Monetary Policy and the Firm: Some Empirical Evidence

Saleem Bahaj Angus Foulis Gabor Pinter*

September 25, 2017

Preliminary – Please do not circulate

Abstract

We use a detailed micro dataset to study the monetary transmission at the firm-level. We identify monthly monetary policy shocks using high frequency surprises around policy announcements. Notably, we exploit variation of firms filing their yearly financial accounts in different months of the year, which allows us to estimate the causal effect of monthly monetary policy shocks on firm-level outcomes. We find that contractionary monetary policy shocks have a persistent effect on the individual firm with the effect on employment peaking as far as 5-7 years after the shock hits. This is in contrast with the more transitory effect of monetary policy shocks on the aggregate economy, suggesting that incumbent firms are affected more by monetary policy shocks and that new entrants promote the recovery. We show that with a delay, rates of entry and size at entry increase following a contractionary shock. Incumbents who entered recently are most sensitive to the shock suggesting monetary policy redistributes activity from the current generation of entrants to the next generation.

Keywords: Firm Dynamics, Monetary Policy, Persistence, Entry-Exit

JEL Classification: D22, E22, E32, E52, L11, M13

*We are grateful for helpful comments to Andy Blake, Giancarlo Corsetti, Fred Malherbe, Silvia Miranda-Agripinno, Michael Mcmahon (discussant), Ricardo Reis, Adam Szeidl and Jasmine Xiao. We also thank discussants and seminar participants at the 2017 RES conference (Bristol), the CCBS Chief Economists' Workshop, and the CCBS-MacCalm Macro-finance Workshop at the Bank of England. Bahaj: saleem.bahaj@bankofengland.co.uk; Foulis: angus.foulis@bankofengland.co.uk; Pinter: gabor.pinter@bankofengland.co.uk, Bank of England and Centre for Macroeconomics.

1 Introduction

What is the effect of monetary policy on the individual firm? Is the effect on the average firm different from that on the aggregate economy? How do firm characteristics such as size and age govern the firm's response? These are fundamental questions in economics, yet we have surprisingly little empirical evidence to answer these questions convincingly.

The empirical literature has faced two main challenges. The first is related to identification: exogenous movements in the policy instrument need to be disentangled from the endogenous reactions of monetary policy to the business cycle. Without careful identification of monetary policy *shocks*, no obtained correlation between the policy instrument and firm-level outcomes could be interpreted as the causal effect of monetary policy on the firm. The second challenge is technical in nature: there is a lack of firm-level microdata which (i) would include detailed accounting information on firms covering the entirety of the size distribution, and (ii) would match the high business cycle frequency of most available monetary policy shock measures. Because of the lack of comprehensive firm-level data, even the most recent empirical literature on monetary policy transmission is restricted to study the effects on the aggregate economy (Ramey, 2016), at industry-level (Ozdagli and Weber, 2017) or on a subset of larger firms (Gorodnichenko and Weber, 2016). We still lack a good understanding of how the average firm responds to monetary policy shocks.

Our main contribution to the literature is to address both challenges in a coherent empirical framework. In particular, we draw on the recent advances of the macroeconometric literature on the identification of monetary policy surprises and obtain monthly series of monetary policy shocks. Moreover, we build on our previous work (Bahaj, Foulis, and Pinter, 2017) and use an unprecedentedly detailed firm-level dataset, which includes more than 20,000,000 relevant firm-account observations for the UK covering the period 1990-2015. The sample is dominated by small and medium sized firms. While the dataset is constructed from financial accounts that are filed at yearly frequency, a key aspect of our dataset is that firms file their accounts in different *months* of the year. Hence, a firm that files its yearly account in January for example will have experienced a different sequence of monetary policy shocks during its accounting window compared to a firm that files in July. The sizable variance of firms' filing dates across the year provides us with a crucial source of variation for identifying the causal effect of monetary policy on the firm.

Our study yields the following robust result: monetary policy shocks have a very persistent effect on the individual firm with the effect on employment peaking as far as 5-7 years after the shock hits. This is in contrast with the more transitory effect of monetary policy on the whole economy with the aggregate effect peaking around 3 years after the shock hits. While the estimated persistence of the aggregate employment response is in line with previous findings

(Christiano, Eichenbaum, and Evans, 1999), our result regarding the persistence of the firm-level response is, to the best of our knowledge, novel in the literature. This finding is robust to various perturbations, in particular to (i) using alternative estimators, (ii) including a rich set of control variables, (iii) expanding the sample size and looking at the effect on total assets, and (iv) using alternative measures of monetary policy shocks.

We provide two pieces of suggestive evidence that can reconcile the firm-level and aggregate responses. First, we show that, while a contractionary policy shock reduces the number of new entrants temporarily, the effect subsequently turns positive thereby contributing to the recovery of aggregate employment. At the same time, the negative response of incumbent firms continues to persist. Second, we find that average birth at entry increases in response to a contractionary policy shock. These two forces may explain the difference between the persistence of the average firm response and the persistence of the aggregate response.

Finally, we explore the role of firm characteristics such as age and size in determining the firm's reaction to monetary policy shocks. We find strong evidence that large firms respond more to a policy shock than small firms. We also find evidence that younger firms respond more than older firms. If we sort firms across both these dimensions it is large recent entrants that are most sensitive to monetary policy shocks.

Overall, our results suggest that monetary policy redistributes economic activity from incumbents, particularly young incumbents, to new entrants. Whether this distributional shifts between firms has important aggregate implications still needs to be assessed. In line with the theories of the scarring and cleansing effects of recessions (Caballero and Hammour (1994); Ouyang (2009)), monetary policy may serve to clear the ground for new entrants by weeding out weaker incumbents or it may damage young incumbents who then fail to realise their potential.

Related Literature Our paper builds on the vast literature on the identification of monetary policy shocks (Sims, 1980; Christiano, Eichenbaum, and Evans, 1999; Ramey, 2016). In particular, we obtain monetary policy shocks using high frequency identification (Gurkaynak, Sack, and Swanson, 2005; Nakamura and Steinsson, 2013; Gertler and Karadi, 2015; Cesa-Bianchi, Thwaites, and Vicondoa, 2016; Miranda-Agrippino, 2016; Mueller, Tahbaz-Salehi, and Vedolin, 2017) as well as using the narrative approach (Romer and Romer, 2004; Cloyne and Hurtgen, 2016).

While the monetary policy literature has traditionally focused on the aggregate effects of policy shocks, a number of recent papers studied how micro-level heterogeneity and distributional effects change the monetary policy transmission (Auclert 2015; Cloyne, Ferreira, and Surico 2016; Kaplan, Moll, and Violante 2016; Sterk and Tenreyro 2016). However, these papers have focused primarily on the household side, and the empirical literature on the firms side remains rather scant. Most of the studies estimating the effect of monetary policy on firms, however, focused on more aggregated levels (Gertler and Gilchrist, 1994; Dedola and Lippi, 2005; Ozdagli and Weber, 2017) or studied a subset of larger firms (Gorodnichenko and Weber, 2016; Ippolito, Ozdagli, and Perez-Orive, 2017), while the effect on the average firm remains largely unknown. As mentioned above, this is mainly due to the lack of comprehensive firm-level micro dataset available to empirical researchers, which is an important gap in the monetary policy literature that our paper and the associated comprehensive data collection work are able to fill in.[1] This is the main technical contribution of our paper to the literature.[2]

Our main conceptual contribution to the literature is to show how much more persistent firm-level responses to monetary policy shocks can be compared to the corresponding aggregate responses. This increased persistence is closely related to the recent work of Moreira (2017), who shows that firms born in recession have permenantly worse outcomes than those born in expansion. A parrallel literature is apparent here at the household side also, the permenant negative consquence of cohorts who enter employment at the time of a recession are documented in Baker, Gibbs, and Holmstrom (1994); Kahn (2010); Oreopoulos, von Wachter, and Heisz (2012). Our work differs from these cohort based studies, first because we can assess how *existing* incumbents, who may have entered sometime ago respond to the shock and, second, our panel setup means that we can see how the same individual frm responds to shocks through time.

The discrepancy between the aggregate response and the individual firm response speaks to the literature on firm entry and exit (Hopenhayn 1992; Samaniego 2008; Bilbiie, Ghironi, and Melitz 2012; Siemer 2014; Clementi and Palazzo 2016). A subset of this literature looked at the relationship between monetary policy and entry-exit dynamics (Bergin and Corsetti 2008; Kobayashi 2011; Bilbiie, Fujiwara, and Ghironi 2014) and found that monetary policy shocks have a procyclical effect on firm entry. Our results corroborate this evidence and extends it by showing (i) that firm entry seems to recover over time and turns positive, and (ii) average size at firm entry seems to increase after a contractionary shock - suggestive of a cleansing effect of monetary policy (Caballero and Hammour 1994).

Our results are also connected with the recent literature on firm dynamics that emphasises the importance of start-ups and young firms for aggregate job creation and the sensitivity of young firms to business cycle fluctuations (Haltiwanger, Jarmin, and Miranda 2013; Ouimet and Zarutskie 2014; Decker, Haltiwanger, Jarmin, and Miranda 2014; Sterk and Sedlacek 2017; Moreira 2017; Schmalz, Sraer, and Thesmar 2017). Our findings are complementary in that we

[1] As will be described in our data section, our comprehensive data collection process is similar to Kalemli-Ozcan, Sorensen, Villegas-Sanchez, Volosovych, and Yesiltas (2015); Gopinath, Kalemli-Ozcan, Karabarbounis, and Villegas-Sanchez (2015). While they study resource misallocation, productivity and capital inflows, our focus is on firms' reactions to monetary policy shocks.

[2] Note that our paper focuses on non-financial firms, and thereby complements the growing empirical literature on the effect of monetary policy on financial intermediaries and credit supply (Kashyap and Stein (2000); Ashcraft and Campello (2007); Jimenez, Ongena, Peydro, and Saurina (2012, 2014)).

also find important role for age in determining monetary policy transmission at the firm-level. However, we find that the relationship between firm responsiveness to aggregate shocks and age is more nuanced than the previous literature might suggest.

Our results relate to the long-standing debate on the role of firm size in affecting the monetary transmission. While the seminal paper of Gertler and Gilchrist (1994) found that small firms respond more to monetary policy shocks, subsequent empirical results found that large firms are indeed more sensitive to business cycles (Chari, Christiano, and Kehoe, 2013; Kudlyak and Sanchez, 2017). While these studies use aggregated data, our paper adds firm-level evidence to this debate.

The remainder of the paper is structured as follows. Section 2 presents our empirical aggregate model to estimate monetary policy shocks; Section 3 provides details on our firm-level data including summary statistics; 4 presents our microeconometric model and our approach to clustering; Section 5 presents our baseline results and robustness checks; Section 6 presents additional results on the role of heterogeneity; Section 7 concludes.

2 Identifying Monetary Policy Shocks

2.1 The External Instrument For Monetary Policy

Our baseline strategy for measuring exogenous fluctuations in UK monetary policy uses the series of Miranda-Agrippino (2016). This series, making use of a high frequency identification strategy, essentially serves as an instrument for monetary policy in our empirical analysis that follows. It is constructed by measuring the reaction in the sterling rate futures market[3] in a 15 minute window on either side of the release of the following UK monetary policy announcements: (i) the outcome of the Monetary Policy Committee meeting; (ii) the publication of the minutes of the meeting; (iii) the publication of the Bank of England's Inflation Report.

The innovation in Miranda-Agrippino (2016), compared to the high frequency identification strategy employed for the US by Gertler and Karadi (2015), is to recognise that these raw market reactions to monetary policy announcements can be predicted both by the central bank's own forecasts and by publicly available information that predates the announcement.[4] The explanation posited for this result is that the market reaction is not a pure monetary shock but may reflect (i) private information available to the central bank regarding either its reaction function or the future path of the economy; or (ii) time varying risk premia that respond to the

[3]The futures contract used embed policy expectations for the next three months. See appendix B of Miranda-Agrippino (2016) for details of the exact contracts and for information about the construction of the series in general.

[4]Miranda-Agrippino (2016) shows that this finding is also true in the US using the Gertler and Karadi (2015) shock series.

Figure 1: The Miranda-Agrippino (2016) instrument for UK monetary policy shocks

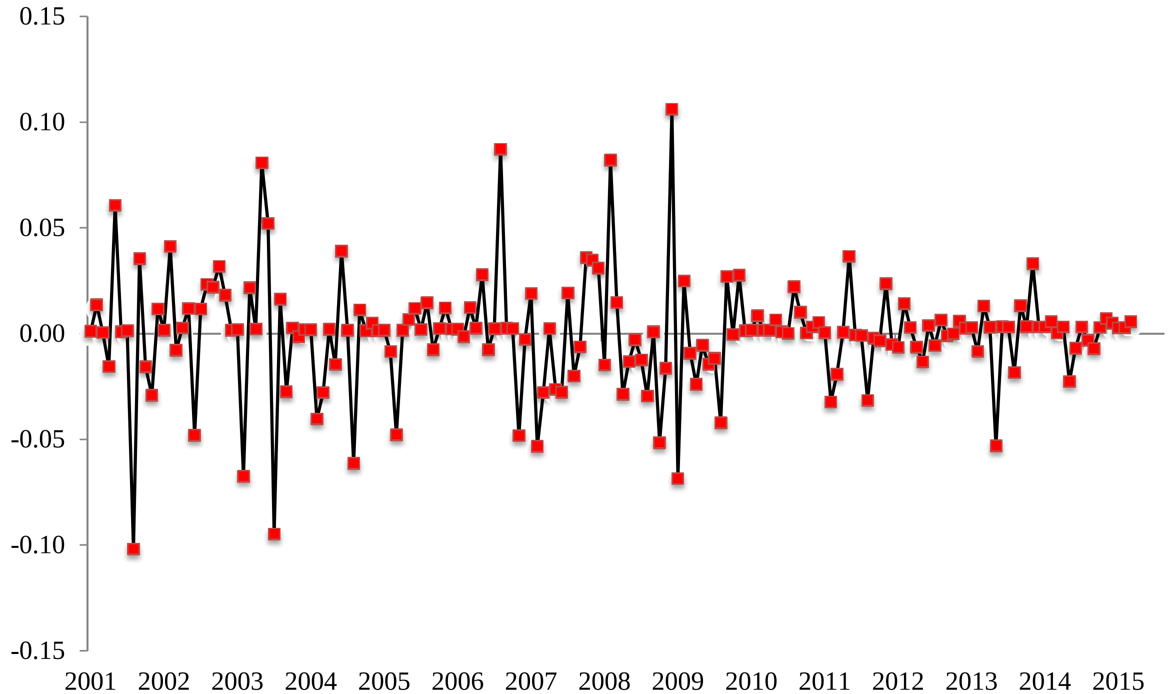

Notes: Instrument for monetary policy shocks from Miranda-Agrippino (2016) (January 2001 to March 2015). The time series is the raw market surprises to monetary policy announcements projected upon (1) CB forecasts and forecast revisions (2) past Bank Rate decisions (3) Libor-OIS spread. The y-axis can be interpreted as changes in an interest rate.

announcement (particularly if some uncertainty is resolved). Failing to control for these factors may cloud the identification of monetary shocks. For example, if the central bank raises rates because it has accurate private information that the economy will expand in future, using the raw market surprise to the monetary policy announcement could attenuate any contractionary impact of policy. To address this, in the spirit of Romer and Romer (2004), Miranda-Agrippino (2016) regresses the raw market reactions on the central bank's own forecasts and on forecast revisions in the lead up to the announcement alongside a set of, predetermined, financial variables designed to capture investor risk appetite. The residuals from these regressions are accumulated across the month and these are what serve as our instrument for monetary policy shocks. The monthly series is plotted in Figure 1 and covers the period January 2001 to March 2015.

2.2 Proxy SVAR

Having obtained a a source of exogenous variation in monetary policy, we use the obtained series as an external instrument in a structural vector autoregression (proxy SVAR) model covering UK aggregate data, and extract the identified monetary policy shock series for use in our firm level regressions. The methodology for proxy SVARs is now relatively standard and we refer readers to Stock and Watson (2012) and Stock and Watson (2012); Mertens and Ravn (2013) for further information about implementation. The methodology only identifies shocks up to a sign and scaling assumption. Hence we scale the identified policy shock by ensuring that a unit change in the shock series brings about a unit increase in the short term interest rate on impact. The F-statistic from the regressions of the residuals on the proxy is 12.2. .

We follow Miranda-Agrippino (2016) by specifying a monthly VAR(12) covering the period January 1981 to March 2015.[5] We include the following time series in the VAR: the UK index of manufacturing production (in logs), one year gilt yields (in percentage points), the UK retail price index (in logs) and aggregate employment (in logs). Note that, in contrast to Miranda-Agrippino (2016), we use employment rather than the unemployment rate as an aggregate labour market variable in order to be consistent with our firm level regressions. We use the industry level break down of UK employment to generate an aggregate series to match our firm level data below. The exact source and construction of all four endogenous VAR series is discussed in Appendix A.

Figure 2 presents the impulse response functions to a contractionary one standard deviation monetary policy shock that emerge from the VAR. The pattern of responses is in line with the monetary policy literature. There is a permanent decrease in prices and a contractionary effect on employment and output. However, at this stage, it is important to note two points: (i) the effect of the shock on aggregate employment is temporary with the effect peaking after 35 months and returning to zero after 60 months. (ii) a one standard deviation monetary policy shock amounts to a 30bp increase to the 1 year yield on impact. Both these findings are almost perfectly in with the Christiano, Eichenbaum, and Evans, 1999 handbook chapter on the effects of monetary policy.

Extracting the identified monetary policy shock series from a VAR as opposed to using the Miranda-Agrippino (2016) series directly as an instrument has three main advantages. First, it allows us to scale the shock series such that we can interpret it in terms of a monthly movement in the interest rate. Second, high frequency market moves are imperfect, noisy measures of the true monetary shock. The narrow window could mean that the measured reaction over or under shoots the true reaction.[6] By projecting the high frequency series on the VAR residuals,

[5] The use of a VAR sample starting 15 years prior to our firm level sample may need explanation. This sample length (i) is comparable to other recent studies, (ii) leads to more precise estimates of aggregate dynamics and, (iii) allows us to use the Cloyne and Hurtgen (2016) series as an alternative instrument for policy shocks as a robustness check. We found that this series requires variation from the 1980s to identify the systematic component of monetary policy hence we use an extended sample in order to focus on one reduced form VAR model. If we restrict our baseline specification to a VAR estimated with data only after the UK exited the ERM in September 1992, the monthly correlation between the identified shocks between the short and long sample is 0.82. The impulse responses are similar, also, albeit with wider error bands.

[6] For instance, this can happen if it takes time for the market participants to fully digest the implications of the announcement. Alternatively, important monetary news could be released outside the three windows considered.

Figure 2: Aggregate Impulse Responses to a Monetary Policy Shock

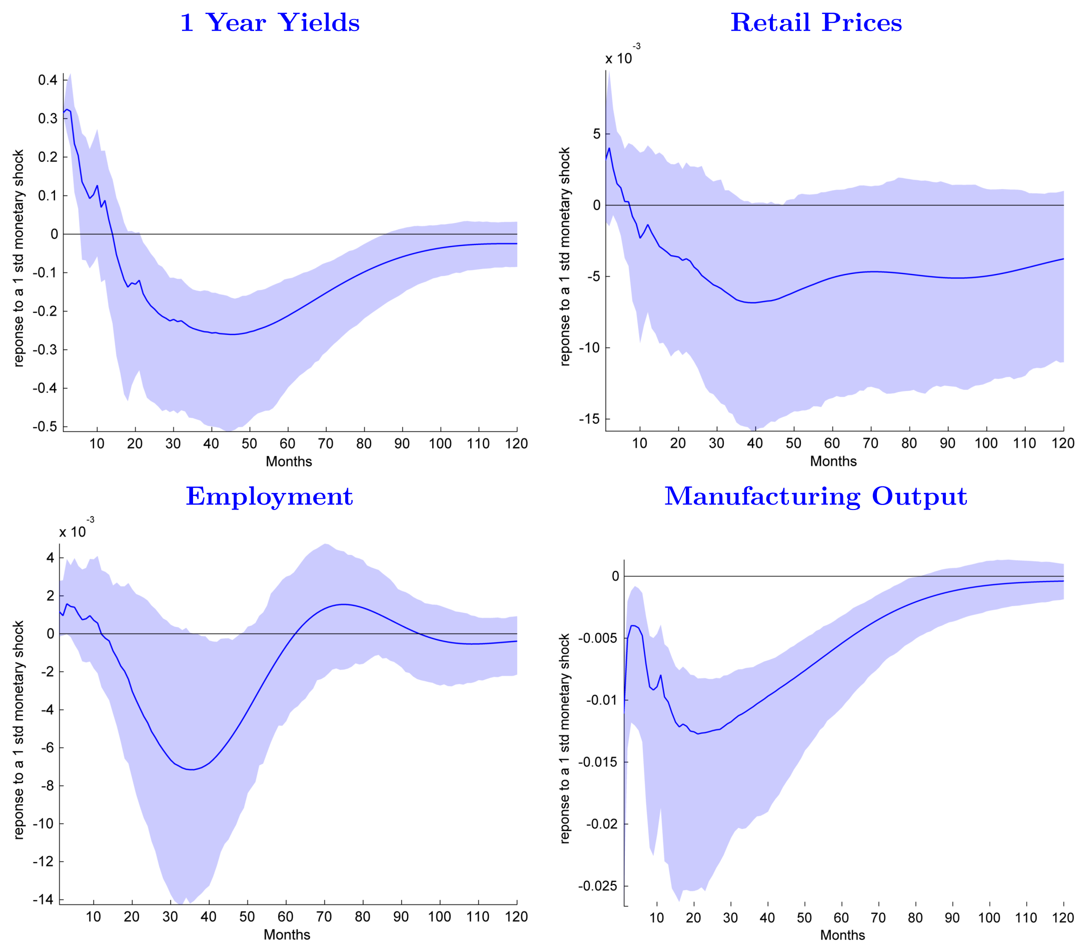

Notes: Estimates are from a proxy SVAR estimated on UK monthly data over the period 1981-2015. Monetary policy shocks are identified using the Miranda-Agrippino (2016) series. The blue solid lines are the point estimates, and the shaded areas are the 90% confidence intervals constructed from a wild recursive bootstrap (Goncalves and Kilian 2004).

we are able to isolate the portion in the forecast error for the interest rate that is caused by a monetary surprise, thereby cleaning out some of the noise in the proxy. Third, we can use the patterns of correlation between the reduced form residuals and the instrument to extend the identified monetary policy shock series back to periods before the instrument was available.[7]

[7]This is a common approach in the proxy SVAR literature (Gertler and Karadi 2015). Specifically, this methodology identifies the contemporaneous coefficients on the reduced form residuals that can be combined to produce the identified shock. Since our reduced form specification extends back beyond 2001, we can use the estimated residuals for the pre-2001 sample, along with the identified coefficients, to extend our shock series prior to 2001. Note that estimating the reduced form portion of the proxy SVAR over a sample that covers a longer period than the external instrument is available.

3 Firm Level Data

Our firm level data for the UK is sourced from a large micro dataset of firms' financial accounts provided by Bureau van Dijk (BVD). This is a commercial dataset based on company filings at Companies House, which is a UK government agency acting as the registrar of companies in accordance with the Companies Acts of 1985 and 2006. The dataset contains information on approximately 4.8 million private and public companies, covering much of the corporate universe of the UK; unincorporated sole traders being a notable missing sector. Companies House filings include the firm's annual accounts, which contain most of the data that is of most interest to us, as well as several other pieces of information about the firm including the firm's location and industry, the individuals responsible for running it, any creditors who have a secured claim on the firms assets etc.

BVD is a live database. This leads to several limitations. First, the company ownership structure is only accurate at the time of access and not for historical observations. Second, companies that die exit the database after four years. Third, the historical information based on past filed accounts has significantly more missing data than the most recent filings. To circumvent these issues, we use archived data sampled at a six monthly frequency to capture information when it was first published. This substantially improves data coverage, allows us to observe the birth and death of companies, provides accurate information on the ownership structure of companies at the time the accounts were filed. As discussed in Kalemli-Ozcan, Sorensen, Villegas-Sanchez, Volosovych, and Yesiltas (2015), the use of archival information and careful treatment of the data is crucial construct an accurate firm-level panel using a data provided by BVD. We discuss our procedure in great detail and the corresponding advantages it brings in terms of coverage in Appendix B.[8]

In effect, our dataset is annual due to the frequency at which firms file their accounts. However, a key feature of the dataset is that firms file their accounts do so at different times during the year. Hence, a firm that files in January for example will have experienced a different sequence of shocks in their accounting window to one that files in July. As shown by Figure 3, filing months are not evenly distributed throughout the year, and firms are more likely to file at the end of the financial or accounting year or at the end of a quarter. Still, the sample size is sufficiently large that every month in the year will have many observations.[9]

One concern is that firms may strategically alter their filing dates in response to monetary

[8]Earlier research using BVD to study the monetary policy propagation at the firm level, such as Bougheas, Mizen, and Yalcin (2006), did not use archival information, and also used the interest rate as a measure of monetary policy (without trying to separate endogenous movements in the policy rate from exogenous movements).

[9]Firms can choose to alter their filing date and file outside an annual cycle (this happens for approximately 4% of firm-account observations). All our regressions focus on time windows where the firm is filing regularly; that is to say, when computing firm growth rates, at any horizon, we exclude any observations when there is an irregular filing period.

Figure 3: Distribution of Filing Dates by Month

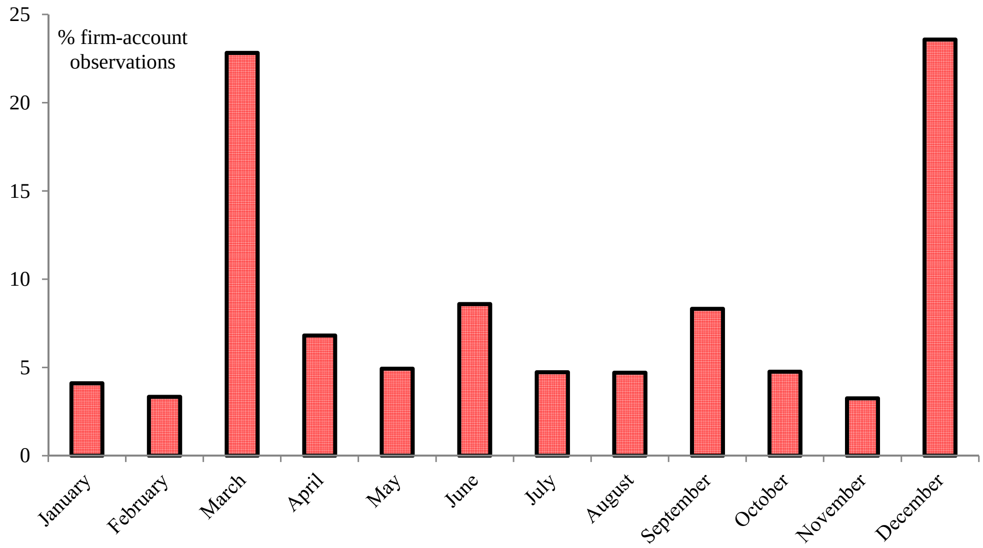

Notes: The Figure presents the distribution across months in which firms file their accounts (all firm account observations 1990-2015).

policy shocks, thereby introducing a selection effect. There are a number of reasons to assume that this sort of behaviour is not distorting our results. First, 91% of firm-account observations are filed in the same month as the first set of the firms' accounts in our sample, and 83% of firms never have an irregular filing window. Second, while firms can always choose to shorten the intervals between account filings, they can only lengthen their accounting period once every 5 years and only to a maximum of 18 months, which limits firms' ability to file strategically. Third, empirically, monetary policy shocks explain a tiny fraction of the time series variation of an individual firm's activity, suggesting the incentive to strategically alter filing dates in response to monetary policy shocks is small.

In our analysis we focus on two measures of firm activity observable in the BVD dataset: *total assets* which captures the size of the firm's balance sheet in nominal terms and *number of employees* which captures the labour force.

3.1 Sample Selection and Summary Statistics

Selection Specific sample selection criteria are laid out in Appendix B.5. In general, we focus on private limited and public quoted firms for whom the UK Companies Act applies. We exclude firms that operate in the financial, public or non-profit sectors.[10] We also exclude

[10]Specifically we exclude companies of the following types: "Economic European Interest Grouping", "Guarantee", "Industrial/Provident", "Limited Liability Partnership", "Not companies Act", "Other", "Royal Charter", "Unlimited", "Public Investment Trust", thereby ensuring that our sample contains only limited liability companies to which the Companies Act applies. In addition, we exclude from the sample firms operating in utilities (2003-UK Standard Industrial Classification [SIC]: 4011-4100); finance and insurance (2003-SIC: 6511-

companies that have a parent with an ownership stake greater than 50%. This is to ensure that the accounts used have the highest degree of consolidation possible, to prevent the double counting of subsidiaries and to ensure that the financial position of the company is correctly accounted for. This leaves us with an initial sample of 20,792,437 firm-account observations covering 3,781,908 unique firms.

Required reporting standards for firm accounts under the Companies Act vary by firm size, and some firms voluntarily provide more information in their statements than is required. This leads to uneven coverage for certain variables, which is an often noted drawback of the BVD database. However, even the smallest firms (defined as micro entities with fewer than 10 employees) must provide a balance sheet to the Registrar. This means that the variable *total assets* is the best reported measure of firm activity, missing for 6.1% of firm-account observations. Hence, we focus on *total assets* as a near representative variable. However, it is worth noting that observations for which *total assets* is missing are still disproportionally located in firms that fail within our sample period.

On the other side of the spectrum, *number of employees* is poorly reported, missing for 96.7% of firm-account observations. However, this series has the advantage that it is directly comparable with the employment series in the VAR. Furthermore, as we will discuss below, firms that do report employment, in total, cover a large share of aggregate employment in the same industries.

Reflecting these differences in reporting, our approach to dealing with missing values is to always use the largest dataset possible for any individual regression model that we estimate, instead of preconditioning on a sample that contains all variables of interest for every model we estimate. This means that when estimating how *total assets* respond to monetary policy shock we are considering a different, broader set of firms than when we estimate how *number of employees* responds.[11] Firms do, however, need to either report the variable of interest for three consecutive years or for two consecutive years twice nonconsecutively to be included in any regression model we estimate.

Our sample period covers firm account filings between 1990 and 2015. However, due to the coverage of our archived data (again, see Appendix B) 96% of firm-account observations are located in the window 1998-2014.

Summary Statistics Table 1 presents summary statistics on the variables of interest for our dataset. We consider these statistics for two subsamples: (I) those regarding the sample of firms able to report sufficient observations on *total assets* (the broader group); and (II) those

6720); real estate (2003-SIC: 7011-7032); public administration (2003-SIC: 7511-7530); education and charity (2003-SIC: 8010-8540); and clubs and organisations (2003-SIC: 9100-9199).

[11] Readers will no doubt be comforted to know that the response of *total assets* for firms who do report *number of employees* is very similar to those firms who do not.

able to report sufficient observations on *number of employees* . These will form the two main subsamples on which we conduct our empirical analysis and by sufficient observations we mean that we have sufficient data on the firm to include it in at least the horizon zero regression as will be described below.

Table 1: Summary Statistics

Variable	**Mean**	**Median**	**25%tile**	**75%tile**	**N**
I: Firms that report *Total Assets* (3,744,718 unique firms)					
Total Assets (£'000s)	2779	55	15	225	12050499
Real Asset Growth (conditional on survival)					
1-year	0.022	0.000	-0.160	0.220	12050499
3-year	0.068	0.031	-0.260	0.430	8072643
5-year	0.160	0.120	-0.310	0.670	4462878
10-year	0.330	0.300	-0.270	0.990	1366788
II: Firms that report *Number of Employees* (105,610 unique firms)					
Total Assets (£'000s)	61718	2326	157	6909	465444
Number of Employees	303	28	4	91	467816
Employment Growth Rates (conditional on survival)					
1-year	0.011	0.000	-.0260	0.065	467816
3-year	0.027	0.000	-0.100	0.190	282028
5-year	0.074	0.013	-0.160	0.340	143259
10-year	0.150	0.130	-0.210	0.560	50047

Note: Firms are counted as reporting total assets/number of employees if they report either for three consecutive years or two consecutive years non-consecutively. Growth rates are calculated for firms who file all accounts in a regular annual pattern (observations for which there is an accounting period that is not annual are excluded). Nominal asset growth is converted into real terms using the UK CPI at the month of filing.

The median firm in sample I has £55,000 in assets with the interquartile range between £15,000 and £75,000. This is in comparison to the UK definition of a "microentity" having assets less than £316,000 and a small company having assets less than £5.1 million.[12] The firms in sample I are therefore capturing the smallest entities (and on a proportionate basis by number of firms). Despite their small size, these firms are not particularly fast growing. The median firm experiences no annual real asset growth and average growth is 2.2% on a year to year basis, roughly in line with real GDP growth. The average growth rate of firms who survive for long time periods is only a little greater: 2.7% on an annualised basis at the 10 year horizon. This evidence is consistent with the findings in Haltiwanger, Jarmin, and Miranda (2013) that firm growth is not necessarily concentrated in small firms but instead it is young firms that

[12] These figures come from technical thresholds used to define company size for the purpose of company reporting rules in the UK. Small companies and microentities have lower reporting burdens than larger companies (e.g. they only need to submit a balance sheet rather then a complete set of accounts to Companies House). A company is defined as "small" ("micro") if it meets two of the following criteria: sales of £10.2 million (£632,000) or less, total assets of £5.1 million (£316,000) or less, 50 (10) employees or less.

grow quickly.

Sample II is biased towards larger firms as they are more likely to report employment. Median total assets are around 40 times at £2.3million. Nonetheless, the median *number of employees* is 22 which is less than the UK small firm threshold of fewer than 50 employees. Employment growth for these firms is more sluggish than real asset growth for the wider sample but is still of a similar order of magnitude and is not dissimilar to average annual UK aggregate employment growth (in the the same industries) of 0.8% between 1996 and 2014. In both samples, the fact that the mean is substantially greater than the median is indicative of the right skew that is typical in firm-level datasets.

Comparisons to the Aggregate Since one of the goals of this paper is to compare aggregate and firm-level behaviour it is useful to ask how the aggregation of our firm level sample compares to the aggregate time series from the UK statistical authority.

The top left panel in Figure 4 compares the sum of employment for all firms that have filed accounts in the prior twelve months with the aggregate employment series for the same industries included in our firm level dataset. The right panel shows the share of aggregate employment covered by our firm level dataset. Prior to 2009, the aggregation of our firm level data tracks aggregate employment quite closely. The correlation between the year over year growth in the aggregate data and the firm level data for the period 2000-2009 is 0.27 allowing for a 6 month phase shift to account for the firm-level series being a rolling 12 month sum. However, once the real implications of the financial crisis manifest in 2009 the relationship breaks down.

An explanation for this is that firm entry and exit became more important drivers of aggregate employment dynamics, relative to the behaviour of incumbent firms. The number of deaths among the firms in our sample increases by 40% over the course of 2008 and 2009 (bottom right panel figure 4).[13] And the recovery from 2012 in employment was accompanied by a surge in entry with 20% more firms incorporated in 2013 as compared to the 2007 pre-crisis peak (bottom left panel figure 4). The aggregation of our firm level data would struggle to detect this, as new entrants are less likely to report employment and firms that exit may cease reporting employment long before they exit or, indeed, be more likely to never report it in the first place. Nonetheless, the share of the aggregate labour market covered by firms in our dataset remains stable between 40-45% of total employment for much of the sample period.

The UK national balance sheet dataset does not have the sectoral detail for us to construct a comparable time series for the *total assets* from aggregate data. However, since *total assets* are near universally reported and it is not unreasonable to think that aggregating over our firms is close to the truth.

[13] Death is defined as the firm's *company status* becoming "dissolved".

Figure 4: Aggregate Behaviour of Firms

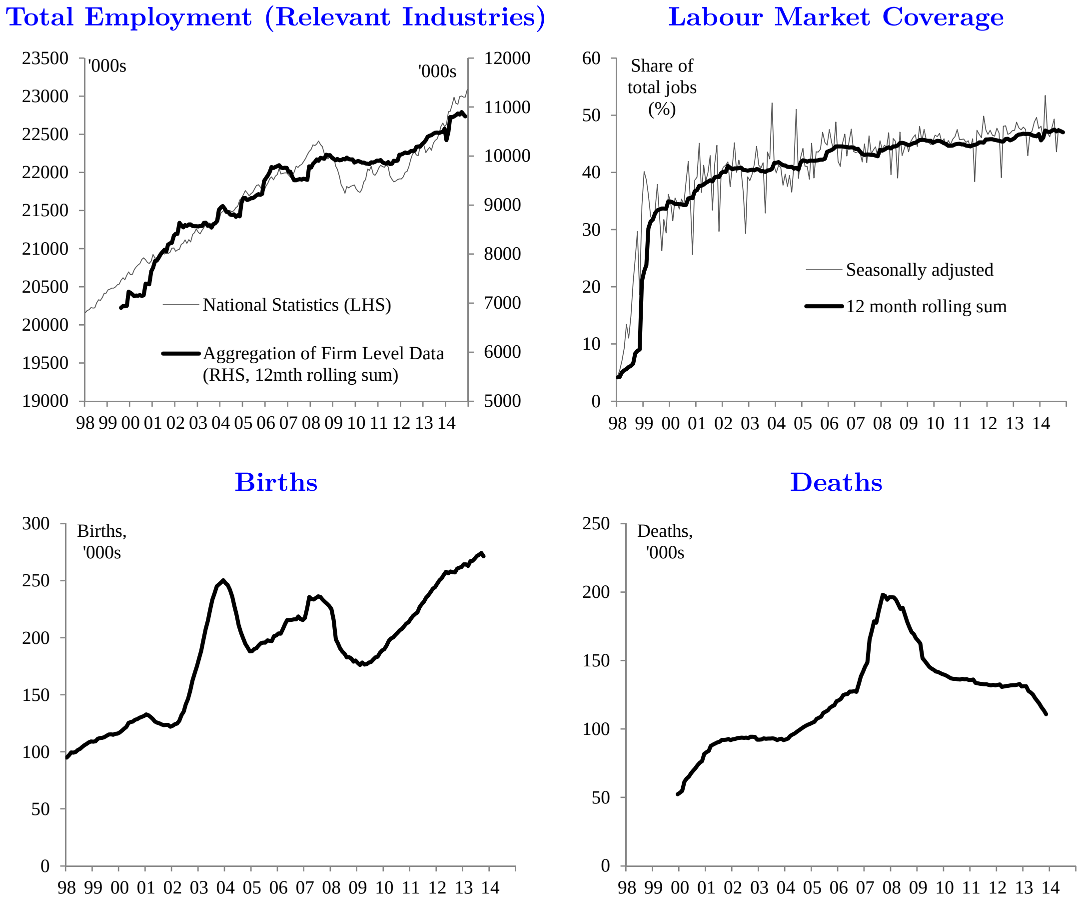

Notes: (i) *top left panel:* Thick black line (aggregation of firm level data) is the sum of the employment of all companies that file in particular month expressed as a 12 month moving sum. Thin black line (national statistics) is employment in the relevant industries as sourced from the UK ONS (see Appendix A). (ii) *top right panel:* thick black line (12 month rolling sum) is the ratio between the two lines in the top left panel. Thin black line (seasonally adjusted) is the constructed by taking the sum of all firms that file in a particular month, seasonally adjusting the time series and multiplying it by 12 dividing by the thin black line in the top left panel. (iii) *bottom left panel:* number of firms with *incorporation date* in a rolling 12 month window. (iv) *bottom right panel:* number of firms a *statement date* where the *company status* was first listed as dissolved in a rolling 12 month window.

4 Empirical Model

4.1 Specification

Let $ACT_{i,t}$ be a measure of firm activity for firm i for accounting period t. Here it is necessary to introduce a brief remark on notation: Our firm level data set is effectively annual so t refers to the year and we use the index $m \in \{1, \dots, 12\}$ to denote months within that accounting year. To ensure no ambiguity, $m = 12$ is the month in which the firm files its accounts within that year, not December.

Our baseline linear specification is specified as a local projection (Jorda 2005) and is a modified version, into a panel setting, of the model discussed in Ramey (2016):

$$log(ACT_{i,t+h}) - log(ACT_{i,t-1}) = \alpha_i^h + \beta^h \sum_{m=1}^{12} w_m e_{m,t} + \gamma^h \times controls_{i,t} + \sum_{j=1}^{4} \phi^h \sum_{m=1}^{12} \tilde{u}_{m,j,t} + \varepsilon_{i,t}^h, \tag{4.1}$$

where $h \in \{0, \ldots, 10\}$ indexes a set of regressions at different growth rate horizons.

The term $e_{m,t}$ denotes the monetary shock for month m of accounting year t as extracted from the VAR described in Section 2. The parameters w_m denotes weights assigned to shocks that occur in different months. If a firm is unable to react immediately to a monetary policy shock one may imagine that shocks that occur earlier in the accounting year have a bigger impact on the firm's behaviour. Empirically, when we estimate the weights, we do not find this and we cannot reject the null that the $w_m = 1$ for all m. For this reason and for the sake of a parsimonious specification with a single coefficient of interest, we impose in our baseline that all the weights are equal.[14] Hence, the sum of the identified monetary policy shocks over the accounting year enters the model as a single regressor so that we can focus on the single parameter β^h and conduct inference over it. As is standard for a local projection, the vector $\beta = \{\beta^0, \ldots, \beta^{10}\}$ is a directly estimated cumulative impulse response function and the standard errors of β^h are used to construct the relevant confidence intervals.

In our baseline specification, the firm level control set ($controls_{i,t}$) is empty as most of the firm level cross sectional heterogeneity is soaked up by the fixed effect α_i^h. We add a rich set of controls as a robustness exercise. The term $\sum_{j=1}^{3} \sum_{m=1}^{12} \tilde{u}_{m,j,t}$ captures other aggregate shocks that occur over the accounting window, and we discuss the construction of these terms and their implications for inference further below.

When we consider how firm characteristics affect the response to monetary policy shocks, we estimate the following non-linear specification:

$$log(ACT_{i,t+h}) - log(ACT_{i,t-1}) = \alpha_i^h + \delta_{m,t}^h + \theta^h X_{i,t} \times \sum_{m=1}^{12} w_m e_{m,t} + \gamma^h \times controls_{i,t} + \varepsilon_{i,t}^h, \tag{4.2}$$

where $\delta_{m,t}$ is a time fixed effect for all firms that file in the same month, which removes the linear effects of all aggregate shocks. The term $X_{i,t}$ is a firm specific interaction variable (or potentially a vector of variables) which can be either continuous or a dummy. When $X_{i,t}$ is a continuous variable we rescale the series by its cross sectional standard deviation.

When constructing cumulative growth rates, $log(ACT_{i,t+h}) - log(ACT_{i,t-1})$, we (i) omit observations in the 99th and 1st percentiles of observations in order to prevent outliers distorting the results, (ii) recast all nominal variables in real terms by dividing through by the seasonally

[14] We present the results with estimated weights in the robustness section.

adjusted UK consumer price index for the month when the account was filed, and (iii) omit observations where any account period in the window between $t-1$ and $t+h$ is irregular.

4.2 Inference

We cluster our error bands at the 4-digit industry level[15] to allow for arbitrary correlation in errors across of firms within the same industry (466 unique clusters). This accounts for all omitted heterogeneity at the firm or sectoral level that may generate correlation in the error terms both through time and in the cross-section.

However, in the linear specification outlined in equation 4.1 there is no time fixed effect; hence, there is the potential for omitted aggregate shocks (excepting monetary policy) to induce cross correlation in the error terms between firms that file in similar time periods. This is of particular concern considering that a critical source of variation that we are exploiting is the different filing dates of firms: failing to account for the fact that firms that filed accounts in February would have received a similar pattern of aggregate shocks to those that file in January may bias down our standard errors. Clustering by firm account filing month would omit potential patterns of correlation that may be present for firms who file their accounts in neighbouring months. Clustering by filing year would also be insufficient as it would omit the correlation between firms that file annual accounts in December of the previous year versus those that file in January of the current year.

Instead, we conduct inference using a multiway clustering approach (Cameron, Gelbach, and Miller 2011) and construct overlapping clusters based around when firms file their accounts. In the ideal case, one would want to allow the error terms of the firms that file in January of year 2 to be correlated with the error terms of all the firms that filed in the past year (i.e. to the previous February of year 1). Similarly, one would want to allow all the firms that file in February in year 1 to be correlated with the firms that filed up to March of year 0 etc. One can implement this by constructing 12 overlapping groups of clusters covering each month of the year and the 11 preceding months. Implementing this ideal case involves clustering across 13 different dimensions (including the industry cluster) which is both very computationally intensive and potentially counter productive given that the number of clusters within each dimension is relatively small. Instead, we limit ourselves to four quarterly clusters and allow firms that filed in Q1 to be correlated with firms that filed in Q4, Q3, Q2 of the previous year. Firms that file in Q2 are correlated with firms that file in Q1 and firms that file in Q3 and Q4 of the previous year; and so on for the other two quarters.[16]

At $h = 0$, this strategy implies that the smallest number of clusters we have is 23. This is

[15] Specifically, we use 2003 UK SIC codes.

[16] With sufficient computational time, estimating our baseline model while clustering in 13 dimensions is possible but it leads to near identical standard errors to the quarterly approximation we are using.

below the common consensus rule of thumb of 50 clusters (for US state-year level data), implying that tests on the significance of our coefficient estimates may be vulnerable to over-rejection. However, the simulations in Bertrand, Duflo, and Mullainathan (2004) and Cameron, Gelbach, and Miller (2008) both suggest that the Wald statistic is only increased by a few percentage points when the cluster is size 20. This is unlikely to be enough to influence the statistical significance of our regression results. Furthermore, we have many thousands of observations within each cluster which helps mitigate the problem induced by the small numbers of clusters. On the other hand, Cameron, Gelbach, and Miller (2008) show that the problem of too few clusters is more likely to emerge if the clusters are unbalanced (as is our situation). Note that multiway clustering can also exacerbate the small cluster problem (Cameron and Miller 2015).

We are not aware of a straightforward correction for multiway clustered errors when the number of clusters are small. Instead, we follow the suggestion of Cameron and Miller (2015). and attempt to directly control for the source of the within cluster correlation. In our case, these are other aggregate shocks that occurred during the firm's accounting period. The term $\sum_{j=1}^{4} \phi^h \sum_{m=1}^{12} \tilde{u}_{m,j,t}$ is included in order to achieve this. The term $\tilde{u}_{m,j,t}$ is the reduced form residual of equation j from our proxy SVAR for month m of the firm's accounting year after having been projected on the identified monetary policy shocks. Hence, including $\tilde{u}_{m,j,t}$ for $j = 1, \ldots, 4$ (alongside $e_{m,t}$) will span all the structural shocks that emerge from our VAR for month m. As also suggested, we present confidence intervals using t_{G-1} where G is the minimum number of clusters. Note also that equation 4.2 does include a time fixed effect so we cluster at the industry level.

5 Firm Level Responses

5.1 Baseline Results

Figure 5 shows baseline responses to a contractionary monetary shock at the firm level for our two dependent variables (and corresponding subsamples) of interest. In terms of scaling, we consider a one standard fluctuation in the accumulated annual shock series $\sum_{m=1}^{12} e_{m,t}$. A one standard deviation positive shock in the monthly shock series generates a 30 basis point increase the interest rate on impact in the month. Ignoring any further dynamics, this equates to an equivalent of a $\sqrt{12} \times 30bp = 104bp$ increase over the course of the year if the accumulated annual shock is one standard deviation.[17]

The left panel in Figure 5 suggests that on impact such a shock brings about a 0.5% decline in firm level employment. The fall in employment continues, however, with the response reaching a trough of -5% after 6 years before a recovery to -3% after 10 years. This implies that 10 years

[17]Simply, $stdev(\sum_{m=1}^{12} e_{m,t}) = \sqrt{12} stdev(e_{m,t})$

after the shock hit, employment in the average firm (incumbent when the shock occurred) is still 3% below the counterfactual level where no shock happened.

Figure 5: Baseline Responses to a Monetary Policy Shock at the Firm Level

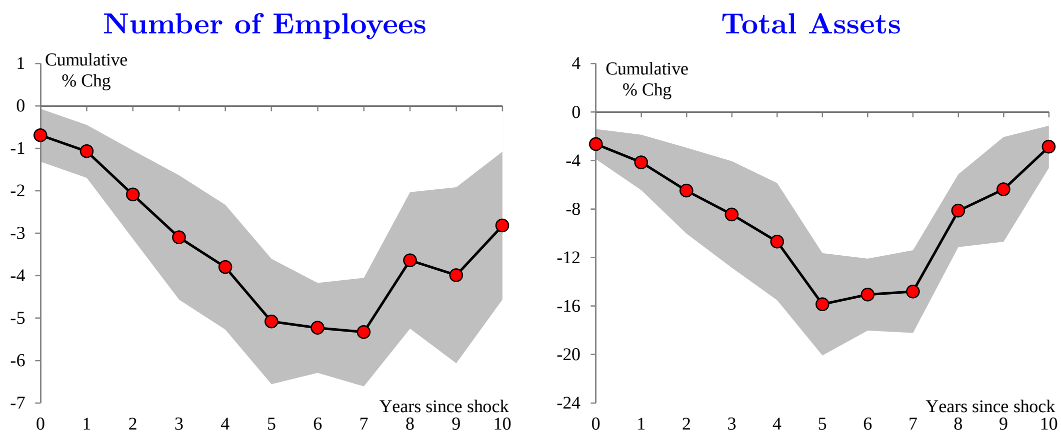

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. Black dotted lines are point estimates. Grey shaded areas are 90% confidence intervals constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. (i) *left panel:* The dependent variable is the cumulative growth rate in log points of employment from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis. Sample is 105,610 unique UK firms over the period 1990-2015 (at $h=0$ zero this corresponds to 467,816 firm-year observations). (i) *right panel:* The dependent variable is the cumulative growth rate in log points of in real total assets from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis Sample is 3,744,718 unique UK firms over the period 1990-2015 (at $h=0$ zero this corresponds to 12,050,499 firm-year observations).

At this point, it is instructive to compare these firm-level employment responses to the aggregate responses in Figure 2. Bear in mind that the aggregate impulse responses are for a one standard deviation shock at a monthly frequency. If we rescale the estimates by $\sqrt{12}$ using the logic above, then the aggregate employment response troughs at -2.5% after 2-3 years.[18] After three years, the average firm (incumbent at the time of the shock) has cut employment by 3% and we cannot reject that the two responses align. Beyond that point the responses diverge. The aggregate returns to zero after 5 years (consistent with Christiano, Eichenbaum, and Evans (1999)), while the firm-level response continues to persist with a recovery that only begins after seven years and is incomplete after ten. Incumbent firms and the aggregate economy have a different dynamic reaction to a monetary policy shock. Given our sample, the responses after 10 years will be driven by observations on firms that file their accounts in the late 1990s and early 2000s. This is a short window that was followed by major financial crisis. The results after 10 years should be interpreted with a degree of caution as a result. However, much of the discrepancy between the aggregate and firm level response emerges after five years and at that horizon we have a large number of observations covering more time periods with different aggregate conditions.

[18]Note that the impulse response functions in Figure 2 come from a VAR. If we recompute the impulses using an equivalent local projection which is in keeping with our firm level estimates, the impulse responses are near identical.

As discussed, we have no direct analogue for *total assets* in aggregate. However, at the firm level the response of *total assets* has similar dynamics to employment in terms, if a little less persistent in the sense that except that assets have largely recovered after 10 years. The major difference, however, is the scale of the response is much larger at -16%. This seems an overly large response to a monetary policy shock. Part of the explanation for this may lie in the fact that the sample is concentrated in very small companies that are more volatile (compare the interquartile range of growth in Table 1 for sample I and sample II). As we will see below, while the dynamic path of *total assets* is robust, the scale of the response shrinks as we perturb the baseline specification. Therefore, what we take from the *total asset* response is that, using an indicator of activity that is near universally reported across firms, we see that monetary policy has a very persistent effect at the firm level.

5.2 Robustness

Given the discrepancy before the aggregate and firm level responses it is first useful to ask how robust the firm level responses are. Figure 6 point estimates for alternative specifications overlayed on the confidence interval for our baseline specification.

Our baseline specification allows firms to exit the sample as they cease reporting. This means that the sample of firms at short horizons is greater for those at longer horizons, potentially distorting the impulse response. However, if we recompute our baseline specification restricting our attention to firms who report 10 years worth of data (Upper panel of 6 – solid line with white square markers) we see a near identical response.

What matters for the aggregate economy is the weighted average growth rate among firms, not the simple average. Our baseline specification is unweighted. If we weight the regressions by initial firm size[19] (Upper panel of 6 – dashed line with black diamond markers) we get similar firm level employment responses. However, the scale the *total asset* response is reduced somewhat, with a peak response at -12%.

In the baseline specification we use a fixed effect estimator. This could be problematic in the sense that estimating fixed effects requires using forward looking information about the firm taken from after the shock hits – this can be a particular problem for firms that we only observe for a short time period. However, our results are robust to using a random effects estimator instead (Upper panel of 6 – solid line with black triangular markers).

[19]Either by the initial *number of employees* or *total assets* depending on the dependent variable.

Figure 6: Firm Level Response to a Monetary Policy Shock: Robustness

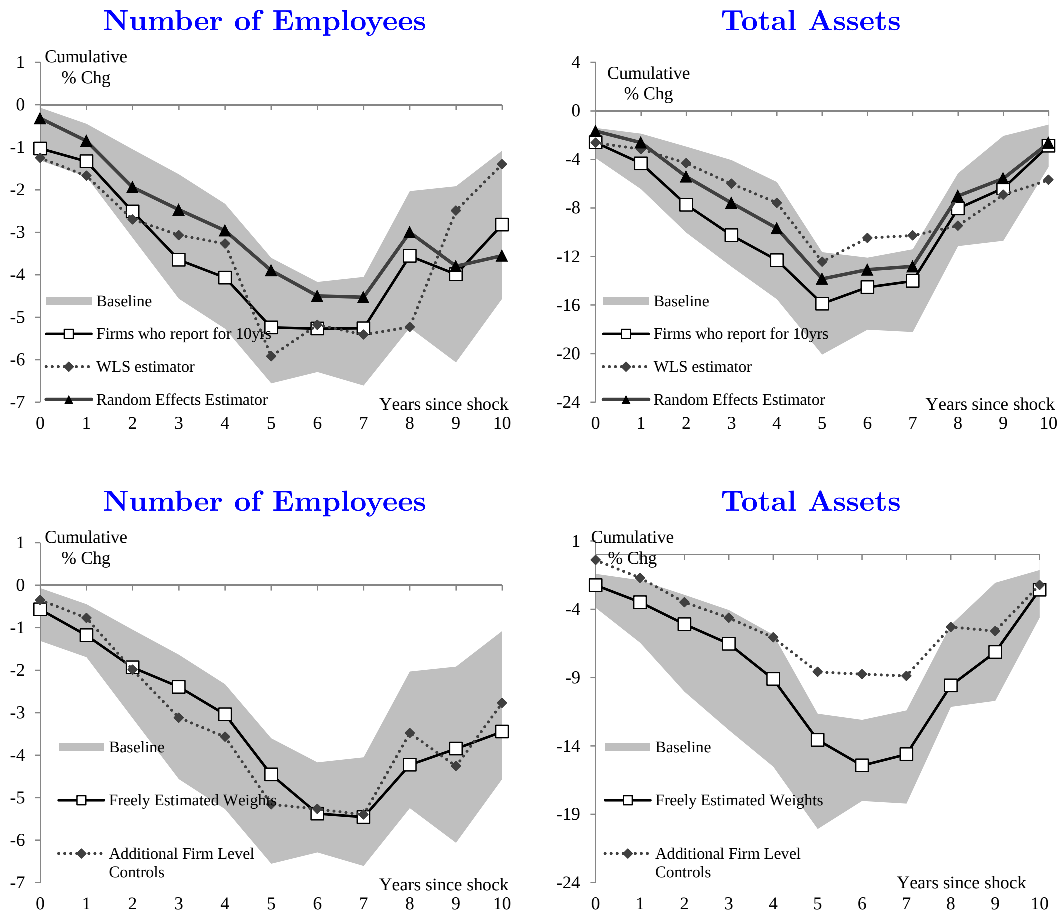

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. Grey shaded areas are 90% confidence intervals from our baseline specification constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. *Upper panels*: (i) black solid lines with square markers are point estimates from a sample where we only include firms who report for 10 years, (ii) black dashed lines with diamond markers are point estimates when we weight our baseline regression model by initial *number of employees* left panel or initial *total assets* right panel, and (iii) black solid lines with triangular markers are point estimates from where we estimate our baseline regressions using a random effects estimator. *Lower panels*: (i) black solid lines with square markers are point estimates from a specification where we allow w_m to be estimated, (ii) black dashed lines with diamond markers are point estimates where we include additional controls in our baseline specification (see main text). *Left panels:* The dependent variable is the cumulative growth rate in log points of employment from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis. *Right panels:* The dependent variable is the cumulative growth rate in log points of in real total assets from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis.

Our baseline specification also included no firm level controls. In Figure (Lower panel of 6 – dashed line with black diamond markers) we show the results when we include a rich control set including the firm's debt to asset ratio, cash to asset ratio, profit to sales ratio, a dummy for whether the firm paid dividend in the accounting year and a dummy for whether the firm increased its fixed assets (invested) in the accounting year (t). These variables enter both linearly and as quadratic terms, and we interact the control set with dummy variables denoting which of the quintiles of the size (measured by *total assets*) and age distribution the firm is located in, at time t. Adding these controls has minimal effects on the results for *number of employees.* The *total asset* response is roughly halved compared to the baseline however. The likely explanation for this is that many of the smallest firms do not report information on some of the control variables used and thus the sample shifts towards larger less volatile firms. Last, we show the results when the weight parameters (w_m) are estimated rather than restricted to be equal (Lower panel of 6 – solid line with white square markers) .

5.3 Alternative Identification Strategies

These findings are, of course, all contingent on the identifications strategy used. There are alternative instruments available for monetary policy for the UK beyond the Miranda-Agrippino (2016) instrument used in our baseline. Cloyne and Hurtgen (2016) provide a narrative based measure of monetary policy shocks in the spirit of Romer and Romer (2004). Cesa-Bianchi, Thwaites, and Vicondoa (2016) provide an alternative high frequency based series which is similar to the "raw" series analysed by Miranda-Agrippino (2016) whereby the market reaction is not projected on other observables.

Figure 7: Firm Level Response to a Monetary Policy Shock: Different Monetary Policy Identification Strategies

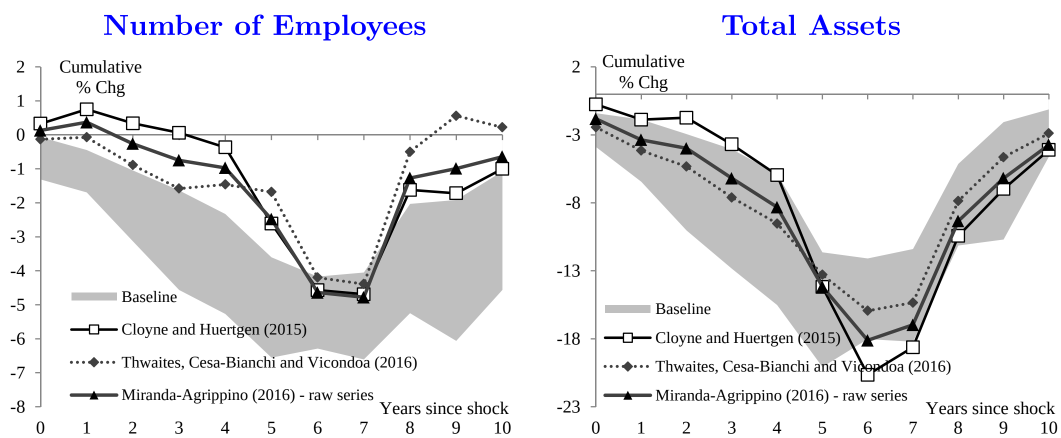

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. Grey shaded areas are 90% confidence intervals from our baseline specification constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. Black solid lines with square markers are point estimates from a model where we identify monetary policy using the Cloyne and Hurtgen (2016) monetary policy series. Black dashed lines with diamond markers monetary policy identified using Cesa-Bianchi, Thwaites, and Vicondoa (2016). Black dashed lines with diamond markers monetary policy identified using Cesa-Bianchi, Thwaites, and Vicondoa (2016). Black solid lines with triangular markers are point estimates from where we estimate our baseline regressions using a random effects estimator. (i) *left panel:* The dependent variable is the cumulative growth rate in log points of employment from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis. (i) *right panel:* The dependent variable is the cumulative growth rate in log points of in real total assets from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis.

Figure 7 presents the results when we identify monetary policy shocks using these three alternative instruments. Specifically, we re-estimate the VAR described in section 2.2 and use each individual series to extract three different monetary policy shock series that form alternative $e_{m,t}$ to be plugged into equation 4.1. All three generate similar responses. The *total assets* responses are consistent across these different series and in line with the baseline. The response of *number of employees* is more muted at short horizons but all three alternative series produce a trough response of around -5% after 6 years. The recovery after that point is also more rapid with the estimates returning to zero after 10 years.

5.4 Reconciling Firm-level and Aggregate Responses

The result is that incumbent firms suffer more persistent decline compared to the aggregate, obviously leads to the question how can we reconcile the two? The missing group from the analysis are new entrants. For aggregate employment to recover while incumbent firms are still contracting, new firms need to fill the gap. This could happen in one of two ways, either more firms are born or the firms that are born are larger, either would mean entrants are hiring more workers.

A lack of a reference point, like an accounting year, makes it difficult to study how individual entrants react to monetary policy shocks that occur before they are born. However, we can investigate how monetary policy affects the behaviour of entrants in aggregate. To do this, we calculate two monthly series from our micro dataset: (i) the number of firms incorporated in a particular month, and (ii) the average size of firms incorporated in a particular month (measured by total assets deflated by the UK CPI).[20] We then manually seasonally adjust both series using an X13 filter.

We explore the effect of a monetary policy shock on these two series using the monthly series $e_{m,t}$, computed as described above, in an aggregate local projection (following equation 2.11 in Ramey (2016)). In addition to the shock series, we include 12 lags of the aggregate series from the VAR as controls.[21]

Figure 8 presents the estimated impulse response function to a monthly one standard deviation contractionary monetary policy shock for the number of entrants and the average size at entry.[22] Several points stand out. First, the monetary policy shock causes a decline in entry – at least initially. This is consistent with other previous empirical evidence (Bergin and Corsetti 2008; Kobayashi 2011). However, this effect troughs after around 30 months, with entry falling by 0.2%. Entry recovers from that point and after around 5 years becomes positive (although this estimate is not significant from zero). On the other hand, size at birth does not respond to monetary policy shocks initially, but after around 30 months starts to increase, at the peak response size at entry increases by 40%.

These responses are consistent with the discrepancy between the behaviour of the aggregate

[20] We use *total assets* measured at the first account published after incorporation. *Employment* among newly founded firms is poorly reported.

[21] As an alternative to including the shock series directly in the local projection, Ramey (2016) suggests that the series could instead be used as an instrument for the endogenous variable of interest (i.e. the interest rate). Taking that approach yields similar results.

[22] We truncate the impulse response at 90 months as the estimates become erratic after this horizon.

Figure 8: The Effect of a Contractionary Monetary Policy Shock on Entrants

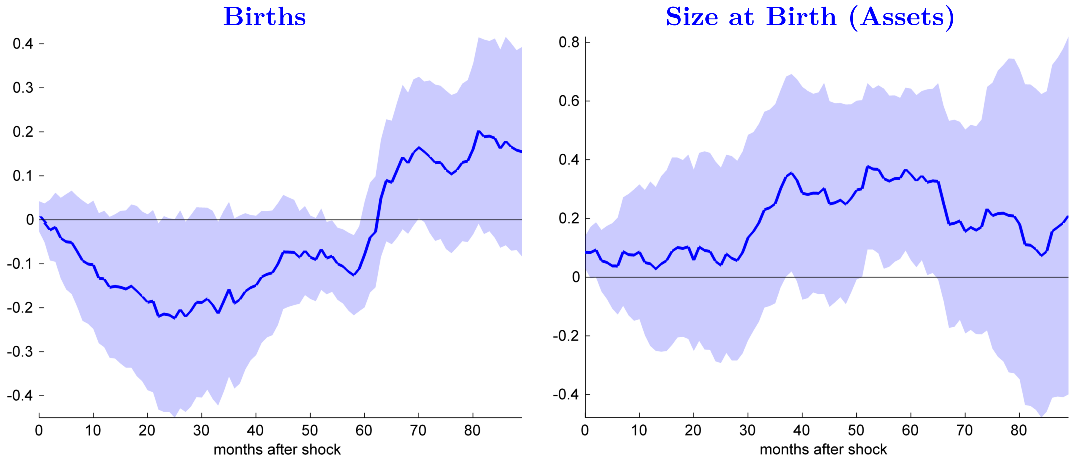

Notes: Aggregate responses to a 1 standard deviation monthly monetary policy shock. Local projections estimated using UK monthly data 1998-2014. Blue solid line is the point estimate, shaded areas are 90% confidence intervals constructed from a HAC adjusted standard errors. *Left panel:* log change in incorporations of new enterprises from $t-1$ to $t+h$ where t is the date of the monetary policy shock. *Right panel:* average birth size measured as real total assets.

economy and incumbent firms. After 30 months when the aggregate response starts the deviate from the response of incumbents is exactly when we see the size of entrants start to increase and rates of entry start to recover. After around 60-70 months (when the deviation between incumbents and the aggregate is the greatest) we see positive rates of entry and the largest firm size at birth.

While encouraging, we view this evidence only as suggestive. Both the numbers of company births and size at birth series are noisy, subject to sampling error and have a short time series (1998-2014) which makes drawing accurate conclusions about responses at long horizons challenging.

What is left is to consider what mechanisms could lead to monetary policy persistently affecting incumbents while encouraging larger new entrants. The literature on the affect of macroeconomic conditions on entrants (e.g. Moreira (2017); Sterk and Sedlacek (2017)) emphasises the importance of demand side frictions in generating persistence. Incumbent firms need to maintain their brand and reputation through continued interactions with their clients. This implies that demand becomes a stock variable (Ravn, Schmitt-Grohe, and Uribe (2006)) and if a contractionary monetary policy shock has a persistent negative effect on the demand for the products of incumbents it could allow new entrants to gain market share increasing their size. Alternatively, financial mechanisms may be a work. If incumbent firms have long term, fixed rate, debt liabilities, a contractionary shock could cause firms to hold back on new investment as the rewards would disproportionately accrue to creditors (Myers (1977)) . New entrants would have the opportunity to choose their liabilities after the shock the hits. There are a number of other mechanisms that could generate persistence at the firm level such as adjustment costs (Caballero (1999)), collateral constraints (Kiyotaki and Moore (1997)) or the real option value of waiting (Dixit (1992)). However, it less clear why these forces would affect incumbents differently from entrants but this is not to say why .

We can exploit cross-sectional heterogeniety in how firms respond to monetary policy to evaluate which of these mechanisms are more relevant and will do so in future work.

6 Heterogeneity

Given that future entrants may behave very differently than incumbents in response to monetary policy shocks, a natural next step is to consider how recently entered incumbents may respond differently compared to firms that have existed for a long time. Figure 9 considers how age affects a firm's response to monetary policy shocks. The upper panel shows how firms in different quintiles of the age distribution respond to the shock.[23] As can be seen, it is typically the oldest firms that respond least to a monetary shock. Once we look beyond the oldest quintile, it is hard to say conclusively that age still matters. Though there is some evidence that firms in the youngest two age quintiles are more responsive than those in third and fourth quintiles. The lower panel of Figure 9 shows the case when we simply include an interaction term (by re-estimating equation 4.2 and setting $X_{i,t} = log(1 + age_{i,t})$) rather than conditioning by quintile. On a linear basis, becoming older seems to attenuate a firm's response to monetary policy shocks. Specifically, the results suggest that for every one standard deviation increase in a firm's age, the given firm would cut employment by 2 percentage point less six years 6 years after the monetary policy shock hit. Overall, our results complement the US evidence on the high procyclicaliy of young firms (Fort, Haltiwanger, Jarmin, and Miranda, 2013; Siemer, 2014), and this finding seems present when conditioning on monetary policy shocks as well.

Of course, young firms are also typically small firms. And, as mentioned, heterogeneity in how small firms and large firms respond to monetary policy has been a focus in the literature. Figure 10 presents equivalent results to Figure 9, for size rather than age. In this case, we construct quintiles of size by conditioning on real total assets at the time the shock hits (and set $X_{i,t} = log(totalassets_{i,t-1})$ in the lower panel). The results here all point to larger firms being *more* responsive to monetary policy (particularly for the largest firms).

[23]Note that we omit estimates of the youngest quintile of firms in the right column, as the these are erratic due to a small sample size problem at long horizons.

Figure 9: Firm Level Response to a Monetary Policy Shock: by Portions of the Age Distribution

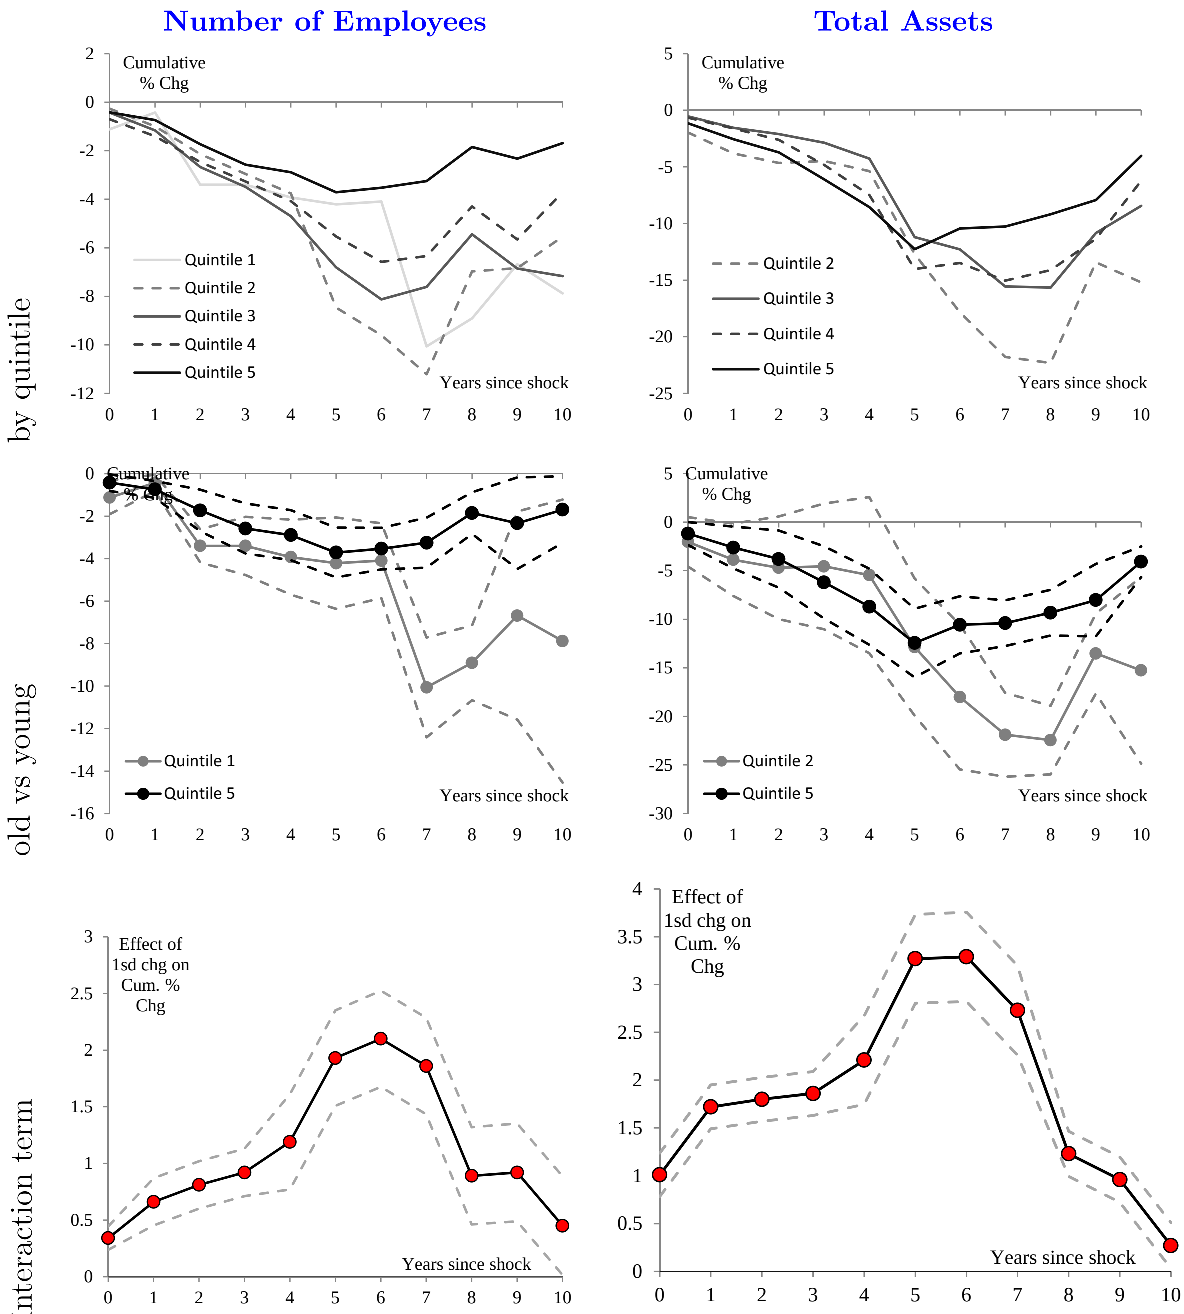

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. (i) *upper left and right panel.* Point estimates for response for firms in the five different portions of the age distribution (when the shock hit). (ii) *middle left and right panel.* Comparing oldest and youngest quintile. Dash lines enclose 90% confidence intervals constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. (iii) *lower left and right panel.* coefficient of interaction term of monetary policy on $log(1 + age_{i,t})$, dash lines enclose 90% confidence intervals.

It is important also to recognise the difference between sample I (firms that report total assets) and sample II (firms that report employment) in this context. The median firm in quintile 5 of the size distribution of sample I has £508,000 in assets that is compared to £2.3m in assets for the median firm in all of sample II (note that £2.3m is roughly the the 98th percentile of *total assets* in sample I). Sample II has a fewer number of much larger firms. Hence, the left column in Figure 10 essentially shows how medium size companies respond less when compared to large companies. The right column in Figure 10 shows that microentities respond less than very small firms. However, there is a non-monotonicity in size at some point. As a whole, firms in sample I are more responsive to monetary policy than those in sample II. If we estimate the equivalent of the right panel of Figure 5, restricting our attention to just firms in sample II, the trough in the response is around -8%.

Given that size and age both seem to matter, how do they interact? Figure 11 presents a double sort where we consider how firms that are both young (in quintiles 1 and 2 of the age distribution) and small (in quintiles 1 and 2 of the size distribution) compare with firms that are old (in quantile 5 of the age distribution) and large (in quantile 5 of the size distribution). We can also look at young, large firms and small, old firms. Note that we just focus on the employment response. The message that emerges from this analysis is that small firms are generally less sensitive to monetary policy shock regardless of their age (although within the group younger small firms are more sensitive). Large firms are sensitive but it is large young firms that respond the most to monetary policy shocks.

There has been debate in the recent literature about the relative cycality of firms by size and age.[24] Overall, our results on size are in contrast with the findings of the seminal paper of (Gertler and Gilchrist, 1994) and more consistent with the subsequent evidence of (Moscarini and Postel-Vinay, 2012; Chari, Christiano, and Kehoe, 2013; Kudlyak and Sanchez, 2017). Fort, Haltiwanger, Jarmin, and Miranda (2013) emphasise the importance of controlling for age as well as size. Our results are in line with theirs in the sense that we find that young firms are more responsive to an aggregate shock. However, they also find that small young firms are more responsive than old, large firms in contrast to our results when we double sort.

Fort, Haltiwanger, Jarmin, and Miranda (2013) do not discuss the large, young firm group in great detail; in our sample this is the group who are most sensitive to monetary policy. This is interesting in the sense that it is the flip side of Figure 8. While contractionary monetary policy may encourage larger entrants, it is also particularly damaging for large incumbent firms who have recently entered. Whether this distributional shift has important consequences for the aggregate economy still needs to be assessed. The importance of new entrants in job creation and reallocation of productive resources is emphasised in Decker, Haltiwanger, Jarmin, and Miranda (2014). But the literature is silent on whether larger entrants are a particularly important part of that process. Conversely, young firms also have a high probably of failure. It may be that contractionary policy weeds out the poor performing younger firms and clears the

[24]With the exception of Gertler and Gilchrist, 1994, this literature is about aggregate conditions in general rather than monetary policy specifically. It could be that there are some discrepancies in the results due to differences in how small versus large or young versus old firms respond to different types of shocks.

ground for a new generation to enter. This would align with the cleansing theory of recessions as in Caballero and Hammour (1994).

Figure 10: Firm Level Response to a 1 Standard Deviation Monetary Policy Shock: by portions of the size distribution

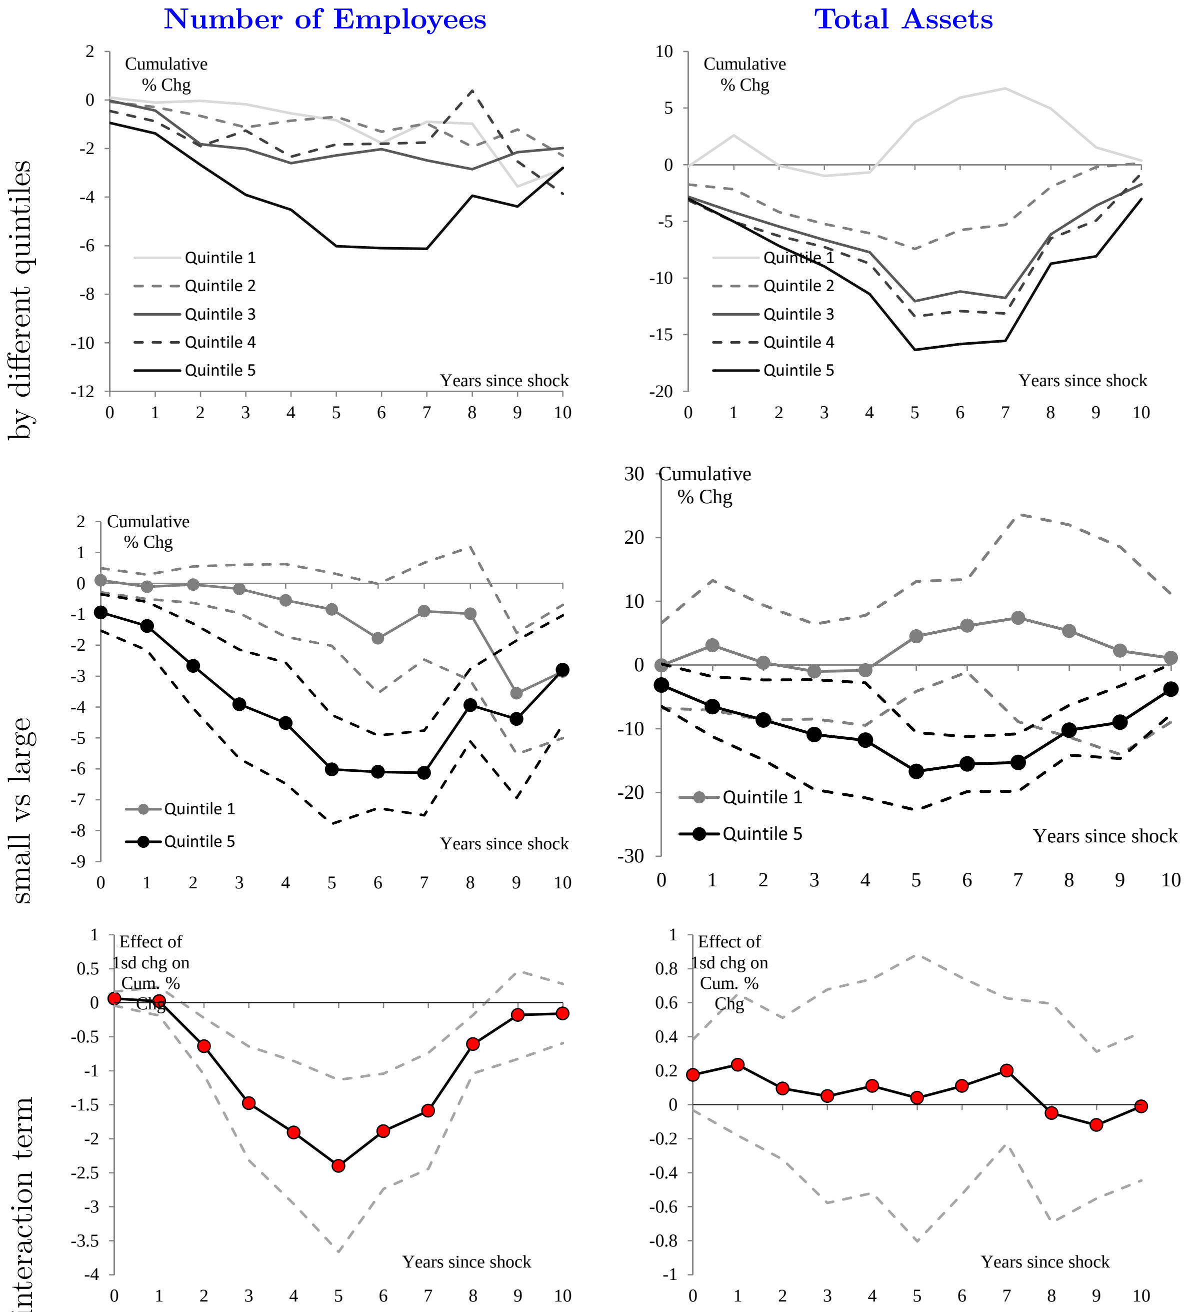

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. (i) *upper left and right panel.* Point estimates for response for firms in the five different portions of the size distribution (when the shock hit). (ii) *middle left and right panel.* Comparing largest and smallest quintile. Dash lines enclose 90% confidence intervals constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. (iii) *lower left and right panel.* coefficient of interaction term of monetary policy on $log(totalassets_{i,t})$, dash lines enclose 90% confidence intervals.

Figure 11: Double Sort: Size and Age

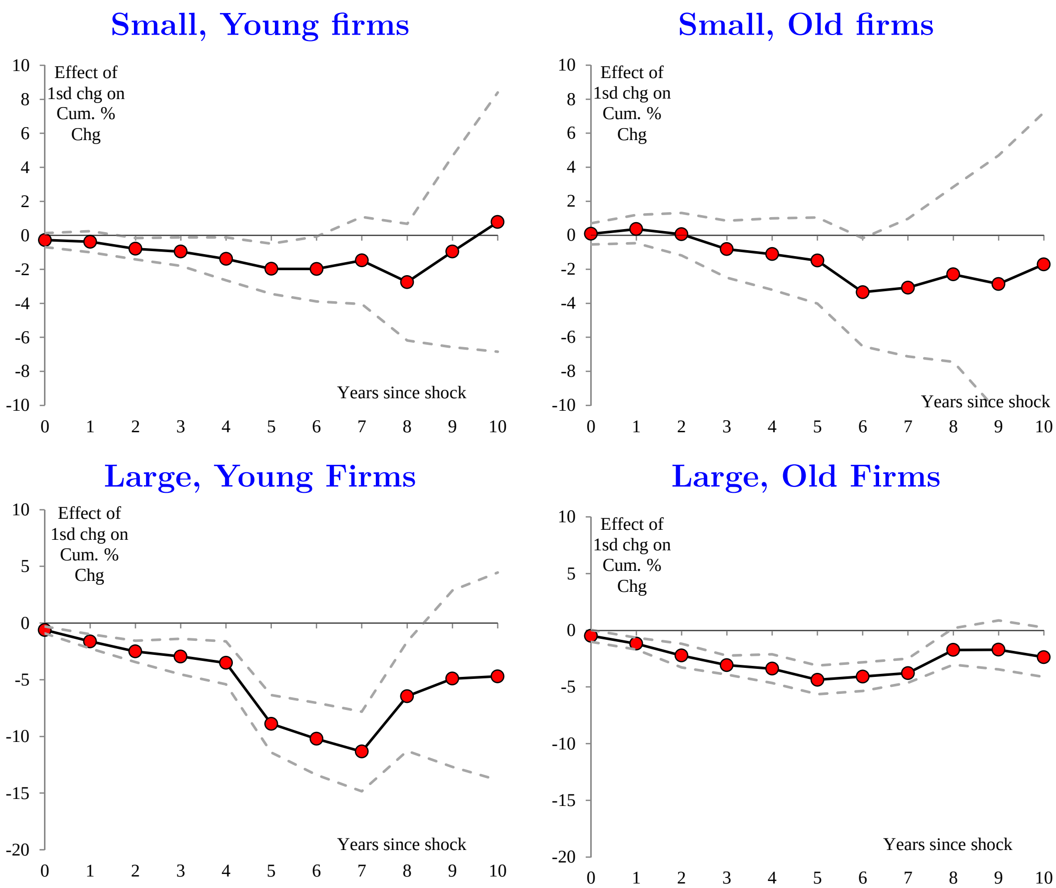

Notes: Firm level responses to a 1 standard deviation contractionary monetary policy shock. The dependent variable is the cumulative growth rate in log points of employment from $t-1$ to $t+h$ where t is the date of the monetary policy shock and h is the x-axis. Black dotted lines are point estimates. Grey dash lines enclose 90% confidence intervals constructed from a t_{G-1} distribution where G is the minimum of clusters in the regression. Small = firm in first two quintiles of the size distribution (measured by assets) when the shock hits. Young = firm in first two quintiles of the age distribution when the shock hits. Large = firm in fifth quintile of the size distribution (measured by assets) when the shock hits. Old = firm in the fifth quintile of the age distribution when the shock hits.

7 Conclusion

The goal of this paper is to consider how incumbent firms react to monetary policy. Our first finding is that monetary policy has a more persistent effect on incumbent firms than on the economy as a whole. The average incumbent firm still has lower employment 10 years after a monetary policy shock; this is long since after the aggregate economy has recovered. Reconciling this with the aggregate is the behaviour on new entrants who, with a delay, enter in larger numbers and average greater size following a contractionary shock. However, among incumbents it is recent entrants, particularly large ones, who are most sensitive to a monetary policy shock. Contractionary policy seems to both shrink recent entrants and make space for new firms to enter the market.

The aggregate consequences of these distributional shifts between firms (or generations of firms) still needs to be assessed. In future work, we plan to explore further how the characteristics of incumbents determine how they respond to monetary policy. Is it the case that the recent entrants that are hit hardest by contractionary policy are productive firms whose prospects are damaged. Or does policy have a cleansing affect and primarily affect less productive recent entrants? The other key unanswered question is what mechanism generates the persistent responses. Digging deeper into differences across industry may allow us to explore the importance of demand side frictions. Our dataset contains comprehensive information on firm balance sheets, so exploring the relevance of financial frictions is also on the agenda.

References

Ashcraft, A. B., and M. Campello (2007): "Firm balance sheets and monetary policy transmission," *Journal of Monetary Economics*, 54(6), 1515–1528.

Auclert, A. (2015): "Monetary Policy and the Redistribution Channel," 2015 Meeting Papers 381, Society for Economic Dynamics.

Bahaj, S., A. Foulis, and G. Pinter (2017): "Home Values and Firm Behaviour," Discussion Papers 1724, Centre for Macroeconomics (CFM).

Baker, G., M. Gibbs, and B. Holmstrom (1994): "The Wage Policy of a Firm," *The Quarterly Journal of Economics*, 109(4), 921–955.

Bergin, P. R., and G. Corsetti (2008): "The extensive margin and monetary policy," *Journal of Monetary Economics*, 55(7), 1222–1237.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 119(1), 249–275.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi (2014): "Optimal monetary policy with endogenous entry and product variety," *Journal of Monetary Economics*, 64(C), 1–20.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz (2012): "Endogenous Entry, Product Variety, and Business Cycles," *Journal of Political Economy*, 120(2), 304–345.

Bougheas, S., P. Mizen, and C. Yalcin (2006): "Access to external finance: Theory and evidence on the impact of monetary policy and firm-specific characteristics," *Journal of Banking and Finance*, 30(1), 199–227.

Caballero, R. J. (1999): "Aggregate investment," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 12, pp. 813–862. Elsevier.

Caballero, R. J., and M. L. Hammour (1994): "The Cleansing Effect of Recessions," *American Economic Review*, 84(5), 1350–1368.

Cameron, A., J. Gelbach, and D. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *The Review of Economics and Statistics*, 90(3), 414–427.

——— (2011): "Robust Inference With Multiway Clustering," *Journal of Business and Economic Statistics*, 29(2), 238–249.

Cameron, A., and D. Miller (2015): "A Practitioners Guide to Cluster-Robust Inference," *Journal of Human Resources*, 50(2), 317–372.

Cesa-Bianchi, A., G. Thwaites, and A. Vicondoa (2016): "Monetary Policy Transmission in an Open Economy: New Data and Evidence from the United Kingdom," Discussion Papers 1612, Centre for Macroeconomics (CFM).

Chari, V. V., L. J. Christiano, and P. J. Kehoe (2013): "The Gertler-Gilchrist Evidence on Small and Large Firm Sales," mimeo, Northwestern University.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 2, pp. 65–148. Elsevier.

Clementi, G. L., and B. Palazzo (2016): "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations," *American Economic Journal: Macroeconomics*, 8(3), 1–41.

Cloyne, J., C. Ferreira, and P. Surico (2016): "Monetary policy when households have debt: new evidence on the transmission mechanism," Bank of England working papers 589, Bank of England.

Cloyne, J., and P. Hurtgen (2016): "The macroeconomic effects of monetary policy: a new measure for the United Kingdom," *American Economic Journal: Macroeconomics*, 8(4), 75–102.

Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014): "The Role of Entrepreneurship in US Job Creation and Economic Dynamism," *Journal of Economic Perspectives*, 28(3), 3–24.

Dedola, L., and F. Lippi (2005): "The monetary transmission mechanism: Evidence from the industries of five OECD countries," *European Economic Review*, 49(6), 1543–1569.

Dixit, A. (1992): "Investment and Hysteresis," *Journal of Economic Perspectives*, 6(1), 107–132.

Fort, T. C., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2013): "How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size," *IMF Economic Review*, 61(3), 520–559.

Gertler, M., and S. Gilchrist (1994): "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *The Quarterly Journal of Economics*, 109(2), 309–40.

Gertler, M., and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76.

Goncalves, S., and L. Kilian (2004): "Bootstrapping autoregressions with conditional heteroskedasticity of unknown form," *Journal of Econometrics*, 123(1), 89–120.

Gopinath, G., S. Kalemli-Ozcan, L. Karabarbounis, and C. Villegas-Sanchez (2015): "Capital Allocation and Productivity in South Europe," Working Paper 21453, National Bureau of Economic Research.

Gorodnichenko, Y., and M. Weber (2016): "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review*, 106(1), 165–199.

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

Haltiwanger, J., R. S. Jarmin, and J. Miranda (2013): "Who Creates Jobs? Small versus Large versus Young," *The Review of Economics and Statistics*, 95(2), 347–361.

Hopenhayn, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–50.

Ippolito, F., A. K. Ozdagli, and A. Perez-Orive (2017): "The Transmission of Monetary Policy through Bank Lending: The Floating Rate Channel," Discussion paper.

Jimenez, G., S. Ongena, J.-L. Peydro, and J. Saurina (2012): "Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications," *American Economic Review*, 102(5), 2301–26.

——— (2014): "Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking?," *Econometrica*, 82(2), 463–505.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

Kahn, L. B. (2010): "The long-term labor market consequences of graduating from college in a bad economy," *Labour Economics*, 17(2), 303–316.

Kalemli-Ozcan, S., B. Sorensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas (2015): "How to Construct Nationally Representative Firm Level data from the ORBIS Global Database," NBER Working Papers 21558, National Bureau of Economic Research, Inc.

Kaplan, G., B. Moll, and G. L. Violante (2016): "Monetary Policy According to HANK," Working Paper 21897, National Bureau of Economic Research.

Kashyap, A. K., and J. C. Stein (2000): "What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?," *American Economic Review*, 90(3), 407–428.

Kiyotaki, N., and J. Moore (1997): "Credit Cycles," *Journal of Political Economy*, 105(2), 211–48.

Kobayashi, T. (2011): "Firm entry, credit availability and monetary policy," *Journal of Economic Dynamics and Control*, 35(8), 1245–1272.

Kudlyak, M., and J. M. Sanchez (2017): "Revisiting the behavior of small and large firms during the 2008 financial crisis," *Journal of Economic Dynamics and Control*, 77, 48 – 69.

Mertens, K., and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103(4), 1212–47.

Miranda-Agrippino, S. (2016): "Unsurprising shocks: information, premia, and the monetary transmission," Bank of England working papers 626, Bank of England.

Mitchell, J., R. J. Smith, M. R. Weale, S. Wright, and E. L. Salazar (2005): "An Indicator of Monthly GDP and an Early Estimate of Quarterly GDP Growth," *The Economic Journal*, 115(501), F108–F129.

Moreira, S. (2017): "Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles," Working Papers 17-29, Center for Economic Studies, U.S. Census Bureau.

Moscarini, G., and F. Postel-Vinay (2012): "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment," *American Economic Review*, 102(6), 2509–2539.

Mueller, P., A. Tahbaz-Salehi, and A. Vedolin (2017): "Exchange Rates and Monetary Policy Uncertainty," *The Journal of Finance*, Forthcoming.

Myers, S. C. (1977): "Determinants of corporate borrowing," *Journal of Financial Economics*, 5(2), 147–175.

Nakamura, E., and J. Steinsson (2013): "High Frequency Identification of Monetary Non-Neutrality: The Information Effect," NBER Working Papers 19260, National Bureau of Economic Research, Inc.

Oreopoulos, P., T. von Wachter, and A. Heisz (2012): "The Short- and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 4(1), 1–29.

Ouimet, P., and R. Zarutskie (2014): "Who works for startups? The relation between firm age, employee age, and growth," *Journal of Financial Economics*, 112(3), 386–407.

Ouyang, M. (2009): "The scarring effect of recessions," *Journal of Monetary Economics*, 56(2), 184–199.

Ozdagli, A., and M. Weber (2017): "Monetary Policy through Production Networks: Evidence from the Stock Market," Working Paper 23424, National Bureau of Economic Research.

Ramey, V. A. (2016): "Macroeconomic Shocks and Their Propagation," vol. 2 of *Handbook of Macroeconomics*, pp. 71 – 162. Elsevier.

Ravn, M., S. Schmitt-Grohe, and M. Uribe (2006): "Deep Habits," *Review of Economic Studies*, 73(1), 195–218.

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

Samaniego, R. M. (2008): "Entry, Exit and Business Cycles in a General Equilibrium Model," *Review of Economic Dynamics*, 11(3), 529–541.

Schmalz, M. C., D. A. Sraer, and D. Thesmar (2017): "Housing Collateral and Entrepreneurship," *The Journal of Finance*, 72(1), 99–132.

Siemer, M. (2014): "Firm Entry and Employment Dynamics in the Great Recession," Finance and Economics Discussion Series 2014-56, Board of Governors of the Federal Reserve System (U.S.).

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

Sterk, V., and P. Sedlacek (2017): "The growth potential of startups over the business cycle," *American Economic Review*, Forthcoming.

Sterk, V., and S. Tenreyro (2016): "The Transmission of Monetary Policy through Redistributions and Durable Purchases," Working Papers 1601, Council on Economic Policies.

Stock, J. H., and M. W. Watson (2012): "Disentangling the Channels of the 2007-2009 Recession," NBER Working Papers 18094, National Bureau of Economic Research, Inc.

Appendix

A VAR Data Sources

Our baseline VAR specification contains four series (all data was accessed either through Bank of England internal databases or downloaded directly from the ONS and is accurate as of 10th May 2017).

Manufacturing output This is the UK index of manufacturing production (ONS code K22A). The raw series is a seasonally adjusted chain value measure and is expressed as an index. Note that the manufacturing index excludes the mining and quarrying, energy supply, and water and waste management industries.

1 year Gilt Yields The average of the yield on 1-year gilts at close for business dates in the month. The raw data provider is the Bank of England.

Employment (workforce jobs in relevant industries) Our measure of aggregate employment comes from the workforce jobs by industry (ONS code JOBS02) database. We take total jobs less the workforce of the following sectors: (i) Electricity, gas, steam & air conditioning supply; (ii) Water supply, sewerage, waste & remediation activities; (iii) Financial & insurance activities; (iv) Real estate activities; (v) Public admin & defense; compulsory social security; (vi) Education; (vii) Human health & social work activities; (viii) Private households. This comes close to matching our selection of firms which excludes the public, non-profit and unincorporated sectors (as well as finance, real estate and utilities etc.).

The jobs series is available quarterly. To convert it into a monthly frequency we use the regression based interpolation technique of Mitchell, Smith, Weale, Wright, and Salazar (2005) using the monthly total employment series (ONS code MGRZ) as an interpoland. We specify the interpolation procedure as an ARX(1,1).

Retail Price Index This is the headline UK retail price index (ONS code CHAW). This series run back to 1987. We extend it back to the start of the sample using the historical RPI series (ONS code CBAB) which runs from 1974-1987 assuming that the two indices grow at the same rate. We then seasonally adjust the spliced series using an X.12 filter.

B Firm Level Data

B.1 Company Reporting Rules in the United Kingdom

The statutory reporting requirements for companies registered in the United Kingdom are mainly governed by the Companies Act 2006 and prior to that the Companies Act 1985. The last provisions of the Company's Act 2006 came into force in 1st October 2009. This means that firms in the United Kingdom operated varying reporting standards during our sample period, the most relevant change in standards for our purposes is the treatment of director's addresses which is discussed in detail below. The Act covers private and public limited companies. Other types of companies, for instance Partnerships or LLPs, are covered by separate legislation but have their own reporting standards and still must file accounts at the registrar. As described below, these are omitted from our analysis to ensure a consistent legal basis for the type of firm under consideration.

Companies House is the Registrar of companies in the United Kingdom. The agency has the responsibility for examining and storing all the statutory information that companies in the United Kingdom are required to supply. Companies House also has the responsibility to make the filed information public; however, there are exceptions to what a company (or individuals that run or exert significant control over a company) have to make publicly available. While Companies House filings often go hand with a companies tax returns (annual accounts can be filed jointly with a tax return), this information held by the Registrar is not directly used for the purposes of calculating corporation tax. Tax returns by companies are dealt with separately by Her Majesty's Revenue and Customs, the United Kingdom Tax Authority.

Reporting Requirements. At the end of the a company's financial year a company must prepare a set of statutory annual accounts that they file with Companies House. These include a version of the firm's balance sheet and profit and loss account. All limited company's are required to report in some way or another to Companies House. However, reporting requirements, particularly over the annual accounts, vary by firm size (see part 1 of the Companies House guide for additional details).[25] Companies House must also be informed of the firms' name, registered office, share capital and charges against the company's assets for the purposes of securing a loan. Companies must also maintain, , a register of directors (including the director's usual residential address). If any of these details change the company must inform Company's House by via an event driven filing.

[25] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/533350/GP2_Life_of_a_company_Pa ver0.1-6.pdf

Time Lags. Companies have 21 months from incorporation to file their first set of accounts with Companies House. Subsequent annual accounts must be filed within 9 months of the company's financial year end for private companies and 6 months of the company's financial year end for public companies. Companies can amend the accounts retrospectively to fix errors and present data revisions. Companies can also amend the end of their accounting year (but not retroactively), which can lead to irregular accounting windows of different lengths than a year. However, companies must file accounts every 18 months. Event driven Companies House filings have shorter time lags. For instance, all appointments, changes to personal details and cessations of a company's directors should be reported to Companies House within 14 days of the changes being made.

B.2 BvD's Collection and Coverage of Firms in the United Kingdom

Companies House is the original source of our data but our direct source is Bureau van Dijk (BvD) who aggregate the data and provides a workable interface to access it. For the United Kingdom and the Republic of Ireland, BvD provides firm-level data through a product called FAME (Financial Analysis Made Easy). This is distinct from the more commonly used Amadeus and Orbis products provided by BvD which cover firms at the European and Global level respectively (although UK firms form a subset in both data sets).

BvD does not source its data from Companies House directly. In between Companies House and BvD is another data provider, Jordans, with whom we have no direct contact. Jordans serve as the direct source for BvD. In the FAME user guide, BvD describes the logistics of the data collection procedure as follows:

Once accounts are filed at Companies House they are processed and checked, put onto microfiche and made available to the public. Companies House aim for a turnaround time of 7-14 days, however this will increase at peak times (October).

Jordans collect data from Companies House daily and transfer it from microfiche to their database with a turnaround time of 3-5 days. This may take longer at peak times of the year (October) and also if figures appear to be incorrect and need to be rechecked with Companies House.

Bureau van Dijk collect data from Jordans on a weekly basis and create the appropriate search indexes to link with the FAME search software. These indexes are then tested and published to the internet server within 48 hours of receiving the data.

In theory, this time frame would imply that most live companies in the Bureau van Dijk database would have their latest accounts filed within the past year (9 months after the firm's financial year plus one-two month's processing time) but lags of two years are not uncommon. Given that lags can occur at four different stages (the filing stage and the three processing stages that follow), we have not been able to determine the root cause of this.

There are four sub-databases in FAME (A,B,C and D) which are ordered by the size of the company as determined by different thresholds in their accounts (e.g. balance sheet size). We have access to and use data from all four databases to achieve the widest possible coverage.

There is conflicting information regarding how long inactive companies remain in the FAME database. When the Bank of England contacted BvD regarding this issue, BvD's claim was that Jordans (their data provider) would only keep inactive companies in the database for five years, so those firms would be lost from the source material. However, BvD would then (on a quarterly cycle) re-upload the missing companies from their own archives ensuring that no data lost from FAME or their other products. However, this claim may not be accurate. From inspecting different vintages of the FAME data set it was that firms did exit the database. For instance, almost 50% of firms in the database in January 2005 were not present 10 years later. Furthermore, some 3 million companies left the database between 2013 and 2014. We discuss the consequences of this in more detail below.

B.3 Treatment of the BvD UK Accounting Data

B.3.1 The Sample of BvD Discs Used

The Bank of England received DVDs and later Blu-Ray discs from BvD on a monthly basis. These discs contained a snapshot of the FAME database for UK firms during the month in question. We refer to these discs as different vintages of the database. From month to month, the database is updated both as companies filed new annual accounts and as companies conducted event driven filings with Company's House. However, for the majority of firms there is no change from one month to the next as no new filings take place.

Our general principle was to sample these discs at a six monthly frequency. We did not pursue a higher frequency as the cost in terms of the amount of time needed to process each disc and the capacity required to store the information was excessive given how little additional information would be gained. The recorded information for an individual company does not change so frequently as to require multiple observations within a six month period. In principle, since accounts are typically filed on an annual basis, we could have also sampled the discs annually and still have guaranteed that for any given company, all the annual accounts filed over our sample period would have appeared as the most recent observation in at least one of the sampled discs. However, we chose biannual sampling for two reasons. First, companies can occasionally have irregular filing periods, if a firm changes its financial year end date, and file twice within a year we would like to have both filings. Second, as described above, other company information can change outside of accounting periods. These are so-called event driven filings. By sampling discs at a biannual frequency we are less likely to have event driven filings causing a deviation between the non-accounting information accurate as of when the disc was produced and the accounting information that is accurate as of the account filing date.

Over the course of the past decade some of the Bank of England's discs have been lost or become damaged so we are not able to pick the same months in every year to conduct our sampling. We chose the last available monthly disc in each half of the year - i.e. June and December are our preferred discs for any given year. If either June or December were not available we substitute in May or November etc. If no disc was available in a half year (for instance, if there are no discs available between January and June) we would use the next available disc in the following half of the year. The complete list of discs used is below:

January 2005, December 2005, June 2006, December 2006, May 2007, December 2007, June 2008, December 2008, May 2009, December 2009, June 2010, September 2010, September 2011, December 2011, April 2012, November 2012, August 2013, December 2013, June 2014, September 2014, August 2015.

B.3.2 Download Strategy

We focus on companies that have either a registered office or primary trading address in England, Wales or Scotland. Our downloads were conducted in regional blocks within each disc and we extracted data for both active and inactive companies. All the data we use is denominated in GBP. The discs have an inbuilt panel structure in the sense that it is possible to download up to 10 years of historical observations for a firm in each vintage of the database. We exploited this by downloading the most recent observation for each firm and two years of lags for discs in the middle of our sample. For the first disc (January 2005) in the sample we downloaded five years of lag data to add additional historical coverage. For the final disc, August 2015, we downloaded the full 10 years of data in order to evaluate the benefits of using the archive discs versus a single snapshot of the database.

B.4 Merging the Discs into a Firm Panel

Each company is the UK is assigned a unique Company's House Registration Number (CRN) upon formation which stays with the company throughout its lifetime. The CRN may change if Companies House chooses to adopt a new numbering format (see Section 1066 of the Companies Act 2006). Fortunately this did not happen over our sample period thus we use the CRN as an identifier to determine the same firms across different vintages of FAME. This allows us to build a firm level panel using information across all vintages.

Beyond expanding the time coverage of the data, forming the firm panel has three benefits over using 10 lags of downloaded data. First, the combined data across vintages contains information on companies who were alive historically but died and are no longer present in the final disc, guarding against survivorship bias. Second, the combined data improves coverage for the companies still present in the final disc, with historical observations sometimes lost from the database between vintages.[26] Third, our approach allows the observation of event-driven variables such as *company status* (whether the company is live, dormant, dissolved etc) when they were first reported, which is important as these variables can change independently of the company accounts. In the case of company status this is particularly important, enabling a more accurate measure of the date when a company dies, as discussed further below.

The benefits the panel structure brings over a single download with 10 lags of data in terms of firm coverage and reporting of variables is set out following the explanation of the firm panel construction.

B.4.1 Treatment of the First Disc

In the first disc (January 2005), we used the additional lagged accounting information as historical observations of the firm accounts. The dates of historical accounts are generated using the *statement date* of the latest set of accounts and the *number of months* covered by the accounting period (12 months in the vast majority of cases), both are variables reported in FAME. For young firms, this process can generate purported accounting periods before the firm was born. To correct for this, all historically generated observations where the *incorporation date* is after the statement date are dropped.

B.4.2 Treatment of Multiple Observations on the Same Firm Accounts

With company data downloaded from BvD at a biannual frequency, the same set of company accounts frequently appear in multiple different BvD discs (up to a maximum of 21 observations on the same accounts). The next step in the formation of the firm panel is to treat these multiple observations on the same set of accounts. At this point the data set is restricted to companies that report the statement date of their accounts, allowing a given set of company accounts to be uniquely identified using the CRN and statement date. The treatment of data is broken up into three groups:

Variables Never Revised by Later Data. As discussed above, information on directors is event-driven, and can change outside of firm accounting periods. To ensure accuracy, for all director variables we retain information from the earliest disc where the accounts are filed. In particular, this process omits information on directors appointed after the disc when the accounts were first published. Multiple trading addresses listed by the firm are treated in the same manner.

[26] For reasons that are unclear, again the information from BvD suggests this should not happen.

Variables Only Revised by Later Data When Initially Missing. A small number of other variables such as the *company status* and the *primary UK sic code* (the primary industry to which the company belongs) can be changed independently of the firm accounts but take a unique value per firm at a given point in time, and are less likely to change over time. For these variables information is used from the earliest disc in which the accounts appear. However, in contrast to director information, as these variables are less likely to change over time, the initial observations on a variable are replaced with subsequent observations if it is initially missing. Table 2 provides a stylised example of this for variables with and without missing data. This treatment also covers lagged accounting information held on the discs.

Table 2: Treatment of Duplicate Accounts:

Variables *Only* Revised When Missing

Firm	BvD disc	Account Date	Variable X	Variable Y
1	A	31/03/2006	x_A	
1	B	31/03/2006	x_B	y_B
1	C	31/03/2006	x_C	y_C

Resolved Accounts

Firm	BvD disc	Account Date	Variable X	Variable Y
1	n.a.	31/03/2006	x_A	y_B

Variables Always Revised by Later Data Unless Subsequently Missing. The remaining data are accounting variables such as *land and buildings* and *number of employees* that are specific to the accounting period in question. Companies revise their historical accounts over time and using the panel structure such revisions are captured. The general principle is to use the latest data on the company's accounting period for these variables, capturing improvements made to the accounts from subsequently filed revisions. Sometimes these data revisions are only filed for the variables that have changed, which can result in missing values on non-revised variables in later discs. To circumvent this problem, the latest non-missing data is taken for this group of variables. Table 3 provides an example of this, for variables with and without missing data. As with the prior group of variables the treatment here is also applied to lagged accounting information.

B.4.3 Treatment of Downloaded Lagged Accounting Information.

Following the data harmonisation in the prior step, for each company statement date there is a unique observation for every variable. This includes the current value of accounting variables at each statement date, as well as two years of lagged accounts. The next step combines this

Table 3: Treatment of Duplicate Accounts:

Variables Revised *Unless* Missing

Firm	BvD disc	Account Date	Variable W	Variable Z
1	A	31/03/2006	w_A	z_A
1	B	31/03/2006	w_B	z_B
1	C	31/03/2006	w_C	

Resolved Accounts

Firm	BvD disc	Account Date	Variable X	Variable Y
1	n.a.	31/03/2006	w_C	z_B

lagged accounting data with data from previous accounts, to incorporate revised accounting data. The first step is to identify and treat missing accounts.

Identifying and Treating Missing Accounts. The firm data is set to panel form using the company CRN and *statement date* of accounts. Before harmonising lagged accounting data with previous accounts it is determined if any firm account observations are missing. Using the *statement date* and *number of months* variables (length of the accounting period) it is determined if successive accounts are the correct number of months apart. Prior to treatment, 97.8% of company observations have no accounts missing, with 1.8% having one set of accounts missing and 0.1% having two accounts missing. Accounts are generated where missing accounts are identified (up to four missing accounts), with the statement date set as the *statement date* of the subsequent accounts less the *number of months* in the accounting period associated with that statement date (taking the last day of the month in question). For the generated accounts, variables without lagged accounting data are assumed to take the same value as at the first statement date after the missing accounts. Following this treatment, 99.81% of company observations have no accounts missing. Variables with lagged accounting data are treated for the missing accounts in the same way as for the rest of the data set, as discussed next.

Harmonisation of Accounting Data. As accounting data can be revised, the general principle for treating it is to use the latest available non-missing data.[27] A stylised set of accounts are presented in Table 4. When there are no accounts missing for a firm and accounting data has not been revised, the diagonal entries in the table will be the same. Thus, for example, the current value of variable x in the 2006 accounts will be the same as the first lag of x in the 2007 accounts, which will in turn be the same as the second lag of x in the 2008 accounts: $x_{C,2006} = x_{L1,2007} = x_{L2,2008}$. Where accounting revisions occur these values will differ.

[27] One exception is the *QuiScore* (a credit score for the firm) which can change outside of the filing of the firm's accounts. For this variable the earliest non-missing value is taken.

Table 4: Treatment of Lagged Accounting Information

Firm	Account Date	No. Months	Variable X, Current	Variable X, Lag 1	Variable X, Lag 2
1	31/03/2006	12	$x_{C,2006}$	$x_{L1,2006}$	$x_{L2,2006}$
1	31/03/2007	12	$x_{C,2007}$	$x_{L1,2007}$	$x_{L2,2007}$
1	31/03/2008	12	$x_{C,2008}$	$x_{L1,2008}$	$x_{L2,2008}$
1	31/03/2009	12	$x_{C,2009}$	$x_{L1,2009}$	

Resolved Accounts

Firm	Account Date	No. Months	Variable X, Current	Variable X, Lag 1	Variable X, Lag 2
1	31/03/2006	12	$x_{L2,2008}$	n.a.	n.a.
1	31/03/2007	12	$x_{L1,2008}$	n.a.	n.a.
1	31/03/2008	12	$x_{L1,2009}$	n.a.	n.a.
1	31/03/2009	12	$x_{C,2009}$	n.a.	n.a.

Consider the 2006 accounts in Table 4. No accounts are missed for the subsequent two accounts between the 2006 accounts and the two that follow (with the time between accounts equal to the number of months covered by each of the accounts that follow) so the current data, $x_{C,2006}$, and the elements running down the diagonal of the table, $x_{L1,2007}$ and $x_{L2,2008}$ refer to the same accounting variable over the same time period. In the first instance, the twice lagged accounts from two periods ahead are used to update variable x as this is non-missing. Thus, in the resolved accounts shown below (for 2007, 2008 only), the value is $x_{L2,2008}$ (which may or may not differ from $x_{C,2006}$). Contrast this with the 2007 accounts. In this case, the twice lagged accounts from two periods ahead has a missing value for x. In this case, the latest available non-missing data is the lagged accounts from one period ahead, and so the resolved value for x in the 2007 accounts is $x_{L1,2008}$. With all accounting variables treated in this way the first and second lags in the accounts are dropped, leaving only the current value of x at the accounting date, as shown in the resolved accounts.

B.4.4 Firm Birth and Death

The way we construct our data also allows for accurate calculation of firm birth and death dates.

Firm Death Date. The advantages of using archival data are particularly pronounced for determining the date at which the firm dies. This is due to two main disadvantages from using a single vintage of BVD. First, within a single vintage only shows the current *company status*, not lagged information on it. Thus, it can reveal that a company has been dissolved, but not *when* they were dissolved. Second, as discussed in the next section, there is a significant survivorship bias in the data with many companies that die exiting the database over time. Thus, when using a single download of the data, a significant number of companies who previously died will be omitted, with no way of knowing when they died, or were even last present in the database.

Using the panel of firm information, we form the following definition of a firm's the death date which we call *statement death date.* This variable records the date of the first set of accounts where *company status* is "dissolved" for the company, and, necessarily, is calculated only for the firms where *statement date* is present. It is important to note that if *company status* is live the first time the firm files their accounts for a given year, but subsequently die and file no more accounts, BvD will update *company status* to reflect the death in subsequent vintages. Our treatment process (as described in section B.4.2) will revert the *company status* as live for that set of accounts, as they were when they were first filed. However, we calculate *statement death date* prior to performing this step meaning that we record the death date accurately and that the company was live when the accounts where filed.

Firm Birth Date. The firm's birth date is given by the variable *incorporation date.* This is seldom absent, but where it is, it can be imputed using its value across different vintages of BvD discs, given that it's value doesn't change over time.

B.4.5 Enhancement of Data Coverage Through Using the Firm Panel

The final combined panel of firms, comprised of companies with non-missing statement dates, contains 28.9 million firm-account observations, with 4.8 million distinct firms. The combination of data across several vintages has significant advantages over data extracted from a single vintage:

- First, and most straightforwardly, with a maximum of 10 sets of accounts being accessible from a given vintage, by using multiple historical discs, a greater time period can be covered.

- Second, even within the time period covered by the 10 lagged accounts, our merged dataset brings significant benefits in terms of coverage of the accounting information firms report.. To demonstrate this, we downloaded 10 accounts for each company from the August 2015 disc and compared the value of firm's *total assets,* a particularly well-reported variable, to the same variable over the same set of 10 accounts using the data as created from the our combined dataset using all 21 discs. The proportion of observations for which *total assets* are missing from each data set is shown in Figure 12. Using the full panel, *total assets* is consistently well-reported, as shown in red, with data missing for only around 3% of firm observations throughout the sample. Data downloaded only from the 2015 disc has similar coverage of *total assets* for the first five accounts, before dropping off substantially, with around a third of observations missing this data by the final lagged accounts.

Figure 12: Proportion of Observations with Total Assets Missing

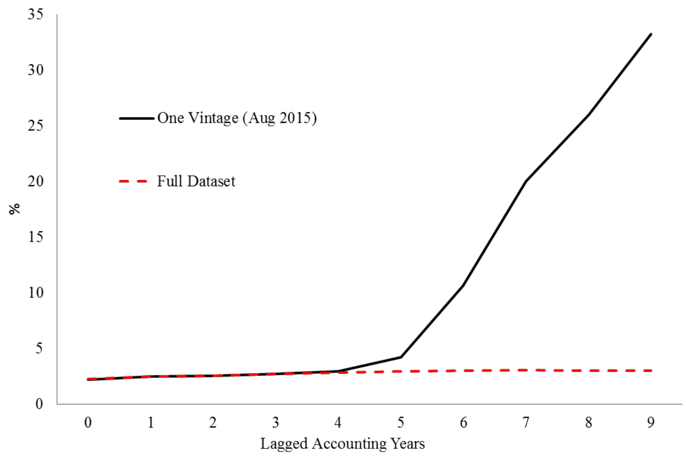

Notes: the figure displays the proportion of total assets missing among companies with a non-missing statement date. *One download* refers to the 10 lagged accounts downloaded for the companies present in the August 2015 disc. *Full panel* refers to the final panel produced from the 21 discs from 2005 to 2015, as described above, covering the same period.

- Third, the combined dataset has significantly greater *firm* coverage. Figure 13 displays the proportion of companies present in each accounting year in the panel that are still present in the August 2015 disc. Only 55% of the companies that filed accounts in 2000 are still present in the August 2015 disc. Note, this is not the requirement that the company *accounts* from 2000 are present in the 2015 disc, only that information on the *company* itself is still present. The difference in asset reporting in Figure 12 is driven largely by firms exiting the database before the 2015 disc. Indeed, 94% of the firm observations where *total assets* is reported in the full panel but not from the 2015 discs have a *statement death date* prior to 2015.

B.5 Sample Selection

Our key sample selection criteria are articulating in the main text; for completeness here we describe the conditions under which companies and observations can enter our sample.

- We restrict our sample to only include limited liability, for profit companies to which the Companies Act applies. Specifically, we include "Private Limited", "Public AIM", "Public Quoted", "Public Not Quoted". This information is contained in the *legal form* field in the FAME database.

- We exclude firms in certain industries based on the *primary UK sic code* field in the FAME database which is available for the 2003 UK Standard Industrial Classification (SIC) codes for all the discs used in our sample. We exclude from the sample firms operating in utilities (2003-SIC: 4011-4100); finance and insurance (2003-SIC: 6511-6720);

Figure 13: Fraction of Companies Present in August 2015 Disc

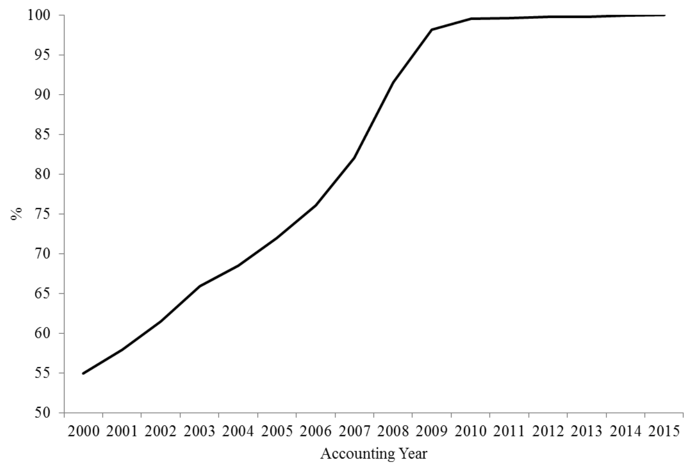

Notes: the figure displays the proportion of companies in each statement year, as derived from the full panel of 21 discs, that are present in the August 2015 disc.

real estate (2003-SIC: 7011-7032); public administration (2003-SIC: 7511-7530); education and charity (2003-SIC: 8010-8540); and clubs and organisations (2003-SIC: 9100-9199). We compute and work with the modal sic code across all of a firm's the observations in the database as a result firms do not switch industry classification throughout our sample.

- We exclude companies that have a parent or are part of a group but not the ultimate owner. Our criteria for doing so is whether the company reports an ultimate owning company on FAME. Those that do not report an ultimate owner company or whose ultimate owning company name is the same as the company name remain in the sample. There is no historical information within FAME about whether or not a company had an ultimate owner. Hence, the ownership information in FAME is only accurate as of the vintage of the database. The use of historical vintages of the database allows us to circumvent this issue and track ownership status through time. We always take data on ownership from the earliest disc available after a company has filed its annual accounts. For simplicity, we exclude all firms who ever have an owner at any point in the sample.

- As our empirical analysis relies upon a mix of flows, stocks and changes in stocks we exclude observations where the accounting period is irregular, e.g. if the company filed two sets of accounts within a year. Specifically, we compute the gap in months between accounts and drop observations where this is greater than 13 or less than 11. In other words, the firm's account window must be between 11 and 13 months to enter our sample. When estimating local projections with dependent variable $log(ACT_{i,t+h})-log(ACT_{i,t-1})$, all accounts filed between $t-1$ and $t+h$ must have a regular accounting period. We do not focus strictly on firms who have a 12 month accounting period as it is fairly common for firms to be recorded as filing at the end of a month in one financial-year and then recorded as filing on the first day of the following month the next financial-year. Observations where there is no information on the filing date are excluded.