firm size and aggregate capital stock data. For this reason the perpetual inventory method likely leads to bias for short firm-level time-series, unless measures of capital stock are available.

[8]ORBIS data are available for a 10 year window. This study employes the 2015 edition of ORBIS, which does include 2014 data. However, ORBIS data are completed with a lag of two years, and the 2014 data had not been fully updated in the version of ORBIS employed in the study, so it was omitted from analysis.

[9]Information on the data cleaning methods used can be found in Section A.

observations for which it is possible to compute a production function: i.e. firm-years where operating revenue, wages, fixed assets and materials are present. It can be seen that in 2013, for the cleaned data sample, we attain 63% coverage by operating revenue in the manufacturing sector. Many of the observations in the services data do not meet the criteria for inclusion in the cleaned data, and coverage in 2013 represents only 23% by operating revenue (having fallen from 61% in the raw data). Coverage by number of employees relative to coverage by revenue is a little lower for manufacturing, at 59% in the clean data in 2013. Coverage by employees is comparable to coverage by revenue for services, at 22% in 2013. The smaller coverage by employees in part reflects the fact that the sample used to compute the table includes many firms with missing data on employment, although this issue is ameliorated in the later years of the sample. It will remain true however that a large body of workers can be found in firms that are sufficiently small to be excluded from the ORBIS dataset, reducing measures of coverage using employment relative to measures of coverage by production. The increase in the percentage of Italian firms contained within the dataset over the sample years is notable, and may reflect an expansion in the data collection methods operated by Bureau van Dijk. An alternative hypothesis would be superior relative performance of firms within the sample in comparison to firms outside the sample.

Data are not from a census, and were not collected as a random sample from the underlying population, and so caution must be exercised while attempting to extrapolate in-sample results to the general economy, given the selected nature of the usable sample. However, the table shows that coverage is sufficiently high that any TFP gains from reallocation computed in-sample will comprise a large enough segment of the Italian economy as to be macro-economically important.

I also examine the representativeness of the sample with respect to the population size distribution using Eurostat data, as can be seen in Table 2. Here we see for manufacturing that the contribution of small firms to the sample, for labour market variables, to be fairly consistent with the contribution of small firms in the population. The relative contribution of medium firms for these variables is somewhat over-represented relative to medium firms in the aggregate data. The relative contribution of small and very large firms are both lower in the ORBIS data than in the aggregate data. The sevices data can only be compared to the aggregate data in the case of gross output. In this case small firms are under-represented and large firms over-represented, relative to the aggregate data. Overall I find that the ORBIS data do an acceptable job at mapping the relative contributions of the three size-classes to their relative contributions in aggregate data. Although (for manufacturing) the data overall account for around 60% of the Italian economy, as was seen in Table 1, it appears from inspection of Table 2 that it remains a representative sub-sample. The table also records that restricting data to a balanced panel would severely reduce the representativeness of the data for smaller firms, which is unsurprising since these firms are likely to be younger, with a shorter time-series.

Basic summary statistics are presented in Table 7 for the manufacturing sector only, since these will be most relevant for the adjustment-cost analysis, which will focus on the manufacturing sample. The means for the production function variables are pooled across years. We can see that the average value added for the services sector is much larger than that of the manufacturing samples. We also see very large standard deviations for these variables, which will likely reflect between-industry variation in the average size of firms. Table 7 also breaks down summary statistics according to size, for reference.

4 Misallocation Analysis

In this section I summarize the nature of misallocation in the Italian data generally, and draw attention to the key features that the model of labour adjustment frictions will need to match. Misallocation of capital is discussed, as well as labour, for completeness.

4.1 Misallocation Framework (Hsieh and Klenow, 2009)

In this section I outline the misallocation framework of Hsieh and Klenow (2009) [henceforth, HK]. Such analyses have been applied to many firm-level datasets in a rapidly growing literature, and an extensive discussion can of course be found in the original paper. I focus the exposition on the relation between industry TFP and the distribution of factor revenue products.

The model presented in HK is a standard static model of monopolistic competition, with firms differing in their idiosyncratic productivities in the manner of Melitz (2003). However, firms also potentially face idiosyncratic "distortions" or "wedges" that are designed to encapsulate the many (un-modeled) reasons why allocations of resources may deviate from efficiency in the static case. These distortions appear as extra payments (or subsidies) that are taken by optimizing firms as exogeneous when choosing how much to produce. The economy is closed, and has two inputs to production (capital and labor). As will be become clear in later discussion, and as emphasized by Asker et al. (2014), the decision of a given firm to deviate from its optimal choice in a frictionless static model can also be rationalized

Table 2: Share of Total Manufacturing Economic Activity By Size Class in 2012 (%)

		MANUFACTURING		*SERVICES*	
		Full	*Balanced*	*Full*	*Balanced*
Employment					
ORBIS	1-19 employees	23	1	20	2
	20-249 employees	51	68	39	33
	250+ employees	25	29	39	64
Eurostat (SBS)	1-19 employees	29		.	
	20-249 employees	41		.	
	250+ employees	27		.	
Wage Bill					
ORBIS	1-19 employees	16	1	18	3
	20-249 employees	49	63	36	42
	250+ employees	30	30	42	52
Eurostat (SBS)	1-19 employees	21		.	
	20-249 employees	43		.	
	250+ employees	34		.	
Gross Output					
ORBIS	1-19 employees	15	3	25	9
	20-249 employees	44	65	39	52
	250+ employees	34	26	32	36
Eurostat (SBS)	1-19 employees	18		44	
	20-249 employees	38		22	
	250+ employees	41		31	

Notes: Presents the percentage share of economic activity (by variable) acounted for by firms belonging in three size categories, for the full and balanced samples. Also presents the corresponding size class shares from Eurostat Structural Business Statistics (SBS). Data are from the year 2012 since disaggregate SBS data for Italy are most complete for this year. Disaggregate data on wages and employment by size are not available for services.

without static wedges, as a temporary deviation from the static first-best in the presence of adjustment costs and stochastic revenue.

Assume there are S different industries, each comprised of M_s firms, indexed by i, each producing differentiated varieties of goods.[10] Output for industry s is given by the CES aggregator:

$$Y_s = \left(\sum_{i=1}^{M_s} Y_{si}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}},$$

where $\epsilon > 1$ is the elasticity of substitution between varieties, Y_{si}.[11]

Firms produce their variety according to the Cobb-Douglas production function

$$Y_{si} = A_{si} K_{si}^{\alpha_{K,s}} L_{si}^{\alpha_{L,s}}, \tag{1}$$

where A_{si} is the firm-specific productivity, K_{si} is capital, L_{si} is labor, and $\alpha_{K,s}$ and $\alpha_{L,s}$ are the capital and labour shares, respectively. Here we have made the assumption that firms within the same industry share the same production technology, and we impose constant returns to scale, so $\alpha_{K,s} + \alpha_{L,s} = 1$.

Firms choose labor and capital to maximize their profits, which are given by:

$$\Pi_{si} = (1 - \tau_{Y,si}) P_{si} Y_{si} - w L_{si} - (1 + \tau_{K,si}) R K_{si},$$

where P_{si} is the price of variety i, w is the wage, and R is the rental rate of capital. Here $\tau_{Y,si}$ and $\tau_{K,si}$ are the firm-specific distortions, which can be respectively termed the "output distortion" and "the capital distortion". The distortions can be understood to be taxes in the case they are positive, or subsidies in the case they are negative. Thus, a positive output distortion represents un-modeled factors that suppress the size of a firm below its efficient level, increasing the marginal products of labor and capital equally. HK give the example of government restrictions on size and transportation costs. A positive capital distortion represents un-modeled factors that increase the costs of capital relative to labor for a particular firm, for example financial frictions or trade union power. The reason the effects of factors affecting the price of labor and capital are incorporated together is that we are only able to identify as many distortions as there are factors of production.

[10]Each firm produces one variety, and each variety is produced by a single firm, so the index i is interchangeable for firms and varieties.

[11]This analysis is entirely static, so the time-subscript is supressed. During the dynamic analysis a time-subscript will be introduced.

Profit maximization implies the following

$$P_{si} = \frac{\epsilon}{\epsilon - 1} \left(\frac{R}{\alpha_{K,s}}\right)^{\alpha_{K,s}} \left(\frac{w}{\alpha_{L,s}}\right)^{\alpha_{L,s}} \frac{(1+\tau_{Ksi})^{\alpha_{K,s}}}{A_{si}(1-\tau_{Ysi})}, \tag{2}$$

where $\epsilon/(\epsilon-1)$ can be understood as the mark-up charged by the firm, over its marginal cost. Thus the prices set by firms are inversely proportional to their productivity, and to the output distortion, and proportional to the capital distortion weighted geometrically by the capital share.

As is described in HK, with no distortions, more productive firms will grow until the the marginal revenue product of labor and capital are respectively equalized within industries to equal factor costs, as can be seen from the following equations:

$$\begin{aligned} MRPL_{si} \equiv& \alpha_{L,s}\frac{\epsilon-1}{\epsilon}\frac{P_{si}Y_{si}}{L_{si}} = w\frac{1}{1-\tau_{Ysi}}, \\ MRPK_{si} \equiv& \alpha_{K,s}\frac{\epsilon-1}{\epsilon}\frac{P_{si}Y_{si}}{K_{si}} = R\frac{1+\tau_{Ksi}}{1-\tau_{Ysi}}. \end{aligned}$$

We note that the presence of distortions will cause allocations of labor and capital to deviate from the efficient level, and that this causes divergences in the marginal revenue products of factors that would not exist otherwise. Thus, if distortions are present, welfare gains from reallocation of factors of production exist, in principle. However, as will be argued later in the paper, they can also be rationalized as *temporary* deviations from frictionless optimality in response to productivity shocks (the argument first made by Asker et al. (2014)). Further, since no attempt will be made to endogenize τ_K or τ_Y, and since they are essentially residuals designed to explain differences in factor revenue products in given industries, the remainder of the paper will summarize misallocation in terms of the distribution of MRPK and MRPL.

HK show that we can then define industry TFP as the following:

$$TFP_s = \left[\sum_{i=1}^{M_s}\left(A_{si}\left(\frac{\overline{MRPK}}{MRPK_{si}}\right)^{\alpha_s^K}\left(\frac{\overline{MRPL}}{MRPL_{si}}\right)^{\alpha_s^L}\right)^{\epsilon-1}\right]^{\frac{1}{\epsilon-1}}, \tag{3}$$

where $\overline{MRPK}$ and $\overline{MRPL}$ are weighted averages of their respective factor productivities.[12]

[12] $\overline{MRPL}_s \equiv w/\left(\sum_{i=1}^{M_s}(1-\tau_{Y,si})\frac{P_{si}Y_{si}}{P_sY_s}\right)$ and $\overline{MRPK}_s \equiv R/\left(\sum_{i=1}^{M_s}\frac{(1-\tau_{Y,si})}{(1+\tau_{K,si})}\frac{P_{si}Y_{si}}{P_sY_s}\right)$. In fact HK display Equation 3 in terms of revenue productivity, $TFPR_{is} = P_{is}A_{is}$. I present their analysis in terms of

It is true from Equation 3 that TFP_s depends negatively on the variance of the marginal productivities of labour and of capital. This can be seen clearest in the case that $\log A_{si}$, $\log MRPK_{si}$, and $\log MRPL_{si}$ is multivariate normal. In this case Equation 3 can be re-expressed:

$$\begin{aligned}\log TFP_s =& \frac{1}{1-\epsilon}\left(\log M_s + \log E[A_{si}^{\epsilon-1}]\right)\\ &+\left(-\frac{\epsilon}{2}+\epsilon\alpha_{K,s}-\left(\frac{\epsilon\alpha_{K,s}^2}{2}+\frac{\alpha_{K,s}\alpha_{L,s}}{2}\right)\right)\sigma_{ML,s}^2\\ &-\left(\frac{\epsilon\alpha_{K,s}}{2}+\frac{\alpha_{K,s}\alpha_{L,s}}{2}\right)\sigma_{MK,s}^2\\ &+(1-\epsilon)\alpha_{K,s}\alpha_{L,s}\sigma_{MK,ML,s},\end{aligned} \tag{4}$$

where $\sigma_{ML,s}^2$ is the industry-level variance of $\log MRPL$, $\sigma_{MK,s}^2$ is the industry-level variance of $\log MRPK$, and $\sigma_{KY,s}$ is their covariance.

We can also see from the coefficients on Equation 4 that, as well as being decreasing in the variance of the two factor productivities, industry TFP is further decreasing in their covariance, given that $\epsilon > 1$. If goods are substitutes, it is especially harmful for aggregate productivity if firms with relatively high marginal productivities to given factors also have relatively high marginal productivity to the other factor.

Aggregation

Throughout the study, in order to aggregate the measures of misallocation from the S industries, we assume the existence of a single final good produced by a representative firm in a perfectly competitive final good market. The firm combines the output Y_s of the S different sectors using a Cobb-Douglas production function:

$$Y = \prod_{s=1}^{S} Y_s^{\theta_s},$$

where $\sum_{s=1}^{S}\theta_s = 1$ are the shares of aggregate output for the S industries. Cost minimization implies that:

$$\theta_s = \frac{P_s Y_s}{PY},$$

marginal revenue products, which is equivalent, and befits my focus on factor revenue productivities.

where P_s is the price index for industry output Y_s, and $P \equiv \prod_{s=1}^{S}(P_s/\theta_s)^{\theta_s}$ is the price index of the final good and is set equal to 1. We should note that the misallocation analysis is conditional on a fixed aggregate stock of capital and labor, and we must make the assumption that the number of firms in each industry is not affected by the extent of misallocation.

4.2 Measurement

With respect to attaining the needed measures of production function variables from ORBIS, the output measure, PY, is added value (computed as operating revenue subtract material expenditure), capital is measured as fixed assets, and the labour input is measured as number of employees. Fixed assets are in book value, which admittedly does not control for the effects of depreciation upon the capital stock.[13] In the case that firms in the same industry have the same capital profile this will not affect measures of misallocation. Further, this is also the same treatment of capital as in the baseline specification of HK.[14]

The measurement of the industry-level capital share, $\alpha_{K,s}$, deserves some discussion. In the presence of distortions, one cannot simply compute them from observed factor shares within industries, since the production function parameters are not separately identifiable from the average levels of the distortion in each sector. One solution is to take factor shares from an economy we believe to be relatively undistorted, and here I follow HK and choose the U.S. value in order to summarize misallocation.[15] The implication of such an approach is that any measures of misallocation that are derived are best understood as measures of misallocation *relative* to those levels that obtain in the U.S., and influence factor share measurements in U.S. data.[16]

[13] Any capital stock correction, i.e. the perpetual inventory method, risks biasing measures if the depreciation rate is improperly calibrated, with the bias increasing over time as mistakes are compounded.

[14] In fact these variables are deflated, although it does not affect any of the dispersion or reallocation results. The manufacturing data are deflated using price indices from Eurostat at the NACE r2 digit level. The services data are deflated using the HICP, since industry-level deflators do not exist for the Italian services sector to my knowledge. The capital stock is deflated using the investment goods deflator, available from the Italian statistical authority (Istat). Details can be found in Section A.

[15] A small discrepency with the later study of adjustment-costs to labour is introduced, where I use a system GMM approach to estimate factor shares. This does not affect the nature of the argument, since the model is designed to match moments that do not depend on factor shares. The use of U.S. shares ensures consistency between my misallocation estimates and those of HK.

[16] I do however, differ insofar as I choose to use BEA data to attain measures of the shares, as opposed to the NBER-CES dataset used in HK. The reason is that the NBER-CES dataset covers only the manufacturing industry, whereas the BEA data covers both manufacturing and services. In order to preserve comparability between the analysis for manufacturing and services, I choose to use the BEA data throughout. Admittedly, here I do sacrifice variability for the manufacturing sector in the sense that the BEA data is less "finely" defined with respect to industries, and while I am able to match at the NACE 2-digit level for services, all

At the baseline, I set ϵ to 4. As detailed in HK, this is relatively "conservative" with respect to this study, in the sense that estimates as high as ten have been recorded in the literature, and the size of computed TFP gains from reallocation are increasing in this parameter. In fact, HK themselves set a value of 3; I use 4 for conformability with the study of Asker et al. (2014).

At the baseline, "industries" are defined according to 4-digit level NACE revision 2 codes. I do delete industries for which we see less than 20 observations during any year in the sample, post-cleaning. One would expect the size of computed gains to be increasing in the "breadth" of industry definitions, since dispersion in outcomes between increasingly diverse firms is attributed to misallocation, as opposed to heterogeneity of production functions and factor costs.

4.3 Analysis

From Table 3, which plots the components of the TFP as expressed in the decomposition of Equation 4, we see that the contribution of the variance of the marginal revenue product of labour is lower than that of capital, which has been shown to be the case in several papers, including Asker et al. (2014) and **?**. We can also see that the misallocation appears to have increased in the years from 2005 to 2013 for manufacturing, and decreased for Services, at least at this second-order approximation. An in-depth discussion of the dynamics of misallocation over the period, as well as a comparison of the scale of Italian estimates with estimates from other countries found in the literature, can be found in Section 7. The main point to conclude thus far is that, while capital misallocation is the more important determinant of productivity in the Italian case, the contribution of labour misallocation is far from irrelevant. Further, the contribution of their covariance is negligible.

4.4 Evidence on the Relation Between Distortions and Size

In order to examine a potential relationship between distortions and firm size, I assess the correlations of the marginal revenue products with size measured in terms of labour and value added. Columns (1) and (2) of Table 4 show the correlation of MRPL and MRPK in deviation from their industry-year means, with firm size. With respect to labour the correlation is almost zero. Interestingly, there is a strong correlation between the misallocation of labour and the value-added of firms. This suggests firms with a large-value added ought

manufacturing firms get the same factor shares.

Table 3: Decomposition of $\log(TFP)$

	Year	Term I $\frac{1}{1-\sigma}\left(\log M_s + \log E[A_{si}^{\sigma-1}]\right)$	Term II $b_1\sigma_{ML}^2$	Term III $b_2\sigma_{MK}^2$	Term IV $b_3\sigma_{MK,ML}$
MANU	2005	1.11	-0.028	-0.080	-0.0048
	2009	1.15	-0.035	-0.12	-0.0044
	2013	1.15	-0.034	-0.12	-0.001
SERV	2005	1.23	-0.080	-0.14	-0.0062
	2009	1.21	-0.080	-0.13	-0.0023
	2013	1.19	-0.082	-0.098	-0.011

Notes: This table shows the contribution of the four components of the log-normal decomposition of $\log(TFP_s)$, expressed relative to the value of $\log(TFP)$ under the approximation. The aggregate contribution of Term x is calculated as $\sum_{s=1}^{S}(\theta_s/\log(TFP_s)) * \text{Term}\, x_s$, $x \in \{I, ..., IV\}$. The coefficient terms are: $b_1 \equiv \frac{1}{2}\{-\epsilon + 2\epsilon\alpha_{K,s} - \left(\epsilon\alpha_{K,s}^2 + \alpha_{K,s}\alpha_{L,s}\right)\}$, $b_2 \equiv -\frac{1}{2}\left(\epsilon\alpha_{K,s}^2 + \alpha_{K,s}\alpha_{L,s}\right)$, $b_3 \equiv (1-\epsilon)\alpha_{K,s}\alpha_{L,s}$.

to take on more employees, in Italian data, and therefore provides suggestive early evidence that the large proportion of small firms in Italy may contribute to its low TFP, as has been hypothesized in papers such as Pellegrino and Zingales (2014).[17] With respect to capital misallocation both values are low. Columns (3) and (4) present correlation of the squared deviation of MRPL and MRPK from their industry-year means. The idea is to capture whether small or large firms are contributiong more to the overall variance of the two factor revenue products, one of the ultimate determinants of misallocation. Again, with respect to labour the correlation is low. It is slightly higher with respect to value added for labour, reflecting the previous result of a large contribution of high-value added with sub-optimal levels of employment, from a welfare perspective. These results are reflected in Columns (5) and (6), which show the absolute deviation of each observation of marginal revenue product from its industry-year mean. Figures 7 and 8 arrange observations by bins in terms of size by labour and value-added, and then computes the average marginal revenue product deviations and absolute deviation from their industry-year mean.

[17] The negative correlation between firm size (measured as labour costs) and measures of misallocation for Italy is also documented in the paper of **?**, using data on manufacturing firms from CERVED.

Table 4: Correlations Between MRP Deviations and Firm Size in Manufacturing

Size Measure	$\widehat{MRPL}$	$\widehat{MRPK}$	$\widehat{MRPL}^2$	$\widehat{MRPK}^2$	$\left\lvert\widehat{MRPL}\right\rvert$	$\left\lvert\widehat{MRPK}\right\rvert$
Labour	-0.000451	-0.00524	-0.00246	-0.00211	-0.00515	-0.00464
Value Added	0.338	-0.00334	0.0870	0.000394	0.258	-0.000695

Notes: Table displays the correlation between MRPL and MRPK in deviation from their respective industry-year means, with Labour and Value-Added. Also displayed are the squared deviation and the absolute deviation. These values are normalized by their industry-year mean.

5 Misallocation and Adjustment Costs to Hiring and Firing

This model presents a model of dynamic labour demand, and uses it to assess the contribution of adjustment costs to labour to observed misallocation of labour in the Italian case, as well as assessing its capability to match its relation with firm size.

5.1 Evidence for "Lumpiness"

Evidence for adjustment costs to changes in factor use has typically come from observed "lumpiness" in the investment and hiring decisions of firms, implying that firms undertake large increases in factor use in infrequent spikes – this would follow from a large fixed-cost to adjustment. I present the histograms of hiring and investment rates, $(E_{t+1} - E_t)/E_t$ and $(K_{t+1} - K_t)/K_t$ and respectively, as can be seen in Figure 6.

We see that the distributions of both factors are non-normal, with mass around zero. Interestingly, in these data, it is the hiring rate that displays the greater evidence for inaction, in the sense that we see a very striking spike in mass around zero for this factor. Thus, it seems that adjustment costs to labour, as opposed to capital, may be the more important when considering factor use by Italian firms. We further see that the distribution of investment has a larger dispersion than that for labour, insofar as we see more mass above even 50% in this factor. Such high levels of factor growth are comparatively rare for labour. We see little difference in the distributions between manufacturing and services. The non-normality of these distributions provides preliminary evidence for the existence of non-differentiable adjustment costs in these data, particularly in the case of labour.

5.2 A Structural Model

In this section I estimate a dynamic model of labour demand in order to assess the ability of adjustment costs to explain observed dispersion in MRPL. Under the first specification of the model, the firm faces adjustment costs only to its labour input. The model is in discrete time. We assume capital to be a rented input, and adjusted frictionlessly.[18]

Each firm i in period t produces a single differentiated output good using a constant-returns to scale production function:

$$F(A_{it}, L_{it}, H_{it}) = A_{it}L_{it}H_{it}, \tag{5}$$

where L_{it} is the number of employees, H_{it} is hours per employee, and A_{it} is productivity (which is potentially stochastic). The output of this function is value added, as in the previous analysis.[19]

Firms face an isolastic demand curve with an elasticity given by ϵ:

$$Q_{it} = B_{it}P_{it}^{-\epsilon}, \tag{6}$$

where Q_{it} is the output of a firm, and B_{it} is a demand shifter. Firms are allowed to face demand shocks, so B_{it} is potentially stochastic. The elasticity of substitution does not vary across firms, and is restricted such that $\epsilon > 1$ (ensuring goods are substitutes).

We can combine the production function and demand curves to attain the revenue function:

$$S(\Omega_{it}, L_{it}, H_{it}) = \Omega_{it}(L_{it}H_{it})^{\bar{\beta}}, \tag{7}$$

where $\bar{\beta} = [1 - (1/\epsilon)]$, and $\Omega_{it} \equiv A_{it}^{1-(1/\epsilon)}B_{it}^{1/\epsilon}$. Here Ω is a composite measure of technical productivity and demand shocks.[20] The TFPR measure will be modelled as an AR(1) stochastic process in logs, representing the composite effect of persistent idiosyncratic innovations to demand and productivity faced by the firm.

As mentioned, at this stage we abstract from capital. Throughout the analysis of labour adjustment costs we maintain the assumption that it is a frictionless input to production.

[18] Here I am broadly following the model of Asker et al. (2014), with alterations to focus on labour adjustment costs instead of those for capital.

[19] Asker et al. (2014) employ a production function in gross output, with materials as a frictionless third input to production. I prefer to omit a study of the role played by material inputs for convenience, studying a production function in value added. This can be understood as equivalent introducing an assumption that the production is Leontief in materials relative to output.

[20] Bloom (2009) refers to this measure as "business conditions".

In fact, it can be show that the reduced form of a model with frictionless capital optimally chosen has the same functional form as 7. This becomes relevant during the estimation of Equation 7, but for now we proceed under the assumption that capital can be frictionlessly adjusted.

We can therefore derive a function for the period profit, which is given by:

$$\Pi(\Omega_{it}, L_{it}, H_{it}) = \Omega_{it}(L_{it}H_{it})^{\bar{\beta}} - w_1(1 + w_2 H_{it}^{\gamma})L_{it}.$$

The monthly wage is parameterized by w_1 and w_2. There is a one period "time-to-build" for labour, which follows Bloom (2009) and is plausible given that the model is simulated at monthly frequency. At the baseline I do not allow an exogenous attrition rate for labour, so workers are assumed not to quit the firm voluntarily.

Under the baseline specification, I assume fixed disruption costs to the inclusion of new workers to production, denoted by C_L^F. This is designed to capture the idea that time and resources must be diverted to interviewing new workers, and training them on arrival, etc. There are also convex adjustment costs to labour, parameterized by C_L^Q. There are also piecewise-linear per employee adjustment costs to hiring and firing workers, denoted by C_L^{P1} and C_L^{P2} respectively. The adjustment costs function therefore takes the following form:

$$C(\Omega_{it}, L_{it}, H_{it}, E_{it}) = C_L^F \mathbf{1}_{\{E \neq 0\}} \Pi(\Omega_{it}, L_{it}, H_{it}) + C_L^Q L_{it} \left(\frac{E_{it}}{L_{it}}\right)^2 + \left(C_L^{P1}\mathbf{1}_{\{E>0\}} - C_L^{P2}\mathbf{1}_{\{E<0\}}\right) E_{it}, \tag{8}$$

where $E_{it} \equiv L_{i,t+1} - L_{it}$ represents hiring. As previously mentioned, I then assume that $\omega_{it} \equiv \ln(\Omega_{it})$ follows an AR(1) process given by

$$\omega_{it} = \rho\omega_{i,t-1} + \sigma v_{it},$$

where $v_{it} \sim N(0, 1)$ is an i.i.d. standard normal random variable. This defines the transition function for $\Omega_{it} : \phi(\Omega_{it+1}|\Omega_{it})$.

It is possible to solve for optimal hours as a function on labour and revenue productivity, thus optimal hours H_{it}^*can be expressed as a function of productivity and labour:

$$H^*(\Omega_{it}, L_{it}) = \left(\frac{\beta}{w_1 w_2 \gamma}\Omega_{it} L_{it}^{\beta-1}\right)^{\frac{1}{\gamma-\beta}}.$$

We can then substitute for optimal hours in the profit and adjustment costs functions as a function of L_{it} and Ω_{it} only, and denote these functions $\Pi^*(\Omega_{it}, L_{it}, H_{it})$ and $C^*(\Omega_{it}, L_{it}, E_{it})$.

This implies the firm's value function $V(\Omega_{it}, L_{it})$ is given by the following Bellman equation:

$$V(\Omega_{it}, L_{it}) = \max_{E_{it}} \ \Pi^*(\Omega_{it}, L_{it}) - C^*(\Omega_{it}, L_{it}, E_{it}) + \beta \int_{\Omega'} V(\Omega'_{it}, L_{it} + E_{it}) \phi(\Omega_{it+1}|\Omega_{it}) d\Omega_{it},$$

which implies that a firm's policy function is the hiring choice that maximizes the firm's continuation value, subtract the cost of hiring. I assume stationarity of the problem. Given that it is not possible to solve the firm's problem using analytical methods, the function is solved for numerically by policy-function iteration; the details of the algorithm are described in Section B.2.

5.3 Estimation

Estimation proceeds in two stages. First I jointly estimate both the production function parameter and the persistence parameter in the AR(1) productivity process using a system GMM approach, following Blundell and Bond (2000). For the next stage, given the estimates attained via GMM, I estimate the adjustment cost function parameters using the method of simulated moments, following Bloom (2009). Certain parameters are assigned values commonly used in the literature.[21]

System GMM

Typically, in order to estimate production functions the literature uses estimation strategies of the Olley and Pakes (1996) and Levinsohn and Petrin (2003) type, which use lags either investment or materials respectively as a proxy for the predictable component of the productivity shock. However, these methods require frictionless adjustment of factors of production, and are therefore not appropriate in the context of this investigation. Asker et al. (2014) assume the frictionless adjustment of labour (and materials), and are then able to consistently back out the parameter on capital from the constant returns to scale assumption they impose. In this study, which is designed to quantify adjustment costs to labour, such an assumption clearly cannot be made. I thus follow the system GMM approach of Blundell and Bond (2000), in a similar manner to Cooper and Haltiwanger (2006).

[21]Such a two-stage approach follows that of Cooper and Haltiwanger (2006).

As mentioned previously, during the estimation of the production function parameter to be incoporated into the baseline model (which is expressed in terms of labour), we in fact employ a two-factor production function in both capital and labour. After estimation, we then back-out the value of the (single) factor elasticity in the baseline production function from the (two) estimated elasticities that result from successful estimation of the two-factor model. Consider the Cobb-Douglas production function, now in both labour and capital:

$$\tilde{F}(\tilde{A}_{it}, K_{it}, L_{it}) = \tilde{A}_{it} K_{it}^{\alpha_{K,s}} L_{it}^{\alpha_{L,s}}.$$

[22]I impose constant returns to scale, so $\alpha_{K,s} + \alpha_{L,s} = 1$. As before, we can re-express in terms of a revenue function in capital in labour, given that firms face a downward facing demand function. Thus:

$$\tilde{S}(A_{it}, K_{it}, L_{it}, H_{it}) = \tilde{\Omega}_{it} L_{it}^{\beta_{L,s}} K_{it}^{\beta_{K,s}}, \tag{9}$$

where $\tilde{\Omega}_{it} = \tilde{A}_{it}^{\phi} B_{it}^{1/\epsilon}$, $\beta_{L,s} = \phi \alpha_{L,s}$, and $\beta_{K,s} = \phi \alpha_{K,s}$, and $\phi \equiv (\epsilon - 1)/\epsilon$.

We will be able to take Equation 9 to the data. In order to estimate the parameters of Equation 9, we take logs, express in terms of observations on firms in given years and industries (indexed by i , t and s respectively), introduce a statistical error υ_{it}, in order to get the following:

$$s_{it} = \beta_{L,s} l_{it} + \beta_{K,s} k_{it} + \tilde{\omega}_{it} + \upsilon_{it}, \tag{10}$$

where y_{it} is log value added of firm i in year t and industry s, l_{it} is log employment, k_{it} is log capital stock, $\tilde{\omega}_{it}$ is the log of each firm's idiosyncratic revenue productivity shock. We then make the following assumptions regarding the revenue productivity process:

$$\begin{aligned} \upsilon_{it} &= \gamma_t + \eta_i + m_{it}, \\ \tilde{\omega}_{it} &= \rho_{\tilde{\omega}} \tilde{\omega}_{i,t-1} + e_{it}, \quad |\rho_{\tilde{\omega}}| < 1, \\ e_{it}, m_{it} &\sim MA(0), \end{aligned} \tag{11}$$

The idiosyncratic productivity shock $\tilde{\omega}_{it}$ is modelled as a first-order autoregressive process with parameter $\rho_{\tilde{\omega}}$. The remaineder of the error term is modelled so that it contains three terms: (1) an aggregate term, or year fixed-effect, γ_t; (2) a firm-level fixed effect which is

[22]Unfortunately I am unable to observe hours, so this variable is omitted from the estimation.

unobservable, η_i; (3) serially uncorrelated measurement error, m_{it}. Note that while the three later terms are modelled in the estimation of Equation 10, only the idiosyncratic shock features in the simulation of the model.

Equation 10 can then be quasi-first differenced, employing autoregressive parameter

$$s_{it} = (\phi - \beta_{K,s}) * l_{it} - \rho_{\tilde{\omega}}(\phi - \beta_{K,s}) * l_{i,t-1} + \beta_{K,s}k_{it} - \rho_{\tilde{\omega}}\beta_{K,s}k_{i,t-1} + \rho_{\tilde{\omega}}s_{i,t-1} \dots \tag{12}$$
$$+ (\gamma_t - \rho_{\tilde{\omega}}\gamma_{t-1}) + (1 - \rho_{\tilde{\omega}})\eta_i + (e_{it} + m_{it} - \rho_{\tilde{\omega}}m_{i,t-1}),$$

embedding the restriction that given the assumption of constant returns to scale, it is true that $\alpha_L + \alpha_K = \phi$. Alternatively, again following Blundell and Bond (2000), we can re-express Equation 12 in terms of a reduced form coefficient vector: $\pi = (\pi_1, \pi_2, \pi_3, \pi_4, \pi_5)$

$$s_{it} = \pi_1 l_{i,t} + \pi_2 l_{i,t-1} + \pi_3 k_{i,t} + \pi_4 k_{i,t-1} + \pi_5 s_{i,t-1} + \gamma_t^* + (\eta_i^* + w_{it}),$$

given $\gamma_t^* \equiv (\gamma_t - \rho_{\tilde{\omega}}\gamma_{t-1})$, and $\eta_i^* \equiv (1 - \rho_{\tilde{\omega}})\eta_i$. We also have three restrictions: $\pi_2 = -\pi_1\pi_5$, $\pi_4 = -\pi_3\pi_5$, under which estimates of $(\alpha_K,\ \rho_{\tilde{\omega}})$ can be attained by minimum-distance. If we were to assume away measurement error ($var(m_{it}) = 0$) we would have $w_{it} = e_{it} \sim MA(0)$, otherwise $w_{it} \sim MA(1)$.

Given an assumption on the initial conditions

$$E[x_{i1}e_{it}] = E[x_{i1}m_{it}] = 0, \quad \text{for } t = 2, ..., T,$$

we attain the following moment conditions

$$E[x_{i,t-j_1}\Delta w_{it}] = 0, \quad \text{where } x_{it} = (l_{it},\ k_{it},\ s_{it}),$$

for $j_1 \geq 2$ when $w_{it} \sim MA(0)$, and for $j_1 \geq 3$ when $w_{it} \sim MA(1)$.

However, Blundell and Bond (2000) note that if instruments are weak, the first-differenced GMM estimator has poor finite sample properties. However, under the assumption that

$$E[\Delta n_{it}\eta_i^*] = E[\Delta k_{it}\eta_i^*] = 0,$$

and given that the initial conditions satisfy

$$E[\Delta s_{i2}\eta_i^*] = 0,$$

it is possible to attain additional moment conditions

$$E[\Delta x_{i,t-j_2}(\eta_i^* + w_{it})] = 0,$$

for $j_2 = 1$ when $w_{it} \sim MA(0)$, and for $j_2 = 2$ when $w_{it} \sim MA(1)$. Thus Blundell and Bond (2000) advise using lagged first differences of the variables as instruments for the equations in levels, and the same approach is followed in this paper.[23]

For the purposes of the SMM part of this paper we do not disaggregate estimation by industry, and pool all manufacturing firms together. In order to estimate an aggregate production function parameter from the industry-level Equation 10, we estimate for given industry sub-samples, and then take a weighted-average (with weights in terms of the average share of value added for a given industry in the sample period). The aggregate labour parameter is then backed out from the constant returns assumption. The dependent variable of sales in the analysis is measured as real operating revenue. Capital is measured as real fixed assets. Labour is measured as the number of employees. I set $j_1 = 3$ and $j_2 = 2$ to account for the possibility of auto-correlated measurement error.

The estimates attained are 0.343 for the capital share of sales (and therefore 0.407 for the labour share), and 0.803 for the for the persistence parameter. The capital share parameter is close to the estimate presented in Cooper and Haltiwanger (2006), of 0.28. The persistence parameter is slightly lower, since they report a value of 0.885. This suggests that the persistence of Italian revenue shocks could be lower than those in U.S. data. The revenue productivity shocks themselves can be extracted by first de-meaning the residuals of Equation 10, to remove the year fixed-effects. Then the residual of Equation 11 is understood to be the revenue productivity shock (for the purposes of shock extraction, industry-level coefficients are used).

The curvature parameter used for the single-factor production function can then be backed out from the estimates of α_L and α_K from the two-factor production function. Following Cooper et al. (2015), the value of the single production function parameter used in the case when the only input is labour follows from optimization of capital for given labour in the two-factor production function, under the assumption that there are no adjustment costs of investment. This implies that the period profit function (ignoring payments to labour):

$$\pi(\tilde{\Omega}, K, L) = \tilde{\Omega} L^{\alpha_L} K^{\alpha_K} - rK.$$

[23] Blundell and Bond (2000) show that joint stationarity of the s_{it} and x_{it} processes is sufficient (but not necessary) for the validity of the additional moment restrictions for the equations in levels.

Given that under this specification capital is a flexible input, maximization with respect to capital leads to the reduced form (Equation 7), where:

$$\beta = \frac{\phi\alpha_L}{1-\phi\alpha_K}. \tag{13}$$

We are then able to re-parameterize Equation 7, substituting $\overline{\beta}$ for β.[24] Since Equation 13 is true even with the inclusion of hours, 7 is re-expressed with β as the exponent on labour. These operations then imply that the coefficient of the labour-only production function is 0.619. This is slightly below the calibration of α in Cooper et al. (2015) as 0.7.

Adjustment Cost Parameters

In order to estimate the parameter vector $\boldsymbol{\theta} = \{C_L^F, C_L^Q, C_L^{P1}, C_L^{P2}, \gamma\}$, I use the method of simulated moments, conditional on the values of β and $\rho_{\tilde{\omega}}$ previously estimated by system GMM. This requires that a vector of moments be selected for our model to match, which are chosen to help identify adjustment costs specifically. The model is then simulated, and selected moments calculated from the simulated dataset. I choose the value of $\boldsymbol{\theta}$ that minimizes a distance criterion between the simulated moments and the moments drom the data.

The simulations are conditional on the choice of several parameters that are calibrated, and are not estimated from the data. The wage equation parameter w_1 is arbitrary, and set to normalize the annual wage bill to equal 1, as in Bloom (2009). The wage equation parameter w_2 is set to minimize wages per hour at a 40 hour working week. As discussed, the parameter ϵ is set to equal 4, matching the value in Bloom (2009) and Asker et al. (2014).

I choose to target the following moments:

1. $\Pr[|\Delta \ln L| < 0.05]$: the proportion of firms with less than a 5 percent year-on-year change in labour;

2. $\Pr[|\Delta \ln L| > 0.2]$: the proportion of firms with more than a 20 percent year-on-year chage in labour;

3. $\mathrm{sd}(\Delta \ln L)$: the standard deviation of the year-on-year change in log labour;

4. $\mathrm{skw}(\Delta \ln L)$: the skewness of the year-on-year change in log labour;

[24] Details of the derivation can be found in Appendix Section 9.

5. $\mathbb{E}[|\Delta \ln L| \mid |L| > 0.05]$: the mean absolute year-on-year change in log labour, conditional on it being greater than 5% in absolute value.

Here I am following the moment selection of Asker et al. (2014), with additional target moments taken from the literature to attain a just-identified setup: the skewness of the hiring distribution is used in Bloom (2009), the conditional expectation of the year-on-year change of log capital is used in Catherine et al. (2017) (in this context I am of course using labour). An examination of the target moments from the Italian data used in this paper and how they compare to target moments computed for other studies in the literature can be found in Table 14, for reference.

Table 8 presents some statistics relating to adjustment costs, which will form the basis for the minimum-distance estimation procedure. We can see that the fraction of firms with a hiring rate below 5% in absolute value is 0.32, the fraction of firms with a hiring rate above 20% in absolute value is 0.43, and that the expected value of hiring conditional on a rate greater than 5% is 0.27.

We can denote the predicted moments from the model as $\Psi_c(\boldsymbol{\theta})$, which I solve by iterating the value function to attain the firms' optimal policies and simulating the model forward at monthly frequency for 30 years for 100,000 firms, and computing moments from the last 20 years of the simulated dataset.

The estimates of the vector are given by

$$\hat{\boldsymbol{\theta}} = \arg\min_{\theta \in \Theta} \mathcal{Q}(\boldsymbol{\theta}) \equiv (\hat{\Psi} - \Psi(\boldsymbol{\theta}))' \boldsymbol{W} (\hat{\Psi} - \Psi(\boldsymbol{\theta})).$$

The weighting matrix is chosen to be identity.[25] I use a simulated annealing algorithm to attain the estimates of $\hat{\boldsymbol{\theta}}$, following Bloom (2009).

5.4 Evaluation

Given the estimated parameter vector, $\hat{\boldsymbol{\theta}}$, the model is then simulated again in order to deliver predictions regarding the following moments of interest. Estimates are displayed in Table 5, and specification (4) is chosen for reasons of superior model fit. The model predictions can be seen in Table 8. We see that the model does a good job of matching targeted moments related to inaction: $\Pr[|\Delta \log L| < 5\%]$ is equal to 0.34 (0.32 in the data),

[25] Here I diverge from Asker et al. (2014) insofar as I use only the aggregate data, and do not solve the model using industry parameters before taking the weighted average of their criterion scores, as they do.

Table 5: Estimates from Baseline Case

	Estimates				
Parameter	(1)	(2)	(3)	(4)	(5)
C_L^F	0	0.0664	0	0.0358	0
C_L^Q	0.4579	5.0640	0.4016	0.1040	0*
C_L^{P1}	0.0725	0	0.0785	0.0421	0
C_L^{P2}	0.0725*	0*	0	0.0230	0.0757
γ	5.1522	5.0041	4.2214	5.7118	3.8478

Notes: The table displays the parameter estimates obtained. An asterisk indicates that the parameter was restricted, and not estimated by SMM. In the case of Column (1) and (2) the piece-wise linear hiring and firing costs are assumed to be identical; only one parameter is estimated here. Under Column (2) and (4) an exogenous quit rate is introduced.

$\Pr[|\Delta \log L| > 20\%]$ is equal to 0.43 (0.43 in the data also). The model somewhat overdelivers on the dispersion of hiring rates at 0.50 (0.34 in the data), and the expected value of hiring conditional on action: $\mathbb{E}[|\Delta \log L| \mid \Delta \log L > 5\%]$ is 0.46 relative to 0.27 in the data. The model generates higher skewness of the hiring rate than in the data: 0.32 in the model relative to 0.04 in the data.

With respect to the target statistic of interest, the standard deviation of the marginal product of labour, we see that the model does an excellent job at accounting for this statistic: predicting 0.57, relative to 0.59 in the data.

Of interest is the failure of the model to match the very low correlation between MRPL and firm size. The correlation of MRPL with respect to labour is effectively zero in the data, and 0.85 in the model. The correlation of MRPL with revenue in the model is 0.34 in the data, and effectively 1 in the model.

Overall, while the model does a good job at matching the variance of MRPL in the data, it does a poor job at accounting for the relation of this statistic with firm-size. This suggests that the model of dynamic labour demand may be missing important features of the determinants of misallocation, in the Italian case at least.

6 Structural Analysis of a Regulatory Threshold

There is a body of research that uses the existence of a threshold in the application of unfair dismissal legislation in Italy, that was in place at least until the reforms of the Renzi administration in 2014, and covers the span of the time-series of data used in this paper. Specifically, under these regulations firms with more than 15 employees found it more difficult to dismiss workers than firms with less than 15 employees or less. This feature of the regulatory environment can be directly incorporated into the model.

6.1 Institutional Background

This section explains the institutional framework in place in Italy prior to the reforms implemented by the Monti administration in 2012, and which had remained essentially the same since the previous reform in 1990.

During this time period, an employer could only terminate the contract of an non-temporary employee legally if there were evidence for "just cause" (*causa giusta*).[26] There were two situations that could allow for such just cause: misbehaving workers could be fired (*giustificato motivo soggettivo*), and contracts could be ended if there was an economic reason that made this necessary (*giustificato motivo oggettivo).*

Workers would always have the option to appeal their dismissal in court. And in the case that the dismissal was ruled as unfair by a judge, the firm would have to pay a sum to the worker that varied with firm size. For the cases where the worker either did not contest his or her dismissal in court, or the judge ruled that the dismissal had just cause, the firm was not required to pay any firing costs, other than the obligatory payment of the *trattamento di fine rapporto.*[27] Further, although of course the law did provide for the legal dismissal of workers in the presence of just cause, Hijzen et al. (2013) argue that the absence of a stringent definition of fair dismissal meant the expected costs of dismissal remained high.

As mentioned, the unfair dismissal costs were increasing in the size of the firm. In fact, there was a discrete threshold at 15 employees, at which firms would face different costs in the case that they were ruled to have dismissed a worker unfairly.[28] Before 1990 firms with less than 15 employees did not operate under employment protection legislation. After the 1990 reform, the small firms were subject to EPL, however the costs were still much lower

[26] See Schivardi and Torrini (2008) for a full discussion.

[27] Another feature of Italian labour regulation was a separation indemnity that depended on tenure (*trattamento di fine rapporto*), and which is a payment from the employer to employee on separation, irrespective of the reason. In practice the employer will retain a set fraction (1/13) of the annual salary of its employees, cumulating the payments, then transferring the sum to the worker at separation. Additionally, it is true that the firm may have faced firing costs for the cases in which they were subject to collective bargaining agreements.

[28] Technically, the threshold is for firms that employ more than 15 employees in the same city, or more than 60 employees in total.

for the smaller firms. In the case of an unfair dismissal, firms with 15 employees or less were required to pay the worker an amount between 2.5 and 6 months of their salary (*tutela obbligatoria*). Firms above the threshold would be required by law to re-instate the worker they had previously dismissed (Article 18, *tutela reale*); in the case that the worker refused to resume their employment (as was the vast majority of cases), the firm was obliged to pay the worker 15 months of their salary.[29]In addition to this, firms would be required to compensate the worker for lost wages from their dismissal until their sentence (along with insurance contributions), with no upper bound.[30] These sums were therefore a function of a variable trial length, which could be as much as 5 years, meaning the costs to the firm would become quite large as a result. Further, the firm was also required to pay an express fine, of up to 200% of the original amount due, to the social security system for the delayed payment of contributions (Ichino and Riphahn, 2005). Schivardi and Torrini (2008) argue that uncertainty about the length of potential trials is a critical variable influencing firing costs.

Further, these regulations were enforced in practice—following the introduction of Article 18 in 1970, a special new court system was was created specifically to deal with claims regarding employment contracts.[31]

It was true that firms could create temporary contracts, which could be ended at no cost. Further, permanent contracts had probationary periods that were limited to 3 months by law. Firms were also required to give notice of 1 month to a worker before they terminated the contract.[32]

[29]The requirement to reinstate workers is unusual, but exists also in the employment law of Austria, the Czech Republic, South Korea, and Portugal. In the Italian case around 60% of the labour disputes went to court, with an average decision time of 23 months. For those cases that did reach court, around 45% are rejected during the first level of judgement and around 63% at the second level. See http://eid.sagepub.com/content/early/2016/03/11/0143831X16635830.full.

Technically, the regulation applies to *establishments* with more than 15 employees. The 15 employees refers to the date in which the firing was intimated, which may be ahead of the dismissal date (Garibaldi, Pacelli, and Borgarello, 2004).

[30]See Schivardi and Torrini (2008).

[31]See Ichino and Riphahn (2005). It is worth noting that, while the threshold principally concerns concerns the costs associated with unfair dismissal of workers, it is true that it also relevant for another piece of legislation affecting firms. Workers with more than 15 employees gained the right to elect trade union representatives, called *Rapresentanze Sindacali Aziendali.* These representatives are allowed to call general meetings, call referenda, and put up posters relating to the union within the establishments. Schivardi and Torrini (2008) argue that the rule is likely of minor impact, as the absence of such a representative does not prevent collective bargaining or trade union membership. Schivardi and Torrini (2008) cite a survey provided by the Metalworking Firm Organization, which shows that the share of firms with a firm-level wage contract was not affected by the 15 employee threshold. Firms also have the ability to vote for a worker representative for safety related issues.

[32]See Leonardi and Pica (2013).

In recent years, there have been two key reform packages that have been passed by the Italian Parliament, which changed employment protection legislation in Italy. The first was passed 2012 under Prime Minister Mario Monti, was called the *Riforma Fornero*, and largely affected temporary contracts.[33] More importantly with respect to the motivation of the present paper, the so-called "Jobs Act" of the Renzi administration was passed in December 2014, and removed Article 18, which provided the regulations governing unfair dismissal costs just described.[34] While the 15 employee threshold studied in this paper is no longer in place for all new contracts formed in Italy, the new set of rules did not apply retroactively to existing contracts. The payment is in principle designed to depend on specific criteria, as opposed to judicial discretion. Specifically, the Jobs Act specifies that firms must pay a single indemnity payment in the case of unfair dismissal, which is a function of the tenure of the employee, and can range from 4 to 24 months worth of the salary of the employee.[35] A key feature is that the "stochastic" component of the previous sanctions has been removed by the new regulations, in the sense that firms are no longer obliged to pay workers forgone wages incurred during trials of variable (and often lengthy) terms.

With respect to the present study only the final year of the data sample is affected by new regulation, and only the Fornero reform, not the Jobs Act. The 15 employee threshold remains in place in the Italian economy of today by virtue of the fact that it applies only to contracts formed post 2015, although the "historical" component to the regulatory environment will of course become of diminishing importance in the future.

The regulatory threshold in Italian labour law has been a subject of many papers; Cingano et al. (2015) use firm-level data to examine the effect of employment protection legislation on capital deepening and productivity. They exploit a change in the law in 1990 that made it more difficult for the firms with fewer than 15 employees to fire workers, relative to those above the threshold. They thus employ a combined regression discontinuity design with difference-in-difference methodology, looking at the effect of the law change on firms above and below the threshold.

[33]Law no. 92 of June 28, 2012, in force since July 18, 2012.

[34]For a full explanation of the Jobs Act, see http://www.portolano.it/2015/03/significant-changes-to-improve-the-italian-labour-market/.

[35]These sanctions also apply in the cases that employees are dismissed within collective dismissal procedures.

Figure 1: Graphical Examination of the Threshold

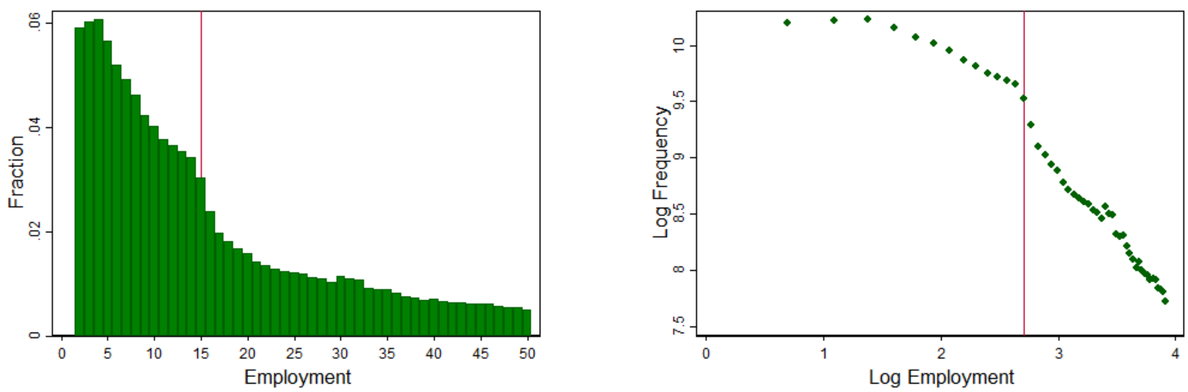

Notes: The left-hand panel shows the histogram of employment. The right-hand panel displays the plot of log frequency against log employment. Figures are displayed for firms with between 1 and 50 employees (firms with only 1 employee have been cleaned from the data). The threshold value of 15 employees is indicated by the vertical red line.

6.2 Graphical Evidence

Figure 1 presents histograms by number of employees for the manufacturing sector separately, pooling observations across years. The threshold value of 15 employees is indicated by a vertical red line. We see initial evidence for bunching of manufacturing firms prior to the threshold in the manufacturing sample.[36]

In order to assess whether the threshold affects the distribution of firms econometrically, I follow the specification of Leonardi and Pica (2013), and estimate the following linear probability model:

$$d_{it} = \delta_1 D13_{it} + \delta_2 D14_{it} + \delta_3 D15_{it} + \boldsymbol{\beta}'\mathbf{X}_{it} + \eta_i + \epsilon_{it}.$$

Here $d_{it} = 1$ if the firm i in year t is has more employees in period $t+1$. The vector X_{it} includes year and industry dummies, and a third-order polynomial in firm employment. Firm fixed effects are also included. The model is estimated only over those firms with between 5 and 25 employees.[37] Results are displayed in Table 9. We can see that the dummy

[36]We do not see similar bunching in the services sample.

[37]I follow the estimation of Leonardi and Pica (2013) broadly. These authors include terms relating to the year 1990, in which the laws were reformed. I omit these since my data are available only from 1995. I do not have a reliable measure of firm age in the ORBIS dataset, so age controls are omitted. I also alter the sequencing, with the dependent variable a growth dummy from t to $t+1$ as opposed to $t-1$ to t, and the employment control being contemporaneous employment as opposed to its lag.

Figure 2: Marginal Revenue Products in Deviation from Industry-Year Means by Employment

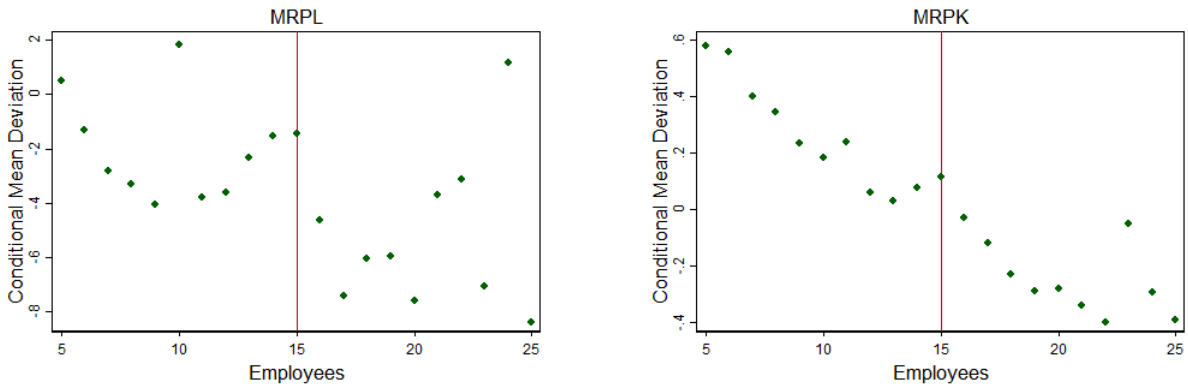

Notes: Displays the marginal revenue products of capital and labour, demeaned by their respective industry and year means, and then subsequently averaged conditional on the number of employees – values are displayed for 5 to 25 employees. The year 2002 is omitted on account of an outlier.

for 15 employees has a negative, and significant effect on the probability of firm growth. This supports the graphical evidence that firms are indeed intentionally holding back their growth so as not to cross the size threshold.

I also examine graphically the distribution of the marginal revenue products of labour and capital around the threshold. To do so, I plot the following function of firm size by labour:

$$f(L) = \frac{1}{N_{\{L_{i,s,t}=L\}}} \sum_{s=1}^{s=S} \sum_{t=1}^{t=T} \sum_{i \in F_s} \mathbf{1}_{\{L_{i,s,t}=L\}} * \left(MRPX_{i,s,t} - \overline{MRPX}_{i,s,t} \right),$$

for $MRPX \in \{MRPL, MRPK\}$, which essentially computes the average of marginal revenue product of each factor from its industry-year mean, conditional on a given value of labour. This function is plotted between 5 and 25 employees in Figure 2. We can see some evidence that the threshold introduced a distortion, in the sense that firms with relatively high marginal revenue products of labour have chosen to remain below the regulatory threshold. We see a smaller, but similar, relationship for the marginal revenue product of capital.

I further investigate whether there could be a discrete change in the *variance* of the factor productivites around the threshold. Under the logic of the paper of Asker et al. (2014), in the presence of higher adjustment costs, we ought to see greater variance of factor revenue

Figure 3: Standard Deviation of (Log) Marginal Revenue Products Conditional on Employment (Unweighted)

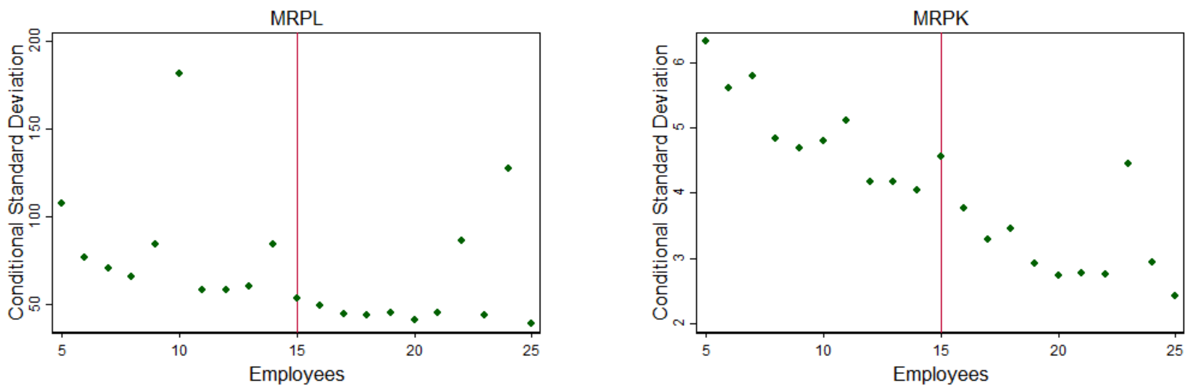

Notes: Figure displays the standard deviation of the marginal revenue products to factors, conditional on the employment level, industry and year. These conditional standard deviations are then averaged over industries and years conditional on the employment level, without applying weights.

product dispersion. One might expect a discrete increase in the variance of factor productivities, conditional on each possible value of employment for the firm. To this end, Figure 3 displays the following function with respect to employment:

$$f(L) = \frac{1}{S * T} \sum_{s=1}^{s=S} \sum_{t=1}^{t=T} \mathrm{Sd}\left(MRPX_{i,s,t} \,|\, L_{i,s,t} = E,\ s_{i,s,t} = s,\ t_{i,s,t} = t \right),$$

for $MRPX \in \{MRPL, MRPK\}$, which simply computes the variance of factor productivities for each employment, industry and year group, before averaging conditional on employment level for a function in L. From visual inspection there is little evidence for any change in the immediate vicinity of the threshold, as can be seen in Figure 3. This suggests that the detrimental effects of the threshold, at least locally, could be more due to firms remaining excessively small to avoid the legislation, than the effect of the legislation on the ability of larger firms to respond to shocks.

6.3 Version of the Model with the Regulatory Threshold Embedded in the Cost Function

The model is identical to that outlined in Section 5.2, other than the cost function (Equation 8), which embeds the discontinuity in the adjustment costs parameters that is associated with Italian employment law. To capture the fact that firms face different adjustment cost parameters for labour as they expand above or below 15 employees, the cost function is specified as follows:

$$\begin{aligned} C(\Omega_{it}, L_{it}, H_{it}, E_{it}) =& 12w\left(C^{P1}E^{+}_{it} + \mathbf{1}_{\{L\leq 15\}}C^{P2}_{S}E^{-}_{it} + \mathbf{1}_{\{L>15\}}C^{P2}_{B}E^{-}_{it}\right) \\ &+ C^{F}_{L}\mathbf{1}_{\{E_{it}\neq 0\}}\Pi(\Omega_{it}, L_{it}, H_{it}) \\ &+ C^{Q}_{L}L_{it}\left(\frac{E_{it}}{L_{it}}\right)^{2} \end{aligned} \tag{14}$$

Equation 14 encapsulates the essence of the relation between labour regulations and firm size in Italy during the sample period. The per-employee cost of firing workers varies according to whether the firm has workers above or below the 15 employee threshold (summarized by the parameters C^{P2}_{S} and C^{P2}_{L}, respectively). Therefore, we can interpret the parameters as follows

$$\begin{aligned} C^{P2}_{S} &= \text{Prob(unfair)}\left[2.5 - 6\,\text{months salary}\right], \\ C^{P2}_{B} &= \text{Prob(unfair)}\left[\text{wages forgone} + \text{social insurance} + \text{fine} + \text{Prob(not reinstate)}[15\,\text{months salary}]\right]. \end{aligned}$$

In fact, I set $\text{Prob(not reinstate)} = 1$, given the low incidence of reinstatement in the data.

6.4 Estimation

In order to identify the effects of the change in the labour law, I match the extended model to the proportion of firms above and below the discontinuity. These values can be seen in Table 11. The same moments from the previous estimation scheme are selected, however this time respectively conditional on the subgroup of firms below the threshold, and the subgroup of firms above. To ensure that the threshold value of labour 15 is meaningful within the model simulation, I also use the model to target a measure of the size distribution of firms. I simply target the fraction of firms below 15 employees, $\Pr[L < 15]$. Further, since the parameter w_1 is an important determinant of the average size of the firms, this parameter is also estimated, in contrast to the estimation strategies of Bloom (2009) and Cooper et al. (2015). Otherwise the approach to estimation remains identical.

I therefore choose the following target moments:

1. $\Pr[\Delta \ln L < 0.05 \mid L \leq 15]$ and $\Pr[\Delta \ln L < 0.05 \mid L > 15]$;
2. $\Pr[\Delta \ln L > 0.2 \mid L \leq 15]$ and $\Pr[\Delta \ln L > 0.2 \mid L > 15]$;

Table 6: Estimates from the Case with the Threshold

	Estimates			
Parameter	(1)	(2)	(3)	(4)
C_L^F	0.0091	0.0074	0.0127	0.0539
C_L^Q	5.4358	3.7925	1.2980	0.2587
C^{P1}	0	0	0.0175	0.0434
C_S^{P2}	0.0077	0.0064	0.0237	0.0183
C_L^{P2}	0.0382	0.0364	0.1194	0.1381
γ	3.8518	2.5485	1.8614	2.7094
w_1	1.3851	1.2804	1.2143	0.3038

Notes: Table shows estimates obtained from the threshold model. In addition to the basic target moments, Column (1) and (3) target $\Pr[L \leq 15]$. Column (2) and (4) target $\Pr[L \leq 15]$, $\Pr[15 < L \leq 30]$, $\Pr[30 < L \leq 50]$, $\Pr[50 < L \leq 100]$, $\Pr[100 < L \leq 150]$. Column (3) and (4) introduce an exogenous quit rate.

3. $\mathrm{sd}(\ln L \mid L \leq 15)$ and $\mathrm{sd}(\ln L \mid L > 15)$;
4. $\mathbb{E}[|\Delta \ln L| \mid |L| > 0.05 \,\&\, L \leq 15]$ and $\mathbb{E}[|\Delta \ln L| \mid |L| > 0.05 \,\&\, L > 15]$;
5. $\Pr[L < 15]$.

6.5 Evaluation

The parameter estimates attained can be seen in Table 6. Specification (3) is chosen for the reasons of model fit. We can see that in all the specifications considered, the estimated model indeed delivers a value of the parameter on firing for firms above 15 employees that exceeds the value for firms below 15 employees. This means that the model matches our prior from knowledge of the institutional background to the Italian case.

The predictions of the model regarding value added and the dispersion of the marginal revenue products of labour are given by Figure 4. The model appears to successfully replicate the increase in the marginal revenue product of labour below the threshold. However, we also see an increase in the marginal revenue product of labour above the threshold, that is not seen in the data. This occurs because there are two effects of the threshold on the firm hiring rule: on the one hand firms are reluctant to hire 16 employees, and become subject to the more expensive regime. On the other hand certain firms who desire to be below the

Figure 4: Threshold Model Predictions Regarding Misallocation

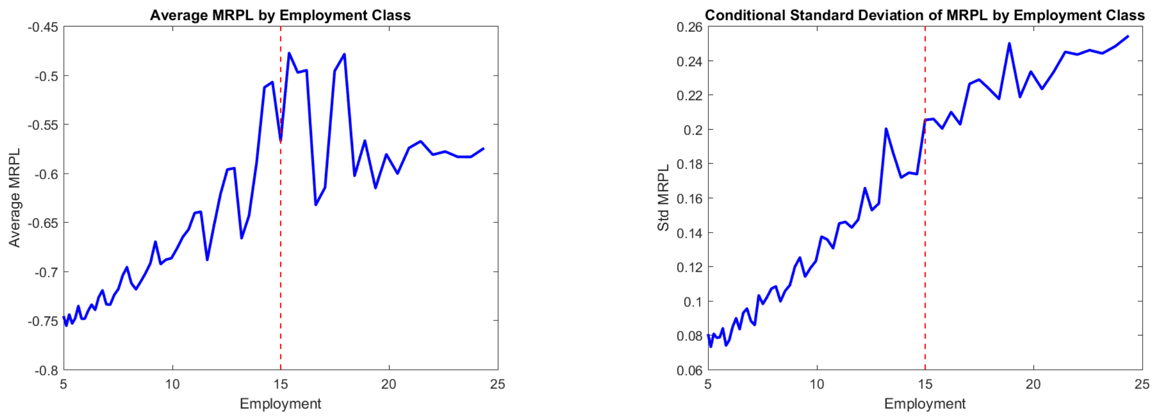

Notes: The left-hand panel shows the average deviation of MRPL (from its overall average), conditional on each labour value in the state-space, from a simulated dataset from the version of the model including a threshold in the costs of dismissal. The right-hand panel shows the standard deviation of MRPL, conditional on each labour value in the state-space. Both charts display output only for simulated firms between 5 and 25 employees.

threshold optimally choose to shed labour in two stages: rather than fire several employees and face the higher firing costs for all terminations, they terminate employees until they are below the threshold, then terminate the remainder to get to their desired size. This implies that firms who wish to shrink below the threshold spend less time above the threshold than they otherwise would, boosting the average marginal revenue product for firms in this range. We can also see some evidence for a level shift in the conditional standard deviation of MRPL after the threshold, induced by the greater firing costs.

Under this version of the model the dispersion of the marginal revenue product of labour is over-predicted, as can be seen from Table 10. The model also fails to match the relation between the dispersion of marginal revenue product and firm size, as we might expect given that the only parameter difference between the two regimes in the model will automatically lead (given estimates attained) to a greater dispersion in the marginal product of labour for large firms. In the data the dispersion of the marginal revenue product of labour for small firms is larger than for large firms. However, as we have seen previously, the overall correlation of the marginal revenue product deviations with firm size is low.

Table 12 shows the effects of the reduction in firing costs on model moments. Given the reduction in firing costs, a lower proportion of firms record hiring rates less than 5%. There is in fact an increase in firms recording values of hiring rates in excess of 20%, as firms are happier to grow in response to a positive shock to revenues, in the knowledge that it is now easier to fire workers. There is an increase in the standard deviation of hiring rates, from 0.27 to 0.39, as firms are more willing to grow and shrink in response to revenue shocks. With respect to the dependent variable of interest, the standard deviation of MRPL, the reform has the effect of reducing the recorded value from 0.90, to 0.78. Although the effect is proportionately greater for those firms who are above the regulatory threshold (falling from 0.91 to 0.71), it is of interest that the standard deviation of MRPL falls also for the firms below 15 employees (from 0.80 to 0.74). This would reflect the increased willingness of firms to grow close to, or even cross, the regulatory threshold (as can be seen from the transition probabilities recorded at the bottom of Table 12).

Whatismore, the reform leads to a 14.44% rise in the labour productivity, measured as the quantity of the aggregate good produced under the model simulation, divided by total labour. This effect is quite large, and there are two caveats of note: (1) the analysis is in partial equilibrium; (2) the sample studied does deviate from the data with respect to an over-representation of large firms, as previously discussed, so admittedly the true predicted effect could be smaller for this reason. Still, the estimated model allows for potentially large gains from reforms similar to those enacted as part of the Jobs Act in 2014.

7 Time Trends

This section discusses observable trends in the data over the time period of study, given the large interest in the nature of trends in misallocation in Southern Europe in recent years, and the fact that the dynamics of misallocation in ORBIS data for services in Italy has not yet been investigated.

In this section I perform the evaluation of potential gains from reallocation for the manufacturing and services sectors. The results are displayed in Table 16. The table shows the gains from reallocation of factors of production for the 18 years in the sample for manufacturing (1995-2013). For services results are displayed from 2005 to 2013.[38] I perform the analysis under 10 specifications in order to assess robustness of results, which are displayed across the columns of the table. The baseline results are displayed in Column (1), under which we set $\sigma = 3$ and use 4-digit NACE industries, for comparability with Hsieh and Klenow (2009). We observe that the gains are 81.34% for manufacturing, and 193.75% for services, for the year 2013. It is thus apparent that the dispersion of marginal revenue

[38]This is on account of conformability issues between industries in the Services sample during the change from NACE 1 to NACE r2 industry classification schemes.

products in the services industry is far larger than that for manufacturing, and thus the gains from reallocation are larger. This is in line with the findings of Dias et al. (2015), the first paper to analyse misallocation in the services industry.

In terms of the magnitude of the gains, relative to other studies, the comparable figure for manufacturing from Hsieh and Klenow is 86% for China in 2005, and 127.5% for India in 1994, and 42.9% for the United States in 1997.[39] Thus the Italian manufacturing gains (81.34% and 193.74%) are somewhat larger than recorded figures for the U.S. In fact, the measure for Italy approximates the size of the measure for Chinese manufacturing. The comparable figures for Portuguese data are reported in Dias et al. (2015), under a two-factor model, for Portuguese data as 37.5% for manufacturing in 2011, which in fact indicates Portuguese misallocation to be *smaller* than that for the U.S. reported in HK. Therefore for manufacturing, at least, we can conclude that the levels of misallocation for Italy is quite high relative to that of the U.S. As mentioned previously, the paper of Gopinath et al. (2017) reports potential reallocation gains of 58% in 1993, 67% in 2006, and 80% in 2011 as baseline estimates using CERVED data on incorporated manufacturing companies.[40]

The magnitude of the gains from reallocation in the services industry seem very large for 2013, with a figure of 193.75% recorded. There are few comparable estimates in the literature, however the analysis of Portuguese data in Dias et al. (2015) attains a value of 56.72%, when they employ a two-factor model. Under their three-factor model the authors attain an estimate of 91.51% in value-added gains in 2011.[41] The larger figure in my data relative to these estimates can partly be accounted for in part by the exclusion of entering firms with less than 20 employees in this study, wheres I do not have a similar restriction for ORBIS data. I will discuss the implications of restricting small firms for robustness. Admittedly, concerns about the measurement error regarding the capital stock in the services industry, where the meaning of capital is somewhat less clear, may pay a role in increasing recorded dispersion in marginal products, biasing estimates for services upwards.

We also note increases in the gains from reallocation, equivalent to increases in mis-

[39]The headline figures from this paper of 30-50% in China and 40-60% in India are calculated relative to 1997 U.S. gains.

[40]In fact, the baseline specification of Gopinath et al. (2017) differs with respect to its use of Italiab labour shares to attain measures of production parameters, as opposed to U.S. shares in this paper. I also choose to use employees as the measure of the labour input, as opposed to labour costs.

[41]Use of a three-factor model would allow for larger estimates of reallocation gains, insofar as materials inputs are subject to distortions that affect their costliness relative to other factors. The two-factor model includes materials expenditure as a component of value added, but only permits distortions to materials within the output distortion, τ_Y. This could be particularly important if we believe energy markets to be distorted, although a three-factor model is not pursued in this paper.

allocation, in both the manufacturing and services industries at the baseline. These gains are presented graphically in Figure 12. In this sense I replicate the findings of increasing misallocation reported in Gopinath et al. (2017) for manufacturing, and show that it is also true for the services industry, admittedly in a smaller time-series of data. Reallocation gains for 2005 are at 48.02% for manufacturing in 2005, and rise to 81.34% by 2013. Reallocation gains for services are at 136.45% for 2005, and rise to 193.75% for 2013. However we also note the size of the gains for services in fact the increases in reallocation gains are much smaller after 2010.

Turning to robustness checks, it can be generally be summarized that the magnitude of the gains from manufacturing are sensitive to the specification chosen, with the various checks employed tending to inflate estimates. This is in line with reports from other robustness excercises undertaken in the literature. The increase in misallocation for manufacturing over the sample period is robust to all specifications, although is smaller after the imposition of a balanced panel. In this sense the firms that are "survivors" in the Italian manufacturing industry seem to be more immune to the forces driving increases in misallocation in sample. The increase in misallocation for services is fairly robust, although not for the case of a wide industry definition (2-digit NACE), or for the case of a balanced panel.

Column (2) of Table 16 includes firms with only one employee, which were omitted at the baseline. We can see that TFP gains are much the same, or perhaps slightly larger.

In Column (3) I restrict the sample to firms with more than 20 employees. We see that the manufacturing and services estimates fall in size markedly. This provides some further evidence that very small firms contribute more to misallocation than average firms in both sectors. Another reason I restrict the analysis to avoid the smallest firms in Column (3) is that it renders my dataset somewhat more comparable to the baseline of Dias et al. (2015) for Portugal, and indeed my estimates for 2011, of 63.22% for manufacturing, and 139.14% for services, are lower, although are still higher than comparable estimates reported in this paper (37.5% for manufacturing and 56.7% for services).[42]

Column (4) sets the measure of employment to wages, as opposed to number of employees. Manufacturing results fall to 81.59% for 2011, and services results fall to 174.48%. It is possible that the relatively larger falls for services can be explained by a strengthening of the relationship between the number of employees in a firm, and the quality of the total labor input, in the services sector. However, it could be the case that the assumption of equal wages for firms within sectors is less appropriate for the services sector.

[42]Strictly, I cut all firm-year observations with less than 20 employees, whereas they do not include entrants with less than 20 employees, but allow incumbents to fall below 20 employees.

Column (5) reports estimates for a pooled trim of the distributions of TFPR and TFPQ across years, to bring my the data cleaning of the study strictly into line with those of HK. Estimates remain close.

Column (6) displays results for a balanced panel, which delivers a markedly smaller sample, and is likely to affect its size distribution also, as older firms tend to be larger. The increases in misallocation for the manufacturing sector are reduced. We also see a large reversal in the increases in misallocation observed for the services sector in the final years, however prior to these movements the upward trend in misallocation of the Services industry still exists. It is possible that these results are correcting for increases in recorded misallocation that are due to expansions in the coverage of the ORBIS dataset. It could also be that, given many missing values, particularly amongst smaller firms, the differences in time trends for the balanced panel are due to reduced representativeness. Further, restricting to a balanced panel also omits legitimate entrants, who may face larger wedges than incumbents.

Columns (7) and (8) show that, predictably, increasing the sizes of industries to NACE 3-digit and 2-digit sectors deliver progressively larger estimates, since increasingly heterogeneous firms are treated as homogeneous by the model, with increased reallocation gains identified in consequence, although the effects are rather muted for manufacturing.

Column (9) shows the effects of increasing the calibrated value of σ to 4, as opposed to 3. This delivers increases in the size of estimates, as was also reported in HK. Generally speaking, the conservative choice of σ is more appropriate, given the sensitivity of results to this parameter, however σ was set to 4 in the analysis of adjustment costs, following Asker et al. (2014) and Bloom (2009).

Column (10) examines the extent to which the assumption of time-varying industry shares, θ_s, may be affecting the variation over time in misallocation measures. Under this specification the industry shares are held constant at their 2005 values, which implies that any observed increases or decreases in misallocation can be ascribed solely to changes in the wedges (as opposed to sectoral composition). We see much larger increases in misallocation over time in manufacturing, while the increases observed in the services sector are comparable to the baseline. This indicates that variation in the sizes of the industries may be alleviating the aggregate consequences of changes in misallocation in the manufacturing industry.

We therefore reach the general conclusion that misallocation has been increasing in ORBIS data, although the time-trend for manufacturing is weaker under a balanced panel.

In the case of the services industry the misallocation gains are reversed in the final few years of the sample under the balanced panel specification.

8 Conclusion

This paper has studied the ability of a structural model of firm dynamics to match the evidence on within-industry misallocation for the case of Italy. The paper therefore extends the analysis of Asker et al. (2014) to the domain of the misallocation of labour, which has been shown to be a non-negligible determinant of resource misallocation for the Italian context. The principal findings are that, while the estimated model performs extremely well at matching the variance of the marginal revenue product of labour, it generates a positive relationship between firm-size and the deviation of marginal revenue products of labour. It can be seen in the data the correlation between MRPL deviation and employment is negative but small, while the correlation with value-added is positive but lower than the model prediction. The diagnosis of the determinant of the relation between measures of misallocation and firm-size is left for future work.

However, given the proven ability of the model to summarize the misallocation of labour, it is then adapted for an analysis of a regulatory threshold. This analysis extends the class of models used in the literature devoted to the estimation of adjustment-costs (that stemming from the contribution of Cooper and Haltiwanger, 2006) to a setting where there is *a priori* knowledge of a regulatory threshold at a given level of employees. This is of interest because it ought to discipline the estimation of the parameters of the model by using features of the data that can only result from a discontinuous increase in the firing costs of labour. Since the usefulness of the SMM estimation approach relies on the ability of the model to match statistics from the data that are known by the researcher to be informative for parameters relating to adjustment costs, the Italian case ought to allow for greater precision in estimation. Given these estimates, one is able to conduct policy-simulations that map closely to the existing regulatory framework, and can inform actual policy-decisions (for example those taken by the Renzi administration at the end of 2014, whose effects are still being played out today).

In future work it would be of interest to extend the policy reform analysis to a general equilibrium model, so as to analyse the effect on wages of such a move. Such an analaysis would be able to fruitfully blend aspects of the micro-econometrics literature surrounding regulatory thresholds with the ability to attain precise local estimates of policy reform, with a structural model capable of making general claims.

References

Asker, J., A. Collard-Wexler, and J. D. Loecker (2014). Dynamic Inputs and Resource (Mis) Allocation. *Journal of Political Economy 122*(5).

Bayer, C., A. M. Mecikovsky, and M. Meier (2015). Productivity Dispersion: Could it simply be technology choice?

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*(3), 623–685.

Blundell, R. and S. Bond (2000). GMM Estimation with persistent panel data: an application to production functions. *Econometric Reviews 3*, 321–340.

Calligaris, S. (2015). Misallocation and Total Factor Productivity in Italy: Evidence from Firm Level Data. *LABOUR 29*(4), 367–393.

Calligaris, S., M. D. Gatto, F. Hassan, G. I. Ottaviano, and F. Schivardi (2016). Italy's Productivity Conundrum: A Study on Resource Misallocation in Italy. Discussion Paper 030, European Commission.

Catherine, S., T. Chaney, Z. Huang, D. Sraer, and D. Thesmar (2017). Aggregate Effects of Collateral Constraints.

Cingano, F., M. Leonardi, J. Messina, and G. Pica (2015). Employment Protection Legislation, Capital Investment and Access to Credit: Evidence from Italy. *Economic Journal 126*(595), 1798–1822.

Cooper, R. and J. Haltiwanger (2006). On the Nature of Capital Adjustment Costs. *Review of Economic Studies 73*, 611–633.

Cooper, R., J. Haltiwanger, and J. L. Willis (2015). Dynamics of labor demand: Evidence from plant-level observations and aggregate implications. *Research in Economics 69*, 37–50.

Dias, D. A., C. R. Marques, and C. Richmond (2015). Misallocation and Productivity in the Lead Up to the Eurozone Crisis. International Finance Discussion Papers 1146, Board of Governors of the Federal Reserve System.

García-Santana, M., E. Moral-Benito, J. Pijoan-Mas, and R. Ramos (2016). Growing like Spain: 1995-2007. Working Paper Series 888, Barcelona GSE.

Garicano, L., C. Lelarge, and J. V. Reenen (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review 106*(11), 3439–79.

Gopinath, G., S. Kalemli-Ozcan, L. Karabarbounis, and C. Villegas-Sanchez (2017). Capital allocation and productivity in south europe. *Quarterly Journal of Economics 132*(4), 1915–1967.

Hawkins, W., R. Michaels, and J. Oh (2015). The Joint Dynamics of Capital and Employment at the Plant Level.

Hijzen, A., L. Mondauto, and S. Scarpetta (2013). The Perverse Effects of Job-Security Provisions on Job Security in Italy: Results from a Regression Discontinuity Design. Discussion Paper Series 7594, IZA.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Ichino, A. and R. T. Riphahn (2005). The Effect of Employment Protection on Worker Effort: Absenteeism during and after Probation. *Journal of the European Economic Association 3*(1), 120–143.

Kalemli-Ozcan, S., B. Sorensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas (2015). How to Construct Nationally Representative Firm Level data from the ORBIS Global Database. Working Paper 21558, NBER.

Lanau, S. and P. Topalova (2016). The Impact of Product Market Reforms on Firm Productivity in Italy. Technical Report 119, IMF.

Leonardi, M. and G. Pica (2013). Who pays for it? The heterogeneous wage effects of employment protection legislation. *The Economic Journal 123*(573), 1236–1278.

Levinsohn, J. and A. Petrin (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *Review of Economic Studies 70*(2), 317–341.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica 71*(6), 1695–1725.

Oberfield, E. (2013). Productivity and misallocation during a crisis: Evidence from the Chilean crisis of 1982. *Review of Economic Dynamics 16*(1), 100–119.

Olley, S. G. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica 64*(6), 1263–1297.

Pellegrino, B. and L. Zingales (2014). Diagnosing the Italian Disease.

Schivardi, F. and R. Torrini (2008). Identifying the effects of firing restrictions through size-contingent differences in regulation. *Labour Economics 15*(3), 482–511.

Socio, A. D. and V. Michelangeli (2015). Modelling Italian firms' financial vulnerability. Occasional Paper 293, Banca D'Italia.

van Ark, B. and E. Monnikhof (1996). Size Distribution of Output and Employment: A Data Set for Manufacturing Industries in Five OECD Countries. OECD Economics Department Working Papers 166, OECD.

Ziebarth, N. L. (2013). Are China and India backward? Evidence from the 19th century U.S. Census of Manufactures. *Review of Economic Dynamics 16*(1), 86–89.

A Data Cleaning

- I drop observations that are missing the unique Bureau van Dijk ID code.
- I drop consolidated firms.
- I drop if missing industry information.
- I restrict industries to: 1) manufacturing (NACE C); 2) services (transportation and storage [NACE H], accommodation and food service activities [NACE I], information and communication [NACE J], real estate activities [NACE L], professional, scientific and technical activities [NACE M], administrative and support service activities [NACE M], repair of computers, personal and household goods [NACE S95]).
- For years prior to 2005 NACE Rev 1.1 is converted to NACE Rev 2 using the correspondence table available here: `http://ec.europa.eu/eurostat/web/nace-rev2/correspondence_tables`. I follow Kalemli-Ozcan et al. (2015)

insofar as I match 4 digit NACE Rev 1.1. codes to the most closely related NACE Rev 2 code by hand, in the case that the mapping is not unique.

- Nominal variables are deflated with producer price indices. Fixed assets are deflated with investment goods deflators.
- Value added is calculated as operating revenue subtract materials expenditure.
- Observations are dropped if the following variables are either missing, zero or negative: operating revenue, wages, fixed assets and materials. This reduces the size of the manufacturing sample from 3,008,425 observations to 897,465 observations, and the size of the services sample from 755,851 observations to 901,732 observations.
- Observations are dropped if the following variables are (strictly) negative: intangible fixed assets, loans, long term debt, depreciation, shareholder funds. This reduces the size of the manufacturing sample to 863,219 observations and the services sample to 820,840 observations.
- Further cleaning operations are conducted based on extreme values in certain ratios. I compute the ratio of the wage bill to tangible fixed assets and drop the full firm if in any year there are extreme values (higher than 1000). I then drop extreme observations at the 0.1% and 99.9% level. I drop firm-year observations when the ratio of tangible fixed assets to total assets is greater than 1. This reduces the size of the manufacturing sample to 840,405 observations and the services sample to 769,911 observations.
- I drop firms with (strictly) negative value added. This reduces the size of the manufacturing sample to 839,215 observations and the size of the services sample to 768,854 observations.
- I construct the ratio of the wage bill to value added and drop extreme observations lower/higher than the 0.1% and 99.99% levels. I drop firm-year observations if this ratio is greater than 1.1. This reduces the size of the manufacturing sample o 835,080 observations and the size of the services sample to 761,812 observations.
- I compute the ratio of tangible fixed assets to shareholder funds and drop values outside the 0.1 and 99.99% percentiles. This reduces the size of the manufacturing sample to 833,217 observations and the size of the services sample to 757,331 observations.

- I compute the ratio of total assets to shareholder funds and drop extreme values outside the 0.1 and 99.99% percentiles. This reduces the size of the manufacturing sample to 832,948 observations and the size of the services sample to 755,851 observations.

- I then winsorize (real) operating revenue, added value, wages, fixed assets and materials at the 1 and 99 percentiles.

- I further impose a minimum on the number of observations an industry must attain in all years in order to be included in the analysis. This is chosen to be 20 observations. This reduces the sample for manufacturing to 809,958 , and the sample for services to 574,599. In fact this omits NACE division S95 from the analysis.

Deflators. Disaggregated manufacturing deflators are available from Eurostat data on producer prices by industry, available here http://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=sts_inpp_a&lang=en. These deflators are at various levels of aggregation, and are converted into Nace Rev.2 codes. I attain complete coverage of manufacturing at the 2-digit NACE R2 level, a couple of exceptions that are assigned by hand. The aggregate value added deflators for services are also taken from the AMECO dataset. This is the measure for price deflator for gross value added by main branch. I use the measure for the manufacturing sector.

The investment goods deflator is downloaded from the AMECO dataset of the European Commission. It corresponds to the measure for capital formation and saving, total economy and sectors. The deflator is re-based to equal 1 in 2010. The deflator is available here: http://ec.europa.eu/economy_finance/db_indicators/ameco/zipped_en.htm.

The CPI measure is taken from the annual series of the harmonized index of consumer prices available from Eurostat. The series was rebased from 2005 to 2010.

B Derivations

B.1 Optimizing Over Flexible Inputs

Firms are facing a downward sloping demand function $Y = P^{-\epsilon}$, with a Cobb-Douglas production function in two factors: $Y = \tilde{A}L^{\alpha_L}K^{\alpha_K}$. Letting $\phi \equiv (\epsilon - 1)/\epsilon$. This means that:

$$\tilde{R}(A, L, K) = PY = Y^{\phi} = \tilde{A}^{\phi}L^{\alpha_L\phi}K^{\alpha_K\phi} - rK.$$

Then we can maximize this function with respect to K.

$$K^* = \phi\left(\frac{\alpha_K}{r}\right)Y^{\phi}$$

Which implies:

$$K^* = \left[\left(\frac{\alpha_K}{r}\right)\phi\tilde{A}^{\phi}L^{\alpha_L\phi}\right]^{\frac{1}{1-\alpha_K\phi}}$$

Substituting back into the revenue function:

$$\left(\frac{\alpha_K\phi}{r}\right)^{\frac{\alpha_K\phi}{1-\alpha_K\phi}}\tilde{A}^{\phi\left(\frac{1}{1-\alpha_K\phi}\right)}L^{\alpha_L\phi\left(\frac{1}{1-\alpha_K\phi}\right)} - r\left[\phi\left(\frac{\alpha_K}{r}\right)\tilde{A}^{\phi}L^{\alpha_L\phi}\right]^{\frac{1}{1-\alpha_K\phi}}$$

Collecting terms:

$$\tilde{A}^{\frac{\phi}{1-\alpha_K\phi}}L^{\frac{\alpha_L\phi}{1-\alpha_K\phi}}\left(\frac{\alpha_K\phi}{r}\right)^{\frac{1}{1-\alpha_K\phi}}\left[\left(\frac{\phi\alpha_K}{r}\right)^{\alpha_K\phi} - r\right]$$

Which can be re-defined:

$$AL^{\beta}$$

Where $A \equiv \tilde{A}^{\frac{\phi}{1-\alpha\phi}}$, and term $\beta \equiv \alpha_L\phi/(1-\alpha_K\phi)$. Note that during the estimation of the full production function by system GMM, the estimates of ρ and $\sigma$are relative to the process $\tilde{A}^{\phi}$, and are both rescaled by the term $\frac{1}{1-\alpha_K\phi}$, to reflect the translation of the revenue product shock estimated controlling for capital to the revenue productivity amalgam of productivity and the associated (flexible) capital response. The persistence and variance of A is therefore larger than that of $\tilde{A}$.

B.2 Solution Method

- The model has a state space (A, L) that is discretized.
- In order to discretize the productivity shock A, the Tauchen method is used. The process is discretized across 30 points.
- The labour grid is geometrically spaced, in order to reflect the log-linearity of the revenue productivity process. I choose a labour grid of 600 points.[43]

[43]For reference, Chaney et al. 2015 have a coarse grid of 10,000 points and a fine grid of 811,440 points, but have four endogenous and two exogenous state variables. Bloom (2009) has 40,000 points, but has two endogeous and two exogenous state variables. Cooper et al. (2015) use 230 points for their labour grid, and

- Profits are computed using the production and cost functions, for all possible values of the coarse grid.
- I initialize a value function at the optimum labour choice given that adjustment costs are zero.
- I perform value function iteration to attain the policy functions of the model.
- I follow Bloom (2009) and employ a simulated annealing algorithm in order to estimate the parameters of the model.

C Additional Tables

21 points for their productivity process. Asker et al. (2014) have 567 points on their capital grid.

Table 7: Basic Summary Statistics for Manufacturing Firms

	Overall		Small ($L \leq 15$)		Large ($L > 15$)		$15 < L \leq 30$	$30 < L \leq 50$	$50 < L \leq 100$	$100 < L \leq 150$	$L > 150$
Variable	Mean	Sd	Mean	Sd	Mean	Sd	Mean	Mean	Mean	Mean	Mean
Added Value	5.59	48.46	0.99	21.48	10.74	66.39	2.77	5.11	8.6	15.06	63.75
Employment	41.75	436.86	7.62	3.95	79.92	633.58	21.98	39.23	70.2	121.59	441.87
Wages	1.71	16.21	0.26	3.39	3.34	23.2	0.79	1.52	2.73	4.88	19.99
Capital	4.05	53.61	0.71	39.87	7.79	65.45	1.71	3.13	5.44	10.01	51.95
$\Delta \log L$	0.03	0.43	0.02	0.47	0.04	0.39	0.05	0.05	0.04	0.04	0.03
$\Delta \log K$	0.02	0.32	-0.01	0.35	0.05	0.29	0.05	0.05	0.05	0.05	0.07
$\Delta \log \epsilon$	-0.01	0.36	-0.01	0.39	-0.01	0.33	-0.01	-0.01	-0.01	-0.01	-0.01
sd MRPL	0.77		0.88		0.69		0.65	0.63	0.72	0.73	0.67
sd MRPK	1.37		1.57		1.26		1.3	1.23	1.09	1.36	0.85
$\mathbb{E}[\|e\| \mid \|e\| > 5\%]$	0.27		0.32		0.22		0.22	0.22	0.21	0.22	0.25
$\Pr[\|e\| < 5\%]$	0.32		0.28		0.36		0.3	0.37	0.4	0.43	0.44
$\Pr[\|e\| > 20\%]$	0.43		0.53		0.31		0.39	0.3	0.25	0.25	0.25
N	553,548		292,189		261,359		98,441	66,125	56,434	17,689	22,670

Notes: Presents basic summary statistics for the cleaned data sample of firms in the manufacturing sector. The summary statistics are separated by size groups: (1) all sizes of firm; (2) small firms below 15 employees; (3) large firms above 15 employees, and 5 sub-divisions of this latter group. The values of the aggregate MRPL and MRPK standard deviations are computed at industry-year level, then averaged across years for given industries. Then a weighted average is taken using industry shares by value-added (the shares themselves represent the average share over time for given industries).

Table 8: Model Evaluation for Baseline Case

	Data	Model
$\Pr[\lvert e\rvert < 5\%]$	0.32	0.34
$\Pr[\lvert e\rvert < 1\%]$	0.22	0.30
$\Pr[\lvert e\rvert > 20\%]$	0.43	0.43
$\mathrm{sd}(e)$	0.34	0.50
$\mathrm{skw}(e)$	0.04	0.32
$\mathbb{E}[\lvert e\rvert \mid e > 5\%]$	0.27	0.46
$\mathrm{corr}(e, A)$	0.20	0.79
$\mathrm{corr}(L, L_{-1})$	-0.01	1.00
$\mathrm{corr}(e, e_{-2})$	-0.02	0.99
$\mathrm{corr}(e, e_{-4})$	0.02	0.98
$\mathrm{corr}(e, q_{-2})$	-0.01	0.90
$\mathrm{corr}(e, q_{-4})$	0.03	0.91
$\mathrm{corr}(q, e_{-2})$	0.04	0.83
$\mathrm{corr}(q, e_{-4})$	0.02	0.80
$\mathrm{corr}(q, q_{-2})$	-0.09	0.95
$\mathrm{corr}(q, q_{-4})$	0.08	0.90
$\mathrm{sd}(r)$	25.50	14.95
$\mathrm{skw}(r)$	0.21	3.64
$\mathrm{sd}(e \mid L \leq 15)$	0.46	0.47
$\mathrm{sd}(e \mid L > 15)$	0.28	0.51
$\Pr[\lvert e\rvert < 1\% \mid L \leq 15]$	0.28	0.38
$\Pr[\lvert e\rvert > 20\% \mid L \leq 15]$	0.53	0.51
$\Pr[\lvert e\rvert < 1\% \mid L > 15]$	0.36	0.32
$\Pr[\lvert e\rvert > 20\% \mid L > 15]$	0.31	0.38
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\, L \leq 15]$	0.32	0.45
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\, L > 15]$	0.22	0.47
$\mathrm{sd}(\mathrm{MRPL})$	0.59	0.57
$\mathrm{sd}(\mathrm{MRPL} \mid L \leq 15)$	0.72	0.51
$\mathrm{sd}(\mathrm{MRPL} \mid L > 15)$	0.51	0.49
$\mathrm{corr}(\mathrm{MRPL}, L)$	-0.00	0.85
$\mathrm{corr}(\mathrm{MRPL}, r)$	0.34	1.00
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, L)$	-0.00	0.62
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, r)$	0.09	0.90
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, L)$	-0.01	0.67
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, r)$	0.26	0.91
$\mathrm{corr}(\mathrm{MRPL}, L)$	0.02	0.04
$\mathrm{corr}(\mathrm{MRPL}, r)$	0.02	0.04
$\Pr[L' > 15\,\&\, L \leq 15]$	0.02	0.04
$\Pr[L' > 15\,\&\, L \leq 15 \mid 10 < L \leq 20]$	0.01	0.04

Notes: Data are from the full sample. Lowercase variables, x, represent $100 * \Delta \ln X$. The fractions are computed using pooled data across years. The dispersion and skewness estimates are computed by year, then averaged across the years in the sample.

Table 9: Firm Sorting Regression

	(1) Manufacturing	(2) Services
Dummy 13	0.007	0.007
	0.00	0.01
Dummy 14	-0.005	-0.012
	0.00	0.01
Dummy 15	-0.023***	-0.024**
	0.00	0.01
R^2	0.297	0.275
N	282,923	92,075

Notes: The dependent variable is a growth dummy, with the value 1 if firm i's employment at $t+1$ exceeds t, otherwise 0. Only firms with between 5 and 25 employees are included in the regression. A third order polynomial in employment, sector dummies, and year dummies are also included, as well as firm-fixed effects. Column (1) displays results for Manufacturing. Column (2) displays results for Services. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 10: Model Evaluation for Case with Threshold

	Data	Model
$\Pr[\lvert e\rvert < 5\%]$	0.32	0.68
$\Pr[\lvert e\rvert < 1\%]$	0.22	0.68
$\Pr[\lvert e\rvert > 20\%]$	0.43	0.18
$\mathrm{sd}(e)$	0.34	0.28
$\mathrm{skw}(e)$	0.04	0.59
$\mathbb{E}[\lvert e\rvert \mid e > 5\%]$	0.27	0.38
$\mathrm{corr}(e, A)$	0.20	0.65
$\mathrm{corr}(L, L_{-1})$	-0.01	1.00
$\mathrm{corr}(e, e_{-2})$	-0.02	1.00
$\mathrm{corr}(e, e_{-4})$	0.02	0.99
$\mathrm{corr}(e, q_{-2})$	-0.01	0.72
$\mathrm{corr}(e, q_{-4})$	0.03	0.73
$\mathrm{corr}(q, e_{-2})$	0.04	0.66
$\mathrm{corr}(q, e_{-4})$	0.02	0.63
$\mathrm{corr}(q, q_{-2})$	-0.09	0.94
$\mathrm{corr}(q, q_{-4})$	0.08	0.88
$\mathrm{sd}(r)$	25.50	20.16
$\mathrm{skw}(r)$	0.21	5.56
$\mathrm{sd}(e \mid L \leq 15)$	0.46	0.38
$\mathrm{sd}(e \mid L > 15)$	0.28	0.26
$\Pr[\lvert e\rvert < 1\% \mid L \leq 15]$	0.28	0.50
$\Pr[\lvert e\rvert > 20\% \mid L \leq 15]$	0.53	0.41
$\Pr[\lvert e\rvert < 1\% \mid L > 15]$	0.36	0.71
$\Pr[\lvert e\rvert > 20\% \mid L > 15]$	0.31	0.14
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\ L \leq 15]$	0.32	0.43
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\ L > 15]$	0.22	0.37
$\mathrm{sd}(\mathrm{MRPL})$	0.59	0.88
$\mathrm{sd}(\mathrm{MRPL} \mid L \leq 15)$	0.72	0.81
$\mathrm{sd}(\mathrm{MRPL} \mid L > 15)$	0.51	0.88
$\mathrm{corr}(\mathrm{MRPL}, L)$	-0.00	0.67
$\mathrm{corr}(\mathrm{MRPL}, r)$	0.34	1.00
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, L)$	-0.00	0.44
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, r)$	0.09	0.88
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, L)$	-0.01	0.51
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, r)$	0.26	0.93
$\mathrm{corr}(\mathrm{MRPL}, L)$	0.02	0.01
$\mathrm{corr}(\mathrm{MRPL}, r)$	0.02	0.01
$\Pr[L' > 15 \,\&\, L \leq 15]$	0.02	0.01
$\Pr[L' > 15 \,\&\, L \leq 15 \mid 10 < L \leq 20]$	0.01	0.01

Notes: Data are from the full sample. The fractions are computed using pooled data across years. The dispersion and skewness estimates are computed by year, then averaged across the years in the sample.

Table 11: Size Distribution for Case with Threshold

	Data	Model
$0 < L \leq 15$	0.53	0.14
$15 < L \leq 30$	0.18	0.14
$30 < L \leq 50$	0.12	0.23
$50 < L \leq 100$	0.10	0.29
$100 < L \leq 150$	0.03	0.09
$150 < L$	0.04	0.12

Notes: Data are from the full sample. The fractions are computed using pooled data across years. The dispersion and skewness estimates are computed by year, then averaged across the years in the sample.

Table 12: The Effects of Structural Reform in Parital Equilibrium

	Data	Pre-Reform	Post-Reform
$\Pr[\lvert e\rvert < 5\%]$	0.32	0.69	0.52
$\Pr[\lvert e\rvert < 1\%]$	0.22	0.69	0.51
$\Pr[\lvert e\rvert > 20\%]$	0.43	0.17	0.30
$\mathrm{sd}(e)$	0.34	0.27	0.39
$\mathrm{skw}(e)$	0.04	0.87	0.83
$\mathbb{E}[\lvert e\rvert \mid e > 5\%]$	0.27	0.38	0.44
$\mathrm{corr}(e, A)$	0.20	0.62	0.75
$\mathrm{corr}(L, L_{-1})$	-0.01	1.00	1.00
$\mathrm{corr}(e, e_{-2})$	-0.02	1.00	0.99
$\mathrm{corr}(e, e_{-4})$	0.02	0.99	0.98
$\mathrm{corr}(e, q_{-2})$	-0.01	0.66	0.79
$\mathrm{corr}(e, q_{-4})$	0.03	0.67	0.80
$\mathrm{corr}(q, e_{-2})$	0.04	0.61	0.73
$\mathrm{corr}(q, e_{-4})$	0.02	0.58	0.69
$\mathrm{corr}(q, q_{-2})$	-0.09	0.92	0.92
$\mathrm{corr}(q, q_{-4})$	0.08	0.85	0.85
$\mathrm{sd}(r)$	25.50	15.17	17.26
$\mathrm{skw}(r)$	0.21	3.85	3.90
$\mathrm{sd}(e \mid L \leq 15)$	0.46	0.35	0.37
$\mathrm{sd}(e \mid L > 15)$	0.28	0.25	0.40
$\Pr[\lvert e\rvert < 1\% \mid L \leq 15]$	0.28	0.55	0.55
$\Pr[\lvert e\rvert > 20\% \mid L \leq 15]$	0.53	0.36	0.36
$\Pr[\lvert e\rvert < 1\% \mid L > 15]$	0.36	0.72	0.50
$\Pr[\lvert e\rvert > 20\% \mid L > 15]$	0.31	0.13	0.27
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\ L \leq 15]$	0.32	0.41	0.43
$\mathbb{E}[\lvert e\rvert \mid \lvert e\rvert > 5\%,\ L > 15]$	0.22	0.37	0.45
$\mathrm{sd}(\mathrm{MRPL})$	0.59	0.90	0.78
$\mathrm{sd}(\mathrm{MRPL} \mid L \leq 15)$	0.72	0.80	0.74
$\mathrm{sd}(\mathrm{MRPL} \mid L > 15)$	0.51	0.91	0.71
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, L)$	0.00	0.61	0.74
$\mathrm{corr}(\widehat{\mathrm{MRPL}}^2, r)$	0.34	1.00	1.00
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, L)$	0.00	0.37	0.49
$\mathrm{corr}(\lvert\widehat{\mathrm{MRPL}}\rvert, r)$	0.09	0.89	0.90
$\mathrm{corr}(\mathrm{MRPL}, L)$	-0.01	0.36	0.52
$\mathrm{corr}(\mathrm{MRPL}, r)$	0.26	0.89	0.91
$\Pr[L' > 15\,\&\,L \leq 15]$	0.02	0.01	0.03
$\Pr[L' > 15\,\&\,L \leq 15 \mid 10 < L \leq 20]$	0.02	0.01	0.03
$\Pr[L' \leq 15\,\&\,L > 15]$	0.02	0.01	0.03
$\Pr[L' \leq 15\,\&\,L > 15 \mid 10 < L \leq 20]$	0.01	0.01	0.03

Notes: Table shows the effects on data moments of reducing the adjustment cost parameter faced by large firms to the level faced by small firms.

Table 13: Updated Bloom (2009) Estimate Comparison Table

	Capital			Labour		
Source	PI (%)	Fixed (%)	Quad	PI (%)	Fixed (%)	Quad
Bloom (2009)	33.9	1.5	0	1.8	2.1	0
Ramey and Shapiro (2001)	40-80					
Caballero and Engel (1999)		16.5				
Hayashi (1982)						
Cooper and Haltiwanger (2006)	2.5	20.4	0.049			
Shapiro (1986)			3			1.33
Merz and Yashiv (2007)			14.2			4.2
Hall (2004)			0			0
Nickell (1986)				8-25		
Cooper, Haltiwanger, and Willis (2004)					1.7	0

Table 14: Comparison of Selected Adjustment Cost Target Moments

		Source / Country										
		CHST (2015)	ACL (2014)								B (2009)	CH(2006)
Moment	**IT**	FR	US	CL	IN	FR	ES	MX	RO	SI	US	US
$\Pr[\|I/K\| \leq 5\%]$	**0.17**		0.39	0.19	0.29	0.13	0.20	0.08	0.08	0.15		
$\Pr[\|I/K\| \leq 1\%]$	**0.03**											0.081
$\Pr[\|I/K\| > 20\%]$	**0.52**		0.09	0.11	0.19	0.57	0.41	0.73	0.61	0.52		0.186 *(+ve)* / 0.018 *(-ve)*
$\mathrm{sd}\,(I/K)$	**0.42**		0.21	0.28	0.30	0.57	0.59	0.66	0.72	0.76	0.139	0.337
$\mathrm{skw}\,(I/K)$	**1.42**										1.789	
$\mathrm{sd}\,(E/L)$	**0.34**										0.189	
$\mathrm{skw}\,(E/L)$	**0.04**										0.445	
$\Pr[I/K > 0]$	**0.55**	0.96										
$\Pr[\|\Delta L/L\| \neq 0]$	**0.79**	0.84										
$\mathbb{E}[I/K \mid I/K > 0]$	**0.31**	0.15										
$\mathbb{E}[\|\Delta L/L\| \mid \|\Delta L/L\| > 0]$	**0.27**	0.14										
$\mathrm{corr}\,(i, i_{-1})$	**0.00**										0.258	0.058
$\mathrm{corr}\,(i, a)$	**0.02**											0.143

Notes: Table displays selected target moments used in the literature estimating structural models of adjustment costs. The column marked IT with bold entries represents the estimates from the Italian dataset. CHST (2015) refers to the paper of Chaney et al. (2015), who use French data. ACL (2014) refers to the paper of Asker et al. (2014), who use U.S., Chilean, Indian, French, Spanish, Mexican, Romanian, and Slovenian datasets. B (2009) refers to Bloom (2009), who uses (quarterly) U.S. data from Compustat. CH (2006) use U.S. data from the Longitudinal Research Database.

Table 15: Comparison of Measures of Reported Model Performance

Moment	**IT**	CHST (2015) FR	B (2009) US	CH(2006) US
$\mathbb{P}[\lvert I/K\rvert > 20\%]$				68.28 (+ve) / 166.67 (-ve)
$\text{sd}\,(I/K)$			92.81	
$\text{skw}\,(I/K)$			100.22	
$\text{sd}\,(E/L)$			88.36	
$\text{skw}\,(E/L)$			69.44	
$\mathbb{P}[I/K > 0]$		112.36		
$\mathbb{P}[\lvert \Delta L/L\rvert \neq 0]$		95.24		
$\mathbb{E}[I/K \mid I/K > 0]$		86.67		
$\mathbb{E}[\lvert \Delta L/L\rvert \mid \lvert \Delta L/L\rvert > 0]$		85.71		
$\text{corr}\,(i, i_{-1})$			118.29	148.28
$\text{corr}\,(i, a)$				216.78

Notes: Table displays model performance of baseline models with respect to (selected) target moments as reported in the literature estimating structural models of adjustment costs. The column marked IT with bold entries represents the estimates from the Italian dataset. CHST (2015) refers to the paper of Chaney et al. (2015), who use French data. ACL (2014) refers to the paper of Asker et al. (2014), who use U.S., Chilean, Indian, French, Spanish, Mexican, Romanian, and Slovenian datasets. B (2009) refers to Bloom (2009), who uses (quarterly) U.S. data from Compustat. Bloom (2009) targets 19 other moments in addition to those 5 reported here. CH (2006) use U.S. data from the Longitudinal Research Database.

Table 16: TFP Gains Over Time

	Year	(1) Baseline	(2) i/c E = 1	(3) $E \geq 20$	(4) $L = E$	(5) Pooled Trim	(6) Balanced	(7) 3-digit	(8) 2-digit	(9) $\sigma = 3$	(10) Constant θ
MANUFACTURING	1995	48.02	48.05	39.71	43.73	49.55	37.4	50.63	54.73	63.06	27.93
	1997	54.92	54.94	44.88	50.42	57.08	40.16	58.19	62.55	74.72	31.54
	1999	63.39	63.49	45.24	53.23	62.92	43.87	66.93	72.6	88.46	36.06
	2001	56.35	56.49	46.99	57.66	56.11	34.26	58.25	65.66	78.48	35.19
	2003	63.94	64.27	47.02	59.84	65.62	36.89	66.31	71.77	86.81	36.74
	2005	59.06	59.5	48.29	54.79	59.93	42.14	63.37	69.13	79.96	59.06
	2007	66.12	68.72	54.8	59.92	68.52	43.35	69.66	73.22	84.39	66.93
	2009	81.54	83.73	62.29	73.17	80.72	52.57	82.19	88.33	106.4	81.41
	2011	81.59	84.28	63.22	77.14	81.59	46.49	83.04	85.32	108.97	79.42
	2013	81.34	84.59	62	76.95	80.73	48.29	83.38	86.5	108.72	78.73
SERVICES	2005	136.45	113.16	117.35	126.88	132.35	113.07	184.86	269.96	161.03	136.45
	2007	156.09	163.27	123.27	136.38	155.83	152.18	201.22	280.94	197.91	152.62
	2009	171.62	187.28	140.06	163.33	174	155.37	204.77	261.88	208.09	165.17
	2011	190.46	203.37	135.81	175.85	195.13	150.15	237.99	249.25	272.31	181.93
	2013	193.75	212.79	139.14	174.48	198.73	112.92	234.25	267.64	271.24	187.95

Notes: Entries are percentage point gains from equalizing TFPR within industries. Entries are calculated as $100(Y_{efficient}/Y-1)$, where $Y/Y_{efficient} = \prod_{s=1}^{S}[\sum_{i=1}^{M_s}(\frac{A_{si}}{\bar{A}_s}\frac{\overline{TFPR}_s}{TFPR_{si}})^{\sigma-1}]^{\theta_s/(\sigma-1)}$. Column (1) is the baseline specification. Column (2) includes firms without employment information. Column (3) includes firms with 1 employee only. Column (4) omits firms with less than 20 employees. Column (5) uses employment as the measure of labor input. Column (6) trims scaled $TFPQ$ and $TFPR$ pooling over years. Column (8) uses 3-digit NACE industries. Column (9) uses 2-digit NACE industries. Column (10) sets $sigma = 5$. Column (11) drops real estate from the service sector sample.

Figure 5: Italian Productivity over Time

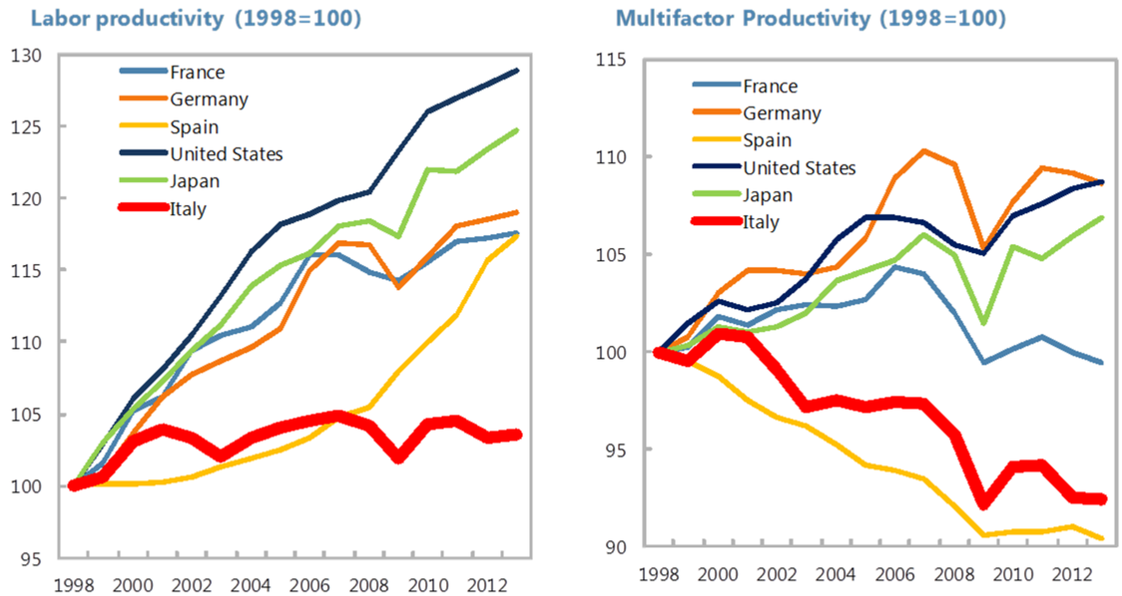

Source: Lanau and Topalova (2016). Labor productivity is measured as GDP per hour worked.

D Additional Figures

Figure 6: Distribution of Investment and Hiring

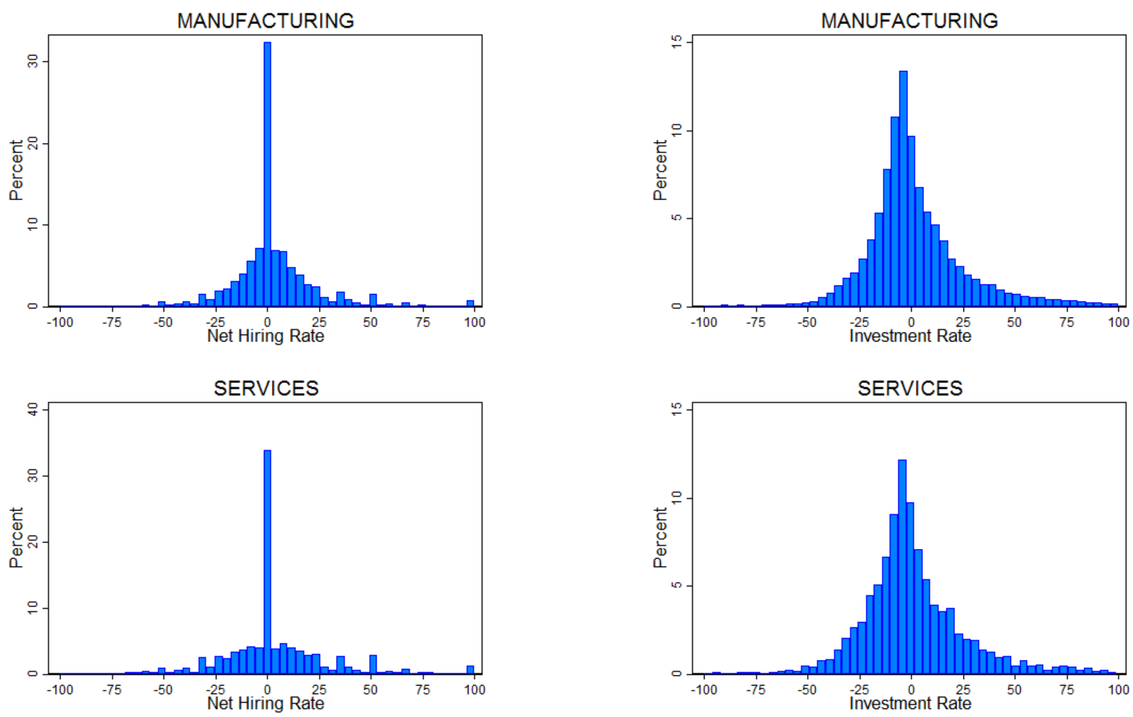

Notes: The hiring rate is defined as $(E' - E)/E$, the investment rate is defined similarly as $(K' - K)/K$. These distributions are calculated from the balanced panel. Distributions are cropped at (-100, +100).

Figure 7: Marginal Revenue Products in Deviation from Industry-Year Means by Size Bin

(a) Bins by Employment

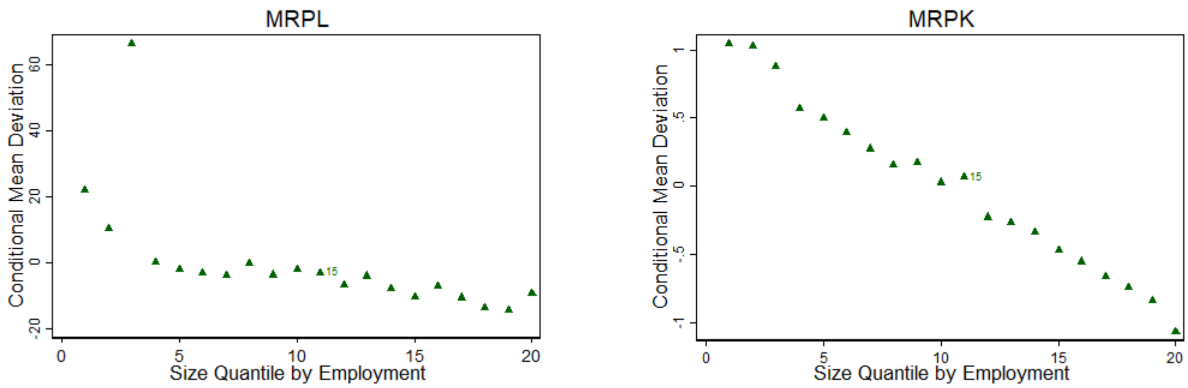

(b) Bins by Value Added

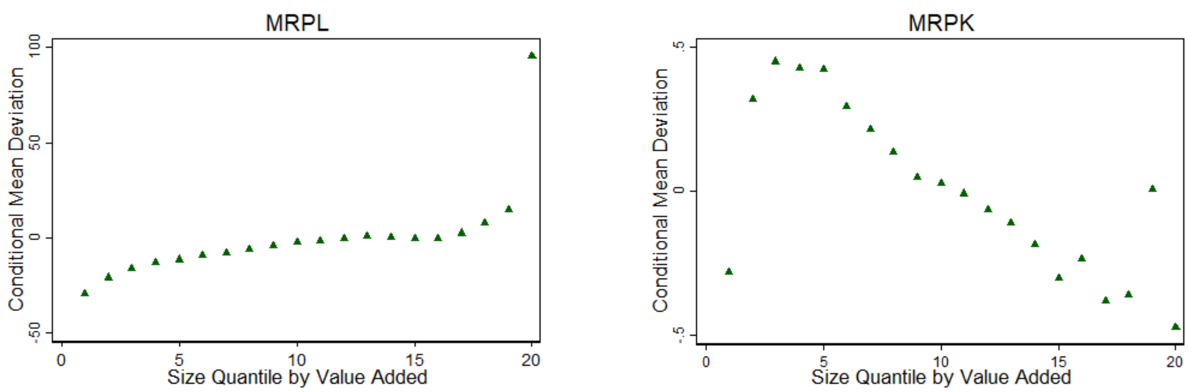

Notes: Figure displays the marginal revenue products of capital and labour, demeaned by their respective industry and year means, and then subsequently averaged conditional on size bin. There are 20 size bins, which are created by ranking observations by employment, displayed for panel (a), and by value added, displayed for panel (b) (unconditional on industry or year). The employment bin containing those firms with 15 employees is indicated. The year 2002 is ommitted on account of an outlier.

$$f(B) = \frac{1}{N_{\{L_{i,s,t} \in B\}}} \sum_{s=1}^{s=S} \sum_{t=1}^{t=T} \sum_{i \in F_{s,t}} \mathbf{1}_{\{L_{i,s,t} \in B\}} \left(MRP_{i,s,t} - \overline{MRP}_{i,s,t} \right)$$

Figure 8: Conditional Mean of Normalized Absolute Deviations by Size Bins

(a) Bins by Employment

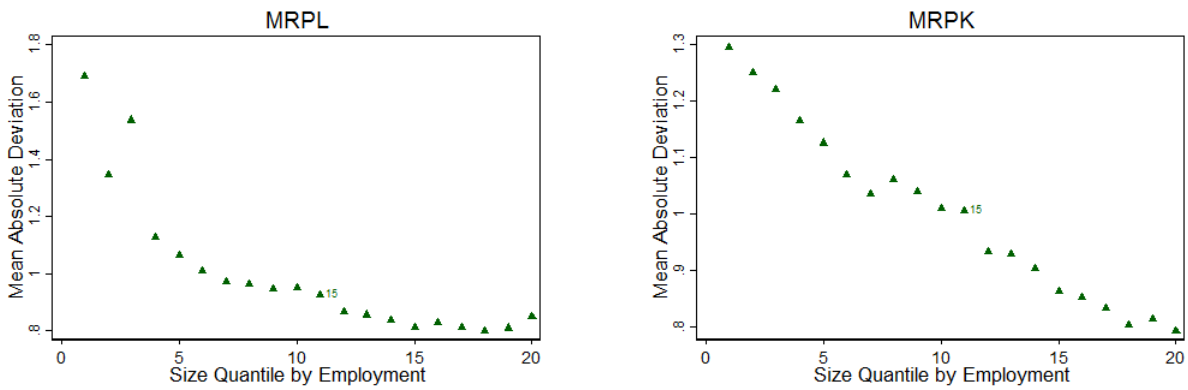

(b) Bins by Value Added

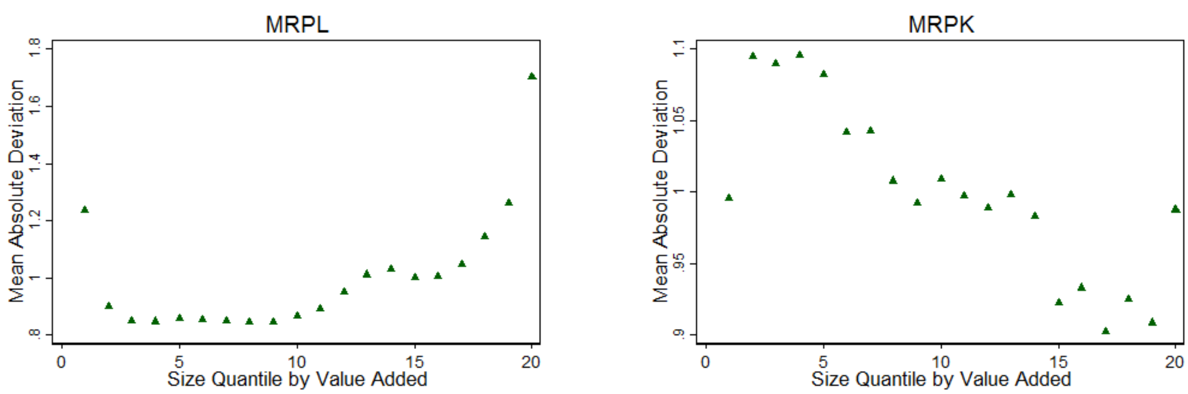

Notes: Figure displays the expected value of the normalized absolute deviation from industry-year mean of the marginal revenue products, conditional on employment level. The absolute deviations are normalized with respect to the industry-year expected absolute deviation. There are 20 size bins, which are created by ranking observations by employment, displayed for panel (a), and by value added, displayed for panel (b) (unconditional on industry or year). The employment bin containing those firms with 15 employees is indicated.

$$f(E) = \sqrt{\frac{\mathbb{E}\left[|MRP_{i,s,t} - \overline{MRP}_{i,s,t}| \mid L_{i,s,t} = E\right]}{\mathbb{E}\left[|MRP_{i,s,t} - \overline{MRP}_{i,s,t}|\right]}}$$

Figure 9: Averages of Key Ratios Conditional on Employment

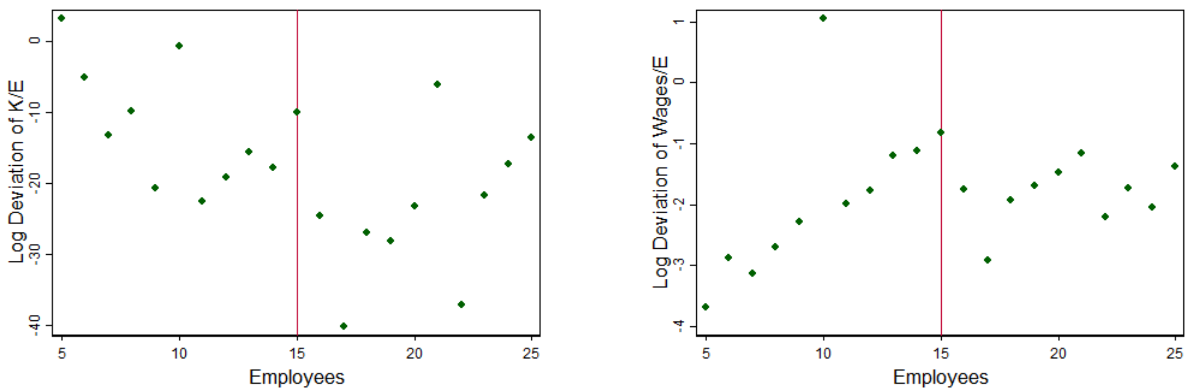

Notes: Figure displays the capital per employee ratio (left-hand panel) and the wage payments to employee ratio (right-hand panel), demeaned by their respective industry and year means, and then subsequently averaged conditional on the number of employees – values are displayed for 5 to 25 employees.

$$f(E) = \frac{1}{N_{\{L_{i,s,t}=E\}}} \sum_{s=1}^{s=S} \sum_{t=1}^{t=T} \sum_{i \in F_{s,t}} \mathbf{1}_{\{L_{i,s,t}=E\}} \left(X_{i,s,t} - \overline{X}_{i,s,t} \right), \quad X_{i,s,t} \in \{ \frac{K_{i,s,t}}{L_{i,s,t}}, \frac{W_{i,s,t}}{L_{i,s,t}} \}$$

Figure 10: Standard Deviation of (Log) Marginal Revenue Products Conditional on Employment (Weighted)

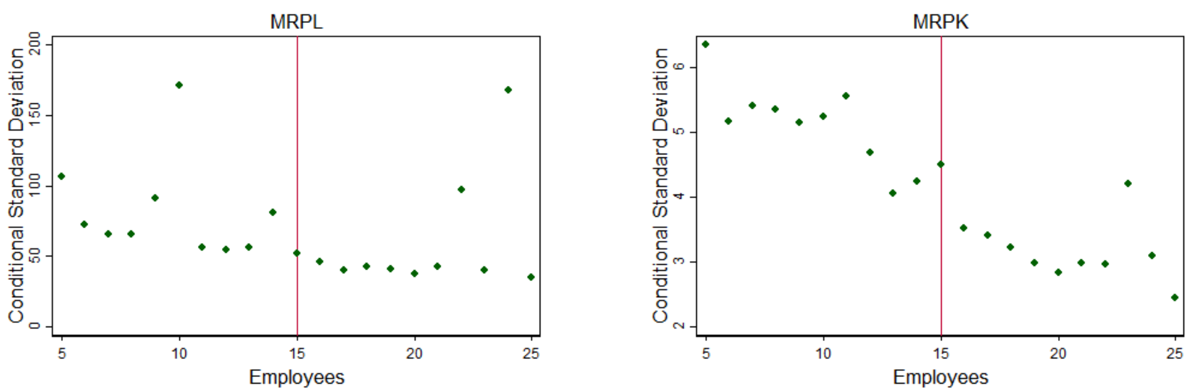

Notes: Figure displays the standard deviation of the marginal revenue products to factors, conditional on the employment level, industry and year. These conditional standard deviations are then averaged across years conditional on the industry and employment level, before then being weighted by industry and averaged over employment group. The industry weights are average value added share of total output for each industry across years (The weights are then normalized to sum to 1).

$$f(E) = \frac{1}{S}\sum_{s=1}^{s=S}\theta_s\left(\frac{1}{T}\sum_{t=1}^{t=T}\mathrm{Sd}\left(MRP_{i,s,t}\,|\,L_{i,s,t}=E,\; s_{i,s,t}=s,\; t_{i,s,t}=t\right)\right),\; \sum_{s=1}^{s=S}\theta_s = 1$$

Figure 11: Conditional Mean of Normalized Squared Deviations by Employment

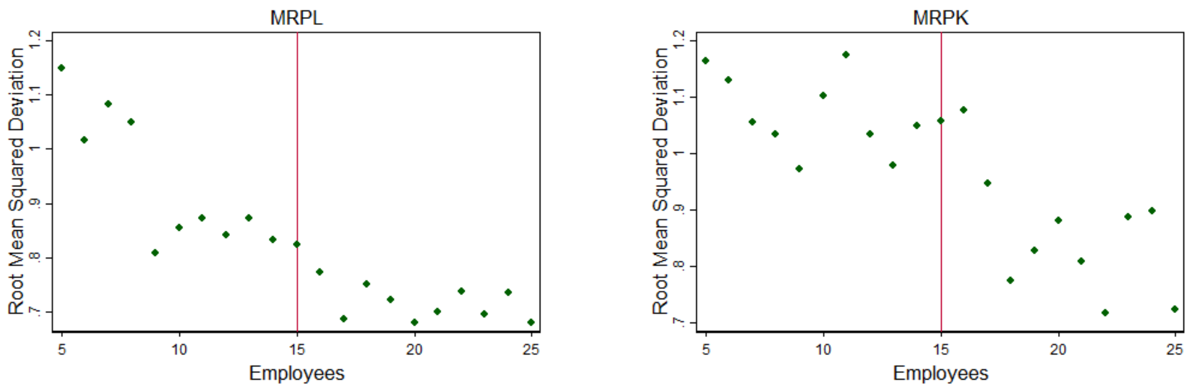

Notes: Figure displays the expected value of the normalized squared deviation from industry-year mean of the marginal revenue products, conditional on employment level. The squared deviations are normalized with respect to the industry-year variance.

$$f(E) = \sqrt{\mathbb{E}\left[\frac{(MRP_{i,s,t} - \overline{MRP}_{i,s,t})^2}{\mathbb{E}\left[(MRP_{i,s,t} - \overline{MRP}_{i,s,t})^2 \mid s_i = s,\, t_i = t\right]} \mid L_{i,s,t} = E\right]}$$

The correlation coefficient between the measure of revenue product deviation and employment is -0.53 for labour, and -0.55 for capital (for the sub-sample between 1 and 100 employees). Both estimates are significant at 1%.

Figure 12: Baseline Reallocation Gains Over Time and Robustness

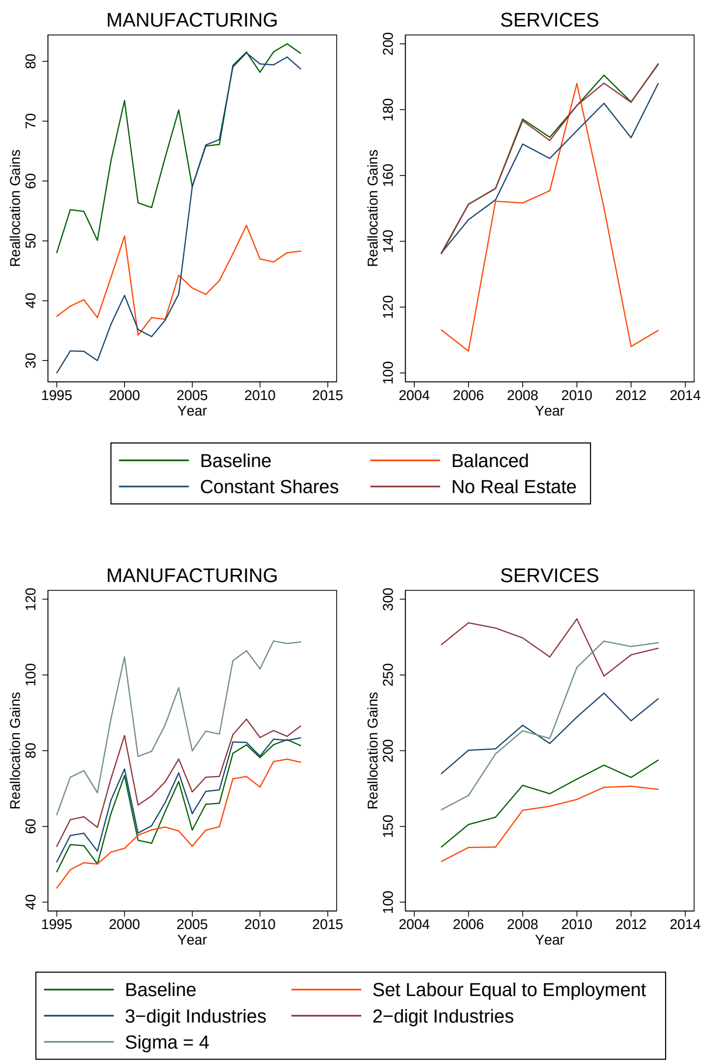

Notes: Different Scales. Upper panel shows misallocation gains under alternate specifications: (1) baseline; (2) version with constant factor shares; (3) version imposing a balanced panel; (4) version excluding real estate (for services only). Lower panel shows misallocation gains under alternate specifications: (1) baseline; (2) industries defined at 3-digit level, as opposed to 4-digit; (3) version with $\epsilon = 5$; (4) industries defined at 2-digit level, as opposed to 4-digit; (4) version with the employment variable as wages.

Chapter 3:

Misallocation and the Financial Crisis

1 Introduction

In this paper I assess the extent to which the efficiency of the allocation of resources across firms within several European economies changed during the 2007-08 crisis and subsequent period of recession. Allocative efficiency in this context refers to an optimal distribution of factors of production across firms of different productivites, within given industries. I also study the relation between efficiency changes and aggregate TFP during the period. I use harmonized European firm-level data, and study 13 countries.

Economists have long considered the possibility that the allocative efficiency of an economy varies with the business cycle. Broadly, there are two competing hypotheses. The first is the "cleansing" hypothesis, which can be understand to state that during recessions allocative efficiency ought to increase as resources are diverted from less productive firms to more productive ones. Cleansing can take place on the extensive and intensive margins. On the extensive margin, we would have cleansing when recessions induce an increased rate of exit for less efficient firms relative to more efficient ones. Along the intensive margin, we would have cleansing when, for a given level of output, recessions induce a shift in production from less efficient firms to more efficient ones. Thus, while production shrinks during a recession, we would see cleansing in the intensive sense if good firms shrink by less than bad ones, on average.[1] Both forms of cleansing would act to mechanically increase the average TFP of incumbents during recessions. These cleansing effects should ameliorate any misallocation of resources across firms within the economy, and potentially dampen output fluctuations over the cycle. To the extent that recessions induce reallocation of resources that are persistent, cleansing effects could further ameliorate the medium-run negative effects of downturns.

The second hypothesis would be that recessions in fact aggravate misallocations of resources across producers – this has been termed the "sullying" or "scarring" effect of recessions. To the extent that the financial sector plays an important role in the reallocation of

[1] Intensive cleansing has also been termed shift-share cleansing.

resources across firms, sullying effects may be particularly pertinent in the case that a recession is accompanied by financial sector breakdown, as was the case in many developed economies during the economic crisis of 2007-08.

Cleansing and sullying arguments both make claims about the allocative efficiency of economies during business cycles, and so as to empirically examine these arguments this paper quantifies the efficiency of 13 European economies during crisis years. Evidence that allocative efficiency increased during the crisis would support the cleansing hypothesis, and evidence to the contrary would support sullying. It must be emphasised that only net changes in efficiency are charted, the gross contributions of cleansing and sullying effects are not computed. Indeed, this study does not take a stand on the differing mechanisms that might lead to changes in efficiency, and instead adopts the Hsieh and Klenow (2009) approach to misallocation, calculating the distribution of wedges in the first-order conditions of optimizing firms. The basic intuition at work is that the greater the dispersion of the marginal products of an input across firms in the same industry, the greater the potential gains from reallocation and the lower the level of allocative efficiency. The exogenous wedges prevent the market from attaining the first-best distribution. [2] The methodology is rich enough to track cleansing and sullying effects in the sense of shifts in the sizes of producers, but also misallocation of the separate factors of production – labour and capital in my set-up. Because the Hsieh and Klenow framework relies on strong functional form assumptions, I also examine reduced form evidence, testing for cleansing effects using a dynamic panel regression approach. It should be stressed that the focus of my paper is on the reallocation of resources across incumbents, and not the effects of firm entry and exit, which are not quantified.[3]

The motivation for this paper is that the overall effect of recessions on allocative efficiency is of interest to economists for a fuller understanding of the dynamics of aggregate productivity during business cycles. It also has macro-policy implications, since cleansing and sullying effects, if present, can alter the strength of the impact of aggregate shocks on output. These effects can also determine the medium-run impact of aggregate shocks. Moreover, cleansing effects are certainly properties of many theoretical models and empirical evidence is relatively sparse, this paper contributes with new quantitative evidence.

The paper contributes to the literature insofar that it is the a cross-country empirical

[2] A key strength of the approach is that under the assumptions of the Hsieh and Klenow methodology we can quantify allocative efficiency having calculated only the dispersion of the marginal *revenue* products, which is all that we are able to calculate using the majority of firm-level datasets, including the data I use for this study.

[3] This is on account of the poor ability of the AMADEUS dataset to track firm entry and exit.

study of changes in the levels misallocation of resources during a financial crisis, while previous research has previously to single-country case studies. The cross-country aspect is important since the strength of cleansing or sullying effects may have structural determinants, and thus it would be of interest to compare how structurally different economies performed with respect to misallocation of resources during the crisis. Since cleansing or sullying effects may determine, in part, the severity of recessions and their after-effects, it would be of interest to see whether they are stronger in countries with particular economic frameworks. A cross-country study of cleansing and sullying could provide economists with clues as to how to affect structural changes to economies in order to make the business cycle less damaging; a worthy goal given the recent experience of developed nations.

The results of the analysis in this paper shows that there is some evidence of downward movements in measures of misallocation during years associated with the financial crisis and subsequent recessions. However, the measures are generally quite volatile, and typically take place around an upward trend, in place prior to the onset of the recession. Descriptive regression analysis is able to detect some evidence for cleansing redistributions of value-added and employment, for the manufacturing sample. The evidence certainly does not seem to favour arguments for the sullying effect of recessions. Corroborating evidence from a separate data-source (the Annual Respondents Database, ARD) is presented for the U.K., given academic interest in the question of this economies recent productivity dynamics. A long-term trend in misallocation is observed for this economy in both the Amadeus and in the ARD dataset, despite the differences in their construction.

2 Literature Review

This paper bridges two literatures which have remained distinct until only relatively recently, although both are broadly aimed at accounting for aggregate productivity. The first is the literature on the cleansing and sullying effects of recessions on aggregate productivity, which has a long history. The second literature has examined the effects of misallocation of resources on TFP, is more recent and has until now largely been confined to static analyses aimed at explaining cross-country differences in levels of productivity, as opposed to studying of its cyclical variation over time. I will discuss the two literatures in turn.

The cleansing view of recessions dates back as far as Schumpeter (1942). One of the first modern formalizations would be that of Caballero and Hammour (1996), in which the authors present a model with exogenous technological progress, and exogenous shifts of demand. The model is a vintage model of technology adoption with costly firm-entry.[4] During periods of low demand, the least productive firms (which are also the oldest) turn unprofitable and exit. However, this cleansing effect is dampened by changes in the level of firm-entry, since the rate of entry also falls with demand, which insulates incumbents from its full effects. Within the modern DSGE framework, Ottaviano (2012) provides a model with firm entry and exit in which cleansing means the propagation effects of these margins are reduced.[5]

However, there are a range of theories on the subject of the effects of recessions on efficiency, and many such theories do not predict cleansing, and propose mechanisms by which recessions might in fact make the distribution of resources across production units less efficient. Barlevy (2002) responds to the stylized fact that jobs created during recessions are likely to be low-paying and temporary in U.S. data, seemingly in contradiction to the theory of cleansing. Barlevy (2002) adds on-the-job search to the standard Mortensen and Pissarides (1994) search-and-matching framework and shows that this leads to an additional "sullying" effect. When the model is calibrated, Barlevy (2002) finds that the offsetting sullying effect of recessions is likely to be much larger than the cleansing effect, accounting for the stylized fact.

More recently, Ouyang (2009) proposes an additional negative effect of recessions on average productivity through a novel mechanism termed as the "scarring effect". Ouyang (2009) takes the Caballero and Hammour (1996) vintage model as a baseline, but makes new firms uncertain as to their own quality. In this environment recessions have a scarring effect on the productivity of incumbent firms in the medium run, because viable firms are destroyed during the recession on account of the truncated learning process. Thus the model contains both the traditional cleansing effect, but also a new scarring effect which works in the opposite direction. The effects of recessions on the productivity of firms therefore depends on which of the two effects is stronger.

A related theoretical literature which also discusses the productivity of firms during business cycle fluctuations are the so-called "pit-stop" theories of firm dynamics. These models focus not on allocative efficiency across firms of given productivities, or on entry and exit, but on the ways in which firms might endogenously change their own productivi-

[4]In a vintage model of technology firms embody the best technology available at the time of creation, and are then stuck with this technology for the rest of their life span.

[5]Most of the theoretical literature has focused on this extensive margin of cleansing, although entry and exit of firms may not be conceptually so different from firm re-sizing (the extensive margin) if one considers firms to be made up of a collection of individual plants or projects, each of which may enter or exit in response to changing business considers according to the decision of the firm's manager.

ties during recessions. The general argument of such theories is that some investments may increase the productivity of the firm in the future but disrupt current production as they are implemented. In these cases, the opportunity cost of the investments is lower during recessions than booms, since profits per unit are reduced. Aghion and Saint-Paul (1998) employ just such a device in their model, and have firms undertake investment counter-cyclically. Like cleansing, this would provide a basis for recessions having an ameliorative effect on the productivity of incumbent firms, although it should occur with a lag.[6] However, despite the predictions of these models, they do seem to run contrary to stylized facts observed in the data, in particular the procyclicality of R&D.[7]

To date the literature has been unable to provide comprehensive empirical evidence on cleansing and sullying effects, largely due to the limited time-series of most firm-level datasets, although studies do exist. Davis et al. (1998) found increased job reallocation during recessions in their study of census data on U.S. firms. Foster et al. (2013) analyse the question of cleansing during the 2007-09 crisis and subsequent recession, again using census data on U.S. firms. They examine the cleansing hypothesis using state-level variation in unemployment as a measure of the business cycle and find cleansing effects. They do however find the cleansing effect to have weakened during the most recent recession. Mustre-del Río (2012) tests the cleansing and sullying effects of recessions with data from the U.S. National Longitudinal Survey of Youth (NLSY) 1979-2006, using the duration of a worker-firm match as a proxy for its quality. The paper finds no systematic evidence for the cleansing effect, and further results tend to support the hypothesis that recessions induce worker-firm matches of shorter than expected duration. The author infers that average match quality is procyclical. High unemployment rates at the start of the match are correlated with low match quality suggesting an active sullying effect of recessions. The cleansing effect is at best unable to offset the sullying effect and at worst not significant.

Outside of U.S. data, Lin and Huang (2012) look at the Taiwanese case and show that a fall in economic activities is associated with a decrease in the fraction of newly created firms accounted for by plants with a high rate of technical change, which they infer indicates that creative destruction is more pronounced during economic contractions. Eslava et al. (2010) study Colombian manufacturing establishments and find that credit-constrained but nevertheless high productivity units may be forced out of the market during recessions, while other productive but unconstrained units may survive. Hallward-Driemeier and Bijkers

[6]These effects are also a property of the machine replacement model of Cooper and Haltiwanger (1993).

[7]See Caballero and Hammour (2005) for a theoretical model which predicts reduced restructuring during recessions, on account of countercyclical contracting expenses associated with R&D.

(2013) study Indonesian manufacturing census data (1991-2001) and rejects the hypothesis that the East Asian crisis unequivocally improved the reallocation process. They find that the correlation between productivity and employment growth did not strengthen and the crisis induced the exit of relatively productive firms. They do find however, that firms that entered during the crisis were relatively more productive, which helped to mitigate the reduction in aggregate productivity.

There has also been treatment in the literature of the way in which financial crises in particular affect the aggregate productivity of the real economy, and we do have evidence that the financial crises in emerging markets of the previous few decades lead to large falls in aggregate TFP. Calvo et al. (2006) show that sudden stop episodes are associated with falls in GDP of 10 percent on average, mainly due to falls in TFP. Meza and Quintin (2006) find that TFP fell by more than two standard deviations in all the cases they study from the Mexican and East Asian crises. Such evidence is certainly consistent with the idea that financial crises might lead to increases in the misallocation of resources, since a poorer distribution of production across firms would lead to falls in aggregate TFP. It may be that these effects are sufficient to counteract cleansing effects, if present.[8] However, without a suitable decomposition, we are unable to link these falls in TFP to decreases in allocation efficiency *per se*.

Of course, the recessions of the developed world during the period following 2008 came at the same time as major financial crises. With respect to the effects of financial crises *specifically* on resource allocation, there is again no theoretical consensus. Barlevy (2003) argues that during recessions few lenders are willing to extend large amounts of credit, and so the projects that survive are those that require less credit, regardless of their efficiency. A model is presented in which the presence of credit frictions mean that resources are directed from more efficient to less efficient uses during recessions. This occurs because more efficient production arrangements are also more vulnerable to credit constraints, since projects which generate more surplus also offer more incentive to deviate, making them fragile. Osotimehin and Pappadà (2016) present a model of firms entering and exiting markets in the presence of credit frictions, and show that, despite their distortionary effect on the selection of existing firms, credit frictions do not reverse the cleansing effect of recessions. Average idiosyncratic productivity rises following an adverse aggregate shock.[9] Saffie and Ateş (2013) develop a model in which sudden stops in the supply of credit to a small open

[8]See Pratap and Urrutia (2012) for a DSGE formulation of these effects.

[9]In their model they also find recessions have only a modest impact on average productivity, irrespective of the level of credit frictions.

economy forces the representative financial intermediary to select only the most promising ideas, which means that firms born during the credit shortage are fewer, but better.[10]

The second main body of research this paper relates to is that focusing on misallocation as an explanation for cross-country differences in TFP levels, as part of the wider development economics literature. The term misallocation is used to refer to deviations from allocative efficiency of factors of production across firms. For examples, papers by Banerjee and Duflo (2005), Restuccia and Rogerson (2008), and Bartelsman et al. (2009) argue that resource misallocation can explain a large part of the differences in total factor productivity between rich and poor countries. As detailed in Restuccia and Rogerson (2013), there are two main approaches to analysing misallocation. The *direct* approach would be for the researcher to select one or more factors or mechanisms thought to be important to misallocation and then proceed to analyse these specific factors analytically or empirically. The *indirect* approach, of which this paper is an example, does not take a stance on which particular factors or mechanisms may lead to misallocation. Since the researcher will never be able to measure all sources of misallocation directly, the indirect approach focuses on the overall net effect of all the various determinants, quantifying the aggregate effect without examining particular factors that could have lead to it. The intuition is that any factors that lead to misallocation can be thought of as generating wedges in the first-order conditions of the optimization problems of firms, and the indirect approach relies on the calculation of these "wedges" for firms. [11] The methodology takes a stand on what represents an efficient allocation of resources, unlike previous methods.[12] The stand-out paper in this field would be Hsieh and Klenow (2009), who develop a methodology to account for misallocation of resources in order to offer an explanation for the large recorded TFP differences between developed and developing countries. They analyse firm-level data from the US, and from India and China. They find larger dispersion of wedges in the marginal products of factors of

[10] They find this result to be supported by Chilean firm-level data. After calibrating the model to the data, they find that failing to account for these selection effects can lead to large overestimation of output losses.

See also: Barlevy (2007) provides an analytical model in an attempt to explain pro-cyclical R&D. Comin and Mulani (2009) present a different theoretical framework which also attains pro-cyclical investment by firms. On R&D Investment: Aghion et al. (2012). Peters (2007) discusses persistence in innovation, and gives potentially: (1) success breeds success, (2) dynamic increasing returns, (3) sunk costs in R&D investments.

[11] As noted in Hsieh and Klenow (2009), in this sense the method represents an example of the approach to business cycle accounting argued for in Chari et al. (2007), in the sense that the wedges will subsume the results of a large class of models, and will, by construction, account for all observed movements of variables of interest in given data.

[12] An alternative approach would be to examine the correlation between firm size and the average product of labour as a measure of allocative efficiency, under the theory that the most productive firms should be the biggest Bartelsman et al. (2009).

production across plants and infer that resources are more misallocated in these countries. They find that were resources to be hypothetically reallocated to U.S. levels of efficiency of allocation, under the assumptions required by the framework TFP would increase by 30-50% in China and 40-60% in India.

More recent papers have studied the dynamics of such misallocation measures. Oberfield (2013) applies a similar approach to the examination of a crisis and business cycle using Chilean data, with the main finding is that misallocation within industries either barely changes, or actually improves. Oberfield (2013) also finds that the measures of misallocation across industries developed in the paper account for about one third of the changes in TFP during the crisis.

This study also closely relates to several recent applications of these methods to misallocation in Southern Europe. Dias et al. (2015) find that within-industry misallocation almost doubled between 1996 and 2011 using Portuguese data, employing a version of the Hsieh and Klenow (2009) framework with a three-factor production function. García-Santana et al. (2016) also find increases in misallocation in Spanish firm-level data for the period 1995-2007, concluding that in the absence of such deterioration, average TFP growth would have been around 0.8% a year, as opposed to the *falls* in Spanish TFP.[13]

The paper of Gamberoni et al. (2016) is the closest paper to this study. These authors study five euro-area countries (Belgium, France, Germany, Italy and Spain), using data from CompNet dataset during the period 2002-2012. They also check robustness using the Amadeus data used in this study, although only for the Spanish sub-sample. They also employ the misallocation framework of Hsieh and Klenow (2009). This paper contributes relative to their study with respect to a larger sample of countries, using a different dataset. I also quantify misallocation trends for the U.K. for the first time, using data from two-different sources. It is reassuring that results are similar, insofar that the paper of Gamberoni et al. (2016) also documents time trends in misallocation, and notable drops in the level of misallocation during the crisis period.

Gopinath et al. (2017), is also a close paper to this study in the sense that they study several Northern and Southern economies, including Italy, France, Germany, and Spain, employing the closely related ORBIS dataset. Their analysis points to large increases in the dispersion of the marginal revenue product of capital in the Spanish manufacturing sector

[13]The paper of Inklaar et al. (2017) studies 52 low- and medium-income countries using World Bank survey data, and finds significant differences in misallocation relative to the U.S., but no correlation between income-level and misallocation. García-Santana and Ramos (2015) study 104 developing countries, and focus on the relation between distortions and firm-size.

between 1999 and 2013, suppressing TFP. They find similar results for the case of Spain and Portugal. Schivardi et al. (2017) undertake a study of credit misallocation for the Italian case using a unique bank-firm linked dataset spanning the period 2004-2013, finding only a modest contribution of capital misallocation to the severity of the great recession.[14]

3 Data

The principal analysis of this study uses data from the Amadeus dataset, which is a commercial database and published by Bureau van Dijk. A total of 41 countries are included in the database, and up to ten years of data are provided for each firm, although coverage varies by country. Amadeus is created by collecting data from 50 vendors across Europe, and the local source for the data is generally the office of the Registrar of Companies. The observations comprise both public and private firms. Amadeus includes consolidated and unconsolidated annual accounts, activities and ownership. I use the 2013 edition of Amadeus, giving observations between 2003 and 2012 inclusive. I have data on Europe's biggest 500,000 public and private companies at the outset.[15]

With respect to country selection, the analysis is applied to the 13 countries with the largest and most complete data for the time-period. The list of countries examined is therefore: Austria, Belgium, Bulgaria, Czech Republic, Germany, Spain, France, Italy, Poland, Romania, Sweden, and the U.K.

One of the more problematic features of the dataset is the problem of survivorship: as companies exit or stop reporting their financial statements, Bureau van Dijk enters a "not available/missing" for four years following the last included filing, and then removes all data from the commercial distribution. Given that the sample ends in 2012, this means that I do not observe firm deaths that occur before 2009. Since the nature of the sample changes in 2009, in that it includes both survivors and firms that will potentially exit, the analysis is restricted to a balanced panel, insofar as balance sheet information for the firm must be observed for dates before 2008. This is equivalent to removing the sub-population of exiting firms from the analysis.[16] This means that cleansing can only be studied in the intensive sense across firms that survive through the sample period, which is an important limitation

[14]Other related studies is that of Kehrig (2015), who studies the cyclicality of TFP dispersion in U.S. data, and finds it to be greater in recessions.

[15]There are several versions of Amadeus, and larger editions with more firms do exist – I have access to the 500,000 firm version. My study treats the data in ways which broadly follow previous studies in the literature which use the Amadeus dataset. See Klapper et al. (2006), Arnold and Schwellnus (2008), Da Rin et al. (2011).

[16]I choose 2004 as the first year of the restriction since there are very few observations in 2003.

of the analysis. However, it is in keeping with the focus on changes in the distribution of resources across incumbent firms.

Amadeus lists both consolidated and unconsolidated filings. Consolidated filings amalgamate the information from the subsidiaries of a firm into a single balance sheet, while unconsolidated filings represent the disaggregated balance sheets in the sense that they are either the information of the parent firms not including their component subsidiaries, or the balance sheets of subsidiaries themselves. All consolidated filings are dropped from the analysis, in order to avoid double-counting. [17]. All legal forms other than the equivalent of public and private limited liability companies are also excluded from the analysis. The majority of listings in Amadeus are of these forms. Minimal further cleaning of the data was undertaken. Observations are required to have non-missing values for production function variables that are strictly positive.This does however, remove a large number of observations from my sample. Wage expenditure is required to be smaller than added value, removing insolvent firms.

The coverage of the cleaned dataset with respect to the underlying population is displayed in Table 1, where employment and revenue are respectively totaled across firm-level observations by sector, year and country, and then compared to the totals reported in data from Eurostat Structural Business Statistics. Admittedly, the coverage is relatively low in all cases. The more representative data samples are to be found in the Belgium, the Czech Republic, Romania and the UK. The data for Germany, and particularly France, have quite low coverage, which suggests that conclusions drawn from these data sub-samples are less likely to be representative.

Table 2 presents the mean and standard deviation of value-added, capital and employees for given countries, for data pooled over years. It is apparent from this table that the German, Spanish, Italian and UK samples are the largest, and incorporate relatively larger firms, on average. The standard deviation of capital in the services industry is much larger than for manufacturing, while the standard deviation of value added is broadly comparable. This could potentially reflect difficulties in measurement of capital for this sector, but could also speak to a weaker relation between capital allocation and profitability in for service-sector firms. Very large standard deviations by employment are observed, which reflects the fact that the data contain both very small firms with one or two employees, and also the largest firms in the given economies. Table 3 displays the same statistics for individual countries, disaggregated by year. Typically, the average size of firms rises over the sample,

[17]Most of the listings in the database are unconsolidated

as well as the standard deviations (although this is not the case for all countries, nor are the increases monotonic). The measures of the standard deviations, in particular, can vary by large amounts over time in the data. Since the measures used in this paper essentially uses dispersion between firm outcomes as a means to identify resource misallocation, this would provide early evidence that misallocation measures are likely to vary substantially over time in the data sample.

4 Misallocation Accounting

This section lays out the methodology employed in the paper, which follows the quantitative model developed by Hsieh and Klenow (2009), henceforth HK. I must emphasize that with respect to the explanation of the model in this section, I largely follow that of HK, and of course invite the reader to refer to this paper for a more thorough description of their methodology. The model presented in HK is a standard static model of monopolistic competition, with the key innovation being that firms face idiosyncratic “distortions”. These distortions prevent the market from removing misallocation, and appear as extra payments (or subsidies) that are taken by optimizing firms as exogeneous when choosing how much to produce. The economy is closed, and has two inputs to production (capital and labour).

Assume there are S different industries, each comprised of M_s firms, each producing differentiated varieties of goods.[18] Output for industry s is given by the CES aggregator:

$$Y_s = \left(\sum_{i=1}^{M_s} Y_{si}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where σ is the elasticity of substitution between varieties, Y_{si}.[19]

Firms produce their variety according to the Cobb-Douglas production function

$$Y_{si} = A_{si} K_{si}^{\alpha_s} L_{si}^{1-\alpha_s}, \tag{1}$$

where A_{si} is the firm-specific productivity, K_{si} is capital, L_{si} is labour, and $\alpha_s \in (0, 1)$ is the capital share. Here we have made the assumption that there are constant returns to scale, and that firms within the same industry share the same production technology.

[18]Each firm produces one good, and each good is produced by only one firm: goods and varieties are both indexed by i.

[19]Throughout the notation I suppress the time subscript.

Firms choose labour and capital to maximize their profits, which are given by:

$$\pi_{si} = (1 - \tau_{Y,si})P_{si}Y_{si} - wL_{si} - (1 + \tau_{K,si})RK_{si},$$

where P_{si} is the price of variety i, w is the wage, and R is the rental rate of capital. Here $\tau_{Y,si}$ and $\tau_{K,si}$ are firm-specific distortions, which can be respectively termed the "output distortion" and "the capital distortion". The distortions can be understood to be taxes in the case they are positive, or subsidies in the case they are negative. Thus, a positive output distortion represents exogeneous factors that suppress the size of a firm below its efficient level, increasing the marginal products of labour and capital equally. HK give the example of government restrictions on size and transportation costs. A positive capital distortion represents exogeneous factors that increase the costs of capital relative to labour, for example financial frictions or trade union power. The reason the effects of factors affecting the price of labour and capital are incorporated together is that we are only able to identify as many distortions as there are factors of production.

Profit maximization implies the following

$$P_{si} = \frac{\sigma}{\sigma - 1}\left(\frac{R}{\alpha_s}\right)^{\alpha_s}\left(\frac{w}{1-\alpha_s}\right)^{1-\alpha_s}\frac{(1+\tau_{Ksi})^{\alpha_s}}{A_{si}(1-\tau_{Ysi})}, \tag{2}$$

where $\frac{\sigma}{\sigma-1}$ can be understood as the mark-up charged by the firm, over its marginal cost. Thus the prices set by firms are inversely proportional to their productivity, and to the output distortion, and proportional to the capital distortion weighted geometrically by the capital share.

The implication of this optimal pricing rule 2, is that the capital-labour ratio, labour allocation, and output can be expressed as follows:

$$\frac{K_{si}}{L_{si}} = \frac{\alpha_s}{1-\alpha_s}\frac{w}{R}\frac{1}{(1+\tau_{K,si})}, \tag{3}$$

$$L_{si} \propto \frac{A_{si}^{\sigma-1}(1-\tau_{Y,si})^{\sigma}}{(1+\tau_{K,si})^{\alpha_s(\sigma-1)}}, \tag{4}$$

$$Y_{si} \propto \frac{A_{si}^{\sigma}(1-\tau_{Y,si})^{\sigma}}{(1+\tau_{K,si})^{\alpha_s\sigma}}. \tag{5}$$

From equations 3-5 we can observe several things about the nature of the efficient economy, where distortions are removed. The first is that all firms in the same industry will have identical capital-labour ratios, given by 3. The second is that the size of the firm is then only determined by its idiosyncratic productivity, as can be seen from equations 4 and 5. In the undistorted case more productive firms are larger: however, the size of the firm is geometrically weighted by the willingness of the consumer to substitute between goods, σ, since there is diminishing marginal utility to the consumer from any one variety. In fact, with no distortions, more productive firms will grow until the the marginal revenue product of labour and capital are respectively equalized within industries, as can be seen from the following equations:

$$\begin{aligned} MRPL_{si} \equiv &(1-\alpha_s)\frac{\sigma-1}{\sigma}\frac{P_{si}Y_{si}}{L_{si}} = w\frac{1}{1-\tau_{Ysi}}, \\ MRPK_{si} \equiv &\alpha_s\frac{\sigma-1}{\sigma}\frac{P_{si}Y_{si}}{K_{si}} = R\frac{1+\tau_{Ksi}}{1-\tau_{Ysi}}. \end{aligned}$$

We note that the presence of distortions will cause allocations of labour and capital to deviate from the efficient level, and that this causes divergences in the marginal revenue products of factors that would not exist otherwise. Thus, if distortions are present, welfare gains from reallocation of factors of production exist, in principle.

4.1 Physical and Revenue Productivity

HK draw a distinction between "physical productivity", which they call TFPQ, and "revenue productivity", which they call TFPR—the expressions for these quantities are given by

$$TFPQ_{si} \equiv A_{si} = \frac{Y_{si}}{K_{si}^{\alpha_s}L_{si}^{1-\alpha_s}}, \tag{6}$$

$$TFPR_{si} \equiv P_{si}A_{si} = \frac{P_{si}Y_{si}}{K_{si}^{\alpha_s}L_{si}^{1-\alpha_s}}. \tag{7}$$

TFPQ is the idiosyncratic productivity of the firm at producing goods from given levels of inputs. TFPR is the idiosyncratic productivity of the firm at producing *revenue* from given levels of inputs. TFPR is of interest to the applied researcher because it is in principle measurable from the majority of firm-level datasets, which contain information on nominal output $P_{si}Y_{si}$, but not physical output Y_{si}, nor the *firm-level* deflators P_{si} required to compute Y_{si} from nominal output.

The HK framework takes the distribution of TFPQ as exogenous. However, in the frictionless case, the model yields the prediction that TFPR ought to be equalized between firms, and in the end, the distribution of TFPR becomes a (measurable) statistic for investigating misallocation. To see this, we insert the expression for the optimal price, (2), into the definition of TFPR, (7), and attain the following:

$$TFPR_{si} = \frac{\sigma}{\sigma - 1} \left(\frac{R}{\alpha_s} \right)^{\alpha_s} \left(\frac{w}{1-\alpha_s} \right)^{1-\alpha_s} \frac{(1+\tau_{K,si})^{\alpha_s}}{1-\tau_{Y,si}}. \tag{8}$$

From equation (8), when $\tau_{K,si}$ and $\tau_{Y,si}$ are zero for all firms in industry s, we see that $TFPR_{si}$ will be equal, and does not depend on $TFPQ_{si}$. The intuition for this result is that firms with higher A will accumulate labour and capital and grow in size, however they will also reduce their prices, until their TFPR is equal to that of less productive ones. Moreover, if we do observe a firm to have relatively high TFPR, we infer that it confronts distortions that render it smaller than is optimal.

We can then define industry TFP as the following:

$$TFP_s = \left[\sum_{i=1}^{M_s} \left(A_{si} \frac{\overline{TFPR}_s}{TFPR_{si}} \right)^{\sigma-1} \right]^{\frac{1}{\sigma-1}}, \tag{9}$$

where $\overline{TFPR}_s$ is a weighted average of revenue total factor productivity.[20] Equation 9 shows that industry TFP is a weighted sum of the individual firm-level productivities, so clearly (and straightforwardly) industry TFP is increasing in the size of firms' $TFPQ$. Moreover, given $\sigma > 1$, TFP_s is a concave function of the summation, which implies that TFP_s increases at a decreasing rate with respect to A_{si} and $\overline{TFPR}_s/TFPR_{si}$.[21] The implication of this is that a greater variance of A_{si} will actually reduce TFP_s, *ceteris paribus*: given that consumers consider the different varieties to be substitutable, increases in the productivity of any one firm is of diminishing marginal benefit, and a situation where one firm increases its productivity while another's productivity falls reduces utility, even if average productivity is held constant. In this sense equation 9 embeds the preferences of the consumer vis-á-vis varieties when aggregating TFP appropriately.

It is also true from (9) that TFP_s depends negatively on the variance of $TFPR_{si}$. This can be seen clearest in the case that $\log A_{si}$, $\log(1-\tau_{Y,si})$, and $\log(1+\tau_{K,si})$ is multivariate normal. Denoting the variances of $\log(1-\tau_{Y,si})$ and $\log(1+\tau_{K,si})$ by σ_Y^2 and σ_K^2, respec-

[20] $\overline{TFPR}_s = \frac{\sigma}{\sigma-1}\left(\frac{\overline{MRPK}_s}{\alpha_s}\right)^{\alpha_s}\left(\frac{\overline{MRPL}_s}{1-\alpha_s}\right)^{1-\alpha_s}$, where $\overline{MRPL}_s \equiv w/\left(\sum_{i=1}^{M_s}(1-\tau_{Y,si})\frac{P_{si}Y_{si}}{P_sY_s}\right)$ and $\overline{MRPK}_s \equiv R/\left(\sum_{i=1}^{M_s}\frac{(1-\tau_{Y,si})}{(1+\tau_{K,si})}\frac{P_{si}Y_{si}}{P_sY_s}\right)$.

[21] See the interesting discussion of the HK paper: https://growthecon.wordpress.com/2014/09/25/measuring-misallocation-across-firms/

tively, and their covariance by σ_{KY}, then,

$$\log TFP_s = \frac{1}{1-\sigma}\log\left(\log M_s + \log E[A_{si}^{\sigma-1}]\right) - \frac{\sigma}{2}\mathrm{var}(\log TFPR_{si}) - \frac{\alpha_s(1-\alpha_s)}{2}\sigma_K^2. \tag{10}$$

We can see in equation (10) the negative relation between the variance of TFPR, and also that its impact is increasing in the parameter σ, i.e. the willingness of consumers to substitute between goods. The intuition for this latter effect is that consumers are less willing to tolerate a given level of dispersion of TFPR as goods become more and more substitutable. We also see an added term relating to the dispersion of the capital wedge, $-\frac{\alpha_s(1-\alpha_s)}{2}\sigma_K^2$, not summarized by TFPR dispersion.

We should note that the above analysis is conditional on a fixed aggregate stock of capital and labour, and we must make the assumption that the number of firms in each industry is not affected by the extent of misallocation.

4.2 TFP Gains from Removing Distortions

The previous section has discussed how the dispersion of TFPR is a useful statistic for evaluating the level of misallocation of resources within industries in the HK framework. It is also possible to quantify industry TFP gains from removing distortions. From equation (9), we can see that in the case that $TFPR_{si} = \overline{TFPR}_s$, the efficient industry productivity, TFP^* is given by:

$$TFP_s^* = \left[\sum_{i=1}^{M_s} A_{si}^{\sigma-1}\right]^{\frac{1}{\sigma-1}}.$$

Provided we have a means to calculate A_{si}, that we know the parameters of the model, and that we have sufficient firm-level data to estimate a production function, we can then compute the gains from removing misallocation for a given industry s are as follows:

$$\frac{TFP_s^*}{TFP_s} = \frac{\left[\sum_{i=1}^{M_s} A_{si}^{\sigma-1}\right]^{\frac{1}{\sigma-1}}}{\left[\sum_{i=1}^{M_s}\left(A_{si}\frac{\overline{TFPR}_s}{TFPR_{si}}\right)^{\sigma-1}\right]^{\frac{1}{\sigma-1}}}. \tag{11}$$

4.3 Aggregation

In order to aggregate the gains from the removal of within-industry misallocation from the S industries, we assume the existence of a single final good produced by a representative firm in a perfectly competitive final good market. The firm combines the output Y_s of the S different sectors using a Cobb-Douglas production function:

$$Y = \prod_{s=1}^{S} Y_s^{\theta_s},$$

where $\sum_{s=1}^{S} \theta_s = 1$ are the shares of aggregate output for the S industries. Cost minimization implies that:

$$\theta_s = \frac{P_s Y_s}{PY},$$

where P_s is the price index for industry output Y_s, and $P \equiv \prod_{s=1}^{S} (P_s/\theta_s)^{\theta_s}$ is the price index of the final good and is set equal to 1. We can express aggregate output as a function of the sector-level production functions:

$$Y = \prod_{s=1}^{S} (TFP_s \cdot K_s^{\alpha_s} \cdot L_s^{1-\alpha_s})^{\theta_s},$$

which implies that

$$\frac{Y^*}{Y} = \prod_{s=1}^{S} \left(\frac{TFP_s^*}{TFP_s} \right)^{\theta_s}. \tag{12}$$

Thus we are able to geometrically average the gains from removing misallocation across the S industries to arrive at an aggregate figure.

4.4 Measurement Issues

As stated, the measures used require data only on nominal output, $P_i Y_i$, not real output, Y_i, and do not require knowledge of price deflators. We infer real output from revenue and our assumed elasticity of demand. Here we follow Hsieh and Klenow and Oberfield and use the wage bill rather than its employment level to measure L_i. This relies on the assumption that firms in the same industry face the same wages for workers of given quality, however, it can be defended on the basis that the measure can account for differences in the quality of each firm's labour force, and also differences in the number of hours worked. Tangible fixed assets are used for the measure of capital, K_i.

The basic approach to the calculation of $\{\alpha_i\}$ is to use a cost shares type method. Cost-minimization under a Cobb-Douglas production function, given perfectly competitive factor markets, implies that factor expenditure will be a constant share of total costs, and that this share will equal the exponent on that factor. In the Amadeus sample, I do see information on total costs, however it is not present for all countries nor for all types of firm, according to the requirements on filing balance sheet information in the respective countries. Therefore, I prefer to calculate cost-shares relative to value-added instead. Though this would be valid under the assumption of perfect competition, I have of course already assumed monopolistically competitive markets. I therefore need the assumption that costs are a constant multiple of revenue along with an assumption that profits are distributed prorata to each of the factors. In my use of value-added cost shares I follow Hsieh and Klenow (2009), and also Arnold and Schwellnus (2008), who compute cost-shares using Amadeus data.

An issue is that we may say that a firm has deviated from optimal production on account of misallocation frictions, simply because we have got the parameters of its production function wrong, and have not accounted for the unique production arrangements that firm represents. Conversely, were we to try to calculate the parameters of the production function by looking at observed factor shares, we might end up with the wrong coefficients because the observed factor shares in the data are distorted by various frictions.

In the original paper of Hsieh and Klenow this identification problem is dealt with under the assumption that α_i is homogeneous across each industry, that it is fixed over time, and further that U.S. factor shares represent the "true" production parameters for given industries – they thus assume that the U.S. data is free from misallocation frictions entirely. The authors are of course not actually committed to the view that the U.S. economy is frictionless, they are merely using the U.S. shares as a benchmark case. And in fact, the benchmark chosen is well suited to the objective of their study: to compare China and India to the U.S. in terms of misallocation. U.S. shares are the appropriate benchmark because they believe that misallocation is more likely to be a problem in China and India than in the U.S. Oberfield also works with U.S. shares in his analysis of the Chilean economy. It is less clear in what sense the U.S. shares represent an appropriate benchmark, since Oberfield does not actually represent a comparison with the U.S. economy.[22] In this study the labour share is simply set equal to that of the industry in question in a given countries.[23] The effect of misspecified production parameters on results ought to be limited, given the focus of the analysis on the variation of misallocation measures over time.

Following Hsieh and Klenow, the elasticity of substitution between firms in the same industry, σ, is set equal to 3. Hsieh and Klenow refer to estimates of this parameter in the

[22] Oberfield does consider alternate assumptions.

[23] The capital share is backed out under the CRS assumption.

trade and industrial organization literatures that range from 3 to 10, and so settled for a conservative measure. Aside from the particular value chosen, the fact that the elasticity of substitution is homogeneous across industries is a strong simplifying assumption. Further, and particularly relevant for any dynamic application of the methodology, it is assumed that the elasticity of substitution is fixed over time. Thus the possibility that consumer demand may change in nature from boom periods to recessions is not accounted for.

I also follow other papers and guard against the influence of measurement error by restricting outliers. I Winsorize the top and bottom 1% of the distributions of capital and labour wedges. This means that, for example, if a firm's capital wedge falls in the top 1% of the distribution, its capital stock is replaced in that year to give it a new capital wedge at the 99th percentile. I choose to Winsorize the data, instead of trimming it, so as not to lose observations.

4.5 Results

Reallocation gains across time for the manufacturing industry, computed according to Equation (12), are presented in Figures 1 and 2. What is apparent is that 8 of the 13 countries display upward movements in the misallocation measures across the sample period (Austria, Belgium, Bulgaria, Czech Republic, Germany, Italy, Ukraine and UK), although these movements can be non-monotonic, and may on occasion partially reverse in the most recent data (Czech Republic, Italy).

In general the time-series are more volatile than one might expect, given that the misallocation measures were developed to quantify the effects of structural characteristics of economies on productivity, as opposed to cyclical phenomena. Of the 13 economies, 9 do appear to show sharp downward movements in the misallocation measures during crisis years of 2008-2010 (Austria, Belgium, Bulgaria, Germany, Spain, Italy, Sweden, Ukraine and the UK). However, similar movements can be observed in other years of the sample, and are often followed by sharp corrections.

The time-series of reallocation gains for the services industry are presented in Figures 3 and 4. Again, we see relatively large movements around an upward trend in most cases. Sweden and Romania both display upward movements for services, while this was not obviously the case for their manufacturing sectors. The services data from the Czech Republic also display a peak midway through the time-series, followed by a reversal, as was the case for their manufacturing sector. The upward movements in the Italian services data do not revert in the final period, as before in the manufacturing case. However, speaking generally, the upward movements in the misallocation measures would seem to hold across both sectors.

In general, from inspection of the graphical evidence, there is less evidence for sharp negative movements in the misallocation series during crisis years. There are some shallow falls in misallocation during the 2008-2010 period in Sweden and Ukraine, and a stronger fall in German data. However, these falls are often preceded by steeper falls in previous periods, after pre-crisis peaks in estimated misallocation (for example the Belgian, Bulgarian, Czech, Spanish, Ukrainian and UK cases). Such movements do not translate obviously into a narrative linking the financial crisis to cyclical movements in misallocation, for the services industry at least.

The relation between the capital wedge and firm size for the manufacturing sample are presented in Figures 5 and 6. Scatter plots are shown for the sample restricted to 2007, and to 2012, in order to assess the possibility that the financial crisis led to a change in the relation between the wedges and firm size. In general smaller manufacturing firms face a higher capital wedge, potentially implying they face difficulties in their ability to rent capital relative to large firms (the relation is less clear in Belgium, Spain, Italy and Romania). This would be consistent with a financial frictions mechanism that requires younger (and most likely smaller) firms to pay higher interest rates, to compensate banks for their inability to supply lengthy credit histories. Admittedly the capital wedge in fact tracks the price of capital relative to labour, it could also be that smaller firms face lower labour costs. From investigation of the graphs, there is little obvious change in the relation between firm-size and the capital wedge in the pre- and post-crisis period.

The relation between the output wedge and firm size for the manufacturing sample is presented in Figures 7 and 8. Since the Figures plot $\log(1 - \tau_Y)$, the dependent variable should be understood to be a "subsidy" to certain firms, relative to the frictionless optimum. The relationship is generally negative, suggesting that smaller firms in these economies benefit from distortions that allow them to exceed their optimal size, relative to larger firms. The implication of this would be that the welfare of these economies would increase were the smaller firms to get smaller, and the larger firms to get larger, *ceteris paribus*. Again, there is little evidence of any changes in the relation between firm-size and the output wedges, in the before- and after- financial crisis samples.

The capital-wedge firm size relation for services is shown in Figures 9 and 10. We no longer see an obvious negative relation between firm size and the capital wedge. In fact the relation is positive for Belgium, Spain, Italy, and Romania. This could reflect the reduced importance of physical capital in the services sector, insofar that differences in the capital wedge between firms would have more to do with differences in the abilities of firms to hire labour. Those countries with a positive relation may reflect economies in which larger firms find it much easier to hire the specialized labour they need, for example. Again we do not see large shifts in the size relation in the pre- and post- crisis sample.

The relation between the capital wedge and firm size for the services sample are presented in Figures 11 and 12, and as in the manufacturing case is generally negative, although this is not robust to all cases. The relation is also typically shallower than that for manufacturing, and again we do not see evidence for a change in the relation between size and firm-level output distortions after the financial crisis in the services data.

5 Panel Analysis

Because of the strong functional form assumptions the misallocation analysis uses, it is also worthwhile to investigate reallocation processes using a reduced form approach which places fewer restrictions on the data. In order to do this, a dynamic regression approach is employed, based on the analysis found in Foster et al. (2013).

The basic intuition behind these regressions is that if recessions do induce cleansing in the intensive (or shift-share) sense, we should see unproductive firms shrinking in size relative to the productive ones. We could examine this in the context of a growth regression by including an interaction term between the productivity of firms and a business cycle indicator, in order to see whether unproductive firms are more responsive to cyclical downturns. I therefore employ a fixed-effects regression model relating growth to productivity, the business cycle and their interaction, of the following form

$$\begin{aligned}\Delta x_{isct} &= \lambda_i + \lambda_t + \beta * \widetilde{tfpr}_{isc,t-1} + \gamma * Cycle_{sct} \\ &\quad +... + \mathbb{1}^E_{t+1} * \delta_E * \widetilde{tfpr}_{isc,t-1} * Cycle_{sct} \\ &\quad +... + \mathbb{1}^C_{t+1} * \delta_C * \widetilde{tfpr}_{isc,t-1} * Cycle_{sct} + \varepsilon_{isct}, \end{aligned} \tag{13}$$

where i is the firm, c is the country, Δx is the change in a measure of firm size from t to $t+1$, $\widetilde{tfpr}_{isct} \equiv tfpr_{isct} - \overline{tfpr}_{sct}$, is the log of firm-level total factor productivity, $tfpr_{isct}$, in deviation from the subsector-by-country-by-year mean, $\overline{tfp}_{sct}$, and $Cycle$ is the change in the aggregate value of output from industry s from t to $t+1$.[24] I employ separate regressions

[24] $Cycle_{ist} \equiv \Delta \ln PY_{sct}$, $\sum_{i \in \mathcal{F}_{sct}} PY_{ist}$, where $\mathcal{F}_{sct}$ is the subset of data on firms in sector s, for country c, in

with the dependent variable, the measure of firm growth Δx, set equal to (the log change in) added value, capital, employment and wage growth. Because I include year dummies the specification exploits cross-country variation in sector-level production levels as opposed to any generalized trend. With respect to timing, here I am examining the determinants of growth from t to $t+1$ based on the productivity at the establishment in period t and the business cycle conditions from t to $t+1$.

The regression also includes two dummies: $\mathbb{1}_E$ is a dummy variable indicating a positive value of $Cycle$ at $t+1$, $\mathbb{1}_C$ is a dummy variable indicating a negative value. Thus, the coefficient on $Cycle$ given that output is contracting (meaning $Cycle$ is negative) is equal to $(\gamma + \delta_C * tfpr)$. One would expect that γ is positive, since firms should grow in expansions and shrink during booms, *ceteris paribus.* As explained in Hsieh and Klenow (2009), revenue productivity provides a summary measure of the extent to which firms are sub-optimally too-large or too-small, were the hidden wedges to be removed. In the absence of the exogenous wedges, firms with high TFPR would grow, and low TFPR firms would shrink, until all firms had the same revenue productivity. We therefore have a cleansing effect in the intensive sense if δ_C is negative, since this implies that $(\gamma + \delta_C * tfpr) * Cycle$ is *less* negative during contractions as the idiosyncratic $tfpr$ of a firm increases. Note that, were I to omit the dummy variable, and include only $\delta * tfp * Cycle$ in my specification, the coefficient on $Cycle$ is $(\gamma + \delta * tfpr)$ and a negative value of δ would imply that that high TFPR firms shrink by less during recessions than low TFPR firms, but also that they grow slower during booms. Contrariwise, a positive value of δ would imply that high TFPR firms grow by more during expansions, but shrink by more than unproductive firms during recessions. The δ term would become a general coefficient indicating sensitivity to cyclical fluctuations, not an indicator of cleansing effects.

With respect to productivity estimation, I continue to use cost shares relative to value-added. The value of productivity itself is can be calculated according to Equation (7). The measure is therefore one of revenue productivity relative to the firm's peer group. With respect to aggregate data, I use data from the United Nations statistics division, which provide GDP data disaggregated by sector. The variable $Cycle$ is deflated according to the respective country's CPI from Datastream.

Regressions are also run where the $\widetilde{tfpr}$ terms are replaced respectively with $\widetilde{mrpk}$ and $\widetilde{mrpl}$, the log deviations of MRPK and MRPL from their industry-country-year means. This allows investigation of the contribution of each factor of production to any cyclical effects

time period t.

on resource misallocation.

See Table 4 for the results of the dynamic panel analysis, for the manufacturing sample. The coefficients on the lagged dependent variables are always negative and highly significant, implying mean reversion. The coefficient on the lag of log TFPR seems to vary in sign depending on whether the dependent variable is added value or capital growth (for which it is negative, and highly significant) or employment growth (for which it is negative, and significant at 5%).

The coefficient on the growth of value-added in the sector is always positive, but is only significant in the case that value-added growth is the dependent variable (where this would follow almost mechanically, given the definitions of the two variables), and for wages. This suggests the effect of cyclical changes in value-added may have a weak relation on factor growth, in the data. The interaction between TFPR and the expansion dummy is positive, and significant, implying resources are reallocated to high TFPR firms during expansions. This should act to actually ameliorate misallocation, by the logic of Hsieh and Klenow (2009). The contraction term is negative for all cases except the capital growth regression, but only significant for the employment growth regression. For the employment growth case, we have partially significant evidence that high TFPR firms are less sensitive to cyclical downturns, implying that resources are diverted from lower- to higher- TFPR firms, and a reduction in misallocation.

Table 6 displays the results incorporating the marginal revenue product of capital as in the interaction term with respect to the expansion/contraction dummy. In this case the coefficient on the contraction term is negative for all cases except capital, and is significant for added value growth, as well as employment growth. This implies that firms with relatively high MRPK for their industry shrink less during contractions, though we do not attain a significant response of the correct sign with respect to capital growth itself.

Table 8 shows the results from regressions employing the marginal revenue product of labour in interaction with the business cycle variable. We see that, with respect to the regressions with interacted variables, there is a significant ameliorating effect of expansions on misallocation via relatively high value added growth from high MRPL firms. We do not attain significant results for the contraction term, and results can change in sign.

Comparable regressions for services are displayed in Tables 9, 7, and 5, though the coefficients on the interaction terms of interest are in all cases insignificant.

Given the diverse responses of efficiency to the recession I have documented by focusing on the individual countries, it is very likely that the functional form of this equation imposes too much homogeneity across the 13 countries. However, it does at least seem to mirror the fact that the misallocation analysis also generated only very limited evidence of cleansing effects, though there is some evidence for ameliorative effects of recessions on misallocation with respect to value-added growth and employment growth (via a shifting of resources to high TFPR firms, and high MRPK firms), for the manufacturing sample only. Certainly the evidence from the regression analysis leans towards cleansing arguments, as opposed to sullying ones.

6 A Case Study of the U.K.

This section applies similar methods to UK micro-data, in order to evaluate the contribution of within-industry resource misallocation to the UK experience. This provides an external check on the results obtained via the Amadeus dataset, since the UK data are taken from a survey, and is thus constructed under different criteria. The UK is also an interesting case, given academic interest in what has been termed the "productivity puzzle" of this economy, whereby labour productivity and TFP growth has been subdued in post-crisis years.

Barnett et al. (2014) argue that increased dispersion in sectoral gross value added and output price deflators is consistent with increased misallocation of capital in the years following the crisis, using a model of perfect competition. Barnett et al. (2014) also supply corroborating evidence by charting a decrease in the strength of the response of investment to the lagged rate of return on capital post-crisis, using a reduced form approach and the same firm-level survey data employed by this study. The authors argue that this suggests a recent increase in frictions that prevent the market reallocating capital to where returns are highest, implying an increase in capital misallocation.

There are many other candidate explanations for the productivity puzzle, including measurement error, the impact of the crisis on a relatively large financial sector, and bank forbearance towards unproductive firms.[25] Barnett et al. (2014) summarize estimates of the importance of these factors, as well as impaired capital reallocation, and conclude that they have the ability to explain 6-9 percentage points of the 12 percentage point productivity gap in evidence by 2013 Q4 (having already corrected for the contribution of measurement issues). The analysis of this paper furthers our ability to explain productivity dynamics in the UK.

[25] See Bryson and Forth (2015) for a discussion.

6.1 Data

This analysis uses firm-level survey data from the Annual Business Survey (2008-2013), and the Annual Business Inquiry (1997-2007).[26] The Annual Business Survey covers the production, construction, distribution and service industries of the UK economy, excluding Northern Ireland. Selected firms are legally obligated to respond to the survey. The population from which the sample is drawn are those firms that either pay VAT, or operate payroll schemes ("pay-as-you-earn", PAYE). Importantly, the survey is a census of larger firms (250+) employees and a stratified sample of smaller ones. The survey asks firms about their value-added, labour costs, and capital expenditure.[27]

The data do not contain information on the capital stocks of firms, and they are therefore estimated using the perpetual inventory method. Thus, capital stocks are calculated according to the law of motion:

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

where K_t is capital in year t, where δ is an assumed depreciation rate, and I_t is the observed value of capital expenditure for the firm. An initial level of capital is apportioned from sectoral capital stock data according to the share of that firm's materials purchases relative to its industry.[28] For full details see Appendix B.

The survey nature of the ABS and ABI datasets presents certain challenges to the researcher. Because the smaller firms are not sampled each year, we cannot attain a panel of observations for this section of the dataset. This is a necessary prerequisite for the perpetual

[26]The survey changed title in 2008, with some changes to the methodology that ought not to greatly affect key variables of interest (questions for key variables used in the study remain the same). An example of such a change would be the transition from SIC92 to SIC07 industry codes, which is corrected using a correspondence helpfully provided by Jennifer Smith at the University of Warwick. However it is true that the employment variable is affected by a discontinuity-–the ONS began to collect information on employment in a separate questionnaire from 2008, as part of the Business Register and Employment Survey. BRES data are not used in this study to avoid this issue.

[27]Data on the number of employees (as well as more detailed employment information) are available from a separate survey called the Business Register Employment Survey for the post-2007 period, whereas previously these questions were contained within the Annual Business Inquiry. Because of concerns regarding a structural break, these data are not employed in the study. Where employment is used (for certain imputations), it is sourced from the administrative data on employment contained within the ABS (which is from the Inter-Departmental Business Survey).

[28]Broadly, the methodology of Gilhooly (2009) was followed closely, with some adjustments regarding apportionment of initial capital stocks. Although the ABS and ARD do contain disaggregate information on capital stocks, only the aggregate measure was used because of concerns about structural breaks in these variables owing to the ABI to ABS transition. Grateful acknowledgement is made of the use of Stata do files provided by Bob Gilhooly.

inventory method, which we need to use since capital stocks are not observed. Because of this issue, large-scale imputation of missing values of capital is necessary if we are to analyse smaller firms (less than 250 employees). However, at the baseline, this study restricts attention to a balanced panel of firms that excludes almost all of these observations.[29]

In order to assess representativeness, Figure 13 compares the growth rate of labour productivity for the sample with that of the raw dataset, and those for the dataset restricted to 250+ firms. Labour productivity in the data is computed as totaled real value added divided by totaled employment, where employment is the point-in-time number of employees.[30] The growth rates for the manufacturing and services sectors are compared to the aggregate growth rates of output per job for the same sectors. We see that the levels of growth in the data are higher than for the ONS aggregate data. The dataset does seem to show relatively lower growth rates in the 2008/09 years, reflecting the onset of the Great Recession. A degree of bounce-back is observed in the aggregate and survey data for 2010, before lower figures are recorded for 2011 and 2012. We do see large growth rates in 2013, which suggests increased divergence between the data and aggregate figures for this year.[31]

An ambiguous association between ARD/ABS sample data aggregates and the associated figures from macro-data, although a little discomfiting, broadly matches the findings of Barnett et al. (2014), who also find "productivity growth in the sample is stronger over the full sample period" when compared both to aggregate data sources (and even to official ONS estimates from the same data). In general, we can tentatively conclude that broad movements in labour productivity are followed by the sample data with some noise, although the divergence of estimates for 2013 from official data is notable.

6.2 Results

The time-series of reallocation gains, under the restriction to a balanced panel, is presented in Figure 14. What is immediately apparent is that we can see a clear positive trend in re-

[29] Imputation from employment is conducted for large firms with missing values, however here the problem is less severe. Firms with more than 50% imputed values are excluded from the analysis.

[30] The point-in-time employment comes from the IDBR dataset, which is an administrative dataset that also forms the population register for the survey (meaning administrative data are available for the firms that were surveyed and the firms that were not). Because the IDBR data come from the version of the dataset at the time the survey was issued, it has the disadvantage of being temporally earlier information than the value added information. While the survey data also asks for employment information, and is thus a more consistent source of employment data for productivity calculations, it is affected by a structural break in 2008, so is not used.

[31] There is a discrepancy in the Services data, insofar as I do not include financial firms in the data, whereas the ONS do, this may affect comparability.

allocation gains over the 2002 to 2013 period, indicating that rises in the misallocation of labour and capital were in place before the financial crisis, and then seem to have accelerated post-crisis somewhat. The result may be surprising, given the relatively robust TFP growth pre-crisis, but would imply that TFP growth in the UK could have been higher, in the absence of factors leading to resource misallocation before 2008. The movements are thus comparable to the experiences of Portugal, Spain and Italy, reported in Dias et al. (2015), García-Santana et al. (2016), and Gopinath et al. (2017) respectively, and less so to the experiences of the Northern European economies reported in Gopinath et al. (2017). This is an interesting result that shows that such movements may not be restricted to Southern European data.

The level of estimates are higher for services than for manufacturing, reflecting previous studies.[32] Other than genuinely higher levels of resource misallocation in the services sector, a reason for this may be increased model misspecification for this case, as well as relatively more measurement error for capital stocks. The reallocation gains in the manufacturing industry are comparable to estimates for the U.S. recorded in the original paper by HK, who report figures of 30-40%.[33] This is to be expected, given the similarities between the U.K. and U.S. economies. Although HK did not examine the services data, the level of reallocation gains for the services sector are also comparable, albeit generally smaller, than those reported in Dias et al. (2015), who give figures between 58 and 91%.

The plotted growth rates in the right panel of Figure 14 look very large, and seem unexpectedly volatile. The final growth rate for manufacturing appears anomolous, and may potentially be driven by the same issues that drive sample productivity growth rates away from aggregate statistics in 2013.

7 Conclusion

This paper has studied changes in measures of misallocation across 13 different economies during the period of the financial crisis and subsequent recessions. Data was used from Amadeus, as well as from the ARD database of the UK, in order to corroborate observed movements.

[32]This is also a finding of Dias et al. (2015), and of García-Santana et al. (2016).

[33]Strictly, the balanced panel restriction would probably lower estimates relative to estimates where this restriction is not imposed, so recorded levels may be higher under a closer comparison of datasets. On the other hand, the use of the perpetual inventory method will almost certainly boost misallocation measures relative those that use capital stock data, as HK do.

Generally the evidence favours the cleansing effect of recessions, as opposed to the sullying effect. However the movements in the misallocation quantified rarely tell a clear or consistent story. Regression analysis is employed to econometrically test the cleansing and sullying hypotheses. Results indicate that changes in value-added and employment were induced by the recession, in a way that reduced the overall level of misallocation *ceteris paribus*, for the manufacturing sample. Future work would benefit from a tighter link between the movements and patterns uncovered, and changes in the distribution of the reduced-form wedges that give cause to misallocation in the Hsieh and Klenow (2009) framework. In particular it would be useful to understand why firms with higher MRPK are able to better maintain value-added and employment growth, while the same is not true for high MRPL firms.

References

Aghion, P., P. Askenazy, N. Berman, G. Cette, and L. Eymard (2012). Cyclicality of R&D Investment: Evidence from France. *Journal of the European Economic Association 10*(5), 1001–1024.

Aghion, P. and G. Saint-Paul (1998). Virtues of Bad Times: Interaction Between Productivity Growth and Economic Fluctuations. *Macroeconomic Dynamics 2*, 322–344.

Arnold, J. and C. Schwellnus (2008). Do Corporate Taxes Reduce Productivity and Investment at the Firm Level? Cross-Country Evidence from the Amadeus Dataset. Technical report, CEPII.

Banerjee, A. V. and E. Duflo (2005). Growth Theory through the Lens of Development Economics. In P. Aghion and S. N. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 1A, Chapter 7, pp. 473–552. Elsevier B.V.

Barlevy, G. (2002). The Sullying Effect of Recessions. *The Review of Economic Studies 69*(1), 65–96.

Barlevy, G. (2003). Credit market frictions and the allocation of resources over the business cycle. *Journal of Monetary Economics 50*(8), 1795–1818.

Barlevy, G. (2007). On the Cyclicality of Research and Development. *American Economic Review 97*(4), 1131–1164.

Barnett, A., S. Batten, A. Chiu, J. Franklin, and M. Sebastia-Barriel (2014). The UK productivity puzzle. Technical Bulletin Q2, Bank of England.

Barnett, A., B. Broadbent, A. Chiu, J. Franklin, and H. Miller (2014). Impaired capital reallocation and productivity. *National Institute Economic Review 228*(1), R35–R48.

Barnett, A., A. Chiu, J. Franklin, and M. Sebastiá-Barriel (2014). The productivity puzzle: a firm-level investigation into employment behavior and resource allocation over the crisis. Technical report, Bank of England.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2009). *Producer Dynamics: New Evidence from Micro Data*, Chapter Measuring and Analyzing Cross-Country Differences in Firm Dynamics, pp. 15–76. NBER Book. University of Chicago Press.

Bryson, A. and J. Forth (2015). The UK's Productivity Puzzle. Discussion Paper 448, NIESR.

Caballero, R. J. and M. L. Hammour (1996). On the Timing and Efficiency of Creative Destruction. *The Quarterly Journal of Economics*.

Caballero, R. J. and M. L. Hammour (2005). The Cost of Recessions Revisited: A Reverse-Liquidationist View. *Review of Economic Studies 72*(2), 313–341.

Calvo, G. A., A. Izquierdo, and E. Talvi (2006). Sudden Stops and Phoenix Miracles in Emerging Markets. *The American Economic Review 96*(2), 405–410.

Chari, V. V., J. Kehoe, and E. R. McGrattan (2007). Business Cycle Accounting. *Econometrica 75*(3), 781–836.

Comin, D. and S. Mulani (2009). A theory of growth and volatility at the aggregate and firm level. *Journal of Monetary Economics 56*(8), 1023–1042.

Cooper, R. and J. Haltiwanger (1993). The Aggregate Implications of Machine Replacement: Theory and Evidence. *The American Economic Review 83*(3), 360–382.

Da Rin, M., M. Di Giacomo, and A. Sembenelli (2011, October). Entrepreneurship, firm entry, and the taxation of corporate income: Evidence from Europe. *Journal of Public Economics 95*(9-10), 1048–1066.

Davis, S. J., J. C. Haltiwanger, and S. Schuh (1998). *Job Creation and Destruction*. MIT Press.

Dias, D. A., C. R. Marques, and C. Richmond (2015). Misallocation and Productivity in the Lead Up to the Eurozone Crisis. International Finance Discussion Papers 1146, Board of Governors of the Federal Reserve System.

Eslava, M., A. Galindo, M. Hofstetter, and A. Izquierdo (2010). Scarring Recessions and Credit Constraints: Evidence from Colombian Firm Dynamics.

Foster, L., C. Grim, and J. Haltiwanger (2013). Reallocation in the Great Recession: Cleansing or Not?

Gamberoni, E., C. Giordano, and P. Lopez-Garcia (2016). Capital and labour (mis)allocation in the euro area: some stylized facts and determinants.

García-Santana, M., E. Moral-Benito, J. Pijoan-Mas, and R. Ramos (2016). Growing like Spain: 1995-2007. Working Paper Series 888, Barcelona GSE.

García-Santana, M. and R. Ramos (2015). Distortions and the size distribution of pplant: evidence from cross-country data. *SERIEs 6*, 279–312.

Gilhooly, B. (2009). Firm-level estimates of capital stock and productivity. *Economic and Labour Market Review 3*(5), 36–41.

Gopinath, G., S. Kalemli-Ozcan, L. Karabarbounis, and C. Villegas-Sanchez (2017). Capital allocation and productivity in south europe. *Quarterly Journal of Economics 132*(4), 1915–1967.

Hallward-Driemeier, M. and B. Bijkers (2013). Do crises catalyze creative destruction. *The Review of Economics and Statistics.*

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Inklaar, R., A. A. Lashitew, and M. P. Timmer (2017). The Role of Resource Misallocation in Cross-Country Differences in Manufacturing Productivity. pp. 733–756.

Kehrig, M. (2015). The cyclical nature of the productivity distribution.

Klapper, L., L. Laeven, and R. Rajan (2006, December). Entry regulation as a barrier to entrepreneurship. *Journal of Financial Economics 82*(3), 591–629.

Lin, Y.-C. and T.-H. Huang (2012). Creative destruction over the business cycle: a stochastic frontier analysis. *Journal of Productivity Analysis 38*(3), 285–302.

Meza, F. and E. Quintin (2006). Financial Crises and Total Factor Productivity. Technical report, Federal Reserve Bank of Dallas.

Mortensen, D. T. and C. A. Pissarides (1994). Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies 61*, 397–415.

Mustre-del Río, J. (2012). Job Duration and the Cleansing and Sullying Effects of Recessions. Research Working Paper 12-08, Federal Reserve Bank of Kansas City.

Oberfield, E. (2013). Productivity and misallocation during a crisis: Evidence from the Chilean crisis of 1982. *Review of Economic Dynamics 16*(1), 100–119.

Osotimehin, S. and F. Pappadà (2016). Credit frictions and the cleansing and sullying effect of recessions. *The Economic Journal*.

Ottaviano, G. I. P. (2012). Firm heterogeneity, endogenous entry, and the business cycle. *NBER International Seminar on Macroeconomics*.

Ouyang, M. (2009). The scarring effect of recessions. *Journal of Monetary Economics 56*(2), 184–199.

Peters, B. (2007). Persistence of innovation: stylised facts and panel data evidence. *The Journal of Technology Transfer 34*(2), 226–243.

Pratap, S. and C. Urrutia (2012). Financial frictions and total factor productivity: Accounting for the real effects of financial crises. *Review of Economic Dynamics 15*(3), 336–358.

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics 11*(4), 707–720.

Restuccia, D. and R. Rogerson (2013). Misallocation and productivity. *Review of Economic Dynamics 16*(1), 1–10.

Saffie, F. E. and S. T. Ateş (2013). Fewer but Better: Sudden Stops, Firm Entry, and Financial Selection.

Schivardi, F., E. Sette, and G. Tabellini (2017). Credit misallocation during the european financial crisis. Technical report, CEPR.

Schumpeter, J. A. (1942). *Capitalism, socialism, and democracy.* New York: Harper.

A Data Cleaning

- I drop entries that are missing the Bureau van Dijk identification code.
- I restrict the analysis to the following sectors: manufacturing (NACE 3), construction (NACE 6), wholesale and retail trade (NACE 7), accommodation and food service activities (NACE 8), information and communication (NACE 9).
- If firms switch industries, I use the modal industry for each plant. If there is a tie, I use the earlier industry.
- I deflate the nominal variables with the consumer price index for each country.
- I drop observations if the following variables are missing: value added, tangible fixed assets, employment.
- I drop observations if the following variables are less than or equal to zero: value added, tangible fixed assets, employment.
- I drop consolidated accounts.
- I keep only public and private legal forms.
- I restrict attention to industries for which there are at least 20 observations for each year in the sample.

B Capital in the ARD

This project uses data from the Annual Business Survey, which was called the Annual Business Database prior to 2008. The ABS does not contain information on the capital stock of firms. In order to attain a measure of firm-level capital stocks, one must estimate them using the Perpetual Inventory Method (PIM). We assume that the firm-level capital stock evolves according to the following equation of motion:

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

where K is capital, I is investment, and δ is the rate of depreciation.

In order to run the PIM, one needs to initialize the calculation with a starting value. Since no information on initial capital stocks is available, a proportion of the level of sectoral capital is allocated to firms. These data come from the Volume Index of Capital Services (VICS). I make the capital stock commensurate with the size of the sample. I broadly follow the procedure detailed in documentation to the ARD Capital Dataset.

1. I first compute the share of investment in the sample for a given sector, out of the aggregate data on investment from Gross Fixed Capital Formation data,

$$\varphi_{jt} = \frac{|\sum_{i \in jt} I_{jt}|}{|I_{jt}|},$$

   where i is the firm, and j is the sector.

2. I then use materials purchases to allocate the scaled measure of sectoral capital. I do this in several steps.

   (a) I calculate the material share for each firm, using data on materials expenditure:

$$M_{ijt} = \frac{m_{ijt}}{\sum_{i \in jt} m_{ijt}},$$

   (b) I then average the material shares across the first three years of the firm's life

$$\overline{M}_{ijt} = \frac{\sum_{i,t \in \{b_i, b_i+2\}} M_{ijt}}{3},$$

   where b_i is the year firm i enters the dataset. The reason I average across years is that I want to reduce noise, but do not want to average across *all* years of a firm's life, since this would only be valid if firms maintain a constant size relative to their industry for their entire lifetime;

   (c) on account of entry and exit from the sample, these average shares do not necessarily sum to 1 in a given year, so it is then necessary to scale them by the summation of material shares:

$$\gamma_{ijt} = \frac{\overline{M}_{ijt}}{\sum_{i \in jt} \overline{M}_{ijt}}.$$

3. Then, the initial value of capital for a firm i born in year 1, is given by the following:

$$K_{ij1} = K_{j1} * \varphi_{j1} * \gamma_{ij1} + I_{j1}.$$

4. The estimate of the capital stock for the rest of the firm's life is given by:

$$K_{ijt} = (1 - \delta)K_{ij,t-1} + I_{ij,t}.$$

It should also be noted that I impute missing values of investment from observations on employment, using the average relationship between investment and employment for each firm. The employment variable was itself linearly interpolated, however the number of interpolations ought to be relatively low, since employment comes from the administrative data in the IDBR. Firms with more than 50% imputations are dropped from the analysis.

Table 1: Coverage Table (%)

		MANUFACTURING					*SERVICES*				
Country	Variable	2008	2009	2010	2011	2012	2008	2009	2010	2011	2012
AT	Employees	24.1	23.89	23.67	28.12	24.23	23.95	23.67	22.57	28.05	24.82
	Revenue	28.61	26.83	29.38	33.18	29.1	25.09	23.89	28.44	30.03	24.95
BE	Employees	45.85	45.02	45.55	47.4	50.46	33.06	34.45	32.18	32.77	34.63
	Revenue	62.95	70.55	62.39	66.23	73.72	44.8	43.33	41.87	40.77	42.73
BG	Employees	16.78	17.41	21.3	30.2	35.89	28.69	29.15	33.82	39.57	49.72
	Revenue	19.79	19.03	23.34	38.95	51.38	49.56	49.74	60.64	63.87	69.27
CZ	Employees	35.91	47.54	49.88	52.3	47.02	34.55	43.23	47.25	51.7	42.61
	Revenue	34.58	45.83	49.05	48.45	46.32	38.55	48.85	51.09	48.23	40.66
DE	Employees	17.57	19.51	19.18	19.27	15.46	12.97	13.03	13.28	13.27	10.54
	Revenue	20.43	21.52	20.87	20.49	17.01	14.53	15.31	16.25	16.16	12.89
ES	Employees	24.49	27.03	28.85	30.13	28.49	28.11	29.64	30.76	31.14	26.81
	Revenue	33.66	36.55	38.24	39.34	37.01	42.55	44.2	44.47	43.3	39.86
FR	Employees	3.04	2.91	2.8	2.88	1.84	8.43	6.44	5.73	5.31	3.89
	Revenue	4.28	3.68	4.07	4.09	2.29	7.61	7.47	7.73	6.14	4.85
IT	Employees	20.38	20.89	23.03	24.74	24.87	18.72	19.71	20.52	21.77	20.02
	Revenue	30.9	31.55	31.71	34.25	34.5	23.05	26.4	23.62	24.99	25.09
PL	Employees	10.61	12.83	11.64	7.78	3.8	21.37	27.01	21.03	15.92	7.65
	Revenue	9.31	16.27	14.59	7.9	4.46	18.88	45.37	21.9	12.67	7.24
RO	Employees	32.69	34.17	37.74	37.96	40.37	45.2	43.62	44.04	45.05	44.85
	Revenue	35.98	38.86	41.96	44.23	46.5	55.92	67.21	62.9	62.56	64.75
SE	Employees	14.17	14.78	15.05	16.11	16.51	22.2	24.37	24.15	24.51	24.6
	Revenue	14.37	17.64	18.27	17.91	17.8	21.57	26.31	28.42	27.62	27.42
UK	Employees	39.91	42.23	41.99	44.66	44.72	34.41	39.11	38.75	39.88	39.19
	Revenue	39.81	46.93	57.82	50.3	48.68	27.19	34.1	33.86	33.59	31.09

Notes: Table displays percentage coverage of totals computed from (cleaned) AMADEUS data relative to the appropriate (NACE r2.) sectoral aggregates available from Eurostat. Eurostat data at the correct level of disaggregation are not available prior to 2008.

Table 2: Summary Statistics

	MANUFACTURING							*SERVICES*						
	Value Added		Capital		Employees			Value Added		Capital		Employees		
Country	Mean	Sd	Mean	Sd	Mean	Sd	N	Mean	Sd	Mean	Sd	Mean	Sd	N
AT	23.59	56.71	18.42	80.34	255.34	422.08	3,041	14.15	36.64	11.15	49.43	207.53	875.88	5,118
BE	23.9	113.73	12.69	46.95	189.28	454.48	12,253	8.28	29.81	3.85	17.32	85.91	224.66	26,733
BG	2.62	10.91	4.77	12.5	292.88	352.38	3,461	2.7	6.28	5.09	23.22	191.96	719.54	5,334
CZ	7.91	21.37	11.79	36.03	321.02	514.13	13,792	5.32	17.69	8.45	86.76	195.97	678.41	13,372
DE	19.74	100.14	11.68	52.12	259.58	851.17	36,464	11.81	35.41	7.02	28.21	182.23	409.17	43,991
ES	12.04	38.23	12.8	44.21	167.66	485.7	32,700	10.43	106.92	11.34	124.65	156.81	1,035.66	69,971
FR	19.21	108.09	13.67	91.39	215.66	479.27	4,667	11.22	66.89	8.81	122.18	157.41	1,084.8	22,234
IT	16.77	48.62	14.83	47.62	211.51	1129.95	34,998	14.27	85.92	28.17	830.57	193.27	1,058.41	32,501
PL	7.36	15.33	11.57	28.06	306.84	429.13	6,451	5.05	61.27	5.81	40.01	200.17	913.99	12,467
RO	4.45	16.57	9.7	38.04	403.27	816.49	9,059	4.25	20.01	8.91	50.7	245.62	878.48	12,881
SE	12.14	39.46	12.73	53.61	139.2	348.82	6,748	6.59	18.5	2.11	9.32	87.91	312.69	19,960
UA	6.63	31.87	11.49	46.77	941.3	2,996.92	8,945	3.59	20.17	7.11	61.83	397.39	1,823.95	7,509
UK	23.76	118.65	17.46	76.92	318.63	914.91	32,102	23.93	147.03	21.32	245.57	476.69	4,053.55	56,808

Notes: Table displays means and standard deviations computed from (cleaned) AMADEUS data for value added, capital (measured by tangible fixed assets), and employees. For each country and sector, statistics are computed on a pooled sample over years.

Table 3: Summary Statistics by Year

		MANUFACTURING							*SERVICES*						
		Value Added		Capital		Employees			Value Added		Capital		Employees		
Country	Year	Mean	Sd	Mean	Sd	Mean	Sd	N	Mean	Sd	Mean	Sd	Mean	Sd	N
AT	2007	20.68	36.67	14.06	56.4	211.27	205.05	190	14.42	34.86	12.53	59.8	192.23	494.67	534
	2009	20.19	48.72	17.59	77.48	256.63	441.55	548	13.41	35.56	11.62	47.86	217.29	965.56	833
	2011	24.01	56.09	18.74	78.55	247.8	404.08	669	14.19	36.38	10.02	40.3	210.69	998.3	1,032
BE	2005	24.41	102.02	12.23	39.12	199.97	458.59	1,107	7.69	25.17	3.49	15.71	82.98	156.1	2,246
	2007	27.67	138.16	13	47.14	195.42	472.24	1,265	9.12	37.38	3.82	16.48	83.33	160.79	2,709
	2009	24.51	156.94	13.05	49.95	183.43	438.77	1,288	8.33	26.15	4.18	18.34	87.3	251.71	2,849
	2011	23.16	107.65	12.73	50.95	178.35	450.41	1,339	8.39	24.38	4.24	20.44	86.83	226.55	3,172
BG	2005	1.67	2.2	2.77	4.01	286.56	324.5	201	1.75	3.06	3.66	9	168.94	294.31	345
	2007	2.41	3.98	3.97	7	293.93	314.94	366	2.75	5.39	4.86	17.58	262.08	1,916.25	580
	2009	1.89	3.35	4.37	11.65	246.72	250.42	392	2.47	5.12	5.12	25.78	172.72	247.19	554
	2011	3.52	9.81	7.02	19	303.16	426.34	514	3	6.88	6.17	27.42	185.57	376.71	772
CZ	2005	6.23	19.15	9.31	30.91	296.02	538.25	1,128	3.43	9.33	3.96	24.21	161.97	610.68	1,065
	2007	8.42	24.55	11.53	34.14	314.28	505.1	1,553	5.51	16.04	7.59	87.63	188.06	582.11	1,468
	2009	7.34	15.59	12.59	38.01	314.33	487.34	1,635	5.95	20.36	10.15	90.52	188.91	578.05	1,605
	2011	9.3	23.74	13.46	38.3	342.04	482.36	1,647	6.18	21.29	9.74	91.04	221.03	630.42	1,664
DE	2005	25.03	104.04	15.74	78.64	325.46	1,082.44	1,864	15.87	48.12	9.45	34.87	221.33	419.57	1,887
	2007	18.39	56.55	10.23	38.72	239.9	568.56	4,610	11.05	35.51	6.27	25.67	171.3	327.86	5,434
	2009	16.03	51.81	10.67	40.26	249.78	1,043.95	5,124	10.89	37.6	6.64	27.35	173.56	346.19	6,207
	2011	18.29	54.27	10.94	46.51	244.47	660.56	5,492	11.3	26.88	6.45	26.47	174.87	358.96	6,926

Table 3: Summary Statistics by Year

		MANUFACTURING							*SERVICES*						
		Value Added		Capital		Employees			Value Added		Capital		Employees		
Country	Year	Mean	Sd	Mean	Sd	Mean	Sd	N	Mean	Sd	Mean	Sd	Mean	Sd	N
ES	2005	11.7	40.38	11.99	44.65	168.59	506.99	3,348	9.57	101.69	10.69	122.47	148.26	1,020.68	7,205
	2007	13.97	45.74	13.03	45.7	176.59	509.24	3,225	11.92	120.01	12.39	125.59	162.16	1,084.18	7,059
	2009	11.1	32.49	13.56	43.14	161.86	452.43	3,321	10.82	114.88	12.54	121.12	161.08	1,024.28	7,251
	2011	12.04	35.79	13.5	42.23	162.01	447.46	3,394	10.57	99.53	9.73	106.83	166.17	1,063.22	7,207
FR	2005	16.36	41.98	9.53	34.35	218.33	492.65	572	10.03	60.11	8.51	131.86	154.34	1117.9	2,655
	2007	24.21	136.85	14.51	98.53	213.2	433.8	527	11.08	46.34	7.69	112.87	153.56	586.93	2,386
	2009	16.45	37.14	14.9	62.82	220.45	466.3	404	12.21	80.74	7.6	83.47	165.96	1238.3	2,056
	2011	21.39	60.86	21.51	122.35	229.1	547.91	374	11.86	70.75	10.94	170.54	160.18	1253.98	1,912
IT	2005	14.95	35.12	12.45	44.76	198.13	562.47	3,257	13.81	83.33	31.59	1002.23	193.01	1,315.99	2,869
	2007	18.33	54.66	13.69	50.63	207.74	563.87	3,580	15.11	86.77	32.07	1011.9	201.32	1,275.9	3,355
	2009	15.6	43.06	16.54	54.05	205.33	554.19	3,646	15.38	117.79	35.73	1016.07	204.12	1,067.76	3,441
	2011	18.25	62.28	16.97	47.88	212.82	660.39	3,906	14.81	81.29	25.7	590.53	203.14	944.35	3,887
PL	2005	6.43	15.11	10	25.69	294.63	424.8	625	3.95	11.64	5.82	40.21	209.76	953.23	1,125
	2007	8.91	17.58	13.21	30.16	332.48	483.25	687	5.6	19.59	6.89	43.18	206.52	980.36	1,340
	2009	7.96	16.26	12.9	28.49	305.79	427.55	935	8.23	151.15	6.97	47.69	214.01	1,049.83	1,911
	2011	7.07	12.45	11.07	25.26	281.52	324.19	605	4.71	21.21	5.63	34.72	212.87	1,154.49	1,164

Table 3: Summary Statistics by Year

		MANUFACTURING							*SERVICES*						
		Value Added		Capital		Employees			Value Added		Capital		Employees		
Country	Year	Mean	Sd	Mean	Sd	Mean	Sd	N	Mean	Sd	Mean	Sd	Mean	Sd	N
RO	2005	2.85	11.87	5.82	35.72	374.8	943.31	720	2.02	5.92	3.38	16.65	172.23	473.27	1,100
	2007	5.06	19.52	10.54	40.41	451.91	843.27	1,031	5.24	25.15	9.85	50.8	282.9	1,146.88	1,416
	2009	4.17	12.52	10.45	32.8	379.57	679.87	1,067	5.29	24.84	10.59	47.69	270.94	884.65	1,492
	2011	5.35	16.22	11.95	38.71	398.98	735.27	1,102	5.17	19.19	12.15	66.31	272.61	923.94	1,558
SE	2005	10.6	25.34	11.49	51.39	125.46	224.77	604	5.24	17.03	1.68	6.46	75.99	325.85	1,788
	2007	13.13	32.43	12.49	49.4	135.25	237.77	638	6.47	18.7	1.84	7.6	85.1	317.91	1,959
	2009	9.75	24.39	11.52	44.61	124.47	217.52	708	6.15	17.48	1.97	8.49	89.4	294.34	2,212
	2011	13.05	32.12	13.07	46.93	129.45	219.45	741	8.23	19.96	2.56	11.54	93.6	274.27	2,301
UA	2005	6.84	31.33	10.43	40.91	1033.74	3,291.22	969	2.49	11.21	6.04	76.17	305.02	1,159.99	722
	2007	8.51	43.82	12.42	49.31	952.7	2,936.6	1,092	3.88	18.6	7.3	64.9	382.74	1,167.33	927
	2009	4.82	16.52	11.27	47.43	790.87	2,503.35	1,063	3.14	11.73	6.67	45.94	423.23	1,355.1	913
	2011	7.98	30.94	14.21	54.57	715.87	1,819.63	1,138	5.64	31.12	8.46	57.21	455.22	2,681.66	1,184
UK	2005	23.46	84.9	20.09	92.99	328.37	886.28	2,918	26.52	161.29	27.06	300.04	486.96	3,626.91	5,026
	2007	24.05	77.66	18.04	75.81	316.38	807.11	3,315	26.59	150.1	23.86	280.63	466.69	3,404.67	5,823
	2009	19.62	72.49	14.59	60.27	307.04	932.57	3,432	21.88	146.49	18.06	206.88	490.7	4,606.31	6,107
	2011	24.58	97.91	16.39	66.75	309.77	926	3,555	23.13	145.68	19.29	230.84	487.36	4760.72	6,457

Notes: Table displays means and standard deviations computed from (cleaned) AMADEUS data for value added, capital (measured by tangible fixed assets), and employees. For each country and sector, statistics are computed on a pooled sample over years.

Table 4: TFPR – Manufacturing

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.530***			
	(0.0384)			
L.Log Capital		-0.407***		
		(0.0289)		
L.Log Labour			-0.531***	
			(0.0490)	
L.Log Wages				-0.435***
				(0.0424)
L.Log TFPR	-0.0285***	0.0166***	-0.00557*	-0.00145
	(0.00431)	(0.00277)	(0.00214)	(0.00229)
Cycle	0.109**	0.0356	0.00880	0.0358*
	(0.0345)	(0.0189)	(0.00751)	(0.0162)
TFPR * Expansion	0.0206*	-0.00366	-0.00599	0.000933
	(0.00687)	(0.0154)	(0.00533)	(0.00828)
TFPR * Contraction	-0.0123	0.0119	-0.0159*	-0.0179
	(0.0145)	(0.0141)	(0.00604)	(0.00904)
N	47361	47361	47361	47360
r2	0.295	0.213	0.326	0.322

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $*** < 0.01,\ ** < 0.05,\ * < 0.1$.

Table 5: TFPR – Services

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.556***			
	(0.0297)			
L.Log Capital		-0.434***		
		(0.0156)		
L.Log Labour			-0.492***	
			(0.0456)	
L.Log Wages				-0.445***
				(0.0354)
L.Log TFPR	-0.0318***	0.0333***	-0.000348	0.00508*
	(0.00459)	(0.00441)	(0.00208)	(0.00191)
Cycle	0.0968**	0.0388	0.0126	0.0404
	(0.0287)	(0.0423)	(0.0101)	(0.0192)
TFPR * Expansion	0.0311	-0.0213	-0.00473	-0.00197
	(0.0183)	(0.0143)	(0.00776)	(0.00924)
TFPR * Contraction	0.0204	0.0218	0.0120	-0.000906
	(0.0274)	(0.0286)	(0.00883)	(0.00912)
N	59755	59755	59755	59754
r2	0.292	0.179	0.294	0.287

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $***<0.01,\ **<0.05,\ *<0.1$.

Table 6: MRPK – Manufacturing

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.587***			
	(0.0308)			
L.Log Capital		-0.330***		
		(0.0261)		
L.Log Labour			-0.477***	
			(0.0420)	
L.Log Wages				-0.419***
				(0.0320)
L.Log MRPK	-0.0841***	0.0471***	-0.00962*	-0.0179
	(0.0128)	(0.00738)	(0.00322)	(0.00825)
Cycle	0.0666**	0.0431*	0.00242	0.0266
	(0.0207)	(0.0167)	(0.00585)	(0.0134)
MRPK * Expansion	-0.0104	-0.0155	-0.00341	-0.00380
	(0.0122)	(0.0139)	(0.00276)	(0.0100)
MRPK * Contraction	-0.0525**	0.0134	-0.0132*	-0.00630
	(0.0157)	(0.0108)	(0.00498)	(0.0145)
N	164436	164436	164436	164429
r2	0.398	0.230	0.315	0.338

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $*** < 0.01,\ ** < 0.05,\ * < 0.1$.

Table 7: MRPK – Services

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.587***			
	(0.0235)			
L.Log Capital		-0.312***		
		(0.0222)		
L.Log Labour			-0.459***	
			(0.0348)	
L.Log Wages				-0.408***
				(0.0312)
L.Log MRPK	-0.0743***	0.0742***	-0.00644**	-0.0169***
	(0.00859)	(0.0161)	(0.00182)	(0.00254)
Cycle	0.0946*	0.0755	0.0180	0.0452
	(0.0345)	(0.0349)	(0.0159)	(0.0216)
MRPK * Expansion	0.00672	-0.0314	-0.00440	0.00113
	(0.00783)	(0.0234)	(0.00360)	(0.00391)
MRPK * Contraction	-0.0177	0.0259	-0.00211	-0.00290
	(0.00947)	(0.0159)	(0.00381)	(0.00501)
N	255875	255875	255875	255873
r2	0.382	0.210	0.292	0.303

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $***< 0.01,\ **< 0.05,\ *< 0.1$.

Table 8: MRPL – Manufacturing

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.506***			
	(0.0321)			
L.Log Capital		-0.365***		
		(0.0266)		
L.Log Labour			-0.477***	
			(0.0423)	
L.Log Wages				-0.417***
				(0.0316)
L.Log MRPL	-0.268***	0.0412***	0.00966	0.0418***
	(0.0295)	(0.00653)	(0.00508)	(0.00581)
Cycle	0.162**	0.0221	-0.00469	0.0181
	(0.0481)	(0.0129)	(0.00691)	(0.0115)
MRPL * Expansion	0.0434*	-0.0246	-0.0258	-0.0152
	(0.0174)	(0.0255)	(0.0137)	(0.0268)
MRPL * Contraction	-0.0443	0.0147	-0.0143	-0.0196
	(0.0460)	(0.0196)	(0.0123)	(0.0188)
N	164594	164594	164594	164590
r2	0.411	0.228	0.315	0.339

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $***<0.01,\ **<0.05,\ *<0.1$.

Table 9: MRPL – Services

	(1) dlnPY	(2) dlnK	(3) dlnemp	(4) dlnwage
L.PY	-0.484***			
	(0.0295)			
L.Log Capital		-0.370***		
		(0.0126)		
L.Log Labour			-0.459***	
			(0.0355)	
L.Log Wages				-0.401***
				(0.0322)
L.Log MRPL	-0.306***	0.0271**	0.00396	0.0442***
	(0.0316)	(0.00678)	(0.00623)	(0.00894)
Cycle	0.182**	0.0842	0.0161	0.0350
	(0.0511)	(0.0636)	(0.0166)	(0.0246)
MRPL * Expansion	0.0376	0.0623	-0.00183	0.00751
	(0.0430)	(0.0676)	(0.0140)	(0.0169)
MRPL * Contraction	0.0209	0.0243	-0.0171	-0.0560
	(0.0628)	(0.0691)	(0.0140)	(0.0316)
N	256172	256172	256172	256171
r2	0.405	0.207	0.292	0.304

Notes: Standard errors clustered at the industry country level, and displayed in parentheses. Statistical significance: $*** < 0.01,\ ** < 0.05,\ * < 0.1$.

Figure 1: Misallocation across Time – Manufacturing [1/2]

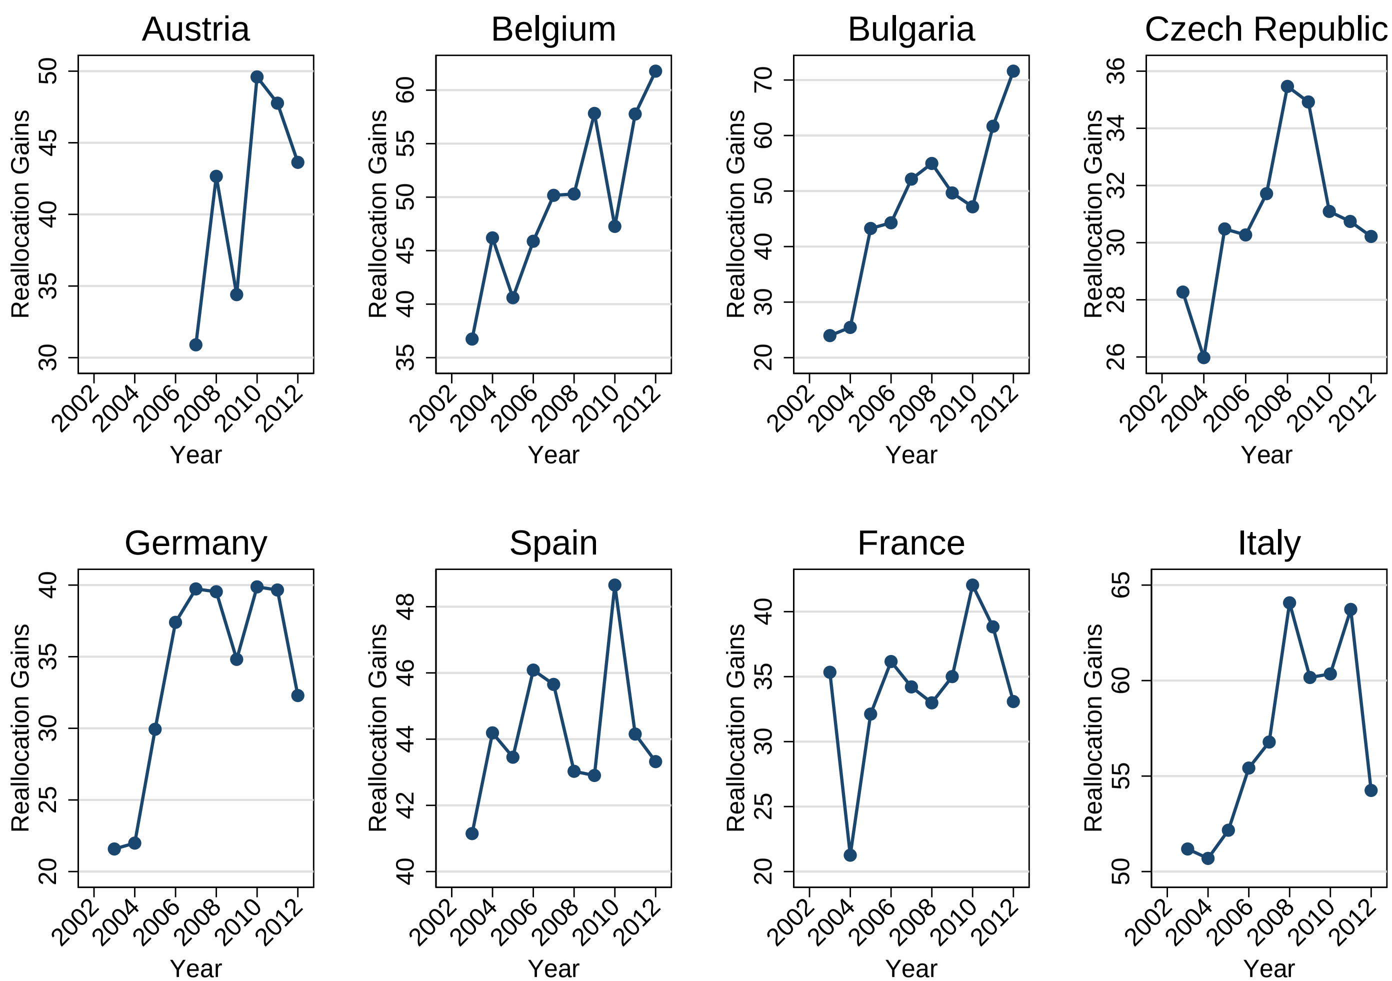

Notes: Y-axis scales differ. Figure shows the time-series of computed potential percentage reallocation gains (misallocation) for the manufacturing sample. Entries are calculated as $100 * (Y_t^*/Y_t - 1)$, where $Y_t^*/Y_t = \prod_{s=1}^{S} \left(TFP_{s,t}^*/TFP_{s,t}\right)^{\theta_s}$.

Figure 2: Misallocation across Time – Manufacturing [2/2]

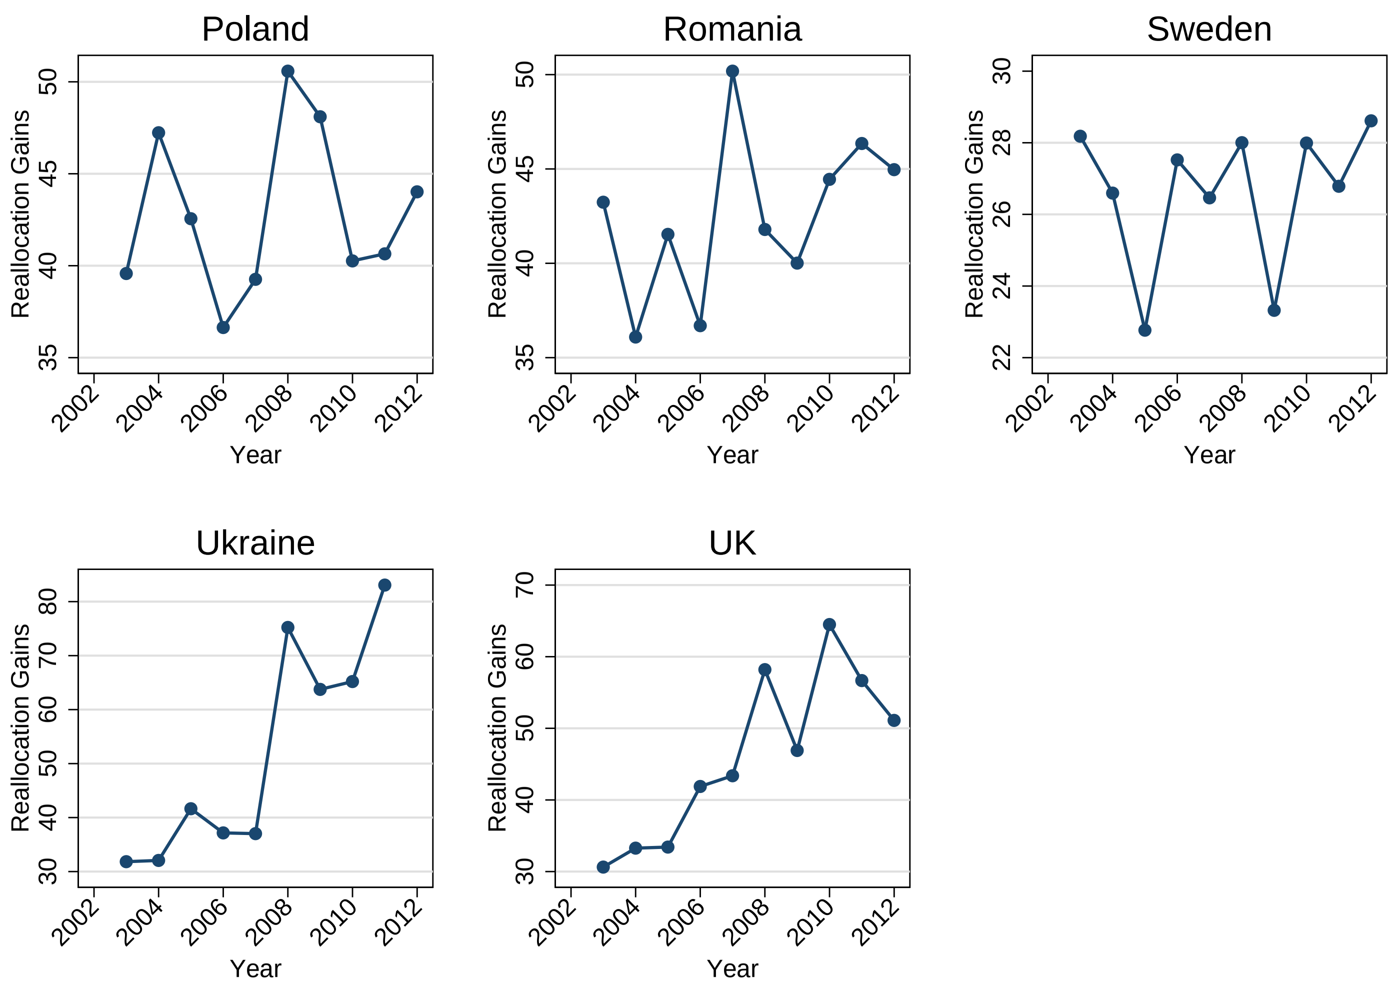

Notes: Y-axis scales differ. Figure shows the time-series of computed potential percentage reallocation gains (misallocation) for the manufacturing sample. Entries are calculated as $100 * (Y_t^*/Y_t - 1)$, where $Y_t^*/Y_t = \prod_{s=1}^{S} \left(TFP_{s,t}^*/TFP_{s,t}\right)^{\theta_s}$.

Figure 3: Misallocation across Time – Services [1/2]

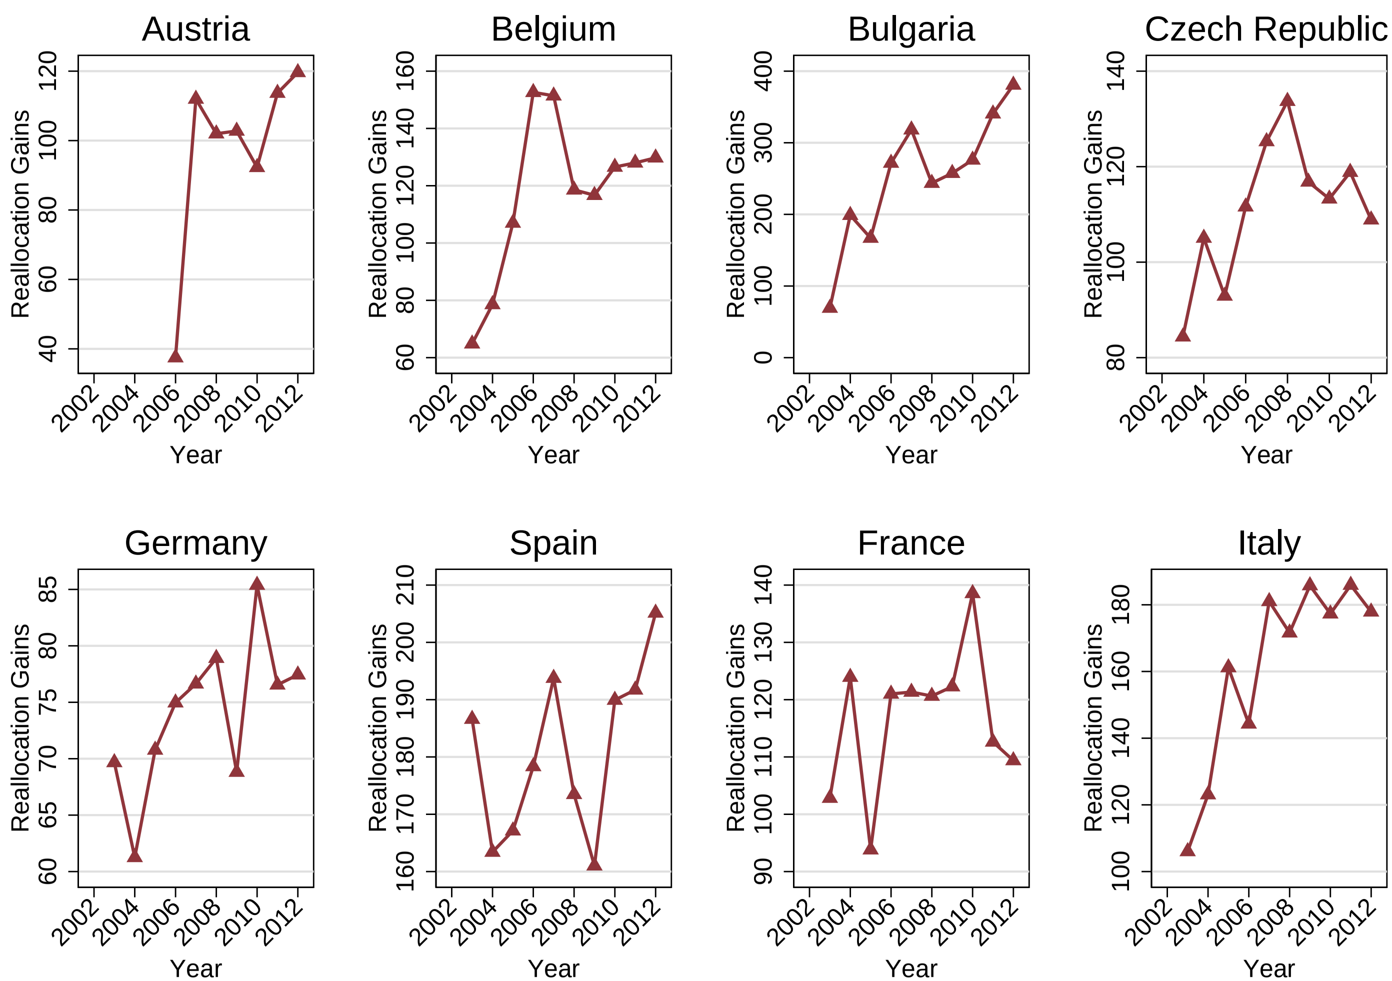

Notes: Y-axis scales differ. Figure shows the time-series of computed potential percentage reallocation gains (misallocation) for the services sample in percentage points. Entries are calculated as $100 * (Y_t^*/Y_t - 1)$, where $Y_t^*/Y_t = \prod_{s=1}^{S} \left(TFP_{s,t}^*/TFP_{s,t}\right)^{\theta_s}$.

Figure 4: Misallocation across Time – Services [2/2]

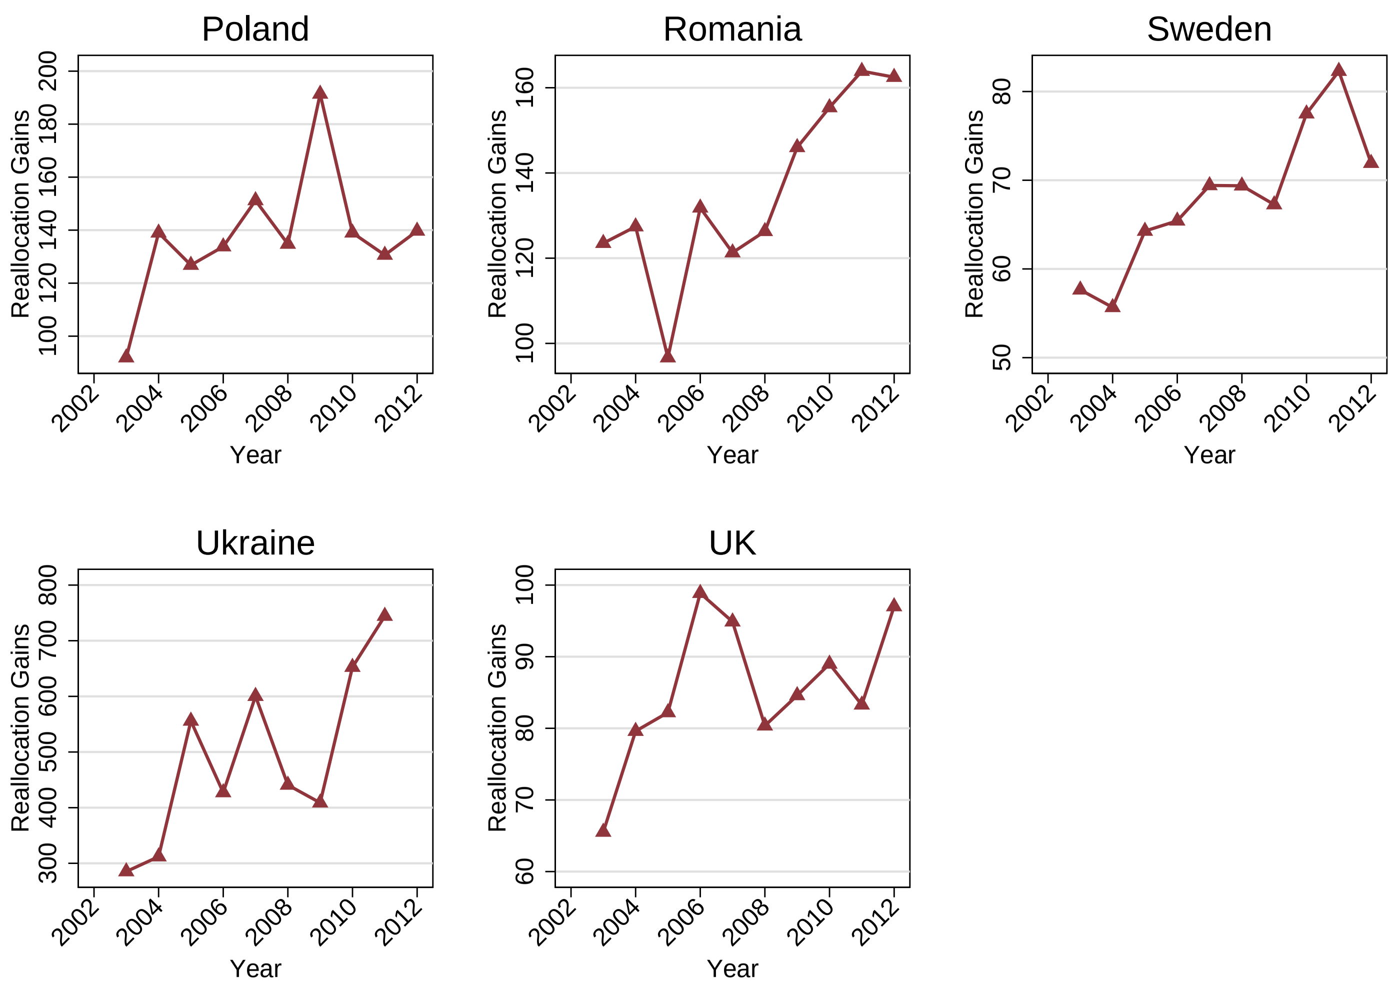

Notes: Y-axis scales differ. Figure shows the time-series of computed potential percentage reallocation gains (misallocation) for the services sample in percentage points. Entries are calculated as $100 * (Y_t^*/Y_t - 1)$, where $Y_t^*/Y_t = \prod_{s=1}^{S} \left(TFP_{s,t}^*/TFP_{s,t}\right)^{\theta_s}$.

Figure 5: Relation between Capital Wedges and Firm Size – Manufacturing [1/2]

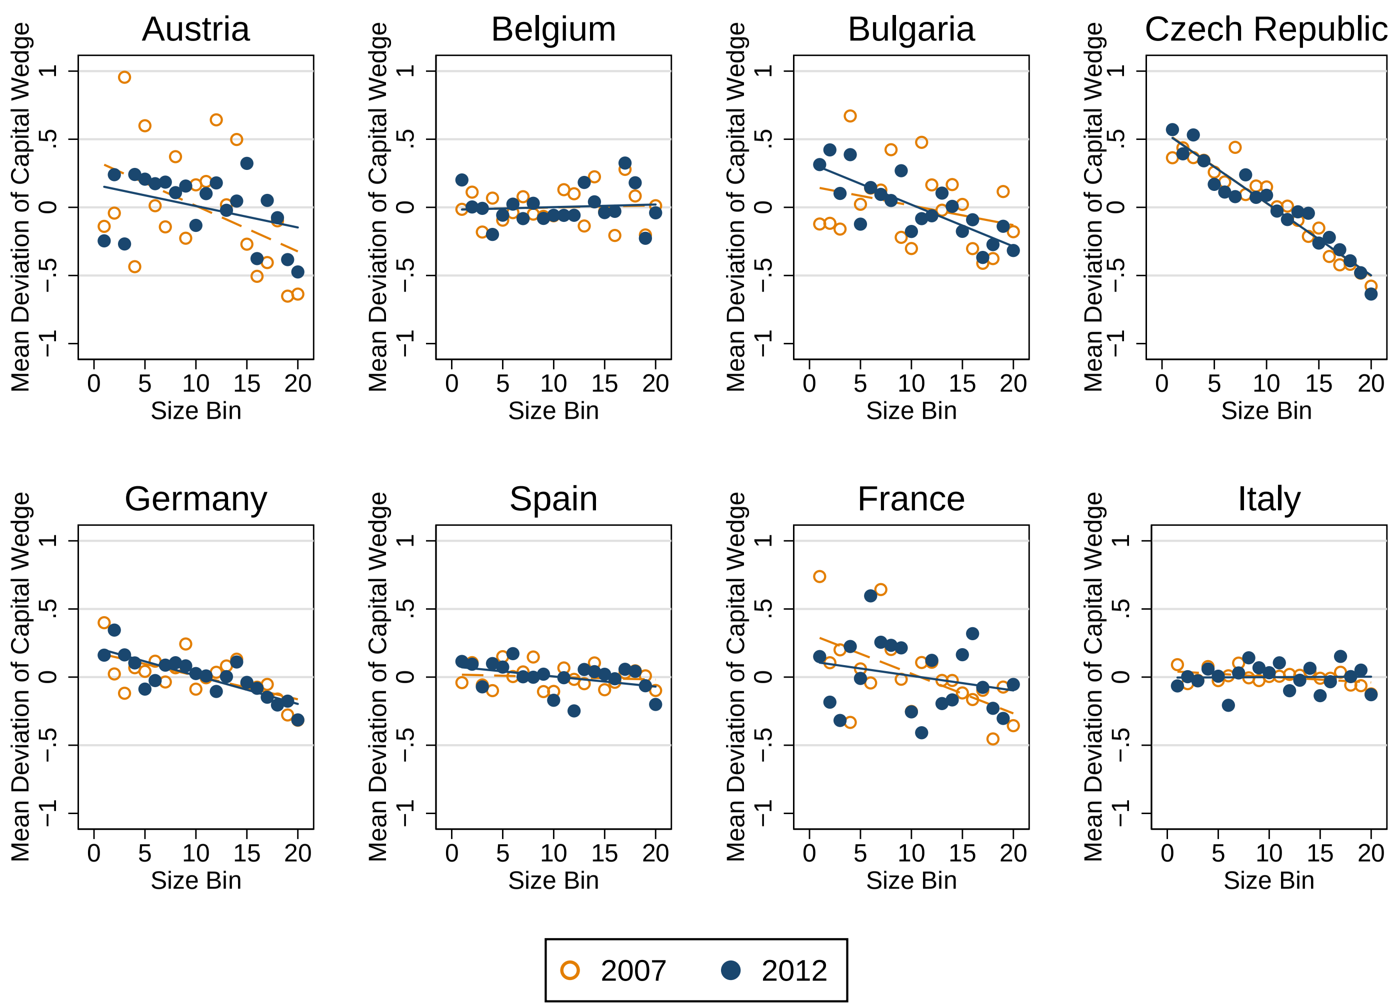

Notes: Figure displays the capital wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the manufacturing sample. The capital wedge is calculated: $1 + \tau_{Ksi} = \frac{\alpha_s}{1-\alpha_s}(wL_{si}/RK_{si})$.

Figure 6: Relation between Capital Wedges and Firm Size – Manufacturing [2/2]

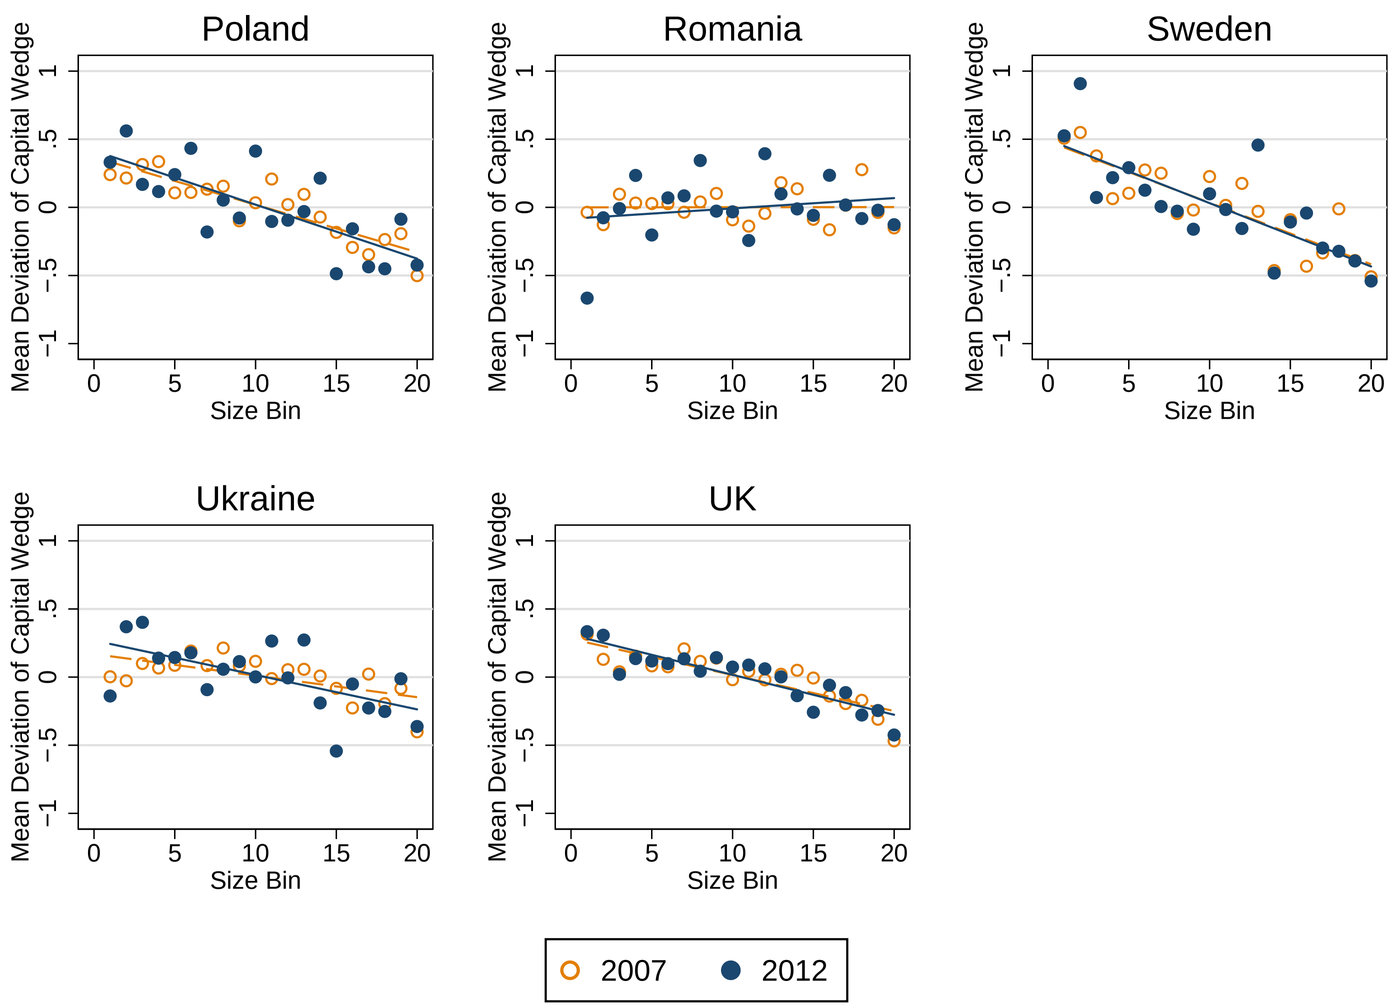

Notes: Figure displays the capital wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the manufacturing sample. The capital wedge is calculated: $1 + \tau_{Ksi} = \frac{\alpha_s}{1-\alpha_s}(wL_{si}/RK_{si})$.

Figure 7: Relation between Output Wedges and Firm Size – Manufacturing [1/2]

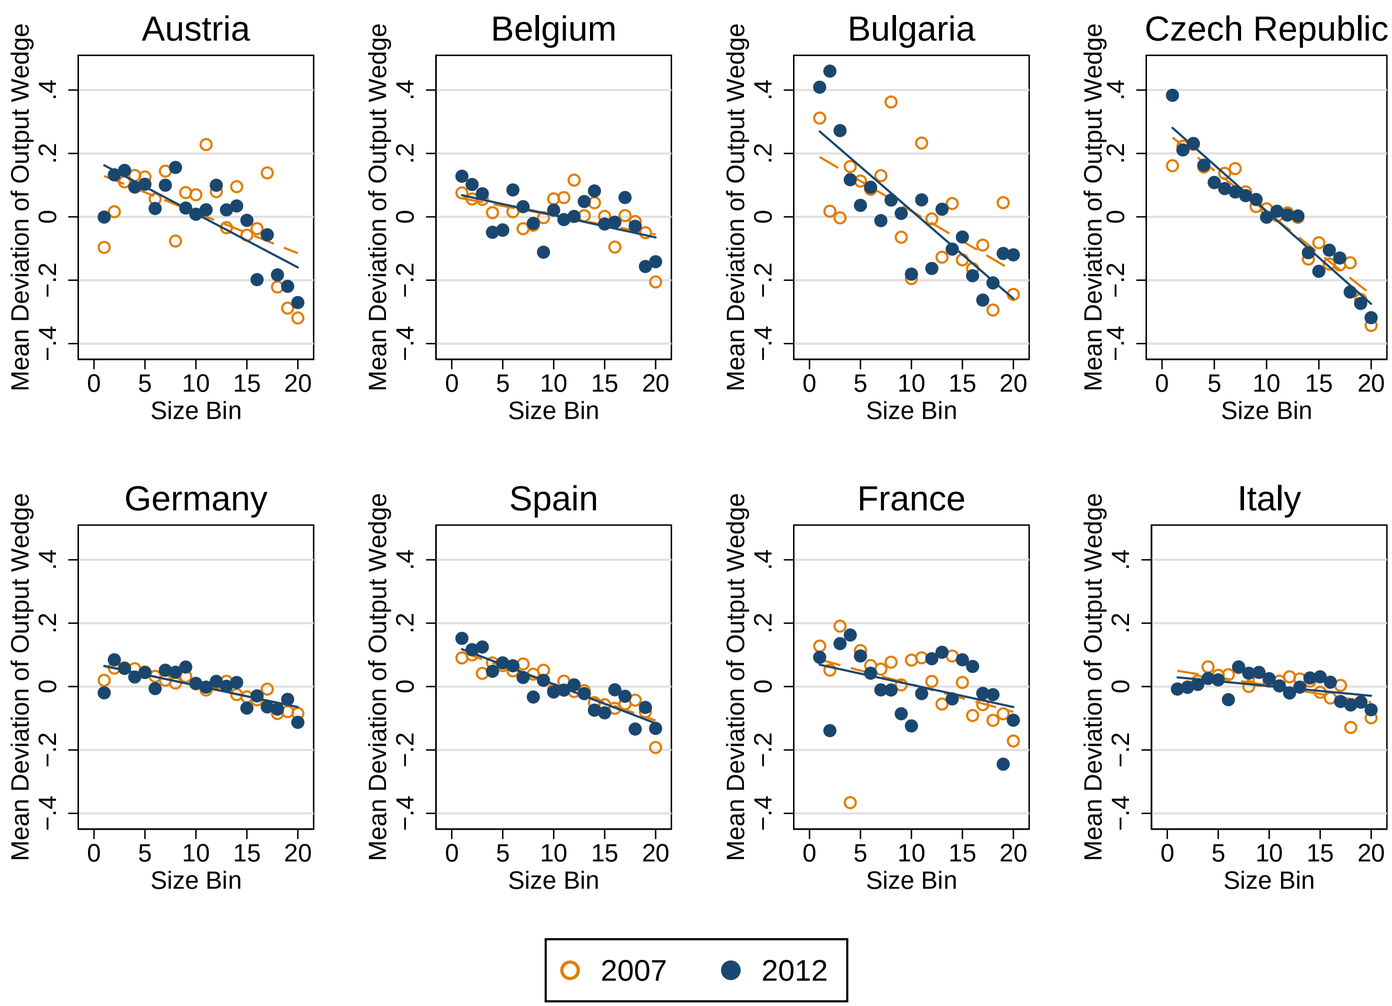

Notes: Figure displays the output wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the manufacturing sample. The output wedge is calculated as: $1 - \tau_{Ysi} = \frac{\sigma}{\sigma-1}(wL_{si}/(1-\alpha_s)P_{si}Y_{si})$.

Figure 8: Relation between Output Wedges and Firm Size – Manufacturing [2/2]

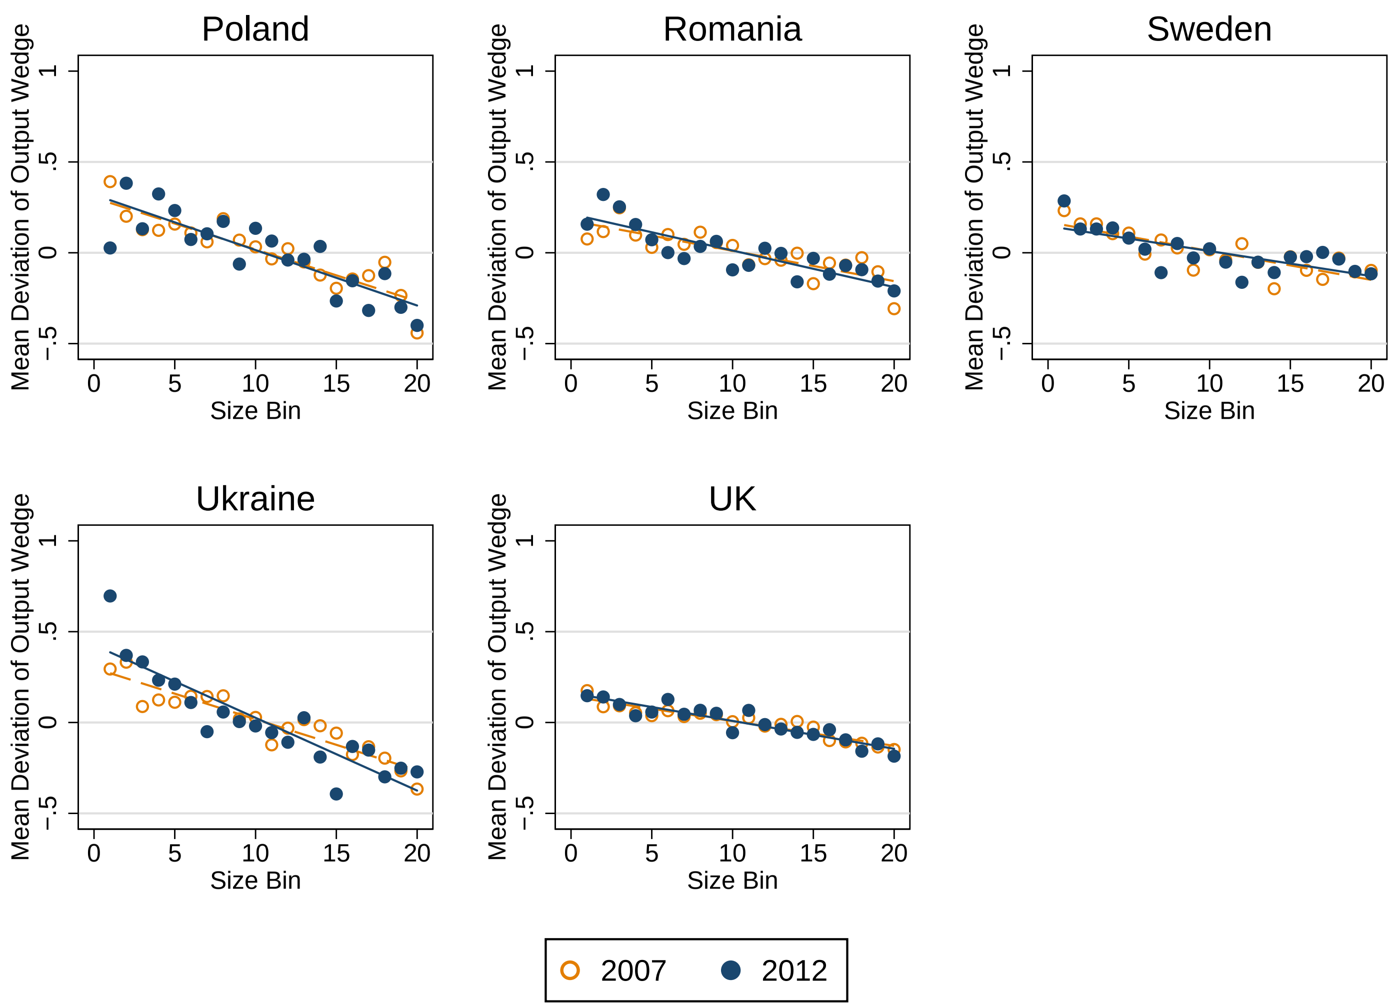

Notes: Figure displays the output wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the manufacturing sample. The output wedge is calculated as: $1 - \tau_{Ysi} = \frac{\sigma}{\sigma-1}(wL_{si}/(1-\alpha_s)P_{si}Y_{si})$.

Figure 9: Relation between Capital Wedges and Firm Size – Services [1/2]

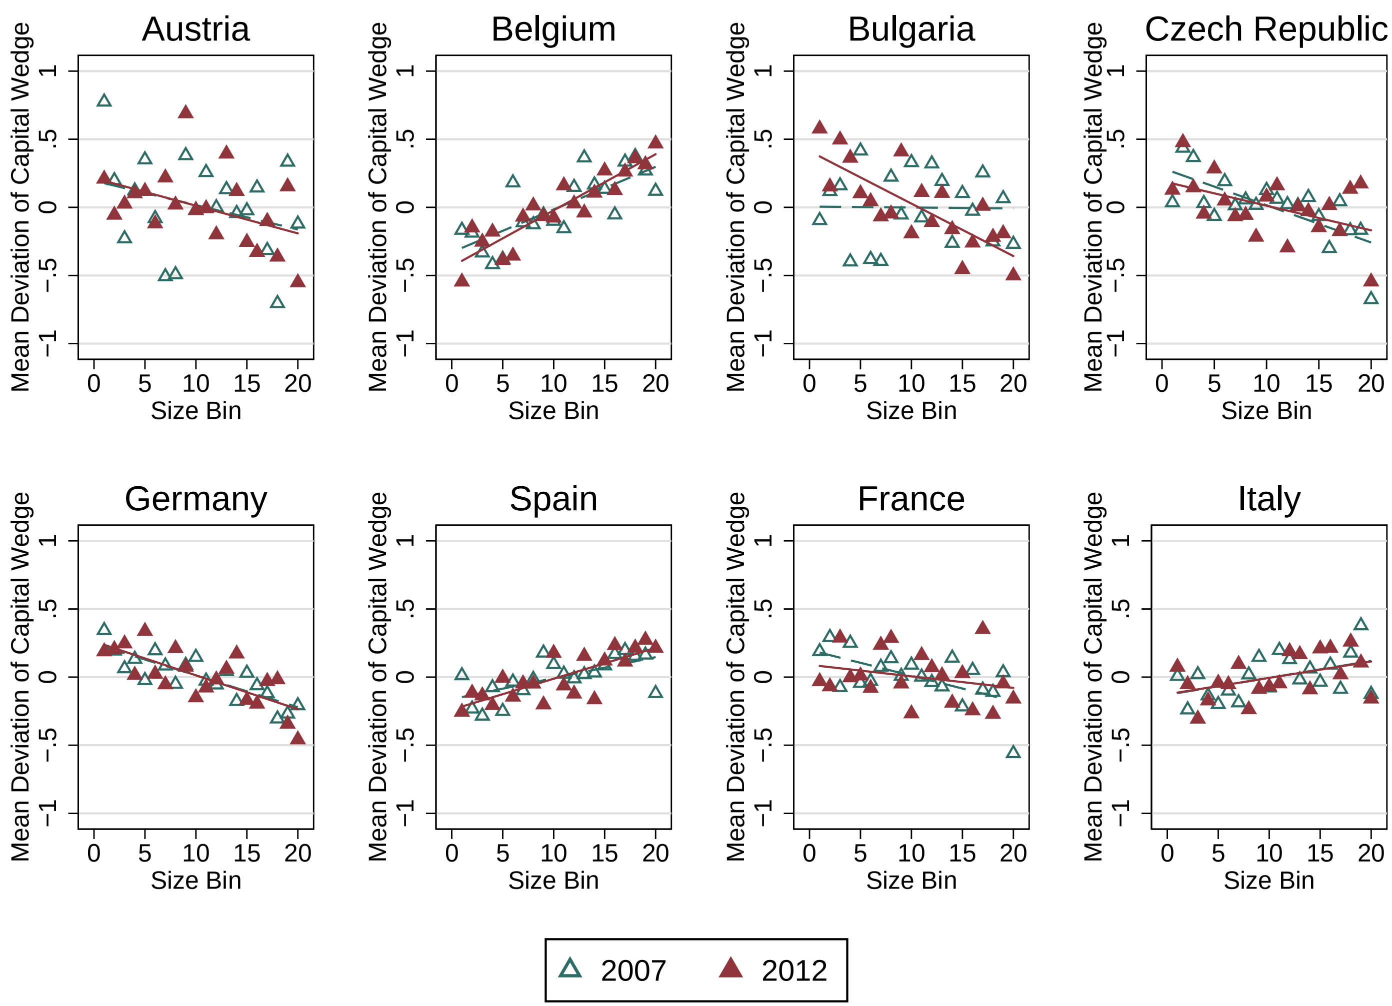

Notes: Figure displays the capital wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the services sample. The capital wedge is calculated: $1 + \tau_{Ksi} = \frac{\alpha_s}{1-\alpha_s}(wL_{si}/RK_{si})$.

Figure 10: Relation between Capital Wedges and Firm Size – Services [2/2]

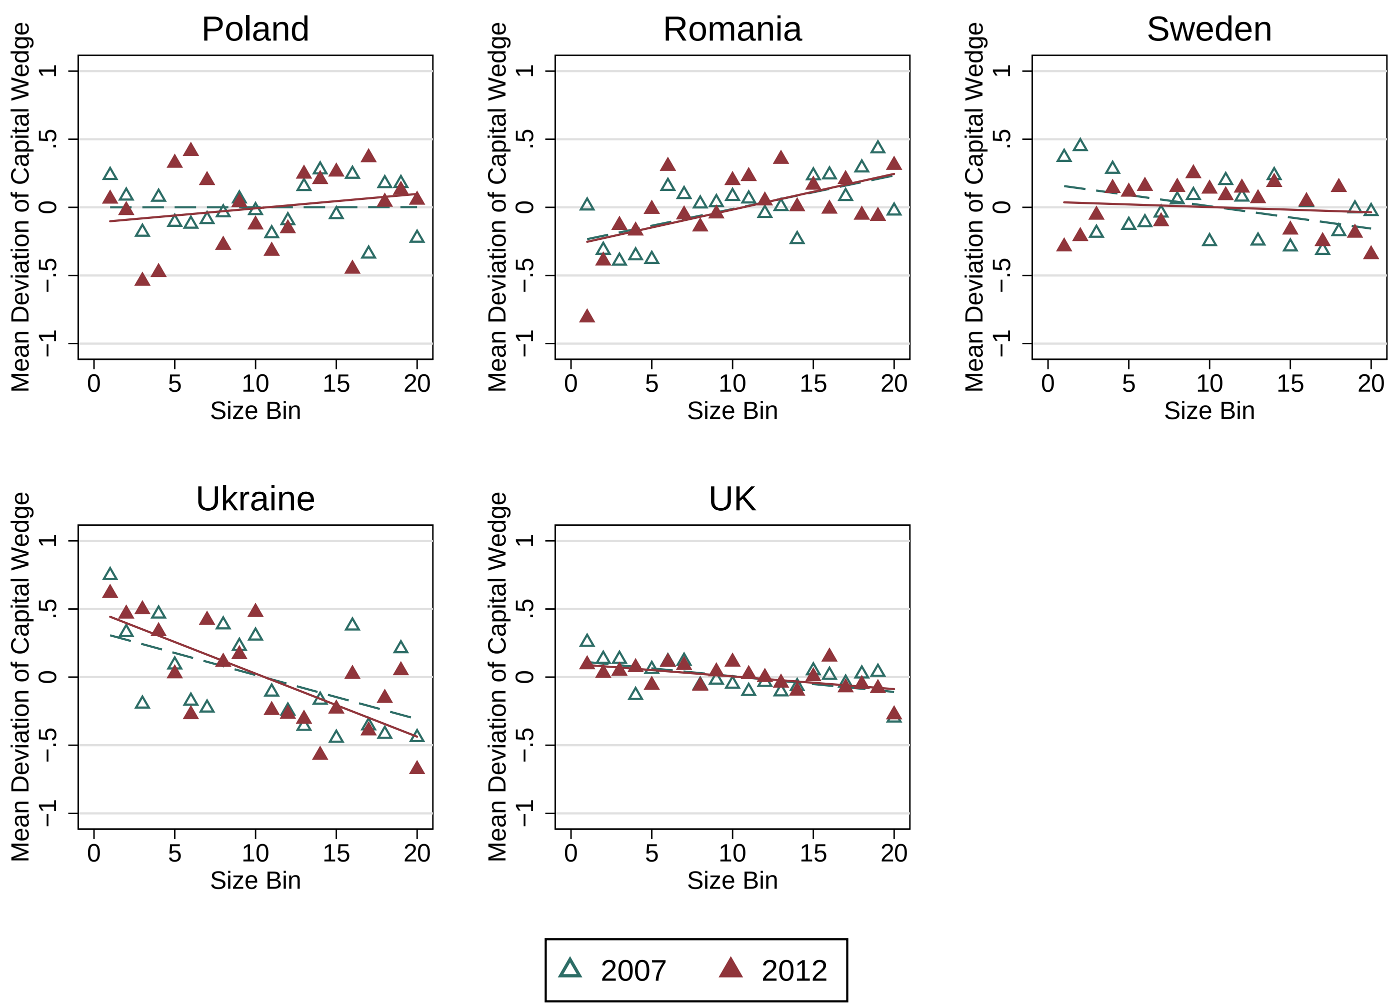

Notes: Figure displays the capital wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the services sample. The capital wedge is calculated: $1 + \tau_{Ksi} = \frac{\alpha_s}{1-\alpha_s}(wL_{si}/RK_{si})$.

Figure 11: Relation between Output Wedges and Firm Size– Services [1/2]

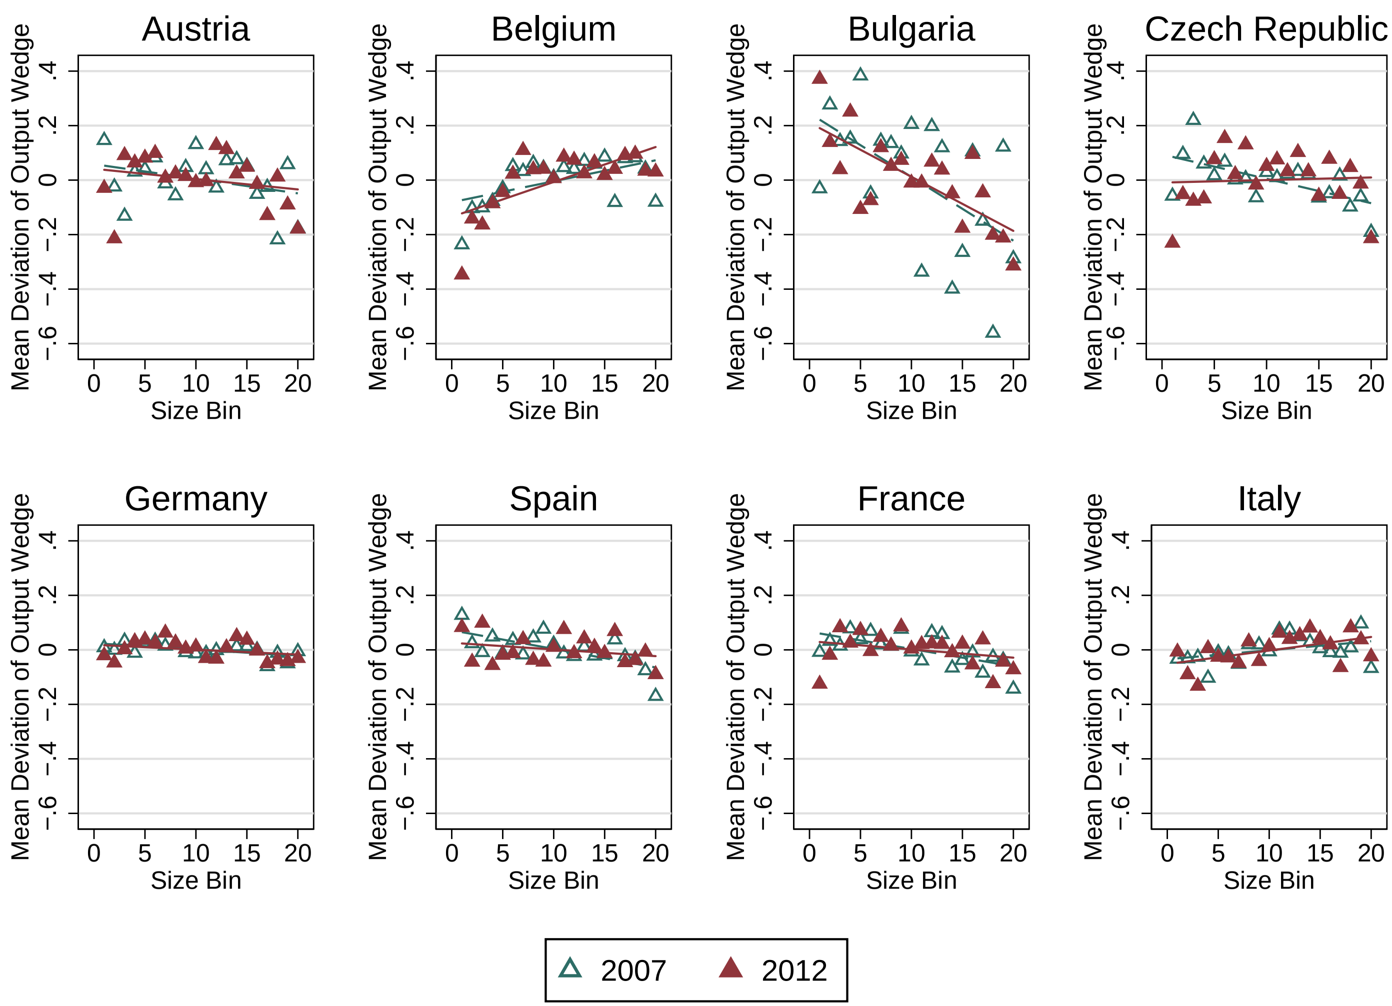

Notes: Figure displays the output wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the services sample. The output wedge is calculated as: $1 - \tau_{Ysi} = \frac{\sigma}{\sigma-1}(wL_{si}/(1-\alpha_s)P_{si}Y_{si})$.

Figure 12: Relation between Output Wedges and Firm Size – Services [2/2]

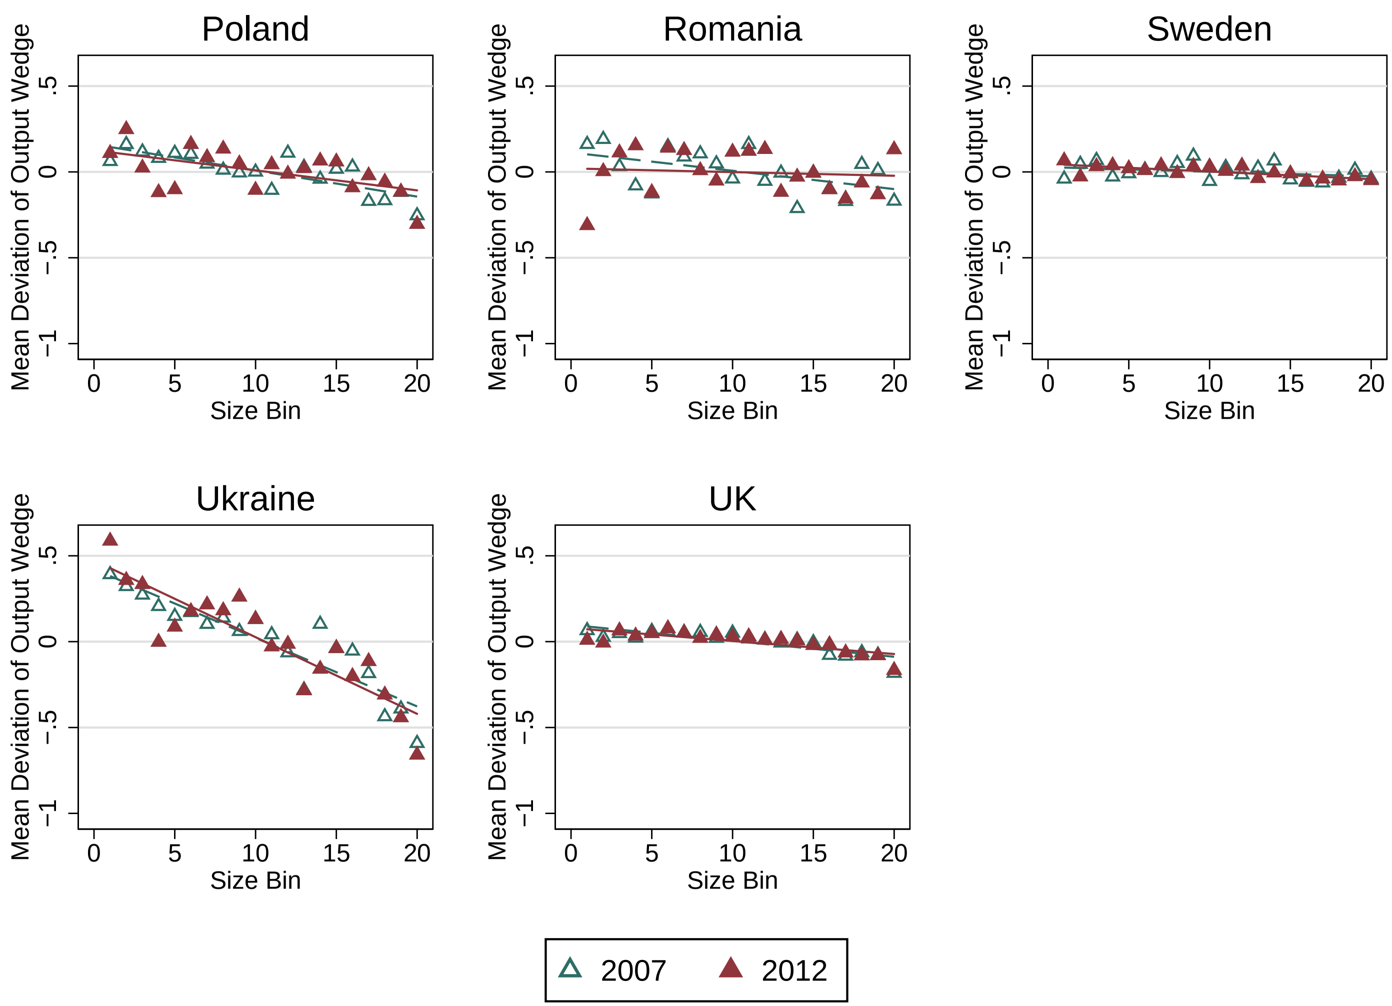

Notes: Figure displays the output wedges, averaged conditional on size bin (and year) and ranked. There are 20 size bins, which are created by ranking observations by value added. Results are displayed for 2007 and for 2012 (with lines of best-fit included). Data from the services sample. The output wedge is calculated as: $1 - \tau_{Ysi} = \frac{\sigma}{\sigma-1}(wL_{si}/(1-\alpha_s)P_{si}Y_{si})$.

Figure 13: Representativeness of Samples

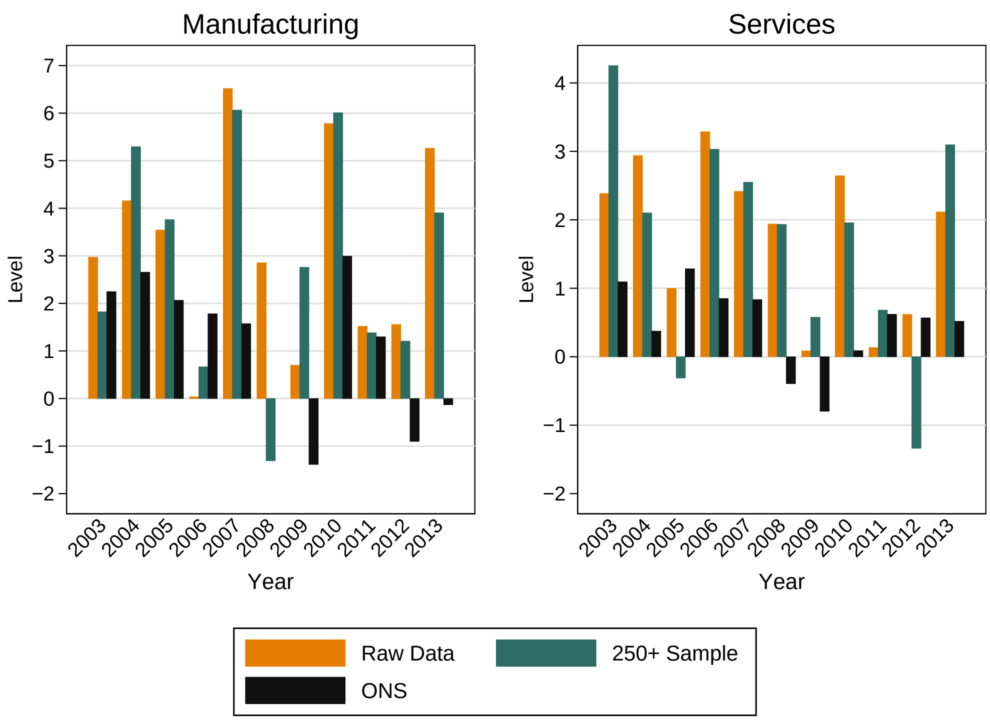

Notes: Aggregate labour productivity is totaled real gross value added at market prices over totalled point-in-time employment from the IDBR data source (point-in-time). The 250+ dataset has been subject to data cleaning, the raw dataset has not. Sampling weights have been used for the raw dataset.

Figure 14: Reallocation Gains

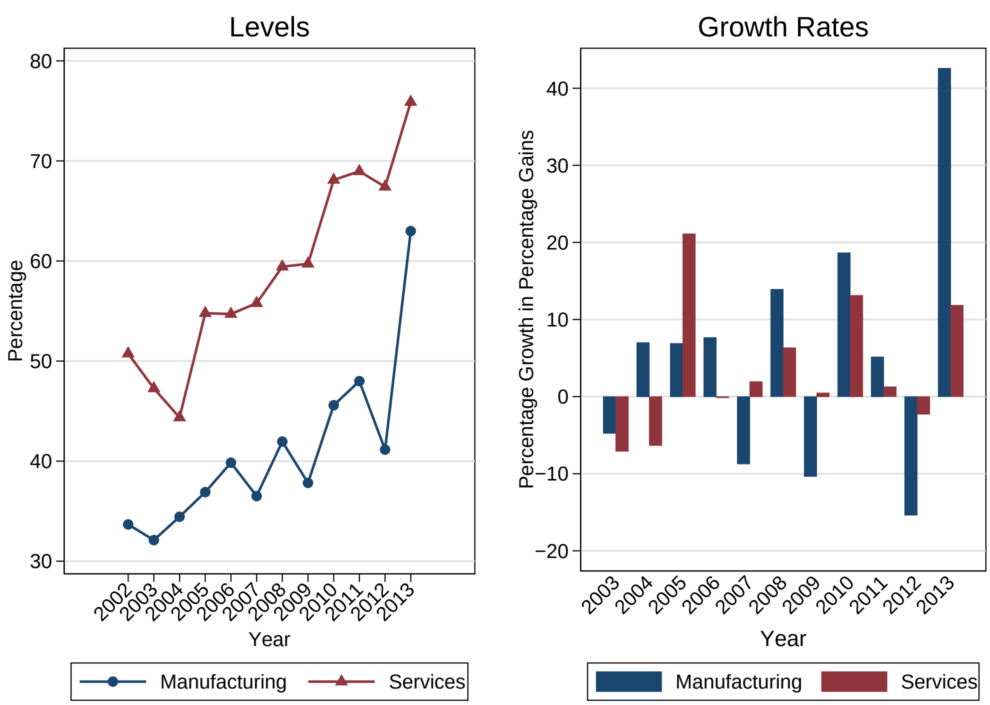

Notes: Reallocation gains for a balanced panel. Entries are percentage point gains from equalizing TFPR within industries. Entries are calculated as $100(Y_t^*/Y_t - 1)$, where $Y_t^*/Y_t = \prod_{s=1}^{S} \left(TFP_{s,t}^*/TFP_{s,t}\right)^{\theta_s}$.