EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Sebastian Hauptmeier, Frédéric Holm-Hadulla,
Katerina Nikalexi

Monetary policy and regional inequality

No 2385 / March 2020

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

We study the impact of monetary policy on regional inequality using granular data on economic activity at the city- and county-level in Europe. We document pronounced heterogeneity in the regional patterns of monetary policy transmission. The output response to monetary policy shocks is stronger and more persistent in poorer regions, with the difference becoming particularly pronounced in the extreme tails of the distribution. Regions in the lower parts of the distribution exhibit hysteresis, consisting of long-lived adjustments in employment and labor productivity in response to the shocks. As a consequence, policy tightening aggravates regional inequality and policy easing mitigates it.

Keywords: monetary policy; regional heterogeneity; local projections; quantile regressions

JEL Classification: C32, E32, E52

Non-technical summary

This paper provides a novel perspective on the regional patterns of monetary policy transmission. Using geographically disaggregated data on economic activity in Europe, we show that the output response to short-term interest rate shocks is significantly more pronounced and persistent in poorer than in richer cities and counties. Moreover, while GDP in the upper part of the distribution returns to its pre-shock level after four to five years, the response in the lower parts does not reverse over this period, thus pointing to pronounced hysteresis in output. The heterogeneous incidence of hysteresis in turn implies that monetary policy has a long-lasting impact on regional inequality, with tightening shocks aggravating and easing shocks mitigating it.

In terms of anatomy, we find hysteresis to originate from long-lived adjustments in both employment and labor productivity. At the same time, employment hysteresis is more pronounced and more broad-based across the distribution. In fact, it even extends to the sample mean, as opposed to productivity hysteresis which concentrates in the lower tails of the distribution. As such, our findings confirm labor markets as an important source of hysteresis in the European context.

Our paper points to the use of geographically disaggregated data as a promising avenue for further insights into how exactly monetary policy shocks propagate to the economy. First, by resorting to information on economic activity at a more granular level than that entering the central bank reaction function, it offers a novel strategy to identify exogenous changes in monetary policy. Second, by providing empirical estimates for the monetary policy impact on regional inequality, it closes an important gap in the large and growing literature on heterogeneity in monetary policy transmission.

From a policy perspective, our findings underscore the challenges of calibrating monetary policy in heterogeneous economies. In the euro area context, the debate has typically interpreted these challenges as a cross-country phenomenon. As such, they attracted particular attention during the euro area sovereign debt crisis from 2010-12, which was marked by strong cross-country divergence in economic performance and raised concerns as to the suitability of a given aggregate monetary policy stance for individual countries. However, our analysis demonstrates that the issue runs deeper: interregional heterogeneity becomes more accentuated at more granular geographical levels and this heterogeneity in turn profoundly alters the impli-

cations of a given monetary policy stance in different parts of the economy. These implications emerge as particularly relevant in view of our finding that monetary policy exerts durable impacts on output and employment – a finding that contrasts with the common notion of stabilization policies merely smoothing out fluctuations in these variables around some natural levels.

1 Introduction

There is an increasingly active debate on whether monetary policy tends to dampen or reinforce economic inequality. Much of this debate has centred on the impact of monetary policy on income, consumption and wealth inequality at the household level (Coibion et al., 2017; Ampudia et al., 2018; Lenza and Slacalek, 2018). An aspect that instead has received little attention so far is that economic inequality typically has a pronounced geographical dimension. In particular, some subnational regions generate significantly and persistently higher per capita incomes than others (Jacobs, 1969; Krugman, 1998); and, in many countries, this aspect of inequality has either intensified over recent decades or, at least, halted its previous declining trend.¹

A natural question is whether these trends have been influenced by macroeconomic policies; and, among the different policy domains, monetary policy emerges as an interesting candidate given its dominant role over the last decade as central banks sought to contain the fallout from the global financial crisis. However, empirical evidence on the link between monetary policy and regional inequality is sparse and, *a priori*, it is ambiguous whether monetary accommodation policy would reduce or aggravate regional inequality.

On the one hand, a large and growing literature has documented that more vulnerable economic agents – *e.g.* in the household-, firm-, or banking-sector – are more sensitive to monetary policy shocks.² To the extent that more vulnerable economic agents cluster geographically, these mechanisms may translate into a stronger response of overall economic activity to monetary policy in less prosperous regions. On the other hand, the findings on different types of individual economic agents are not uncontested, in that recent literature has also highlighted channels that might induce a *lower* responsiveness in the weaker parts of the respective distributions.³ Moreover, there is ample evidence suggesting that the impact of monetary policy depends

¹See Ganong and Shoag (2017) and Austin, Glaeser and Summers (2018), Iammarino, Rodriguez-Pose and Storper (2018), and Martin et al. (2018) for recent evidence on the United States, the euro area, and the UK, respectively; and Bluedorn et al. (2019) for recent IMF work on this issue for a broad set of advanced economies.

²For evidence on households, see above; on firms, see *e.g.* Kashyap, Lamont and Stein (1994), Gertler and Gilchrist (1994); and on banks, see *e.g.* Peek and Rosengren (1995), Kashyap and Stein (2000), Altavilla, Canova and Ciccarelli (2019).

³For instance, Ottonello and Winberry (2018) find investment of weaker firms to be less responsive to monetary policy shocks; Boeckx, Dossche and Peersman (2017) find that the macroeconomic impact of monetary policy measures is weaker in countries with a relatively poorly capitalized banking system; and the burgeoning literature on heterogeneous agent DSGE models, has identified a differentiated set of mechanisms that could make consumption of poorer households more or less responsive to monetary policy shocks than that of richer ones (Kaplan, Moll and Violante, 2018; Auclert, 2019).

on the sectoral composition of output (Carlino and DeFina, 1998; Peersman and Smets, 2005; Dedola and Lippi, 2005); and this composition in turn may be less interest-sensitive in poorer than in richer regions.⁴ Hence, the regional incidence of monetary policy is an interesting testing ground for competing theoretical predictions.

Against this background, the current paper studies the impact of monetary policy on regional inequality, exploiting granular data at the city- and county-level in the euro area. The use of city- and county-level data, as the lowest aggregation level at which official GDP statistics are available, has two key advantages for this type of analysis.

The first advantage pertains to the identification of exogenous variation in monetary policy, which is the key challenge at the heart of empirical monetary economics: as central banks are equipped with explicit macroeconomic stabilization mandates, monetary policy is by construction endogenous to the state and prospects of the economy. As a consequence, it is inherently difficult to disentangle the cause and effect of observed co-movements between policy indicators and macroeconomic aggregates. While the literature has proposed a host of strategies to solve this identification problem, consensus on the actual impact of monetary policy shocks on the economy has remained elusive (e.g. Ramey (2016) documents the wide dispersion in the peak and persistence of the estimated output response to monetary policy shocks in the US).

Hence, there is value in exploiting further, complementary identification strategies to progress towards a consensus on the effect of monetary policy on the economy. And the use of disaggregated GDP data offers a promising avenue in this regard. In particular, our identification strategy allows us to make use of the fact that the mandate of the ECB refers to the euro area aggregate level and the ECB's Governing Council has, since its onset, emphasized that: “[its] single monetary policy will adopt a euro area-wide perspective; it will not react to specific regional or national developments” (ECB Governing Council Press Release, 13 October 1998). This euro-area wide focus of monetary policy implies that, controlling for aggregate conditions, variation in short-term interest rates is exogenous to GDP at the city- and county-level. Thus, exploiting economic information at a more granular level than that entering central bank reaction functions allows us

⁴For instance, this constellation should apply to the classic dichotomy between a (prosperous) manufacturing core and a (less prosperous) agricultural periphery e.g. in the United States and Italy; see Krugman (1998). Since agriculture should be fairly interest insensitive, this in turn would render output in poorer regions less responsive to monetary policy shocks.

to address this type of simultaneity bias. Similar identification strategies based on sub-national data have been put to productive use in estimating fiscal multipliers (see e.g. [Auerbach and Gorodnichenko \(2012\)](#), [Clemens and Miran \(2012\)](#), and [Nakamura and Steinsson \(2014\)](#), as well as [Chodorow-Reich \(2019a\)](#) for a review). But they have received little attention in the monetary economics literature so far.

The second advantage of using data at this maximum degree of disaggregation derives from the particularly pronounced disparities arising at this level. For instance, the dispersion of per capita GDP at the city- and county-level is more than 50% higher than that between countries in the euro area (see Section 3 for details). Accordingly, the disaggregated data provide for much a richer information set to be exploited in analysing the regional transmission of monetary policy. And it allows us to assess whether monetary policy tends to mitigate and accentuate interregional income disparities – an important question in view of the increased scrutiny that the impact of monetary policy on economic inequality has received in public discourse and the increased recognition that rising inequality in many advanced economies exhibits a pronounced interregional dimension ([The Economist, 2016](#); [Coeuré, 2018](#)).

To estimate the dynamic impact of exogenous changes in monetary policy across the entire regional GDP distribution, we combine [Jordà \(2005\)](#)’s local projections method with two complementary approaches for modeling heterogeneity. The first consists of quantile estimation techniques and is based on the novel approach proposed by [Machado and Silva \(2019\)](#) which, building on [Chernozhukov and Hansen \(2008\)](#), provides an avenue to overcome the well-known problems of standard quantile regressions in the presence of fixed effects. The second consists of sub-sample analyses, as suggested *e.g.* by [Crouzet and Mehrotra \(2017\)](#) in the context of firm size classes, and is based on a break-down of the regional distribution into separate quantile-ranges. In terms of policy indicators, we focus on standard monetary policy, implemented via policy-controlled short-term interest rates, since this type of policy dominated over our sample period. We do however confirm the robustness of our findings to the use of shadow rates that also capture the impact of non-standard monetary policy on financial conditions. The data are based on Eurostat’s Nomenclature on Territorial Units for Statistics (NUTS) and broken down to the most disaggregated geographical level at which information on economic activity is available (NUTS3).⁵

⁵The subnational administrative setup differs substantially across euro area countries and not all countries have a local layer of government equivalent to what is referred to as a “county” for instance in the United States or in Germany. As we discuss

Our estimates document a significant and relevant impact of changes in short-term interest rates on regional GDP that strongly intensifies towards the lower end of the distribution. At the sample mean, the response closely resembles the patterns typically found in the literature. After an initial transmission lag, the interest rate coefficient turns significant at a one-year horizon. The impact then further builds up and peaks at the two-year horizon, with the estimates implying a decline in average regional GDP of around 2% in response to a 100 basis point exogenous interest rate increase, before fading out over the remainder of the projection horizon. Moving beyond the sample mean, the quantile estimations point to pronounced heterogeneity in the regional patterns of monetary policy transmission. In particular, output responds more strongly in regions with low versus high per capita GDP and this difference becomes particularly pronounced in the extreme tails of the distribution. Further, the same qualitative patterns emerge when, instead of total GDP, we focus on the capital- and labor-intensive sub-sectors of the regional economies while, among these sub-sectors, capital-intensive production is more responsive, in line with standard notions of the interest rate channel of monetary policy.

The most striking feature of our results is that substantial parts of the regional distribution experience long-lived effects of monetary policy shocks on output. In particular, while GDP in the upper part of the distribution returns to its pre-shock level after four to five years, the response in the lower parts does not reverse over this period, thus pointing to pronounced hysteresis in output. The heterogeneous incidence of hysteresis in turn implies that monetary policy has a long-lasting impact on regional inequality, with tightening shocks aggravating and easing shocks mitigating it.

In terms of anatomy, we find hysteresis to originate from long-lived adjustments in both employment and labor productivity. At the same time, employment hysteresis is more pronounced and more broad-based across the distribution. In fact, it even extends to the sample mean, as opposed to productivity hysteresis which concentrates in the lower tails of the distribution. As such, our findings confirm labor markets as an important source of hysteresis in the European context. And they add to a growing literature arguing that monetary policy may exert durable impacts on output and employment – in contrast to the common notion

in greater detail in Section 3, the NUTS classification tackles this issue by creating a system of economically coherent regions following harmonized standards across countries. For ease of exposition, we refer to NUTS3 regions as the “city- and county-level” throughout this paper but note that this wording is approximate and needs to be interpreted in conjunction with the more precise definitions presented in Section 3.

of stabilization policies merely smoothing out fluctuations in these variables around some natural levels (Dupraz, Nakamura and Steinsson, 2019; Jordà, Singh and Taylor, 2020).

Finally, our findings add a novel perspective to the literature on the geographical incidence of monetary policy. To date, most of the literature on the regional patterns of monetary policy transmission has focused on less granular levels of aggregation, such as (clusters of) US states and similar geographical units in European countries.⁶ By contrast, and notwithstanding the aforementioned benefits, only few papers have studied monetary policy transmission at a richer regional disaggregation level. Important exceptions include Francis, Owyang and Sekhposyan (2012) who estimate city-level responses to monetary policy shocks in the US, as well as Fratantoni and Schuh (2003), Di Maggio et al. (2017), Beraja et al. (2018), and Eichenbaum, Rebelo and Wong (2018) who examine the transmission of monetary policy via mortgage markets at similar aggregation levels as the current paper. Our paper is the first to model the impact of monetary policy across the entire regional GDP distribution of a major advanced economy at the city- and county level; to explicitly tackle the question of how the response differs in richer and poorer places; and thus to shed light on how monetary policy affects regional inequality.

The remainder of the paper proceeds as follows. The next section presents the empirical model and identification strategy. Section 3 describes the data and highlights some salient stylized facts regarding economic activity at the regional level. Section 4 presents our baseline results and a broad range of robustness checks. Section 6 concludes.

⁶For literature on the US, see *e.g.* Carlino and DeFina (1998, 1999); Di Giacinto (2003); Owyang and Wall (2009); Beckworth (2010); Furceri et al. (2019); Leahy and Thapar (2019). For literature on Europe, see *e.g.* Arnold (2001); Arnold and Vrugt (2002); Rodriguez-Fuentes and Dow (2003); Dow and Montagnoli (2007).

2 Empirical setup

2.1 Baseline model and identification

To study the dynamic effects of monetary policy on regional output, we apply [Jordà \(2005\)](#)’s local projections method consisting in a set of regressions of the form:

$$y_{i,t+h} = \alpha_i + \beta_h i_t + \gamma_h \mathbf{X}_{i,t} + \delta_h \mathbf{X}_{j,t} + \theta_h \mathbf{X}_{k,t} + \varepsilon_{i,t+h} \quad (1)$$

where the dependent variable $y_{i,t+h}$ denotes real GDP in jurisdiction i and year $t + h$; α_i is a set of region-fixed effects; i_t is the monetary policy-controlled short-term interest rate in year t ; $\mathbf{X}_{i,t}$, $\mathbf{X}_{j,t}$, and $\mathbf{X}_{k,t}$ are vectors of time-variant control variables at the local-, country- and euro area-level, respectively; and $\varepsilon_{i,t+h}$ is an error term. All variables enter equation 1 in log-levels, except for the short-term interest rate which enters as a percentage per annum. In the most parsimonious version of the model, the local control variables include the population of each jurisdiction; the country controls include GDP of the country j in which jurisdiction i is located; and the euro area controls include GDP and HICP at the euro area level. In the below analysis, we augment this baseline model with a host of further covariates to test the robustness of our findings (see Section 5). Further, we test whether our main results carry over to the use of alternative monetary policy indicators, including a set of “shadow interest rates”. The latter help us capture the impact of non-standard measures affecting the policy stance above and beyond the effect of observed short-term interest rate changes, which may be of particular relevance for the final years of our estimation sample.

Our main interest is in the impulse response of local output to a change in the short-term interest rate in year t , as captured by the coefficient β_h for each horizon h . The interpretation of the β_h -coefficients as the causal effect of short-term interest rate changes on local output at horizon h relies on the identifying assumption that, controlling for macroeconomic conditions, monetary policy does not respond to economic activity at the city- and county-level.

This assumption is supported by three considerations. The first relates to the ECB’s monetary policy reaction function. As highlighted in Section 1, the ECB’s monetary policy mandate, with its primary objective to maintain price stability, refers to the euro area as a whole. In keeping with this euro-area wide mandate,

the ECB's Governing Council has formally defined its price stability objective in terms of the aggregate euro area inflation rate. Vice versa, it has explicitly ruled out regional developments as a determinant of its monetary policy conduct (ECB, 1998). Accordingly, by controlling for euro area aggregate inflation and activity in the ECB's reaction function, we are able to partial out the variation in short-term rates that does not reflect the systematic central bank response to the economy. Based on this remaining variation, which effectively amounts to using Taylor-rule residuals to identify exogenous changes in policy-controlled short-term interest rates, we can then estimate the impact of monetary policy shocks on regional GDP. As such, the granular data on economic activity helps us overcome the risk of simultaneity bias, which constitutes one of the key challenges in estimating monetary policy effects on the economy.⁷

A second consideration supporting the identification strategy relates to the information set available to the Governing Council in its monetary policy deliberations. As a key element of this information set, ECB staff provide a comprehensive analysis of prevailing and prospective economic conditions (ECB, 2019a). However, most of this economic analysis focuses on the euro area level and, to the extent that it does consider more granular information, the maximum degree of disaggregation is at the level of individual countries. Therefore, even if policy-makers, in explicit deviation from the stated ECB strategy, wanted to target regional economic developments, they would lack crucial ingredients to taking an informed decision in this regard; and, given the absence of systematic region-specific economic analyses and forecasts by other major policy institutions, this gap would also not be possible to close via alternative information sources.

A third and closely related consideration pertains to the long lags in the publication of regional data. In particular, information on economic activity at the NUTS3 level usually becomes available only around two to three years after the period they refer to. This very long lag implies that, at the time policy rates are being set, even the raw data that would allow decision makers to consider regional economic developments are missing.

The remaining explanatory variables serve to further sharpen this identification approach. The inclusion of country-level GDP further severs the link between the monetary policy reaction function and regional

⁷The assumption that monetary policy responds to aggregate but not to regional shocks also features prominently in the recent model-based analysis of the aggregate implications of regional business cycles by Beraja et al. (2019), as well as in the conceptual discussion of the increasingly widespread use of regional data in macroeconomics by Chodorow-Reich (2019b).

GDP as the dependent variable. The population variable accounts for the large cross-sectional heterogeneity in the size of the regions.⁸ And the region-fixed effects allow us to control for a host of other, unobserved, factors that may confound the causal interpretation of our estimates.

To assess the response of regional economic activity to monetary policy shocks at the sample mean, we estimate equation 1 via ordinary least squares (OLS). Inference is based on [Driscoll and Kraay \(1998\)](#) standard errors that account for cross-sectional and temporal dependencies in the data. In terms of lag-length, we follow the heuristic from the first step of the [Newey and West \(1994\)](#) plug-in procedure, which in our data set implies two lags (see [Hoechle \(2007\)](#) for details). When experimenting with higher lag-lengths, however, we found the estimated standard errors to remain largely unaffected and the confidence intervals, if anything, to become narrower.

2.2 Modeling heterogeneity

Besides estimating the mean response of regional GDP to monetary policy shocks, our aim is to also shed light on how this response differs across the (conditional) GDP distribution. To this end, we rely on two complementary approaches, the first consisting of quantile regressions and the second of sub-sample analysis.

As regards quantile estimation, the seminal approach proposed by [Koenker and Bassett \(1978\)](#) in principle provides a flexible way to model the conditional distribution of the dependent variable. At the same time, it is not ideally suited for panel data models with fixed effects when the cross-section is large relative to the time dimension, in which case the estimates are prone to incidental parameter problems (see, e.g., [Lancaster, 2000](#)).⁹ To address this issue, we therefore employ the quantile regression approach recently proposed by [Machado and Silva \(2019\)](#), which enables us to control for unobserved heterogeneity while estimating

⁸An alternative would be to directly specify the dependent variables in per-capita terms. Compared to this approach, our choice of specification is somewhat more flexible as it does not restrict the coefficient on the population variable to be equal to one. At the same time, when experimenting with log per-capita GDP as dependent variable, we obtained very similar results.

⁹A related issue arises with regard to the interpretation of the coefficients. In our application, the regions populating the upper and lower parts of the GDP distributions *conditional on the fixed effects* may be very different than those in the respective unconditional distributions. For instance, it is possible for relatively prosperous regions to temporarily end up in the bottom part of the distribution in years in which their GDP exhibits a transitory drop relative to its sample mean. However, our intention is to sort the cross-sectional units according to the more persistent aspect of regional inequality, which requires an alternative approach.

quantile-specific coefficients of the covariates in our model via location- and scale-functions.¹⁰ Overall, the choice of methodology places our paper in a growing literature using quantile estimation techniques in macroeconomic applications, as recently popularized by [Adrian, Boyarchenko and Giannone \(2019\)](#) and applied to a local projections model estimated on a country panel by [Adrian, Grinberg, Liang and Malik \(2018\)](#). The latter setting closely corresponds to the approach we pursue in this paper.

For the sub-sample analysis, we group the regions according to their position in the per-capita GDP distribution and allow the coefficients in equation 1 to differ across groups. Specifically, we define a dummy variable D_{it}^d for each decile d that is 1 for all regions i whose per capita GDP falls within this decile in year t and zero otherwise. We then interact this dummy with all explanatory variables in equation 1, including the fixed effects, to obtain decile-specific coefficient estimates (for a similar approach in the context of firm size distributions, see [Crouzet and Mehrotra \(2017\)](#)). Further, we adopt an analogous approach for other quantiles, such as the interquartile range, in the results presented in section 4.

As pointed out by [Koenker and Hallock \(2001\)](#), this type of sub-sample analysis is conceptually distinct from quantile regressions. However, as we show in section 4, the two methods yield mutually consistent and qualitatively similar conclusions in the application considered in the current paper. Further, besides offering an informative cross-check, the sub-sample analysis provides a more flexible framework to characterize differences in the adjustment to policy shocks across regions. In particular, quantile regressions, as an inherent feature, model heterogeneity in relation to the conditional distribution of the dependent variable. But our aim is to not only study heterogeneity in the response of economic activity but also to understand its anatomy, *inter alia* by testing for differential responses in a broader set of dependent variables, such as productivity and employment. While quantile regressions would only allow us to check for differential responses across the distribution for the respective choice of alternative dependent variables, the sub-sample analysis is amenable to testing whether the response of these alternative dependent variables also differs across the per-capita GDP distribution.

¹⁰In the quantile regressions, we again account for potential error correlation across space and time. To this end, we resort to two-way clustering, given the Driscoll-Kraay correction used for the mean regressions is not available for the quantile estimator. As suggested by [Machado and Silva \(2019\)](#), we also tested whether our conclusions may be biased due to the large ratio of cross-sectional units relative to time periods via the split-sample jackknife estimator developed by [Dhaene and Jochmans \(2015\)](#); this also left our results qualitatively unaffected and, if anything, led to somewhat more accentuated heterogeneity in the responses across regions.

3 Data and stylized facts

3.1 Sources and definitions

Our analysis relies on a rich dataset of economic and demographic indicators at the subnational level in Europe. The data are based on Eurostat's Nomenclature of Territorial Units for Statistics (NUTS), which is a hierarchical system for dividing up the economic territory of the EU into four levels. The highest level (NUTS0) corresponds to the nation-state and the lowest (NUTS3), which we use in the ensuing analysis, roughly corresponds to the city- and county-level. In this section, we describe the main features of our data, whereas further information on their construction and sources is available in [A.1](#).

The source of these data is the European Regional Database (ERD) by Cambridge Econometrics. The ERD is based on Eurostat's REGIO database, but closes certain gaps, especially as regards the period prior to 2000, using national statistics from European Commission's AMECO database and interpolation methods.¹¹ In terms of regional variables, our main focus is on gross domestic product, deflated to 2005 price levels, and population size. The NUTS3 data offer the maximum degree of geographical disaggregation for which information on economic activity is available. At the same time, the coverage in terms of economic variables is relatively limited at the NUTS3 level compared to more aggregated data sets.¹²

Our main regional variables are gross domestic product and gross value added, both deflated to 2005 Euros, as well as population and employment. Further, we make use of the breakdown of regional gross value added into six sectors of the economy, corresponding to the disaggregation in NACE Rev.2 as: agriculture, forestry and fishing; industry less construction; construction; financial and business services; wholesale, retail, transport, accommodation and food services, information and communication; and last, non-market services. We merge the sectors of industry and construction to gauge the GVA of the manufacturing sector typically studied in the monetary policy literature. We call this the capital-intensive sector. Similarly, we merge the sectors of financial and business services, wholesale, retail, transport and food services, information and communication and non-market services to gauge the GVA of the services sectors. We call this the

¹¹For further detail, see [Cambridge Econometrics \(2017\)](#).

¹²For instance, data on inflation and unemployment are available only at higher regional aggregation levels; see *e.g.*, [Beck, Hubrich and Marcellino \(2009\)](#), [Özyurt and Dees \(2015\)](#), and [Belke, Haskamp and Setzer \(2016\)](#) for recent studies using these variables.

labor-intensive sector.

We complement this regionally disaggregated information with euro-area and country-level variables for real GDP and HICP from the AMECO database of the European Commission. Following the bulk of the related literature, the 3-month Euribor serves as our measure of the policy-controlled short-term interest rate (see, *e.g.*, [Coenen et al. \(2018\)](#), [Smets and Wouters \(2003\)](#), [Faust et al. \(2003\)](#)). We source this variable from the Area Wide Model (AWM) database. Further, we also rely on the AWM for information on oil prices (expressed in US Dollar per barrel) and on country-specific long-term sovereign interest rates. The nominal effective exchange rate, considering 67 trading partners, is taken from Bruegel.

The sample includes all NUTS3 regions from the eleven initial euro area member states, excluding Luxembourg, over the period 1999-2015 and from Greece over the period 2001-2015. The starting point of the sample corresponds to the year the euro currency was introduced and the end-point to the last year for which NUTS3-level data are currently available.¹³ Our final sample thus consists of 886 NUTS3 regions.

3.2 Heterogeneity across space and time

A comparison of per capita GDP distributions at the regional and national level demonstrates how the degree of heterogeneity intensifies with greater disaggregation (see Table 1).¹⁴ For instance, in 2015, per capita GDP at the country level ranged from €16,623 in Portugal to €51,789 in Ireland.¹⁵ While substantial, this difference is dwarfed by the dispersion in regional per capita GDP, which ranged from €8,455 for Serres, a region in Northern Greece, to €127,390 in Wolfsburg, a large city in the German state of Lower Saxony. Moreover, this pattern of greater heterogeneity at more granular aggregation levels is not just confined to a few extreme outliers, but it is a general feature of the respective distributions. For instance, the coefficient of variation (CV), computed as the ratio of the standard deviation over the mean of per capita GDP, is around 50% higher at the regional than at the national level (Figure 1).

¹³Luxembourg is excluded because it consists of just one NUTS3 region. The sample for Greece starts in 2001 because this is when the country introduced the euro. We also follow the literature in excluding the five NUTS3 regions of the French overseas territories of Guadeloupe, Martinique, French Guiana, La Réunion and Mayotte.

¹⁴Unless otherwise noted, per capita GDP refers to the ratio of real GDP over population throughout the remainder of the paper.

¹⁵As is well-known, the Irish GDP figures tend to be distorted upwards by the activities of multinational companies. Hence, the difference between the national and regional distribution becomes even more striking when considering the Netherlands, which recorded the second highest per capita GDP in 2015 (of €39,184), as the upper bound for the national distribution.

Table 1: Heterogeneity at different aggregation levels

		Mean	St. Dev.	Minimum	Maximum
GDP	<i>Country</i>	31273	10252	16623	51789
	<i>Region</i>	26250	11929	8455	127390
Capital-intensive GVA	<i>Country</i>	6503	2666	2096	10502
	<i>Region</i>	7371	6302	819	91688
Labor-intensive GVA	<i>Country</i>	18580	5913	10087	30807
	<i>Region</i>	15457	7164	4733	70126
Observations	<i>Country</i>	11			
	<i>Region</i>	886			

Notes: Figures refer to real GDP and real GVA per capita in 2015 at the NUTS3 (region) or NUTS0 (country) level. Capital-intensive GVA is calculated as the ratio of GVA in the sectors of industry and construction over total population. Labor-intensive GVA is calculated as the ratio of GVA in the sectors of financial and business services, wholesale, retail, transport, accommodation and food services, information and communication and non-market services over total population.

Further, an even more pronounced degree of dispersion emerges for gross value added (GVA), and in particular for the capital-intensive part of GVA, whose CV at the regional level is twice as high as that at the national level. Thus, the dispersion not only in overall activity but also in its composition become more accentuated at the disaggregated level, consistent with the typical patterns of spatial concentration of different types of productive activity (Krugman, 1991; Glaeser et al., 1995).

Importantly, these disparities also arise within the countries. In fact, within-country dispersion in per capita GDP in some cases reaches similar levels as for the euro area as a whole; and the average within-country dispersion in 2015 was almost two thirds of that observed for the euro area (Figure 1; similar patterns emerge for the country-specific time-series averages as visible from Figure 15 in A.1). Accordingly, the analysis of regional disparities in the euro area is not just another way of looking at cross-country differences (and putting them under a magnifying glass); instead, it addresses an important aspect of the unequal geography of economic activity that transcends the well-known cross-country perspective.

Moreover, notwithstanding some instances of geographical clustering of more and less prosperous regions, the patterns of within-country inequality are fairly nuanced. For instance, in Italy poorer regions tend to cluster in the South, in Germany in the East, and in Spain in the West (Figure 2). As such, some form of core-periphery divide, which is often emphasized in the context of different countries in the euro area,

Figure 1: Coefficient of variation of regional GDP per capita (2015)

Notes: The coefficient of variation (CV) is computed as the ratio of the standard deviation to the mean of all NUTS3 regions within each country in 2015, except for: the bar denoted EA, which refers to the CV over all NUTS3 regions in the sample; and the bar denoted Mean, which refers to the unweighted average of the eleven within-country coefficients of variation displayed in the graph.

also exists at the regional level within some of the countries. However, there are also instances of strong disparities arising between regions with a close geographical proximity, as visible from regions in the lower quartiles punctuating clusters of regions in the upper quartiles, and vice versa. Overall, the patterns of heterogeneity within countries are thus fairly diverse, which implies that a less granular subnational disaggregation of the data, such as by provinces or states, would miss important aspects of regional inequality.

Beyond this static perspective on regional inequality, also its dynamics offer interesting observations, especially in the aftermath of the global financial and economic crisis. This crisis triggered a steep fall in euro area economic activity, followed by a double-dip recession, both of which is clearly visible at the mean of the regional per capita GDP distribution (Figure 3). Similar patterns emerge for the median, although activity in 2011-12 merely stagnates here, rather than contracting as it did at the mean. For the outer parts of the distribution, the dynamics again show nuanced patterns. While the 25th and 75th percentiles have mostly moved in lockstep, and regions in the 25th percentile even weathered the second dip in economic activity over the 2011-12 period better than those in the 75th percentile, vastly different trajectories emerge for the more

Figure 2: Heterogeneity in GDP per capita

Notes: Figures refer to real GDP per capita in 2015 at the NUTS3 level. Different shadings refer to quartiles.

extreme parts of the distribution. For instance, regions in the 90th percentile on average experienced a solid recovery after the 2009 recession, whereas per capita GDP in the 10th percentile just continued drifting down after the crisis and only showed a mild turn-around in the last two sample years. Moreover, this drifting-apart of poorer and richer regions even intensifies when considering the wedge between the 5th and 95th percentile.

Taken together, these stylized facts confirm that the city- and county-level data used in the current paper provide a very interesting setting to study heterogeneity in monetary policy transmission. Of course, the basic fact that disparities in economic structures and performance become particularly accentuated at this level is well-established in the urban economics literature ([Glaeser, Scheinkman and Shleifer, 1995](#)). But it has so far received little attention in the monetary economics literature. The next section closes this gap by presenting evidence on the effects of monetary policy on regional inequality.

Figure 3: Evolution of average per capita GDP in selected percentiles

Notes: The lines show the normalized percentiles of regional GDP per capita in the total sample of EA11. The percentiles have been normalized to 100 in 2008.

4 The heterogeneous impact of monetary policy on regional output

The impulse response functions (IRFs) point to a significant and economically relevant impact of exogenous changes in short-term interest rates on regional GDP; and this impact becomes stronger and more persistent when moving towards the lower parts of the distribution. We next describe these findings in greater detail.

4.1 Monetary policy responses at the mean

As a starting point, and to benchmark our estimates against the related literature, we first focus on the mean response estimated via OLS. The following patterns emerge for a 100 basis point hike in short-term interest rates (Figure 4).¹⁶ After an initial transmission lag, the downward impact on economic activity turns significant at a one-year horizon. This impact further builds up and peaks at the two-year horizon, with the impulse response functions implying a contraction of around 2% in regional GDP in year $t + 2$, and then gradually fades out over the remainder of the IRF horizon. As evident from the tight confidence intervals, the IRFs are estimated with a higher degree of precision than is typically the case at a more aggregate level, which *inter alia* reflects the large number of observations at this granular aggregation level and the

¹⁶Analogous findings obtain for rate cuts as the model is symmetric.

Figure 4: Impact of monetary policy on regional output at the mean

Notes: Vertical axis refers to impact of 100 basis point rate hike on regional GDP (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands.

possibility to control for various sources of heterogeneity, including via fixed effects, which absorb a lot of the cross-sectional variation in the data.

Overall, the estimates closely resemble the familiar patterns for the output response to monetary policy shocks found in the macroeconomic literature for higher levels of aggregation. For instance, the peak effect of a 2% output contraction in response to an exogenous 100 basis point short-term interest rate hike clearly falls within the range of estimates emerging from prominent contributions to the literature studying the US economy, as reviewed in [Ramey \(2016\)](#) (see Table 1 in that paper). The peak effect is also broadly similar to that deriving from the estimated DSGE model of [Smets and Wouters \(2003\)](#), which often serves as a benchmark reference for the literature on the euro area economy. Moreover, the two-year transmission lag between the policy rate hike and its peak impact, as well as the gradual, but incomplete, fading out of the output response in subsequent years is comparable to the patterns in [Smets and Wouters \(2003\)](#). This broad match between the estimated responses is reassuring as it allays potential concerns of time aggregation bias

arising from the annual frequency of our data (Marcet, 1991; Hansen and Sargent, 1991).

4.2 Heterogeneity across quantiles

Moving beyond the mean, the quantile estimates point to pronounced heterogeneity in the regional patterns of monetary policy transmission. Already for the interquartile range, notable differences emerge in terms of both, the peak and the persistence of the impacts (Figure 5). Output, in both parts of the distribution, hits its trough in the second year after the shock, but the point estimates point to a somewhat deeper contraction at the lower than at the upper quartile. Moreover, this gap widens over the horizon, as output in the upper quartile rebounds, with the point estimate becoming statistically insignificant in year $t + 5$, whereas output in the lower quartile recovers at a much slower pace and fails to return to its initial level by the end of the horizon.

Moreover, these differences in the output response become markedly more accentuated when considering more extreme parts of the distribution. For example, for the lower decile the contraction in GDP reaches a trough of -2.5%, compared to -2.0% for the upper decile (Figure 6); and the maximum contraction at the 5th percentile, standing at -2.6%, is more than one third deeper than at the 95th percentile (see Figure 16 in A.2). Also, the contrast between the quick recovery in the upper part of the distribution and the persistent contractionary effect in the lower part of the distribution becomes even starker for these percentile pairs: while output recovers to its pre-shock levels by the end of the horizon for the upper percentiles, the contraction does not reverse for the lower percentiles; for the latter, instead, the point estimates essentially move sideways from horizon $t + 1$ on and remain close to the trough by the end of the horizon. Finally, while the confidence intervals in the interquartile comparison are fairly close together, and over large parts of the horizon overlap, the difference in the outer tails is clearly statistically significant over the later years of the horizon.

Further, these patterns remain intact when relying on sub-sample analysis instead of quantile regressions (Figures 7 and 8, as well as Figure 17 in A.2). In particular, both richer and poorer regions experience a contraction in GDP in response to the monetary policy tightening shock, but the maximum impact and its persistence are substantially higher in latter. Moreover, the estimated extent of contraction in the sub-

Figure 5: Impact of monetary policy on regional output:
25th versus 75th percentile

Figure 6: Impact of monetary policy on regional output:
10th versus 90th percentile

Figure 7: Impact of monetary policy on regional output:
upper vs. lower quartile

Figure 8: Impact of monetary policy on regional output:
bottom versus top decile

Notes: Vertical axis refers to impact of 100 basis point rate hike on regional GDP (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands. Top row refers to estimates from quantile regressions, with grey (blue) lines depicting the estimates for the higher (lower) percentile. Bottom row refers to estimates from the sub-sample analysis, with grey (blue) lines depicting the estimates for the higher (lower) quantile-range.

sample analysis is similar to the quantile regressions, albeit being somewhat more pronounced for poorer regions.¹⁷ As discussed in Section 2.2, this close empirical match between the two approaches is convenient for estimating differential responses across different types of industry (Section 4.3), as well as for exploring the origins of cross-regional differences in the output effects of monetary policy shocks (Section 4.4).

4.3 The role of industry structures

Subnational jurisdictions typically differ markedly in their industry mix and a comprehensive literature has documented that the responsiveness of output to monetary policy shocks tends to differ across different types of industry (see, e.g., [Dedola and Lippi \(2005\)](#); [Peersman and Smets \(2005\)](#); and [Carlino and DeFina \(1998\)](#) for representative studies). Our data and research design allow us to explore whether these key insights from the related literature are confirmed in our more granular setting and whether regional heterogeneity in industry structures reinforces or mitigates the differential responsiveness to monetary policy shocks.

To answer these questions, we again resort to the subdivision of GVA into its capital-intensive and its labor-intensive sub-sectors, with the former consisting of construction and industry and the latter of services (see Section 3.2). Based on this breakdown, we conduct two exercises. First, we estimate industry-specific output responses to monetary policy shocks for the sample mean so as to cross-check our priors on the relative interest-rate sensitivity across industry-types. Second, we repeat this exercise for the cross-regional GDP distribution via sub-sample analysis so as to test whether differential responses also arise *for a given type of industry*.

Consistent with standard notions of the interest rate channel of monetary policy, the impulse responses point to a significantly greater sensitivity of gross value added (GVA) in the capital-intensive sub-sectors (Figure 9, LHS panel). In particular, capital-intensive output contracts more quickly (reaching its trough already in year $t + 1$) and more strongly (with a trough impact roughly double that estimated for labor-intensive output). At the same time, capital-intensive output displays a fairly dynamic V-shaped recovery,

¹⁷The resultant higher degree of heterogeneity in impact estimates in the sub-sample analysis relative to the quantile regressions is consistent with the specific features of these different estimation approaches. While the quantile regressions estimate the impact at specific percentiles, the sub-sample analysis estimates average effects within percentile-buckets. Given the impact of policy on regional GDP strengthens in monotonous fashion, the contraction estimated at the 10th percentile in the quantile regressions tends to be less pronounced than that estimated for the average of the lowest decile as per the sub-sample analysis.

whereas labor-intensive output remains near its trough for longer and still remains below its initial level through year $t = 5$, thus pointing to some mild hysteresis, which we study in greater detail below.

Figure 9: Capital-intensive vs. labor-intensive sectors

Notes: In the left-hand side (LHS) panel, vertical axis refers to impact of 100 basis point rate hike on regional output in the capital-intensive sector (in blue) and in the labor-intensive output (in grey) at the sample mean (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands. In the right-hand side (RHS) panel, vertical axis refers to the ratio of output in the capital-intensive sector over the output in the labor-intensive sector, averaged over the sample period. Horizontal axis refers to quantiles of the regional per-capita GDP distribution.

A natural question is how the differential responsiveness across industry types interacts with our previous results. For instance, if regions in the lower part of the GDP distribution were to also specialize in capital-intensive production, we could not be sure whether their stronger sensitivity to monetary policy shocks reflects the former or the latter characteristic. However, two pieces of evidence speak against industry structure as an explanation for the greater sensitivity of poorer regions. First, they in fact exhibit a lower share of capital-intensive production (Figure 9, RHS panel). Second, even within each type of industry, the differential response across the distribution remains intact: both the capital- and the labor-intensive production contracts more strongly and more persistently in the lower than in the upper part of the distribution (Figure 10); and, within each group of regions, the relative responsiveness across capital- and labor-intensive production is qualitatively similar. Against this background, the inherent differences across more and less prosperous regions also carry over to the industry-breakdown and so does the finding of monetary policy

Figure 10: Impact of monetary policy on regional output:
bottom versus top decile

Notes: Vertical axis refers to impact of 100 basis point rate hike (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands. The LHS (RHS) chart refers to the GVA of capital-intensive (labor-intensive) sub-sectors. Grey (blue) lines refer to average effects in top (bottom) deciles.

exerting persistent effects on output.

4.4 Sources of hysteresis

The most striking feature of the results so far is that substantial parts of the regional distribution experience long-lived effects of monetary policy shocks on output. This result clearly contrasts with the common notion of monetary policy as merely causing transitory adjustments in the real economy and, as such, it links to a long-standing debate on potential sources of long-term monetary non-neutrality. This debate has enjoyed a revival in the aftermath of the 2007/2008 financial crisis, which brought renewed urgency to the question whether contractions in aggregate demand may give rise to hysteresis, consisting in lasting declines in the productive capacity of the economy (Yellen, 2016). If so, monetary-policy induced changes in activity may also prove persistent and several recent papers support this conjecture. Most closely related to our paper, Jordà et al. (2020) provide empirical evidence of long-lived effects of monetary policy on output in a panel

of advanced economies using a local projections framework.¹⁸

An aspect common to these papers is that they tend to adopt a broad perspective regarding the sources of hysteresis, moving beyond the initial emphasis on the labor market in the seminal contribution by [Blanchard and Summers \(1986\)](#) to also consider productivity and capital accumulation as potential candidates. The debate on which of these potential sources of hysteresis matters (most) is far from settled however. For instance, [Jordà et al. \(2020\)](#) find hysteresis in the impact of monetary policy on the capital stock and on total factor productivity, but not on labor, whereas [Blanchard \(2018\)](#) points to labor markets as an important origin of hysteresis. This question has important policy implications as to how hysteresis should be modeled and what policy options appear best suited to address it.¹⁹

Here, we seek to add to this debate by again exploiting the granularity of our data to study the sources of hysteresis. To this end, we re-run our model separately for employment and labor productivity (with the latter being defined as GDP over employment). We first focus on the sub-sample analysis for the bottom decile, which we know from previous exercises to exhibit pronounced hysteresis in the response of output to monetary policy, and then broaden the analysis to further parts of the distribution.

The results point to strong hysteresis in employment, but also labor productivity displays a lasting downward adjustment in response to monetary policy shocks (Figure 11). Employment monotonously falls in response to the shock and, after hitting a trough in year $t + 4$, only marginally recovers to still stand around 2% below its initial level by the end of the horizon. At the same time, also labor productivity contracts in a persistent manner, albeit less sharply and with a somewhat more pronounced rebound at the end of the horizon. Since our data cover a narrower set of variables than those of [Jordà et al. \(2020\)](#), we cannot disentangle the other factors into whether they reflect a lasting erosion of the capital stock or of total factor productivity (TFP). However, to be consistent with the estimated responses of employment and labor productivity, the combined adjustment in the capital stock and TFP has to also consist in a long-lived contraction (conversely, if capital and TFP would remain constant, labor productivity would rise given the drop in employment would

¹⁸For further empirical evidence see, *e.g.*, [Blanchard et al. \(2015\)](#) and for model-based analysis see, *e.g.*, [Reifschneider et al. \(2015\)](#).

¹⁹For instance, consistent with their emphasis on labor market hysteresis, [Blanchard and Summers \(1986\)](#) focus on the wage-setting process in modeling hysteresis. [Reifschneider et al. \(2015\)](#) instead also consider hysteresis in capital deepening and multi-factor productivity in their extension of the FRB/US model of the Federal Reserve Board staff.

raise the capital intensity).

Further zooming in on employment, the result of strong hysteresis emerges as a fairly broad-based phenomenon, applying to large parts of the distribution. In fact, also at the sample mean employment remains almost 1% below its initial level by the end of the horizon and only shows tentative signs of recovery after hitting its trough in the preceding year (Figure 11). Finally, long-lasting employment effects even materialise in the upper half of the distribution, with the coefficient still remaining significant for regions in the sevenths decile and only fully reversing in the outer tails (Figure 13). Hysteresis in labor productivity by contrast is less pervasive, as already the third decile displays a full recovery over the horizon (Figure 14).

Taken together, these findings yield several important insights. First, the result of monetary policy exerting long-lived effects on output and employment adds to a growing literature that challenges the dominant paradigm of stabilization policies as merely smoothing out fluctuations in these variables around some natural levels and that, more broadly, casts doubt on the sharp distinction between cyclical and structural drivers of economic outcomes that acts as a mainstay of this paradigm (see, *e.g.*, [Dupraz, Nakamura and Steinsson \(2019\)](#) for a recent contribution). Second, the heterogeneous incidence of output hysteresis, with monetary policy exerting persistent effects in the lower parts and only transitory ones in the upper parts of the distribution, implies that monetary policy may have a long-lasting impact on regional inequality. In particular, it implies that monetary policy tightening aggravates regional inequality and policy easing mitigates it. Third, the findings favour theories of hysteresis that extend beyond the labor market (see, *e.g.*, [Moran and Queralto \(2018\)](#) and [Queralto \(2019\)](#) for recent examples). At the same time, the persistent response of employment, even at the sample mean, still points to labor market hysteresis as an important factor in the euro area context.²⁰

²⁰Different sources of hysteresis may also interact. For instance, endogenous growth theory emphasises human capital spillovers that enhance productivity ([Lucas, 1988](#)). These spillovers are likely to decline when labor market hysteresis sets in and, therefore, fewer workers have mutual work-related exposures that may allow them to learn from each other. These mechanisms, in turn, are likely to be particularly relevant at more granular aggregation levels, such as the cities and countries considered in the current paper, given the geographic proximity of economic agents in these settings ([Glaeser, Scheinkman and Shleifer, 1995](#)).

Figure 11: Impact of monetary policy on employment and labor productivity in the bottom decile

Figure 12: Impact of monetary policy on employment at the mean

Figure 13: Impact on employment at horizon $h = 5$ across quantile-ranges

Figure 14: Impact on labor productivity at horizon $h = 5$ across quantile-ranges

Notes: Vertical axis refers to impact of a 100 basis point rate hike on the respective dependent variable (in %), horizontal axis refers to horizon of IRF (in years). In top-LHS panel, the IRFs are for regional employment (in grey) and labor productivity (in blue) in the bottom decile of the per-capita GDP distribution; in top-RHS panel, the IRFs are for regional employment at the sample mean. Solid lines in the top row denote point estimates and shaded areas denote 90% confidence bands. Bottom-LHS panel shows impacts of a 100 basis point rate hike on employment at horizon $h = 5$ (in %) for the decile ranges from: 0-10; 20-30; 60-70; and 90-100 (in that order from left to right); bottom-RHS panel shows corresponding impacts on labor productivity. Diamonds indicate point estimates and bars indicate 90% confidence intervals.

5 Robustness

To check the robustness of these findings, we re-estimated the quantile and sub-sample regressions with the following modifications to our baseline specification: first, we extend the list of covariates in equation 1 with forward-looking indicators of economic activity and prices; second, we replace the 3-month Euribor with shadow rates as an alternative indicator of the monetary policy stance; and, third, we enrich the model with a set of additional asset- and commodity-price variables. As we report in greater detail below, these modifications leave the main conclusions from our baseline estimations intact.

5.1 Forward-looking indicators of activity and prices

By including aggregate measures of economic activity and prices in equation 1, our identification strategy essentially amounts to using Taylor-rule residuals as a measure of exogenous changes in policy-controlled interest rates. A potential concern regarding this choice of specification is that the central bank reaction function may not (only) depend on the realized but (also) on the expected evolution of these variables (Sims, 1992; Romer and Romer, 2004). In fact, the ECB has emphasized that its monetary policy has a “medium-term orientation”, which means that “monetary policy needs to act in a forward-looking manner” (Issing, 1998; ECB, 1998, 2019b). Accordingly, a natural robustness check is to augment our baseline specification with forward-looking indicators of economic activity and prices that are likely to shape the ECB Governing Council’s monetary policy deliberations. In terms of implementation, we construct these variables from the quarterly ECB staff macroeconomic projections, as a rule using the final year of the respective projection horizon as the relevant reference point and computing a weighted average across the four projection vintages of each year (see A.3 for additional detail).

The resultant estimates are very similar to our baseline (Figure 18 in A.4). For the sample mean, the confidence intervals of the two specifications overlap throughout the IRF horizon; and, for the quantile regressions and sub-sample analysis, the alternative specification confirms the weaker and less persistent response of economic activity in the upper than in the lower parts of the distribution. The upper percentiles, in the alternative specification, stage a somewhat more vigorous recovery over the last years of the IRF horizons than in the baseline; and the lower percentiles experience an even stronger contraction, with the point

estimate for the last year in the IRF horizon remaining close to the trough for the quantile regressions. These differences to the baseline are statistically insignificant, however. Against this background, our baseline findings prove robust to the inclusion of forward-looking control variables for prices and activity.

5.2 Shadow rates

A second set of potential concerns relates to the use of short-term interest rates as our indicator of the monetary policy stance. While standard in the related literature, this choice may prove overly narrow when short-term interest rates approach their lower bound and central banks resort to other instruments to inject additional accommodation. This constellation also applies to our sample period, especially towards its end, as the ECB has expanded its monetary policy toolkit by a set of non-standard measures, such as quantitative easing, which it explicitly motivated by a reduced scope for further rate cuts (ECB, 2015; Rostagno et al., 2019). Accordingly, it appears worthwhile testing whether our findings also hold for broader measures of the monetary policy stance.

To this end, we re-estimate our baseline model replacing the 3-month Euribor with a set of shadow short-term interest rate measures, which serve as proxies of the monetary policy stance in the presence of lower bound constraints and non-standard measures (Krippner, 2013; Wu and Xia, 2016; Krippner, 2019). In terms of specific measures, we rely on two variants, the first based on Lemke and Vladu (2017) and the second on Krippner (2013). This specific choice allows us to span the very wide range of estimates provided in the related literature (see, for instance, Figure 26 in Hartmann and Smets (2018) for an overview). Among these estimates, the Lemke and Vladu (2017) shadow rates fall below the actual short-term interest rate only from 2015 onward and the absolute difference to the actual short-term rate over our sample period peaks at around 1 percentage point. The shadow rate by Krippner instead starts falling significantly below the EONIA rate already towards end-2011/early-2012 and the maximum distance over our sample period almost reaches 4.5 percentage points. By considering both of these measures in our robustness checks, we thus accommodate the pronounced model- and estimation-uncertainty that surrounds the construction of shadow-rate measures.

Again our baseline findings prove robust to the change in specification. In the mean regressions, the IRFs based on the Lemke-Vladu shadow rate are almost identical to the baseline and the confidence intervals

overlap throughout all horizons (Figure 19). As this shadow rate deviates from actual short-term rates only in the last sample year, however, the limited relevance for our estimates is not particularly surprising. The specification including the Krippner shadow rate points to a slightly deeper and more protracted trough in the impact of policy rate changes on regional GDP than the baseline (Figure 20); but the confidence intervals between the two specifications overlap over the entire horizon. For both shadow-rate measures, the quantile estimations and sub-sample analysis in turn confirm the familiar pattern of a stronger and more persistent response of economic activity in the lower than in the upper parts of the distribution (Figure 19 and 20).

5.3 Additional covariates

Finally, to guard against omitted variable bias, we extend the fairly parsimonious baseline model with a set of factors that may correlate with both, regional activity and policy rates. Prominent examples include global oil prices and exchange rates against major trading partners. These variables routinely feature in ECB policy communication;²¹ and they appear as potentially important controls at a disaggregated level where the typical spatial concentration of different types of economic activity may render certain regions more and others less responsive to these types of shocks (see, *e.g.* House, Proebsting and Tesar (2019) for a recent analysis documenting differential effects of exchange rate fluctuations across US regions).²² Further, we include the long-term government bond yield spread vis-à-vis Germany in the list of explanatory variables to control for differences in country-specific risk premia, which played a dominant role especially during the euro area sovereign debt crisis.

With some nuances, the main conclusions from our baseline estimations also carry over when including these additional regressors. For the mean, the IRFs display very similar patterns over a horizon of up to year

²¹In particular oil price developments are a regular element in the economic assessment communicated via the Introductory Statements to the Governing Council's monetary policy press conferences. Exchange rates feature less regularly, but do occasionally enter the Introductory Statements, such as on 25 January 2018, when the Governing Council communicated that: "*recent volatility in the exchange rate represents a source of uncertainty which requires monitoring with regard to its possible implications for the medium-term outlook for price stability*". Moreover, the relevance of exchange rates has, at various occasions, been emphasized in the question-and-answer sessions during these press conferences by ECB Presidents. For instance, President Trichet, on 7 October 2010, pointed out that "*excess volatility and disorderly movements in exchange rates have adverse implications for economic and financial stability*"; and President Draghi, on 25 January 2018, stated that "*exchange rates are important for growth and for price stability*".

²²Similar effects appear plausible also for the euro area; for instance, Lane and Stracca (2018) document heterogeneity in exchange rate pass through across euro area countries and, given the large differences in economic structures at the NUTS3 level, such heterogeneity is likely to arise also here.

$t + 3$. After that point, the model including the additional variables yields a somewhat faster recovery of the output losses, but the difference to the baseline is again statistically insignificant. The quantile estimates and sub-sample analysis, in turn, confirm the differential response across the GDP-distribution, characterized by a significantly stronger and more persistent contraction in economic activity in the lower than in the upper parts.

6 Conclusion

This paper has provided a novel perspective on the regional patterns of monetary policy transmission. Using geographically disaggregated data on economic activity in Europe, we have shown that the output response to short-term interest rate shocks is significantly more pronounced and persistent in poorer than in richer cities and counties. Moreover, while GDP in the upper part of the distribution returns to its pre-shock level after four to five years, the response in the lower parts does not reverse over this period, thus pointing to pronounced hysteresis in output. The heterogeneous incidence of hysteresis in turn implies that monetary policy has a long-lasting impact on regional inequality, with tightening shocks aggravating and easing shocks mitigating it.

In terms of anatomy, we find hysteresis to originate from long-lived adjustments in both employment and labor productivity. At the same time, employment hysteresis is more pronounced and more broad-based across the distribution. In fact, it even extends to the sample mean, as opposed to productivity hysteresis which concentrates in the lower tails of the distribution. As such, our findings confirm labor markets as an important source of hysteresis in the European context.

Our paper points to the use of geographically disaggregated data as a promising avenue for further insights into how exactly monetary policy shocks propagate to the economy. First, by resorting to information on economic activity at a more granular level than that entering the central bank reaction function, it offers a novel strategy to identify exogenous changes in monetary policy. Second, by providing empirical estimates for the monetary policy impact on regional inequality, it closes an important gap in the large and growing literature on heterogeneity in monetary policy transmission.

From a policy perspective, our findings underscore the challenges of calibrating monetary policy in heterogeneous economies. In the euro area context, the debate has typically interpreted these challenges as a cross-country phenomenon. As such, they attracted particular attention during the euro area sovereign debt crisis from 2010-12, which was marked by strong cross-country divergence in economic performance and raised concerns as to the suitability of a given aggregate monetary policy stance for individual countries. However, our analysis demonstrates that the issue runs deeper: interregional heterogeneity becomes more accentuated at more granular geographical levels and this heterogeneity in turn profoundly alters the implications of a given monetary policy stance in different parts of the economy. These implications emerge as particularly relevant in view of our finding that monetary policy exerts durable impacts on output and employment – a finding that contrasts with the common notion of stabilization policies merely smoothing out fluctuations in these variables around some natural levels.

References

- Adrian, T., Boyarchenko, N. and Giannone, D. (2019). Vulnerable Growth, *American Economic Review* **109**(4): 1263–1289.
- Adrian, T., Grinberg, F., Liang, N. and Malik, S. (2018). The Term Structure of Growth-at-Risk, *IMF Working Papers 18/180*, International Monetary Fund.
- Altavilla, C., Canova, F. and Ciccarelli, M. (2019). Mending the Broken Link: Heterogeneous Bank Lending Rates and Monetary Policy Pass-Through. *Journal of Monetary Economics* (forthcoming).
- Ampudia, M., Georgarakos, D., Slacalek, J., Tristani, O., Vermeulen, P. and Violante, G. (2018). Monetary Policy and Household Inequality. *ECB Working Paper No. 2170*.
- Arnold, I. J. (2001). The Regional Effects of Monetary Policy in Europe, *Journal of Economic Integration* **16**(3): 399–420.
- Arnold, I. J. and Vrugt, E. B. (2002). Regional Effects of Monetary Policy in the Netherlands, *International Journal of Business and Economics* **1**(2): 123–134.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel, *American Economic Review* **109**(6): 2333–67.
- Auerbach, A. J. and Gorodnichenko, Y. (2012). Fiscal Multipliers in Recession and Expansion, in A. Alesina and F. Giavazzi (eds), *Fiscal Policy After the Financial Crisis*, University of Chicago Press, pp. 63–98.
- Austin, B., Glaeser, E. and Summers, L. H. (2018). Saving the Heartland: Place-Based Policies in 21st Century America. *Brookings Papers on Economic Activity*, Spring 2018.
- Beck, G. W., Hubrich, K. and Marcellino, M. (2009). Regional Inflation Dynamics Within and Across Euro Area Countries and a Comparison with the United States, *Economic Policy* **24**(57): 142–184.
- Beckworth, D. (2010). One Nation Under the Fed? The Asymmetric Effects of US Monetary Policy and its

- Implications for the United States as an Optimal Currency Area, *Journal of Macroeconomics* **32**(3): 732–746.
- Belke, A., Haskamp, U. and Setzer, R. (2016). Regional Bank Efficiency and its Effect on Regional Growth in “Normal” and “Bad” Times, *Economic Modelling* **58**: 413–426.
- Beraja, M., Fuster, A., Hurst, E. and Vavra, J. (2018). Regional Heterogeneity and the Refinancing Channel of Monetary Policy, *Quarterly Journal of Economics* **134**(1): 109–183.
- Beraja, M., Hurst, E. and Ospina, J. (2019). The Aggregate Implications of Regional Business Cycles, *Econometrica* **87**(6): 1789–1833.
- Blanchard, O. (2018). Should We Reject the Natural Rate Hypothesis?, *Journal of Economic Perspectives* **32**(1): 97–120.
- Blanchard, O., Cerutti, E. and Summers, L. (2015). Inflation and Activity—Two Explorations and their Monetary Policy Implications, *NBER Working Paper No. 21726*.
- Blanchard, O. J. and Summers, L. H. (1986). Hysteresis and the European Unemployment Problem, *NBER Macroeconomics Annual* **1**: 15–78.
- Bluedorn, J., Kóczán, Z., Lian, W., Novta, N. and Timmer, Y. (2019). Closer Together or Further Apart? Subnational Regional Disparities and Adjustment in Advanced Economies, *World Economic Outlook, October 2019*, International Monetary Fund.
- Boeckx, J., Dossche, M. and Peersman, G. (2017). Effectiveness and Transmission of the ECB’s Balance Sheet Policies, *International Journal of Central Banking* **2013**(1): 297–333.
- Cambridge Econometrics (2017). *European Regional Database: Technical Note*.
- Carlino, G. and DeFina, R. (1998). The Differential Regional Effects of Monetary Policy, *Review of Economics and Statistics* **80**(4): 572–587.

- Carlino, G. and DeFina, R. (1999). The Differential Regional Effects of Monetary Policy: Evidence from the US, *Journal of Regional Science* **39**(2): 339–358.
- Chernozhukov, V. and Hansen, C. (2008). Instrumental Variable Quantile Regression: A Robust Inference Approach, *Journal of Econometrics* **142**(1): 379–398.
- Chodorow-Reich, G. (2019a). Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?, *American Economic Journal: Economic Policy* **11**(2): 1–34.
- Chodorow-Reich, G. (2019b). Regional Data in Macroeconomics: Some Advice for Practitioners. *NBER Working Paper No. 26501*.
- Clemens, J. and Miran, S. (2012). Fiscal Policy Multipliers on Subnational Government Spending, *American Economic Journal: Economic Policy* **4**(2): 46–68.
- Coenen, G., Karadi, P., Schmidt, S. and Warne, A. (2018). The New Area-Wide Model II: an Extended Version of the ECB’s Micro-Founded Model for Forecasting and Policy Analysis with a Financial Sector. *ECB Working Paper No. 2200*.
- Coeuré, B. (2018). The Local Impact of the ECB’s Monetary Policy. *European Central Bank*. Retrieved on 14 November 2019 at <<https://www.ecb.europa.eu/press/key/date/2018/html/ecb.sp181004.en.html>>.
- Coibion, O., Gorodnichenko, Y., Kueng, L. and Silvia, J. (2017). Innocent Bystanders? Monetary Policy and Inequality, *Journal of Monetary Economics* **88**: 70–89.
- Crouzet, N. and Mehrotra, N. (2017). Small and Large Firms over the Business Cycle. *Federal Reserve Bank of Minneapolis Working Paper No. 741*.
- Dedola, L. and Lippi, F. (2005). The Monetary Transmission Mechanism: Evidence from the Industries of Five OECD Countries, *European Economic Review* **49**(6): 1543–1569.
- Dhaene, G. and Jochmans, K. (2015). Split-Panel Jackknife Estimation of Fixed-Effect Models, *The Review of Economic Studies* **82**(3): 991–1030.

- Di Giacinto, V. (2003). Differential Regional Effects of Monetary Policy: a Geographical SVAR Approach, *International Regional Science Review* **26**(3): 313–341.
- Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A. and Yao, V. (2017). Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging, *American Economic Review* **107**(11): 3550–88.
- Dow, S. C. and Montagnoli, A. (2007). The Regional Transmission of UK Monetary Policy, *Regional Studies* **41**(6): 797–808.
- Driscoll, J. C. and Kraay, A. C. (1998). Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data, *Review of Economics and Statistics* **80**(4): 549–560.
- Dupraz, S., Nakamura, E. and Steinsson, J. (2019). A Plucking Model of Business Cycles. *NBER Working Paper No. 26351*.
- ECB (1998). A Stability-Oriented Monetary Policy Strategy for the ESCB. *European Central Bank*. Retrieved on 14 November 2019 on <<https://www.ecb.europa.eu/press/key/date/1998/html/sp981110.en.html>>.
- ECB (2015). The Governing Council’s Expanded Asset Purchase Programme. *European Central Bank Economic Bulletin, Issue 1/2015, Box 1*.
- ECB (2019a). Macroeconomic Projections. *European Central Bank*. Retrieved on 14 November 2019 at <<https://www.ecb.europa.eu/pub/projections/html/index.en.html>>.
- ECB (2019b). Medium-Term Orientation. *European Central Bank*. Retrieved on 14 November 2019 at <<https://www.ecb.europa.eu/mopo/strategy/princ/html/orientation.en.html>>.
- Eichenbaum, M., Rebelo, S. and Wong, A. (2018). State-Dependent Effects of Monetary Policy: The Refinancing Channel. *NBER Working Paper No. 25152*.

European Commission (2019). *Urban Data Platform Plus*.

Retrieved on 14 November 2019 from <urban.jrc.ec.europa.eu>.

Faust, J., Rogers, J. H., Swanson, E. and Wright, J. H. (2003). Identifying the Effects of Monetary Policy Shocks on Exchange Rates using High Frequency Data, *Journal of the European Economic Association* **1**(5): 1031–1057.

Francis, N., Owyang, M. T. and Sekhposyan, T. (2012). The Local Effects of Monetary Policy, *The BE Journal of Macroeconomics* **12**(2): 1–38.

Fratantoni, M. and Schuh, S. (2003). Monetary Policy, Housing, and Heterogeneous Regional Markets, *Journal of Money, Credit and Banking* **35**(4): 557–589.

Furceri, D., Mazzola, F. and Pizzuto, P. (2019). Asymmetric Effects of Monetary Policy Shocks Across US States, *Papers in Regional Science* **98**(5): 1861–1891.

Ganong, P. and Shoag, D. (2017). Why Has Regional Income Convergence in the U.S. Declined?, *Journal of Urban Economics* **102**: 76 – 90.

Gertler, M. and Gilchrist, S. (1994). Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms, *The Quarterly Journal of Economics* **109**(2): 309–340.

Glaeser, E. L., Scheinkman, J. and Shleifer, A. (1995). Economic Growth in a Cross-Section of Cities, *Journal of Monetary Economics* **36**(1): 117–143.

Hansen, L. and Sargent, T. (1991). Two Difficulties in Interpreting Vector Autoregressions, in L. Hansen and T. Sargent (eds), *Rational Expectations Econometrics*, Westview Press, Boulder London.

Hartmann, P. and Smets, F. (2018). The First Twenty Years of the European Central Bank: Monetary Policy. *ECB Working Paper No. 2219*.

Hoechle, D. (2007). Robust Standard Errors for Panel Regressions With Cross-Sectional Dependence, *The Stata Journal* **7**(3): 281–312.

- House, C. L., Proebsting, C. and Tesar, L. L. (2019). Regional Effects of Exchange Rate Fluctuations. *NBER Working Paper No. 26071*.
- Iammarino, S., Rodriguez-Pose, A. and Storper, M. (2018). Regional Inequality in Europe: Evidence, Theory and Policy Implications, *Journal of Economic Geography* **19**(2): 273–298.
- Issing, O. (1998). The European Central Bank on the Eve of EMU. *European Central Bank Speeches*.
- Jacobs, J. (1969). *The Economy of Cities*, Vintage Books.
- Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections, *American Economic Review* **95**(1): 161–182.
- Jordà, Ò., Singh, S. R. and Taylor, A. M. (2020). The Long-Run Effects of Monetary Policy, Federal Reserve Bank of San Francisco.
- Kaplan, G., Moll, B. and Violante, G. L. (2018). Monetary Policy according to HANK, *American Economic Review* **108**(3): 697–743.
- Kashyap, A. K., Lamont, O. A. and Stein, J. C. (1994). Credit Conditions and the Cyclical Behavior of Inventories, *The Quarterly Journal of Economics* **109**(3): 565–592.
- Kashyap, A. K. and Stein, J. C. (2000). What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?, *American Economic Review* **90**(3): 407–428.
- Koenker, R. and Hallock, K. F. (2001). Quantile Regression, *Journal of Economic Perspectives* **15**(4): 143–156.
- Koenker, R. W. and Bassett, G. (1978). Regression Quantiles, *Econometrica* **46**(1): 33–50.
- Krippner, L. (2013). Measuring the Stance of Monetary Policy in Zero Lower Bound Environments, *Economics Letters* **118**(1): 135–138.
- Krippner, L. (2019). A Note of Caution on Shadow Rate Estimates. *Journal of Money, Credit and Banking*.

- Krugman, P. (1991). Increasing Returns and Economic Geography, *Journal of Political Economy* **99**(3): 483–499.
- Krugman, P. (1998). Space: the Final Frontier, *Journal of Economic Perspectives* **12**(2): 161–174.
- Lancaster, T. (2000). The Incidental Parameter Problem Since 1948, *Journal of Econometrics* **95**(2): 391–413.
- Lane, P. R. and Stracca, L. (2018). Can Appreciation be Expansionary? Evidence from the Euro Area, *Economic Policy* **33**(94): 225–264.
- Leahy, J. V. and Thapar, A. (2019). Demographic Effects on the Impact of Monetary Policy. *NBER Working Paper No. 26324*.
- Lemke, W. and Vladu, A. L. (2017). Below the Zero Lower Bound: a Shadow-Rate Term Structure Model for the Euro Area. *ECB Working Paper No. 1991*.
- Lenza, M. and Slacalek, J. (2018). How Does Monetary Policy Affect Income and Wealth Inequality? Evidence from Quantitative Easing in the Euro Area. *ECB Working Paper No. 2190*.
- Lucas, R. (1988). On the Mechanics of Economic Development, *Journal of Monetary Economics* **22**(1): 3–42.
- Machado, J. A. and Silva, J. S. (2019). Quantiles via Moments, *Journal of Econometrics* **213**(1): 145–173.
- Marcet, A. (1991). Time Aggregation of Econometric Time Series, in L. Hansen and T. Sargent (eds), *Rational Expectations Econometrics*, Westview Press, Boulder and London.
- Martin, R., Sunley, P., Gardiner, B., Evenhuis, E. and Tyler, P. (2018). The City Dimension of the Productivity Growth Puzzle: the Relative Role of Structural Change and Within-Sector Slowdown, *Journal of Economic Geography* **18**(3): 539–570.
- Moran, P. and Queralto, A. (2018). Innovation, Productivity, and Monetary Policy, *Journal of Monetary Economics* **93**: 24–41.

- Nakamura, E. and Steinsson, J. (2014). Fiscal Stimulus in a Monetary Union: Evidence from US Regions, *American Economic Review* **104**(3): 753–92.
- Newey, W. K. and West, K. D. (1994). Automatic Lag Selection in Covariance Matrix Estimation, *The Review of Economic Studies* **61**(4): 631–653.
- Ottonello, P. and Winberry, T. (2018). Financial Heterogeneity and the Investment Channel of Monetary Policy. *NBER Working Paper No. 24221*.
- Owyang, M. T. and Wall, H. J. (2009). Regional VARs and the Channels of Monetary Policy, *Applied Economics Letters* **16**(12): 1191–1194.
- Özyurt, S. and Dees, S. (2015). Regional Dynamics of Economic Performance in the EU: to What Extent Spatial Spillovers Matter? *ECB Working Paper No. 1870*.
- Peek, J. and Rosengren, E. S. (1995). Bank Lending and the Transmission of Monetary Policy, *Conference Series: Federal Reserve Bank of Boston*, Vol. 39, Federal Reserve Bank of Boston, pp. 47–68.
- Peersman, G. and Smets, F. (2005). The Industry Effects of Monetary Policy in the Euro Area, *The Economic Journal* **115**(503): 319–342.
- Queralto, A. (2019). A Model of Slow Recoveries from Financial Crises. *Journal of Monetary Economics*.
- Ramey, V. (2016). Macroeconomic Shocks and Their Propagation, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, pp. 71–162.
- Regulation (EU) No 549/2013* (2013). *Official Journal of the European Union*.
- Reifschneider, D., Wascher, W. and Wilcox, D. (2015). Aggregate Supply in the United States: Recent Developments and Implications for the Conduct of Monetary Policy, *IMF Economic Review* **63**(1): 71–109.
- Rodriguez-Fuentes, C. and Dow, S. (2003). EMU and the Regional Impact of Monetary Policy, *Regional Studies* **37**(9): 969–980.

- Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications, *American Economic Review* **94**(4): 1055–1084.
- Rostagno, M., Altavilla, C., Carboni, G., Lemke, W., Motto, R., Saint-Guilhem, A. and Yiangou, J. (2019). A Tale of Two Decades: The ECB's Monetary Policy at 20. *ECB Working Paper No. 2346*.
- Sims, C. A. (1992). Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy, *European Economic Review* **36**(5): 975–1000.
- Smets, F. and Wouters, R. (2003). An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area, *Journal of the European Economic Association* **1**(5): 1123–1175.
- The Economist* (2016). Place-Based Economic Policies as a Response to Populism. Print edition, 15 December 2016.
- Retrieved on 14 November 2019 at <<https://www.economist.com/finance-and-economics/2016/12/15/place-based-economic-policies-as-a-response-to-populism>>.
- Wu, J. C. and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound, *Journal of Money, Credit and Banking* **48**(2-3): 253–291.
- Yellen, J. L. (2016). Macroeconomic Research After the Crisis. *The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics*, 60th Annual Economic Conference, Board of Governors of the Federal Reserve System Speeches.

A Appendix

A.1 Additional detail on the regional data

The regionally disaggregated data are based on a harmonized and integrated breakdown of territorial units established by a European Parliament and Council Regulation ([*Regulation \(EU\) No 549/2013*, 2013](#)). This regulation ensures comparability of classifications across countries and follows two principles. First, it sets a range for the number of inhabitants that each NUTS region should comprise; NUTS3 regions, for instance, should be drawn up in a way that their population falls between 150,000 and 800,000 inhabitants. Second, it relies on existing administrative units, where available; for instance, if the administrative structure of a country includes counties, then these are used to construct NUTS3 regions for municipalities whose population falls below the minimum number of inhabitants; in countries without counties, the NUTS3 regions are created by aggregating smaller administrative units such that they meet the population requirement. The regulation is legally binding, so member states have to comply as a matter of Union law. Moreover, regular revisions to the NUTS classification ensure that any changes in national administrative regions are reflected in the disaggregation without compromising the comparability of subnational indicators.

The data are sourced from the European Regional Database of Cambridge Econometrics, which is based on Eurostat's REGIO database, and are supplemented with data from the European Commission's AMECO database. The database uses the NUTS 2013 classification, and covers the statistical territory of EU27 and Norway for the period between between 1980, or 1990 for the new member states, and 2015. Nominal measures of gross value added (GVA), gross domestic product (GDP), compensation of employees, and gross fixed capital formation (GFCF) are deflated to 2005 constant price euros using sectoral price deflators obtained from AMECO. Employment and population are also available at the highest level of regional disaggregation. Sectoral disaggregation refers to the sectors of agriculture, forestry and fishing; industry less construction; wholesale, retail, transport; accommodation and food services; information and communication; financial and business services; and finally, non-market services.

An important adjustment we have made to the Cambridge Econometrics dataset refers to the regions that belong to more than one NUTS territorial level. By construction, Cambridge Econometrics have ex-

Table 2: Summary statistics across full sample period (1999 – 2015)

	Mean	Standard deviation	Minimum	Maximum
GDP	25079	10534	7900	127391
Population	349.9	466.8	8.44	6418
HICP rate	1.78	1.04	-1.69	5.28
Short-term rate	2.22	1.55	-0.02	4.64
NEER 67	96.57	5.21	83.21	105.7
Non-oil commodity prices	128.3	38.78	75.93	201.0
Oil prices	63.65	32.31	17.70	110.0
Observations	14958	14958	14958	14958

Notes: All GDP and GVA figures are in real per capita terms. Population is in thousands of people. Non-oil commodity prices are in US dollars and oil prices are UK Brent in US dollars per barrel.

cluded the NUTS3 regions that do not provide further disaggregation. Otherwise put, these NUTS3 regions simultaneously and on their own also form a NUTS2 region, as for example Madrid which is both a NUTS2 region, with the region identifier ES30, and a NUTS3 region, with the region identifier ES300. Similarly, Vienna constitutes both a NUTS2 region (AT13) and a NUTS3 region (AT130). In the NUTS classification, these are denoted by a zero suffix in their region identifier. Omitting those from our sample, as Cambridge Econometrics do, would result in an incomplete representation of Spain or Austria, leading us to exclude very relevant locations of economic activity in these countries. In light of this, we have introduced those NUTS3 regions to the Cambridge Econometrics database, imputing their values using the figures of their NUTS2 regions. This follows the practice done by Eurostat, where all economic and demographic accounts of NUTS2 regions equal the values of NUTS3 regions when the latter does not provide further regional disaggregation than the former.

Cambridge Econometrics have constructed the European Regional Database based on data from the regional database of Eurostat. Since 2019, the ERD has been made available from the territorial dashboard of the European Commission's Joint Research Centre ISPRA ([European Commission, 2019](#)). The primary advantages of using the Cambridge Econometrics database over downloading the data directly from Eurostat are two. First, the sample period for the Cambridge Econometrics database is 1980-2015 for most countries

and 1990-2015 for the rest, while for Eurostat regional economic accounts start only in 2000, and a consistent back series of data is not available for the full series. Cambridge Econometrics splices the earlier series to match more recent series. Splicing the data involves using an overlapping period between the European System of Regional and National Accounts (ESA) 79, ESA 95, or ESA 2010 series to extend the ESA 2010 series backwards using ESA 95 and ESA 79 growth rates. The second advantage of Cambridge Econometrics is that they deal with missing data. Even for the period after 2000 that is covered by Eurostat, there is a considerable number of missing values. Cambridge Econometrics have filled in those gaps by scaling up data from sub-regions, extrapolation and interpolation. They additionally implement manual fixes, and scale up the data to AMECO totals. As a result, their final product is a consistent and complete dataset that makes use of and extends upon information provided by Eurostat or AMECO.

Figure 15: Coefficient of variation of regional GDP per capita for the full sample period

Notes: The coefficient of variation (CV) is computed as the ratio of the standard deviation to the mean of all NUTS3 regions within each country for the sample period 1999-2015, except for: the bar denoted EA, which refers to the CV over all NUTS3 regions in the sample; and the bar denoted Mean, which refers to the unweighted average of the eleven within-country coefficients of variation displayed in the graph.

A.2 Additional charts for Section 4.2

Impact of monetary policy on regional output in the tails of the distribution

Figure 16: Quantile regressions

Figure 17: Sub-sample analysis

Notes: Vertical axis refers to impact of 100 basis point rate hike on regional GDP (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands. LHS-panel shows the 5th percentile (in blue) and the 95th percentile (in grey) of the baseline model estimated via quantile regressions; RHS-panel shows the corresponding average estimates from the sub-sample analysis for the bottom five percent (in blue) and top five percent (in grey) of the per-capita GDP distribution.

A.3 Construction of forward-looking variables

As part of the robustness checks in Section 5, we augment the baseline specification with forward-looking variables for GDP and HICP. We construct these variables from the quarterly ECB staff macroeconomic projections. The ECB staff projections are the natural choice among a variety of alternative sources that provide publicly available macroeconomic projections for the euro area (including for instance other policy institutions or private forecasters). While these projections are a staff-level exercise that does not necessarily have to fully match the Governing Council's assessment of the economic outlook, they are an integral element of the Governing Council's information set. This prominent role becomes visible for instance from the fact that the staff projections are reported in the Introductory Statements following the Governing Council's monetary policy meetings, where they provide the motivation and context for the decisions taken at these meetings. We choose the last projection year as the reference point because we consider it as the most

suitable approximation of the medium-term orientation: while the ECB has clarified that the definition of the “medium-term” may vary depending on the state of the economy and has refrained from defining it in terms of a specific time horizon, existing communication indicates that the outlook for the final projection year carries particular weight in its policy deliberations.

As a rule, we use the final year of the respective projection horizon as the relevant reference point and compute a weighted average across the four projection vintages of each year. For most of the sample period, the horizon of the March, June, and September projection vintages of year t stretched until year $t + 1$ and that of the December vintage until year $t + 2$. The ECB changed this convention in 2014, when it extended the projection horizon by another year. For the sake of consistency, and given this change took place only in the penultimate year of our sample, we abstract from it in constructing the forward-looking activity and price level variables. So, in summary, we proceed as follows. We first calculate a forward-looking inflation variable as:

$$\pi_{t+1|t}^e = \frac{1}{4}(\pi_{t+1|DEC,t-1}^e + \pi_{t+1|MAR,t}^e + \pi_{t+1|JUN,t}^e + \pi_{t+1|SEP,t}^e)$$

where $\pi_{t+1|t}^e$ is the expected inflation rate for year $t + 1$ entering the central bank information set in year t , $\pi_{t+1|DEC,t-1}^e$ is the expected annual inflation rate in year $t + 1$ according to the December ECB staff macroeconomic projections of year $t - 1$, $\pi_{t+1|MAR,t}^e$ is the expected annual inflation rate in year $t + 1$ according to the March ECB staff macroeconomic projections of year t , and $\pi_{t+1|JUN,t}^e$ ($\pi_{t+1|SEP,t}^e$) is the corresponding value according to the June (September) projections of year t . The forward-looking variable for the rate of real economic growth is calculated in analogous fashion. Finally, we apply the respective inflation (GDP growth) rate to the level of HICP (GDP) in year t to back out the expected levels for these variables. Consistent with the baseline specification, expected HICP and GDP are expressed in 100 times their log-levels. In terms of aggregation, we construct the EA figures as the GDP-weighted average of the country-level forecasts for the eleven Euro area countries we include in our sample. For 1999 and 2000, we do not include Greece in the weighted average given it had not adopted the euro at this stage. The weights for each country are based on nominal GDP from the AMECO database.

A.4 Results from robustness checks in Section 5

Figure 18: Inclusion of expected HICP and GDP

Notes: Vertical axis refers to impact of 100 basis point rate hike on regional GDP (in %). Horizontal axis refers to horizon of IRF (in years). Solid lines denote point estimates and shaded areas denote 90% confidence bands. Top panel shows mean estimates for the baseline (in grey) and the modified model (in blue). Bottom LHS-panel shows the 5th percentile (in blue) and the 95th percentile (in grey) of the modified model estimated via quantile regressions; RHS-panel shows the corresponding average estimates from the sub-sample analysis for the bottom five percent (in blue) and top five percent (in grey) of the per-capita GDP distribution.

Figure 19: Use of Lemke/Vladu shadow rate

See notes to Figure 18.

Figure 20: Use of Krippner shadow rate

See notes to Figure 18.

Figure 21: Additional covariates

See notes to Figure 18.

Acknowledgements

The analysis presented in this paper has benefited from discussions with Carlo Altavilla, Fabio Canova, Nicola Fuchs-Schündeln, Christophe Kamps, Maarten Dossche, Giorgio Primiceri, Morten O. Ravn, Hélène Rey, Joao M.C. Santos Silva, Frank Smets, Thomas Vlassopoulos and seminar participants at the ECB. The views expressed here are those of the authors and do not necessarily reflect those of the European Central Bank.

Sebastian Hauptmeier

European Central Bank, Frankfurt am Main, Germany; email: sebastian.hauptmeier@ecb.europa.eu

Fédéric Holm-Hadulla

European Central Bank, Frankfurt am Main, Germany; email: federic.holm-hadulla@ecb.europa.eu

Katerina Nikalexi

European Central Bank, Frankfurt am Main, Germany; email: katerina.nikalexi@ecb.europa.eu

© European Central Bank, 2020

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-4028-3

ISSN 1725-2806

doi:10.2866/98413

QB-AR-20-037-EN-N