NBER WORKING PAPER SERIES

BIG G

Lydia Cox
Gernot Müller
Ernesto Pastén
Raphael Schoenle
Michael Weber

Working Paper 27034
http://www.nber.org/papers/w27034

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

We acknowledge outstanding research assistance by Jiacheng Feng. We thank Klaus Adam, Regis Barnichon, Christian Bayer, Benjamin Born, Gabriel Chodorow-Reich, Zeno Enders, Emmanuel Farhi, Yuriy Gorodnichenko, Edward Knotek, Ludwig Straub and conference and seminar participants at the Banco Central de Chile Heterogeneity in Macroeconomics Conference, Federal Reserve Bank of Cleveland, PUC Santiago, the University of Tuebingen, and the VfS Monetary Meeting Frankfurt for valuable comments. Weber also thanks the Fama-Miller Center at the University of Chicago Booth School of Business for financial support. Schoenle and Weber also gratefully acknowledge support from the National Science Foundation under grant number 1756997. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Central Bank of Chile, the Federal Reserve Bank of Cleveland, the Federal Reserve System, or the National Bureau of Economic Research.

Big G
Lydia Cox, Gernot Müller, Ernesto Pastén, Raphael Schoenle, and Michael Weber
NBER Working Paper No. 27034
April 2020
JEL No. E32,E62

ABSTRACT

Big G typically refers to aggregate government spending on a homogeneous good. In this paper, we open up this construct by analyzing the entire universe of procurement contracts of the US government and establish five facts. First, government spending is granular, that is, it is concentrated in relatively few firms and sectors. Second, relative to private expenditures its composition is biased. Third, procurement contracts are short-lived. Fourth, idiosyncratic variation dominates the fluctuation of spending. Last, government spending is concentrated in sectors with relatively sticky prices. Accounting for these facts within a stylized New Keynesian model offers new insights into the fiscal transmission mechanism: fiscal shocks hardly impact inflation, little crowding out of private expenditure exists, and the multiplier tends to be larger compared to a one-sector benchmark aligning the model with the empirical evidence.

Lydia Cox
Harvard University
lcox@nber.org

Gernot Müller
University of Tuebingen and CEPR
Mohlstr. 36,
72074 Tuebingen,
Germany
gernot.mueller@uni-tuebingen.de

Ernesto Pastén
Central Bank of Chile
Agustinas 1180
Santiago
Chile
and Toulouse School of Economics
epasten@bcentral.cl

Raphael Schoenle
Brandeis University
Department of Economics
415 South Street
Waltham, MA 02454
and CEPR
schoenle@brandeis.edu

Michael Weber
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
michael.weber@chicagobooth.edu

An online appendix is available at http://www.nber.org/data-appendix/w27034

1 Introduction

What is “Big G”? In the national accounts G refers to “government spending”—the part of GDP that comprises government expenditures. This convention possibly helps explain why research on fiscal policy typically entertains a somewhat crude notion of government spending as a homogeneous good, isomorphic to GDP. In empirical and theoretical work, we frequently refer to it as Big G, and the literature assumes policy makers can freely adjust it over time—in response to the business cycle, or for other reasons. The recent “renaissance of fiscal research” survey by Ramey (2019) has changed little in this regard. A number of recent contributions have started to study the role of heterogeneity for the fiscal transmission mechanism but focus exclusively on heterogeneity on the household side (McKay and Reis, 2016; Auclert et al., 2018; Hagedorn et al., 2019).

Starting point of our paper is the observation that Big G itself is fundamentally heterogeneous. Government spending is not simply one large transaction. It is composed of a large number of smaller transactions whose composition differs from the other components of aggregate demand. Empirically, we first establish five facts about government spending by characterizing the underlying components of Big G. In the second part of the paper, we then study the role of these facts for the fiscal transmission mechanism through the lens of a stylized two-sector New Keynesian model. Accounting for the heterogeneity of government spending has first-order effects on the transmission mechanism aligning the model prediction with empirical evidence.

We construct our empirical anatomy of Big G using a database that has only recently become accessible: USASpending.gov. The database provides detailed information on the entire universe of procurement contracts by the US federal government since 2001. These data capture about half of federal consumption expenditures which, in turn, account for about one half of general government spending. For each year, the database records several million government procurement transactions. We establish five facts on the basis of detailed analysis of these data.[1]

The first of our five facts is government spending is granular in the sense of Gabaix (2011). Few firms and sectors supply a large share of government consumption: (i) The largest 20% of suppliers supply 99% of government consumption. (ii) The top 10 firms—or the top 0.01% — among all firms supplying goods and services to the federal government receive more than 35% of all procurement contracts by value and the top 0.1% of firms receive almost 60%. (iii) The most important suppliers to the federal government are concentrated in few sectors: firms in the largest three out of the roughly 20 two-digit NAICS industries supply more than 60% of all government

[1]Defense spending accounts for more than one half of the transactions by value in our data set. We replicate the five facts separately for defense and non-defense spending in Online Appendix A.5.

Figure 1: Consumption Shares: Government vs Household

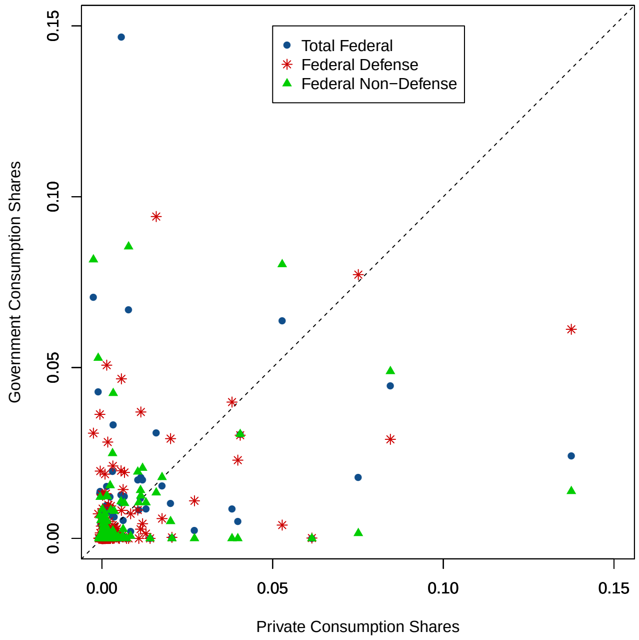

Notes. This figure shows the fraction of government spending in a certain sector on the y-axis and the fraction of private spending in a certain sector over total private spending on the x-axis. We separate total federal spending (blue circles) into defense spending (red asterisks) and non-defense spending (green triangles). We use the BEA *Use Table* to calculate private consumption shares. The sample represents averages over the period between 2001 and 2018.

contracts and the top 1 percent of just over 1,000 six-digit NAICS industries make up around 40% of all government spending. Decomposing the total cross-sectional variance of government contracts, we find almost 100% of the variation is "within" firms or sectors and almost none is "across" firms or sectors providing further evidence of the granular nature of government spending. The underlying cross-sectional size distribution of contracts is characterized by fat tails, providing the basis for these facts.

The second fact is the existence of a sectoral bias in government spending: the share of government spending in each sector differs substantially from the share consumers spend on the goods and services of that sector. Figure 1 shows the share of government spending in a certain sector on the vertical axis, and the same ratio for private consumption on the horizontal axis. We separate total federal spending (blue circles) into defense spending (red asterisks) and non-defense spending (green triangles). Both for overall federal spending, but also for both subcomponents, a substantial difference in the spending patterns of consumers and the government exists. Some sectors that are negligible for the government make up about 14% of private consumption, whereas sectors that are unimportant for consumers are big suppliers to the government. This sectoral bias also holds for non-defense spending. Hence, government spending varies across sectors and does not purely mimic consumer spending. In earlier work, Ramey and Shapiro (1998) stressed the importance of sectoral bias for the fiscal transmission mechanism. Until now, however, no data was available to establish the sectoral bias of government spending systematically.

Third, we show government contracts have a short duration and are often modified. The median contract has a duration of 36 days, 80 percent of contracts last less than one year, and about 30 percent of contract transactions represent modifications to initial contracts. Hence, the government does not tend to enter long-term contracts with suppliers. Only few contracts lasts very long. The median firm supplying goods to the government is in the dataset for 2 years, while the firm with the median value of average annual obligations is in the dataset for only 1 year.

Fourth, idiosyncratic shocks dominate the fluctuations in government spending over time – rather than in the cross section (which the first fact studies). To establish this fact, we decompose growth rates following Gabaix (2011) and Foerster et al. (2011) and find idiosyncratic shocks, both at the firm and the sectoral level, are again key drivers of variation in government spending over time. In addition, when we estimate AR(1) processes for government spending at the sectoral level and study the correlations in the residual spending across sectors, we find aggregate shocks play a negligible role for changes in sectoral government spending over time. Hence, large variation of government spending over time exists that variation in Big G cannot account for. Instead, if an innovation occurs, it is idiosyncratic at the sectoral level. These innovations have both large negative as well as large positive correlations for many sector pairs. Overall, sectoral government spending is generally relatively persistent, consistent with our cross-sectional variance decomposition.

Fifth, government consumption tends to be concentrated in sectors with a relatively high-degree of price stickiness. The frequency of price changes in the top two two-digit NAICS sectors is 9% while it is on average 20% for the remaining sectors in the economy. We use the micro data underlying the producer price index at the Bureau of Labor Statistics to estimate these frequencies. The average frequency of price adjustment overall is 16% which corresponds to prices adjusting approximately every 6 months. Our detailed contract data allow us to further characterize the way in which prices are sticky. The contract data contain information on the types of contracts between buyer (the government) and sellers. The majority of contracts—over 85 percent— are "fixed-price" in nature.

The facts we establish might not appear surprising but so far, no systematic evaluation exists. To better understand whether they matter from a macro perspective, we feed these facts into a two-sector New Keynesian model with government spending a la Woodford (2011) and compare the implications of the model to a one-sector benchmark. The model is deliberately stylized in order to account for the five facts as clearly as possible while only minimally departing from the conventional one-sector model. Importantly, rather than postulating a process for Big G, as is commonly done, we model government spending in each sector as a distinct variable.

Sectoral heterogeneity induces profound changes in the fiscal transmission mechanism in our two-sector model relative to the benchmark economy. We derive a number of closed-form results for the limiting case in which prices are completely flexible in one sector. If government spending is biased towards the flex-price sector, crowding out of private expenditure can be infinite. Empirically, however, government spending is biased towards the sticky-price sector. An increase of government spending in the sticky-price sector induces little crowding out of private expenditure, and hence the output multiplier is considerably larger relative to the one-sector benchmark.

We also run model simulations and show the sectoral heterogeneity of government spending matters quantitatively. Specially, we calibrate the model to capture key features of the data, including the actual degree of price rigidities as well as the sectoral composition of government spending and the relative size of sectors. A fiscal shock in the relatively small and sticky-price sector towards which government spending is biased induces a multiplier effect about three times larger than a shock in the other sector. Moreover, the multiplier becomes even bigger if prices are more flexible in the sector in which private expenditure is concentrated. Hence, just like Barsky et al. (2007) and Barsky et al. (2016) show for the transmission of monetary policy, we find the degree of price stickiness in the sector in which government spending is concentrated is essential and not the economy-wide stickiness.

In the New Keynesian model, monetary policy is also key for the fiscal transmission mechanism (Woodford, 2011; Christiano et al., 2011; Farhi and Werning, 2016). Government spending is inflationary and thus (generally) triggers a response of the central bank. The resulting interest rate increase crowds out private expenditure because of intertemporal substitution. As a result, the multiplier is smaller than unity—in contrast to the textbook IS-LM model in which no intertemporal substitution takes place. In related work, Boehm (2019) distinguishes between government consumption and government investment and finds the multiplier is particularly small for government investment precisely because the intertemporal elasticity of substitution of investment demand tends to be high. Auerbach and Gorodnichenko (2012a,b) and Ramey and Zubairy (2018) discuss empirically whether government spending multipliers are larger in recessions and periods of low interest rates when monetary policy might be less responsive to government spending. A number of recent contributions introduce household heterogeneity and credit frictions in New Keynesian models in order to limit intertemporal substitution (Galí et al., 2007; McKay et al., 2016; Kaplan et al., 2018). As a result, multipliers tend to be larger.

We show sectoral heterogeneity in spending by households and the government combined with sectoral heterogeneity in pricing frictions has a similar effect and as a result, the New Keynesian model becomes "more Keynesian." In a nutshell, since the government spends in relatively sticky-price sectors and the private sector spends in relatively flexible-price sectors, inflation is more responsive to private expenditure than to government spending. Such differential heterogeneity in turn dampens the monetary response to a fiscal impulse: less intertemporal substitution occurs, less crowding out, and the multiplier is larger. We also show, however, things are turned up-side down at the zero lower bound (ZLB). The ranking of multipliers across sectors flips: raising government spending in the relatively flexible sector has now a larger impact because no interest rate response occurs to curb the larger inflationary pressure and hence more crowding in of private consumption happens.

Empirically, accounting for sectoral heterogeneity helps the model to generate predictions that align better with the time series evidence than predictions from the conventional one-sector model. Many studies have established the response of interest rates, both nominal and real, as well as the response of inflation to fiscal shocks tends to be weak or even negative (Mountford and Uhlig, 2009; Corsetti et al., 2012; Ramey, 2016) which is exactly what our model generates. Once we modify the model to account for the evidence on government spending at the micro level, the model also gets the macro evidence right. Our data also allow studying whether fiscal multipliers differ for defense versus non-defense spending but we leave a systematic analysis of fiscal multipliers for future work.

Our paper is related to recent work on the effect of regional fiscal policies in monetary unions (Galí and Monacelli, 2008; Nakamura and Steinsson, 2014; Hettig and Müller, 2018). In this literature, government spending is concentrated in some spatial partition of the economy, and its composition is biased relative to the composition of private expenditures. Just like in our analysis, the effects of fiscal policy turn out to be highly sensitive to the conduct of monetary policy. In contrast to this earlier work, we model private expenditure as being determined at the aggregate level rather than at the regional/ sectoral level. Chodorow-Reich (2019) surveys the recent empirical work on government spending multipliers based on cross-sectional data. Last, we also share modeling features with a number of recent papers that account for heterogeneity on the production side across sectors and firms, tracing out the implications for the business cycle (Acemoglu et al., 2012; Pasten et al., 2019a,b; Baqaee and Farhi, 2019; Ozdagli and Weber, 2017). Bouakez et al. (2018), in particular, study theoretically the transmission of fiscal policy shocks in a rich model featuring heterogeneity in sector size and input-output structure.

2 Data

2.1 Background on USASpending

In the first part of this paper, we undertake a comprehensive analysis of the USASpending.gov database—the official source for federal spending data.[2] We first detail and define several fundamental concepts before we move on to analyze the data. The database was created in response to the Federal Funding Accountability and Transparency Act (FFATA), which was signed into law on September 26, 2006. FFATA requires federal contract, grant, loan, and other financial assistance awards of more than $25,000 are publicly accessible on a searchable website, in an effort to provide transparency to the American people on how the government spends their tax dollars. In accordance with FFATA, federal agencies are required to collect and report data on federal procurement. Agencies must report award data—contracts, grants, loans, and other financial assistance— on a monthly basis through various government systems such as the Federal Procurement Data System (FPDS-NG) for contract data and the Data Act Broker for grant, loan, and other financial assistance data. Some agencies report frequently during a month, while others report once a month or even less frequently if they do not issue awards on a monthly basis. The USASpending.gov database which the Treasury Department hosts, compiles the data from the various government reporting systems. In addition to directly uploading the information which the federal agencies report to systems like the FPDS-NG, the site also utilizes information collected from the recipients of the awards themselves. Though FFATA was not signed into law until 2006, data are available back to 2001 through an external organization.

2.2 What are Government Contracts?

Our data focus on a subset of federal spending—spending on goods and services via government contracts. The Federal Acquisition Regulation (FAR) defines "Contract actions" as "any oral or written action that results in the purchase, rent, or lease of supplies or equipment, services, or construction using appropriated dollars over the micro-purchase threshold, or modifications to these actions regardless of dollar value." The micro-purchase threshold is in general $3,500. As the definition suggests, the goods and services that the government consumes through contracts span a wide range, from janitorial services for federal buildings to IT support services to airplanes and rockets. Contracts can be short-term—e.g., a one-month contact awarded by the Department of Agriculture Rural Housing Service to Sikes Property And Appraisal Service

[2] Demyanyk et al. (2019) and Auerbach et al. (2019) also rely on this database but their focus is on estimating fiscal multipliers. They rely exclusively on contracts awarded by the Department of Defense. On average, these account for about half of the transactions in the database by count and for about two thirds by value. Appendix A.2 provides an overview of other similar data that have been used in the literature.

for single family housing appraisals in September 2008—or longer-term relationships—e.g., the 43 year and 10 month contract awarded by the Department of Energy to Leland Stanford Junior University for the operation and management of the SLAC National Accelerator Laboratory. In awarding contracts, federal agencies must abide by the guiding principles set forth in the FAR. The FAR includes directives on every aspect of contracting, from how contracts should be structured and priced, to how they should be solicited to promote competition and encourage small business participation.

2.2.1 Type of Awards

The government can use different types of awards to procure services. The majority of federal spending through contracts is done through either a definitive contract action (DCA) or a delivery order. A DCA is a legally binding agreement obligating the seller to furnish certain supplies or services and the buyer to pay for them. For example, on April 27, 2018, Lockheed Martin was awarded an $828,724,214 contract to build Guided Multiple Launch Rocket Systems for the Department of the Army. Funds for the project were obligated at the time of the award, and the expected time of completion is September, 2021.

A delivery order, on the other hand, is a contract that does not specify a firm quantity, but provides issuance of orders for the delivery of goods or services during the period of the contract. For example, on January 21, 2015, a company called Ace Maintenance & Services, Inc. was awarded a $13,663,688 contract for janitorial services at Naval Support Activity Bethesda. The work to be performed under the contract included all labor, supervision, management, tools, materials, equipment, facilities, incidental engineering, and other items necessary to provide janitorial services. The initial contract action was for a base period of one year and one month, with the option of four additional years. The contract stipulated a maximum dollar amount for the base period and four option years of $69,698,540. DCAs tend to be used for larger, one-time purchases, while delivery orders are used for smaller and more frequent purchases. We provide additional details on different types of awards see in Online Appendix Section A.3.1.

2.2.2 Type of Contract Pricing

In addition to the type of award, a wide selection of contract pricing are available to the Government and contractors. Contract types are grouped into two broad categories: fixed-price contracts and cost-reimbursement contracts. Within those categories, specific contract types vary according to the degree of risk placed on the contractor for the execution costs of the contract, and the nature of the incentives offered to the contractor for their performance. The most common type of contract is a firm-fixed-price contract, which details a price that is not subject to any adjustment, regardless of the contractor's actual cost experience in executing the contract. Fixed-price contracts can also include provisions for economic adjustment or incentive payments, somewhat reducing the risk placed on the contractor.

Cost-reimbursement contracts are also frequent, and typically include a negotiated fixed fee or an award amount on top of the reimbursement payment. We discuss in further detail what the data on contract pricing look like when we discuss our fifth fact in the next section.

The pricing structure of a contract depends on many factors—price competition, the complexity and urgency of the requirement, and the length of the contract, to name a few. Many contracts are complex and require hybrid pricing structures—the Multiple Launch Rocket System contract mentioned above, for example, is a "cost-plus-fixed-fee, firm-fixed-price, and fixed-price-incentive" hybrid. According to the FAR, "the objective is to negotiate a contract type and price (or estimated cost and fee) that will result in reasonable contractor risk and provide the contractor with the greatest incentive for efficient and economical performance."

2.2.3 Competition

Federal regulations generally require contracting officers to promote full and open competition in soliciting offers and awarding government contracts. In most cases, agencies are directed to use sealed bids, competitive proposals or some combination of competitive procedures to solicit and issue awards. Ultimately, about half of the awarded contracts were fully and openly competitive via negotiated proposals. The Ace Maintenance & Services, Inc. contract for janitorial services, for example, was solicited using the Navy Electronic Commerce Online website, and seven proposals were received.

A number of cases, however, exist, in which full and open competition is not required. Some contracts are deemed "not available for competition," in which case agencies are authorized by statute to solicit bids from only one source. Solicitation from one source is authorized if, for example, the supplies or services required by the agency are available from only one responsible source or, for the Department of Defense, NASA, and the Coast Guard, from only one or a limited number of responsible sources. Supplies can also be deemed available from only the original source if the contract is a follow-on to an existing contract for the continued development or production of a major system or highly specialized equipment. The Lockheed Martin contract for Guided Multiple Launch Rocket Systems is an example for the latter.

Finally, for smaller awards—those below a certain dollar threshold—federal agencies are required to use "Simplified Acquisition Procedures (SAPs)," which reduce administrative costs, increase efficiency, and improve opportunities for small and minority-owned businesses.

2.3 Scope of the Dataset

The dataset we use includes the universe of federal government contract transactions from fiscal years 2001 through 2018.[3] On average, 3.2 million individual contract transaction records exist each year—with almost 5 million annual contracts toward the end of the sample period. The contracts are awarded to over 160 thousand recipient parent companies each year, spanning over 1000 six-digit NAICS sectors. The median transaction value is just under $2,300, while the mean transaction value is just under $140,000, suggesting the distribution is heavily right skewed. The majority of contract transactions (82 percent) represent positive obligations from the government to firms, but also transactions with negative value exist, or deobligations, which occur when a modification to an initial contract is performed (see section 3.3 for details). Figure 2a shows contract obligations are roughly equivalent to total federal government purchases of intermediate goods and services plus gross investment (from the National Income and Product Accounts).[4] Contract obligations represents 12 to 18 percent of general government spending, or about 2 to 4 percent of GDP over the sample period, see Figure 2b.

Each observation in the data traces a contract action from its origin (the parent agency) to the recipient firm (which can be a subsidiary of a parent firm) and the sector and zip code within which the award is executed (see figure A.3 in Appendix **??** for a schematic representation of the data). Six variables uniquely identify each observation: (1) an award identification number, (2) a modification number, (3) a transaction number, (4) a parent award identification number, (5) an awarding sub-agency code, and (6) a parent award modification number.

In our analysis, we outline a number of facts of what we refer to as individual transactions (the observation level of the data), firm-level statistics—for which we aggregate by the recipient parent firm, and sector-level statistics—for which we aggregate by NAICS sectors. The value of each contract transaction, or obligation, is given by the "federal action obligation"—the government's liability for an award transaction. Each transaction is associated with a start and end date for the period of performance of that transaction (barring any subsequent modifications), which we use to calculate "duration." Finally, a transaction will have a "modification number" if it represents an action that makes a change to an initial award. In Section 3, we use these detailed data to document five facts about the nature of this portion of government spending.

[3] Data for fiscal years 2008-2018 can be downloaded from the "Award Data Archive" on the USASpending.gov site. Prior to fiscal year 2008, we use the Custom Award Data download to obtain all prime award contracts for fiscal years 2001 through 2007.

[4] Federal government purchases of intermediate goods and services are equal to federal government consumption expenditures minus compensation of employees and consumption of fixed capital. This number is also equivalent to the gross output of general government minus value added (NIPA Tables 3.9.5 and 3.10.5). Government gross investment consists of spending by the government for fixed assets that directly benefit the public (e.g., highway construction) or that assist government agencies in production activities (e.g., purchases of military hardware).

3 Facts on Government Spending

Government spending is conventionally viewed as a homogeneous good—a relatively constant fraction of GDP that is determined by an ethereal government entity, "G." In this section, we describe five facts about government spending that illustrate government spending is in fact heterogeneous in nature. The granularity in government spending echoes the recent focus on granularity in the firm-size distribution and the input-output structure of the economy but is distinct from it as we show below (Gabaix (2011) and Acemoglu et al. (2012)).

3.1 Granularity

This subsection presents our first and most fundamental fact, government spending is "granular". We use different methods to illustrate this fact. A common definition of granularity proposes a few sectors or firms are disproportionately larger than others. A stricter definition of granularity is in terms of fat tails (see for example Gabaix (2011)): When the size distribution of sectors or firms exhibits fat tails, then some firms or sectors are disproportionately large and granular at any level of disaggregation.

Government spending is granular according to two definitions. First, it is concentrated among a few firms and sectors. Second, a log-normal distribution approximates the government spending distribution well at the transaction level.

Fact 1 *Government spending is "granular:"*

1. *The top 1% of firms receive 80% of all contract obligations and the top 1% of six-digit sectors receive 40% of all contract obligations (where we define rank in terms of firm or sector sales). The top 0.01% of firms receive 30% of contract obligations.*

2. *Nearly 100% of cross-sectional variation in contract spending is within firms or sectors, rather than across.*

3. *The size distribution of contracts has fat tails — in particular, it is approximately log-normal.*

3.1.1 Spending is Concentrated Among Few Firms and Sectors

The first sense in which government consumption is granular is that it is highly concentrated among few firms and sectors. The ten largest suppliers of goods and services to the government (or top 0.01%) account for about one third of total government consumption, and the top 0.1% of firms account for just under one half of total government consumption. Figure 3 illustrates this unequal distribution in the left panel. To put this into perspective, we note on average some 140,000 firms exists in our sample.

A similar spending concentration exists among sectors. The right panel of Figure 3 shows over 60% of contract obligations are directed toward the top three (of roughly 25) two-digit NAICS sectors: 33—manufacturing; 54—professional, scientific, and technical services; and 56—administrative and waste management. The middle panel of Figure 3 shows similar patterns at the more disaggregated sector level—the top 1% (of roughly 1200) six-digit sectors account for about 40% of government consumption, while the top 10% of six-digit sectors account for over 80% of government consumption. Figure 3 also shows the concentration of spending among firms and sectors has been fairly stable over time.

3.1.2 Large Contracts and Firms Drive Cross-sectional Variance

Another way to look at the granular nature of government spending is by looking at a decomposition of the variance of government contracts into the variation that occurs within firms, and the variation that occurs across firms and similarly for sectors. The first decomposition starts with the contract transaction level as the smallest unit of observations, the second with the firm. Specifically, we first calculate:

$$\sum_f \sum_{i \in f} (g_{if,t} - \bar{g}_t)^2 = \underbrace{\sum_f \sum_{i \in f} (g_{if,t} - \bar{g}_{f,t})^2}_{\text{Within Firm}} + \underbrace{\sum_f \sum_{i \in f} (\bar{g}_{f,t} - \bar{g}_t)^2}_{\text{Across Firm}}$$

where $g_{if,t}$ is the total spending amount on individual contract i at firm f in year t, $\bar{g}_{f,t}$ is the firm average in year t, and $\bar{g}_t$ is the overall average in year t. Figure 4 shows this decomposition for all contracts in the left panel, for the top 20% of contracts (which represent 97% of the total value of contracts) in the middle panel, and for the bottom 80% of contracts in the right panel. When we look at the within-firm versus across-firm breakdown for all firms, almost 100% of the variation is "within"—meaning substantial variation exists in the range of contract sizes that an individual firm receives, which completely outweighs any variation in the size of contracts across different firms. The fat right tail of the contracts data fully drives this result.

The empirical variance at the granular level is large and dominates the decomposition. The left panel of Figure 5 shows the density of the log of individual contracts, the density of the log of the average contract amount by firm, and the log of the average contract amount overall. The fat right tail of individual contracts is apparent, and is averaged out at the firm level, creating the high within-firm variation. Looking at the middle and right panels of Figure 5, the top 20% of contracts fully determine this within result. When we restrict our attention to the bottom 80% of contracts, the fat tails are absent and both within- and across-firm variations are present.

Granularity across firms also has implications for the variance decomposition within and across sectors. Instead of looking at the variance of the size of individual contracts within and across firms, we can sum contracts up to the firm level, and decompose the overall variance into the within-sector and across-sector components. Specifically, we calculate:

$$\sum_{s}\sum_{f\in s}(g_{fs,t}-\bar{g}_t)^2=\underbrace{\sum_{s}\sum_{f\in s}(g_{fs,t}-\bar{g}_{s,t})^2}_{\text{Within Sector}}+\underbrace{\sum_{s}\sum_{f\in s}(\bar{g}_{s,t}-\bar{g}_t)^2}_{\text{Across Sector}}$$

where $g_{fs,t}$ is the total amount given to firm f in sector s in year t, $\bar{g}_{s,t}$ is the sector average in year t, and $\bar{g}_t$ is the overall average in year t. Figure A.4 in the Online Appendix shows "within sector" variation dominates across all parts of the distribution. Hence, larger variation exists across firms within a sector, than across sectors within the economy.

Just as in the firm-level exercise, the fat right tail in the data again drives this result. Figure A.5 in the Online Appendix shows the density of firm size has a fat right tail in the case of the full dataset (left panel) and top 20% of firms (middle panel) but a fat left tail in the case of the bottom 80% of firms. In all cases, the fat tail is averaged out at the sector level, creating the high within-sector variation that we see across the board.

3.1.3 The Size Distribution of Contracts Has Fat Tails

Government spending is granular in a statistical sense: The distribution of government contracts is fat-tailed and, in particular, well approximated by a log-normal distribution. A simple way to illustrate this point is to look at a Q-Q (quantile-quantile) plot, in which we plot the actual quantiles of the log transaction values against a set of quantiles from a simulated log-normal distribution with the same mean and variance. If both sets of quantiles come from the same distribution, the plotted points should line up along the 45-degree line. Figure 6 shows that this is the case — the scatter points roughly follow the 45-degree line across the entire distribution. Figure 7 shows the actual density of transaction values and the density of a simulated variable that is log-normally distributed with the same mean and variance, confirming the log-normal distribution appears to be a good fit in the tails. While a log-normal distribution is the best fitting fat-tailed distribution for the full sample of government contracts, we show in Appendix A.4.1 that a Pareto distribution, as in Gabaix (2011), also provides a good approximation to the right tail of the distribution.

3.2 Sectoral Bias: Differential Granularity

The second fact we present establishes government consumption is special compared to household consumption. The composition of government spending across sectors is distinct from the composition of the private consumer basket, which is a natural benchmark for economic activity. The most important firms and sectors as suppliers to the government differ substantially from the most important firms and sectors supplying to private households.

Fact 2 *Government spending is "sectorally biased."*

1. *The top 0.01% of recipients of government obligations account for 17% of average annual government consumption, but only 2% of average annual sales.*[5]

2. *The sector with the largest share in government spending (NAICS 33 — manufacturing) receives 31% of government obligations, but accounts for only 6% of value added. The sector with the largest share in private consumption (NAICS 53 — real estate, rental, and leasing) accounts for 13% of value added, but less than 1% of government obligations.*

We illustrate this fact in Figure 1. The vertical axis measures the share of a (six-digit) sector k, in total government spending, $\frac{G_k}{G}$. The horizontal axis measures the share of the same sector in private consumption $\frac{C_k}{C}$. In the figure we also distinguish between total federal spending (blue circles), defense spending (red asterisks) and non-defense spending (green triangles).

For overall federal spending and each of its subcomponents, the public and private sectoral spending shares differ substantially, that is, $\frac{G_k}{G} \neq \frac{C_k}{C}$. Some sectors that are big suppliers to the government are almost negligible for private consumers. Sector 541300—Architectural, Engineering, and Related Services, for example, accounts for 15% of government consumption but less than one percent of private consumption. The converse is also true.

In a similar vein, Appendix Table A.1 shows for 2017 as an example that the bias in government spending runs both ways.[6] Manufacturing (NAICS 33), for example, accounted for over 30% of government consumption in 2017, but under 7% of value added. Conversely, in the same year, real estate, rental, and leasing (NAICS 53) accounted for 13% of value added, but less than 1% of government consumption. Finally, this feature holds at the firm level as well—Table A.2 compares the top 35 firms in terms of average annual contract obligations to the top 35 non-oil firms from Compustat in terms of average annual sales between 2001 and 2018. Little overlap exists in the firm lists, with only a few firms like Boeing and General Electric showing up in both lists, albeit in very different orders. Taken together, the evidence indicates

[5]Based on sales of all firms in Compustat.

[6]Value added shares come from the National Income and Product Accounts.

the government spending varies across sectors and its composition does not mimic that of private consumption.

3.3 Short Duration

We now turn to the variation of spending over time. The third fact we establish is that government contracts tend to be relatively short lived. Moreover, they are frequently modified.

Fact 3 *Government spending is characterized by short contract durations:*

1. *The median contract has a duration of 36 days.*
2. *80% of contracts last less than 1 year.*
3. *The firm with the median value of average annual obligations is in the dataset for only 1 year.*

To arrive at the first two results, we study the difference between transactions and the overarching contract structure. Each "contract" can consist of a bundle of transactions—the observation level of the data. Some simple purchases may be made with a single transaction, while others may have hundreds of transactions over the life of the contract that continuously modify the order or relationship.[7] Each transaction is associated with an action date—the date when the reported action was issued, a period of performance start date—the date that the transaction begins, and a period of performance end date—the current date when the transaction ends (barring subsequent modifications). We calculate the "duration" of a transaction as the difference between the period of performance start date and the current end date. Similarly, we calculate the duration of a contract as the difference between the period of performance start date of the earliest underlying transaction and the current end date of the latest underlying transaction.

Durations of transactions and contracts can range from 0 days—this might be a transaction that makes an administrative change, closes out an order, or represents a one time purchase of a commercially manufactured good—to over a decade—a contract funding research and development, for example. The length of the transaction depends entirely on the nature of the relationship and the provided product or service. Overall, however, contracts tend to have short lifespans. Over the entire sample, the median contract has a duration of only 36 days.[8]

Figure 8 shows in each year, about 80% of contracts have durations of less than one year. The figure and Appendix Table A.3 also show the distribution at the transaction level is almost

[7] About 80 percent of "contracts" are made up of a single transaction.

[8] Note, for this analysis we keep transactions with durations between 1 and 5500 days (15 years). These contracts represent more than 95% of the total value of obligations.

identical.[9] The similar statistics at the transaction and contract level implies that only very large contracts include multiple transactions.

In addition to being relatively short, contracts are frequently modified. An observation in the data will have a "modification number" if it represents a transaction that makes a change to an initial award. 20 different types of modifications exist, some of which reflect no change to the initial value of the contract, like a change of address, and some of which reflect either additional obligations or deobligations, like an order for additional work. Figure 9 shows the time series of spending summing only modification spending, as well as the series of spending summing the disjoint non-modification spending. Spending through modification transactions is substantial, and is in fact higher than spending from initial (non-modification) transactions.

Occasionally, modifications are used to correct data entry errors. In these cases, we see (sometimes large) obligations that are almost immediately followed by a de-obligation of similar magnitude, under the same award identification number and directed to the same recipient. For example, Figure A.6 in the Online Appendix shows the individual transactions that made up a contract given by the Department of the Army to Emerson Construction Company for the construction of the Army Reserve Center in Fort Worth, Texas. Line 1 shows an obligation of $13,917,176,427 was made on September 29, 2008. Line 2, however, shows that on January 7, 2009, most of this was offset by a $13,901,924,427 de-obligation. The description of the de-obligation transaction says the modification was made to "correct subclins", or sub-contract line item numbers. In other words, it sounds like an administrative error was made. Netting these two roughly $13 billion transactions, and combining the sum with the rest of the transactions associated with that contract, it appears that a total of $16.3 million was ultimately obligated to Emerson Construction Company. The Emerson Construction Company website advertises that they completed the construction of the Army Reserve Center for exactly this amount.[10] Over the entire sample period, there are about 1 million observations (or less than 2% of observations) that are part of these "offsetting transaction pairs."

Our general approach to deal with such errors is the following: in cases in which two potentially offsetting contracts are within 0.5% of each other, we combine the two transactions into one, and apply the net amount to the date of the earlier of the two offsetting transactions.[11]

[9] Though it may be surprising that contracts appear to be shorter than transactions based on the summary statistics, a simple example can explain the discrepancy. Consider a contract that is made up of 3 individual transactions. The initial transaction begins on January 1 and lasts for a period of 1 year, ending on December 31. A modification is made to the contract on April 1, in a transaction that still has an end-date of December 31 (a duration of 274 days). A final modification is made on September 1, in a transaction that still has an end-date of December 31 (a duration of 121 days). The duration of this contract is 364 days. However, when we look at summary statistics of the transactions, we include transactions with durations of 274 days and 121 days which drives up the moments of the transaction-level distribution relative to the contract-level distribution.

[10] http://www.eccinc.com/projects/army-reserve-center-fort-worth

[11] Applying the net amount to the earlier or later action date makes no difference.

The short persistence of firms in the dataset further highlights the short durations of government spending. We illustrate these results in Figure A.7 in the Online Appendix. The figure shows the fraction of firms in the data for a certain number of years. Looking at the entire dataset, one can clearly see that most firms are in the data for periods of time. In fact, the median number of years a firm is in the dataset is 2 years while the firm with the median value of average annual obligations is in the dataset for only 1 year. Among the large firms, such as the top 0.1% of firms, firms tend to be in the data throughout the sample. A handful of such firms exists, and very few of their contracts last very long. They are mostly related to facilities management and investment around the government. They span information technology, professional, scientific, and technical services, administrative and support and waste management and remediation services, as well as manufacturing. The appendix provides details on the identity of these firms and sectors.

3.4 Idiosyncratic Shocks Drive Aggregate Variation over Time

The fourth fact we establish is idiosyncratic (rather than aggregate) shocks drive the variation in spending over time. At the aggregate level, we show that granularity of firms and sectors is an important origin for the growth rate of aggregate government spending, consistent with our previous fact on granularity: A few firms or sectors drive the dynamics of aggregate government spending.

Fact 4 *Idiosyncratic shocks drive aggregate variation over time:*

1. *The "granular residual" explains more than 50% of aggregate government spending growth.*

3.4.1 Granular Origin of Government Spending Fluctuations

We use the notion of granularity to show idiosyncratic shocks matter if we want to account for the growth of aggregate government spending, instead of getting washed out in the aggregate. We follow Gabaix (2011) and Foerster et al. (2011) to establish this fact.

Granular Residual Approach First, as in Gabaix (2011), we calculate the "granular residual", Γ_t, to show shocks to the top suppliers of government consumption drive the fluctuations in aggregate government spending. To see this, let $g_{i,t}$ be the total obligations to recipient firm i in year t. Then, the growth rate of obligations is given by:

$$z_{i,t} = \ln(g_{i,t}) - \ln(g_{i,t-1})$$

The granular residual is then given by:

$$\Gamma_t = \sum_{i=1}^{K} \frac{g_{i,t-1}}{G_{t-1}} (z_{i,t} - \bar{z}_t) \tag{1}$$

where G_t is aggregate government consumption in year t, and $\bar{z}_t = Q^{-1} \sum_{i=1}^{Q} z_{i,t}$ is the average growth rate over the top Q firms. In other words, the granular residual is the weighted difference in growth rates for the top K firms relative to the average growth rate for the top Q firms, where $Q \geq K$.

As in Gabaix (2011), we run a regression of aggregate growth—$Z_t = \ln(G_t) - \ln(G_{t-1})$—on the granular residual and its lags. The granular hypothesis suggest idiosyncratic shocks, captured by the granular residual, account for a large part of the aggregate movement of government spending. Specifically, we estimate:

$$Z_t = \beta_0 + \beta_1 \Gamma_t + \beta_2 \Gamma_{t-1} + \beta_3 \Gamma_{t-2}$$

We estimate this specification for $K = 100$, $Q = 100$ and $Q = 1000$ firms, and on one and two lags of the granular residual term. We see in Table 1 the granular residual explains about 50% of the variation in aggregate government consumption across specifications. These results are in line with the estimates of Gabaix (2011) for the explanatory power of the granular residual for the top firms on GDP growth.

Decomposition of Government Consumption Growth Second, as in Foerster et al. (2011), we perform a different set of exercises to decompose changes in aggregate government spending growth into components arising from aggregate and idiosyncratic (sector-specific) shocks. This second approach delivers results that are consistent with the results we find using the granular residual approach of Gabaix (2011). Using the methodology of Foerster et al. (2011), we decompose aggregate government consumption growth, Z_t, as follows:

$$Z_t = \underbrace{\sum_{i=1}^{N} \omega_{i,t} z_{i,t}}_{\text{(1) Actual}} = \underbrace{\frac{1}{N} \sum_{i=1}^{N} z_{i,t}}_{\text{(2) Equal Weights}} + \underbrace{\sum_{i=1}^{N} \left(\bar{\omega}_i - \frac{1}{N} \right) z_{i,t}}_{\text{(3) Granular Residual}} + \underbrace{\sum_{i=1}^{N} (\omega_{i,t} - \bar{\omega}_i) z_{i,t}}_{\text{(4) Share Deviation}} \tag{2}$$

where i denotes firms or sectors. The term $(1/N) \sum_{i=1}^{N} z_{it}$ weights each sector equally. If z_{it} are uncorrelated, this component has a variance proportional to N^{-1}. The second term, the "granular residual term," $\sum_{i=1}^{N} [\omega_{it} - (1/N)]_{it}$ will be large if the cross-sectional variance of sectoral shares is large at date t.

Figure 10 plots the individual components of equation (2) over time. In Foerster et al.

(2011), the equally weighted component tracks the series for aggregate industrial production growth more closely than the granular residual term. In our case, both series exhibit fluctuations of a similar magnitude to the aggregate growth rate, indicating that both idiosyncratic shocks and covariance across sectors are important drivers of aggregate growth.

Furthermore, we show in Online Appendix Section A.4.2 at the sectoral level that aggregate time fixed effects explain little of sectoral government spending dynamics. Instead, idiosyncratic innovations drive changes in sectoral spending which can have large positive and negative correlations across sectors.

3.5 Government Consumption is Concentrated in Sticky Sectors

The fifth fact documents a new fact about government consumption and pricing frictions: Government consumption tends to be concentrated in "sticky" sectors—that is, sectors in which price changes are relatively less frequent. We document this result in two complementary ways. We use micro data underlying the producer price data from the Bureau of Labor Statistics (BLS) to construct frequencies of price adjustments for the sectors from which the government purchases. An important caveat of this analysis is the assumption that the frequency of price adjustment for private and government consumption are identical. Therefore, we also study the pricing structure of government contracts directly. "Fixed-price" contracts are dominant and reflect the stickiness at the micro level of the individual contracts.

Fact 5 *Government spending is concentrated in sticky sectors*

1. *The monthly frequency of price changes in the top two supplying sectors to the government is 9% while it is 20%, on average, for the remaining sectors.*

2. *80% of all contracts are fixed-price in nature.*

Our main result for this fifth fact is government spending is concentrated in sticky-price sectors. Figure 11 shows the average annual share of government spending in each two-digit sector (x-axis) plotted against the frequency of price changes in those sectors from the BLS. The size of the bubble corresponds to the average sectoral share of annual aggregate spending—a larger bubble means the sector supplies a larger proportion of government consumption. The figure shows the government spends the vast majority of dollars in sectors with low frequencies of price adjustment. The frequency of price changes in the largest 2 sectors is 9% while it is 20% on average.

This finding is consistent with the type of contracts firms use to set their prices. Table **??** summarizes our findings. It shows the distribution by both count and value of pricing types for government contracts. The first two columns show the distribution for all firms, while the last two columns show the distribution of pricing type for the top 10 firms.

By count, the majority of contracts are "firm fixed price" contracts—the pricing type that places all of the risk on the contractor. Fixed price contracts with economic adjustment follow. The total share of contracts that are fixed-price is over 85%. No comparable benchmark for the private sector exists to the best of our knowledge.

By value, a similar picture emerges: a somewhat larger share of contracted funds are cost-reimbursement contracts (cost plus an award fee or cost plus a fixed fee), but fixed-price type contracts still dominate the contracting environment. Larger transactions are relatively more likely to be awarded under a cost-reimbursement contract, while smaller award transactions are relatively more likely to be fixed price. Still, the total share of spending under some form of fixed price agreement amounts to over 62%. This finding justifies using a sticky-price setting to model the effect of government spending.

4 A New Keynesian Model with Sectoral Government Spending

We now develop a two-sector New Keynesian model to assess the relevance of the five facts we document. The model is deliberately stylized departing as little as possible from the one-sector textbook model. Sectors potentially differ along three dimensions: First, the shares of private and public spending and hence, their size; second, the degree of price rigidity; third, the incidence of shocks. Rather than postulating a process for "big G", we model government spending in each sector as distinct variables. In what follows we outline the setup in general terms and derive a number of theoretical results. We then calibrate the model to capture the five facts we established above and study the quantitative importance.

4.1 Setup

We focus on the key equations of the model because it is a simple extension of the textbook version of the New Keynesian model (Woodford, 2003; Galí, 2015). A representative household chooses consumption and labor effort in order to solve an infinite horizon problem subject to a budget constraint and the labor endowment which we normalize to unity:

$$\max_{\{C_{1t},C_{2t},L_{1t},L_{2t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\ln \left(\frac{C_{1,t}^{\omega} C_{2,t}^{1-\omega}}{\omega^{\omega}(1-\omega)^{1-\omega}} \right) - \xi_1 \frac{L_{1t}^{1+\varphi}}{1+\varphi} - \xi_2 \frac{L_{2t}^{1+\varphi}}{1+\varphi} + f(G_{1t}, G_{2t}) \right),$$

subject to

$$\begin{aligned} W_{1t}L_{1t} + W_{2t}L_{2t} + \Pi_t + I_{t-1}B_{t-1} &= B_t + P_{1t}C_{1t} + P_{2t}C_{2t} + P_{1t}G_{1t} + P_{2t}G_{2t} \\ L_{1t} + L_{2t} &\leq 1. \end{aligned}$$

Here, C_{kt} and G_{kt} denote private and public consumption of sector-k goods, with $k = \{1, 2\}$, respectively. G_{kt} is determined exogeneously. Lump-sum taxes finance government consumption for which we substitute in the household budget constraint. Government spending provides utility, but independently of private consumption and leisure. P_{kt} is the price index in sector k. L_{kt} and W_{kt} are labor employed and wages paid in sector k. Our specification assumes sectoral segmentation of labor markets. Below, we set parameters ξ_k to ensure a symmetric steady state across all firms. Households own firms and receive net income, Π_t, as dividends. Bonds, B_{t-1}, pay a nominal gross interest rate of I_{t-1} and we rule out Ponzi schemes.

The optimal allocation of consumption expenditures across sectors requires:

$$C_{1t} = \omega \left(\frac{P_{1t}}{P_{Ct}} \right)^{-1} C_t \text{ and } C_{2t} = (1 - \omega) \left(\frac{P_{2t}}{P_{Ct}} \right)^{-1} C_t, \tag{3}$$

where $P_{Ct} = P_{1t}^{\omega} P_{2t}^{1-\omega}$ is the consumer price index.

The household first-order conditions determine labor supply and define the Euler equation:

$$\frac{W_{kt}}{P_{Ct}} = \xi_k L_{kt}^{\varphi} C_t \text{ for } k = \{1, 2\}, \tag{4}$$

$$1 = \mathbb{E}_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1} I_t \frac{P_{Ct}}{P_{Ct+1}} \right]. \tag{5}$$

Total demand for sectoral output is:

$$Y_{kt} = C_{kt} + G_{kt}. \tag{6}$$

Sectoral output, in turn, is defined as a CES aggregate of differentiated goods indexed by $j \in [0, n]$ in sector 1 and $j \in (n, 1]$ for sector 2:

$$Y_{1t} \equiv \left[n^{-1/\theta} \int_0^n Y_{j1t}^{1-\frac{1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}, \quad Y_{2t} \equiv \left[(1-n)^{-1/\theta} \int_n^1 Y_{j2t}^{1-\frac{1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}. \tag{7}$$

Cost minimization implies the demand for differentiated goods:

$$Y_{j1t} = n \left(\frac{P_{j1t}}{P_{1t}} \right)^{-\theta} Y_{1t}, \quad Y_{j2t} = (1-n) \left(\frac{P_{j2t}}{P_{2t}} \right)^{-\theta} Y_{2t} \tag{8}$$

and defines the sectoral price indices:

$$P_{1t} = \left[\frac{1}{n}\int_0^n P_{j1t}^{1-\theta} dj\right]^{\frac{1}{1-\theta}}, \quad P_{2t} = \left[\frac{1}{1-n}\int_n^1 P_{j1t}^{1-\theta} dj\right]^{\frac{1}{1-\theta}}. \tag{9}$$

Differentiated goods are produced according to: $Y_{jkt} = L_{jkt}$. Firms are constrained in their ability to set prices. With probability α_k, which may differ across sectors, a firm may not adjust its price in the next period. The pricing problem of firm j in sector k is:

$$\max_{P_{jkt}} \mathbb{E}_t \sum_{s=0}^{\infty} Q_{t,t+s}\alpha_k^s \left[P_{jkt}Y_{jkt+s} - \mathcal{C}_{t+s}(Y_{t+s|t})\right].$$

Here $Q_{t,t+s}$ is the stochastic discount factor between periods t and $t+s$ and $\mathcal{C}_{t+k}(\cdot)$ are costs of production. The first order condition is:

$$\sum_{\tau=0}^{\infty} Q_{t,t+\tau}\alpha_k^{\tau} Y_{jkt+\tau}\left[P_{kt}^* - \mathcal{M}\Psi_{kt+\tau}\right] = 0, \tag{10}$$

where $Y_{jkt+\tau}$ is the total production of firm jk in period $t+\tau$, $\mathcal{M} \equiv \frac{\theta}{\theta-1}$ denotes the desired markup and $\Psi_{t+k} = \mathcal{C}'_{t+k}(Y_{t+k})$ are marginal costs. The optimal price, P_{kt}^*, is the same for all firms in a given sector. Thus, aggregating all prices within a sector yields:

$$P_{kt} = \left[(1-\alpha_k)\, P_{kt}^{*1-\theta} + \alpha_k P_{kt-1}^{1-\theta}\right]^{\frac{1}{1-\theta}}. \tag{11}$$

We define (nominal) GDP as follows:

$$P_{Yt}Y_t \equiv P_{1t}Y_{1t} + P_{2t}Y_{2t}, \tag{12}$$

where $P_{Yt} \equiv P_1^n P_2^{1-n}$ is the GDP deflator.

Analogously, we define aggregate government spending ("Big G") as:

$$P_{Gt}G_t \equiv P_{1t}G_{1t} + P_{2t}G_{2t}. \tag{13}$$

Assuming the average weight of sector 1 in total government spending is γ, we define the deflator for government spending as $P_{Gt} \equiv P_1^{\gamma} P_2^{1-\gamma}$.

Lastly, we close the model by specifying an inflation target (of zero):

$$\Pi_{Y_t} = \frac{P_{Y_t}}{P_{Y_{t-1}}} = 0. \tag{14}$$

In the spirit of Svensson (2003), we assume monetary policy adjusts short term interest rates to meet the inflation target at all times.

4.2 Approximate Equilibrium Conditions

We now consider an approximation of the equilibrium conditions around a symmetric steady state in which relative prices are unity and inflation is zero (see Appendix A.1.1 for details). Let γ denote the fraction of government spending in sector 1 in the steady state: $\gamma = G_1/G$. ζ, in turn, is the steady-state ratio of consumption to output: $\zeta = C/Y$, such that $1-\zeta = G/Y$.

The steady-state sizes of sector 1 and 2 are then given by the weighted average of each sector's share in private and public spending, respectively:

$$n = \omega\zeta + \gamma(1-\zeta) \tag{15}$$

$$1-n = (1-\omega)\zeta + (1-\gamma)(1-\zeta). \tag{16}$$

We state the equilibrium conditions in terms of deviations from steady state with lowercase letters denoting percentage deviations from steady state. Market clearing in each sector implies:

$$ny_{1,t} = -\omega\zeta(1-\omega)\tau_t + \omega\zeta c_t + (1-\zeta)\gamma g_{1,t} \tag{17}$$

$$(1-n)y_{2,t} = (1-\omega)\zeta\omega\tau_t + (1-\omega)\zeta c_t + (1-\zeta)(1-\gamma)g_{2,t}. \tag{18}$$

where $\tau_t = p_{1,t} - p_{2,t}$ are the terms of trade. In deriving the expressions we use $p_{1,t} - p_t = (1-\omega)\tau_t$.

The New Keynesian Phillips curves in each sector are given by:

$$\alpha_1\pi_{1,t} = \alpha_1\beta E_t\pi_{1,t+1} + (1-\alpha_1)(1-\beta\alpha_1)\psi_{1t} \tag{19}$$

$$\alpha_2\pi_{2,t} = \alpha_2\beta E_t\pi_{2,t+1} + (1-\alpha_2)(1-\beta\alpha_2)\psi_{2t}, \tag{20}$$

where *marginal costs*, ψ_{kt}, are in real terms (deflated with the producer price in each sector). After substituting for the real wage we have:

$$\psi_{1t} = c_t + \varphi y_{1,t} - (1-\omega)\tau_t \tag{21}$$

$$\psi_{2t} = c_t + \varphi y_{2,t} + \omega\tau_t. \tag{22}$$

An approximation of the Euler equation yields:

$$c_t = E_t c_{t+1} - (i_t - E_t\pi_{c_{t+1}}) \tag{23}$$

$$\pi_{c_t} = \omega\pi_{1,t} + (1-\omega)\pi_{2,t}, \tag{24}$$

where the second equation is consumer price inflation. Equations (12) and (13) and the definition of the deflators for GDP and government spending imply the following equations for real GDP and real aggregate government spending:

$$y_t = ny_{1,t} + (1-n)y_{2,t} \tag{25}$$

$$g_t = \gamma g_{1,t} + (1-\gamma)g_{2,t}. \tag{26}$$

Regarding monetary policy, the inflation target (14) requires:

$$\pi_{y_t} = 0. \tag{27}$$

For government spending we assume an exogenous AR(1) process for both sectors:

$$g_{1,t} = \rho_1 g_{1,t-1} + \varepsilon_{1,t} \tag{28}$$

$$g_{2,t} = \rho_2 g_{2,t-1} + \varepsilon_{2,t}, \tag{29}$$

where $\varepsilon_{k,t}$ are sector specific spending shocks and parameters $\rho_k \in [0,1)$ capture the persistence of the spending processes.

4.3 Results

In this section, we derive a number of closed-form results. We focus on the effect of an exogenous variation in government spending, that is, we state the solution in terms of g_{1t} and g_{2t}. The goal is to illustrate how idiosyncratic variation in government spending at the sectoral level impacts the aggregate economy. To facilitate the algebra we assume $\alpha_1 = 0$, that is, prices are fully flexible in sector 1. We relax this assumption when we present numerical results in Section 4.4 below. We do not restrict the extent of price rigidity in sector 2. Instead we have $\alpha_2 \in [0,1]$.

To solve the model, we first derive the solution for the terms of trade that are the only endogenous state variable in the model. Intuitively, since prices are (potentially) sticky in sector 2, the adjustment to even purely transitory shocks takes time and dynamics of the terms of trade govern the adjustment process. Inflation targeting implies $\pi_{y_t} = n\pi_{1,t} + (1-n)\pi_{2,t} = 0$. This equations allows us to rewrite the Phillips curve in sector 2 (equation (20)) as:

$$n(\tau_t - \tau_{t-1}) = n\beta(E_t(\tau_{t+1} - \tau_t)) - \kappa_2 \psi_{2t}. \tag{30}$$

Marginal costs in sector 2 drive the dynamics of the terms of trade, which are:

$$\psi_{2t} = \left(1+\frac{\zeta\varphi(1-\omega)}{1-n}\right)c_t + \left(1+\frac{\zeta\varphi(1-\omega)}{1-n}\right)\omega\tau_t + (1-\zeta)(1-\gamma)\frac{\varphi}{1-n}g_{2,t}. \quad (31)$$

We use the market clearing condition in equation (18) to substitute for output in equation (22). To substitute for consumption, we exploit the fact that firms in sector 1 are fully flexible in setting their prices and hence, charge a constant markup over marginal costs. As a result, marginal costs are constant in real terms and we obtain the following expression for consumption:

$$c_t = (1-\omega)\tau_t - \left(1+\frac{\zeta\varphi\omega}{n}\right)^{-1}(1-\zeta)\gamma\frac{\varphi}{n}g_{1t}. \quad (32)$$

Intuitively, this expression captures the dynamics of the labor market. Higher government spending induces upward pressure on wages as production and the demand for labor rise. For real marginal costs to remain constant in equilibrium, labor supply must also increase. An increase in the marginal utility of wealth, or equivalently, a drop in consumption delivers the increase in labor supply. This effect accounts for the negative impact of government spending on consumption in expression (32) for given terms of trade.

Using equation (32) in equation (31) and substituting in equation (30) we obtain a second-order difference equation in the terms of trade:

$$\{(1+\beta)+\kappa A_2\}\tau_t - \tau_{t-1} - \beta E_t\tau_{t+1} = \kappa\frac{A_2}{A_1}\frac{\varphi}{n}(1-\zeta)\gamma g_{1t} - \frac{\kappa\varphi}{1-n}(1-\zeta)(1-\gamma)g_{2,t}, \quad (33)$$

where $\kappa \equiv \kappa_2/n$, $A_2 = 1+\frac{\zeta\varphi(1-\omega)}{1-n}$ and $A_1 = 1+\frac{\zeta\varphi\omega}{n}$. The A coefficients increase with the weight of a sectors' share in private consumption as well as with its size. We solve equation (33) to obtain a solution for the terms of trade in government spending. The following proposition summarizes our first result.

Proposition 1 (Solution for terms of trade) *Assuming prices in sector 1 are fully flexible ($\alpha_1 = 0$) and monetary policy targets producer price inflation ($\pi_{y_t} = 0$), the solution for the terms of trade is given by:*

$$\tau_t = \Lambda_0\tau_{t-1} + \Lambda_1(1-\zeta)\gamma g_{1,t} - \Lambda_2(1-\zeta)(1-\gamma)g_{2,t}, \quad (34)$$

where $\Lambda_0 \in (0,1)$ and $\Lambda_1, \Lambda_2 \geq 0$.

Proof. See Appendix A.1.3 ■

The intuition for this case is straightforward: government spending in sector 1 increases the prices of sector 1, thereby raising the terms of trade, while spending in sector 2 reduces the terms of trade.

We now substitute in expression (32) for the terms of trade using equation (34) and obtain our second result. Government spending crowds out private consumption—independently of the sector in which spending occurs.

Proposition 2 (Crowding out of consumption) *Assuming prices in sector 1 are fully flexible ($\alpha_1 = 0$) and monetary policy targets producer price inflation ($\pi_{y_t} = 0$),*
(1) *the solution for consumption is given by*

$$c_t = \Theta_0 \tau_{t-1} - \Theta_1 (1-\zeta)\gamma g_{1,t} - \Theta_2 (1-\zeta)(1-\gamma) g_{2,t} \tag{35}$$

where $\Theta_0 \in (0,1)$;
(2) *$\Theta_1 \in [0,\infty)$, and $\Theta_2 \in [0,\zeta^{-1})$, that is, government spending in either sector crowds out private consumption. The limiting case $\Theta_1 \to \infty$ occurs if $n \to 0$, while $\Theta_2 \to \zeta^{-1}$ obtains if $1-n \to 0$;*
(3) *if $\omega \geq \gamma$, then $\Theta_1 > \Theta_2$, that is, crowding out is stronger in response to sector 1 spending. Also, if $\kappa \to 0$, $\Theta_1 > 0$ and $\Theta_2 = 0$.*

Proof. See Appendix A.1.3 ■

Expression (35) shows, all else equal, higher terms of trade imply higher consumption ($\Theta_0 > 0$). Intuitively, since the terms of trade reduce marginal costs in the flex-price sector, constant marginal costs require consumption to go up in order to put upward pressure on the real wage.

Next, we observe from expression (35) government spending tends to crowd out private consumption, since both Θ_1 and Θ_2 are non-negative. To understand this result, note an increase in government spending in either of the two sectors raises production and employment as well as marginal costs in the sector.[12] As a result, upward pressure on inflation occurs, which induces monetary policy to raise interest rates and, in turn, induces households to reduce their current consumption in both sectors. Put differently, a shock in one sector spills over to the other sector because monetary policy can only manage aggregate demand rather than demand in a specific sector.

This result clarifies why $\Theta_2 \to 0$ as $\kappa \to 0$. $\kappa \to 0$ implies prices are completely sticky in the limit. Hence, government spending in sector 2 does not generate inflationary pressure. Monetary policy remains unresponsive, and private consumption is invariant to the fiscal impulse. For the

[12] Utility that is linear in labor ($\varphi = 0$) is an exception. In this case, marginal costs are independent of the level of production and $\Theta_1 = \Theta_2 = 0$.

same reason, crowding out is stronger in the flex-price sector ($\Theta_1 > \Theta_2$) provided $\omega \geq \gamma$, that is, whenever private consumption is relatively concentrated in the flex-price sector which is the empirically relevant case (facts #2 and #5). This condition also holds in the absence of sectoral bias, that is, when $\omega = \gamma$. Intuitively, consumption drops more in response to an increase of government spending in sector 1, because the higher flexibility in prices results in stronger inflationary pressure. Hence, the monetary authority has to raise interest rates by more to keep inflation in check.

Theoretically, crowding out of consumption in response to sector 1 spending can be arbitrarily large. Specifically, we find $\Theta_1 \to \infty$ if $n \to 0$ (which also implies that $\omega \to 0$). Assuming the weight of sector 1 approaches zero, private consumption is concentrated in the sticky sector. Given the inflationary impact of government spending in the flex-price sector, the reduction in consumption necessary to offset the impact on inflation becomes arbitrarily large because inflation is relatively inelastic to changes in sticky-sector consumption, both public and private. Instead, when $1 - n \to 0$, the crowding out of consumption in response to sector 2 spending, captured by Θ_2, does not exceed ζ^{-1}. At this point, the drop in consumption matches the increase of government spending such that marginal costs, and hence inflation, remains constant. This happens for a relatively modest reduction of consumption because inflation is very elastic to changes in both public and private spending in the flexible sector.

Finally, we establish the effect of government spending on output.

Proposition 3 (Output multipliers) *Assuming prices in sector 1 are fully flexible ($\alpha_1 = 0$) and monetary policy targets producer price inflation ($\pi_{y_t} = 0$), the solution for output is given by*

$$y_t = \Gamma_0 \tau_{t-1} + \Gamma_1 (1-\zeta)\gamma g_{1,t} + \Gamma_2 (1-\zeta)(1-\gamma) g_{2,t} \tag{36}$$

where $\Gamma_0 \in (0,1)$, *and*

$$\Gamma_1 = 1 - \zeta\Theta_1 \ \textit{and}\ \Gamma_2 = 1 - \zeta\Theta_2. \tag{37}$$

Moreover, we find $\Gamma_0 \in (0,1)$, Γ_1 *has full support in* $(-\infty, 1]$, *and* Γ_2 *has full support in* $(0,1]$.

In expression (36), the effect of the lagged terms of trade on output is positive ($\Gamma_0 > 0$) because, as discussed above, their effect on consumption is also positive (see Proposition 2). The coefficients Γ_1 and Γ_2 directly capture the impact multiplier of government spending on output, that is, the change in output divided by the change in government spending.[13] Also, equation (37) shows the sum of the direct effect of higher spending on output and the indirect effect on private consumption, which is negative, determine the overall multiplier.

[13]While $g_{k,t}$ measures the percentage deviation of government spending from its steady-state level, multiplying with $(1-\zeta)\gamma$ and $(1-\zeta)(1-\gamma)$, in turn, transforms this into units of steady-state output.

Given our results regarding Θ_1 and Θ_2, stated in Proposition 2, it follows immediately Γ_1 may actually be negative, while Γ_2 is bounded from below by zero, just as in the one-sector New Keynesian model. Moreover, we also stress the multiplier may not exceed unity, again, just like in the one-sector model unless the zero lower bound on interest rates binds (Woodford, 2011).

4.4 The Quantitative Relevance of Granular Government

In this section, we explore to what extent the departure from the one-sector model matters quantitatively for the fiscal transmission mechanism. We calibrate the model to capture the five facts in a stylized way. Before doing so, we fix three parameters that are independent of the five facts. Specifically, we assume that a period in the model corresponds to one month and set $\beta = 0.997$. Next, we set $\varphi = 4$, in line with estimates for the Frisch elasticity of labor supply (Chetty et al., 2011). Lastly, we assume government spending accounts for 20 percent of GDP and set $\zeta = 0.8$.

We proceed as follows. First, we account for granularity by assuming that government spending is concentrated in sector 2. To determine the steady-state weight of sector 1 in the government's consumption basket, we turn to Table 2 which reports the relevance of sectors for government consumption, starting with the most important sector. We assume sector 2 in the model represents the three most important suppliers to the government: they account for approximately 69 percent of government consumption. Hence, we set the weight of sector 1 in government spending to $\gamma = 0.31$. Next, Table 2 also shows these sectors account for only about 16 percent of value added. Hence, we set the size of sector 1 to $n = 0.84$. Given these parameter values, restriction (15) implies $\omega = 0.9725$. As a result, we account for fact #2 (sectoral bias): private spending is concentrated in sector 1, while public spending is concentrated in sector 2.

We capture the third fact ("Short durations") as follows. To pin down the shock process for government spending we estimate AR(1) processes (28) and (29) at the sectoral level and report results again in Table 2. We find a value for ρ of approximately 0.3 if we look at the three most important sectors (captured by sector 2 in the model) and thus set $\rho_2 = 0.3$. The value for the other sectors (reported in column "average rho out") is also about 0.3, and hence we fix ρ_1 accordingly. Fact #4 documents idiosyncratic shocks drive the aggregate variation in government spending which we account for by modelling distinct shocks in the two sectors rather than a shock to Big G itself.

Lastly, we account for fact #5: the government mainly spends in sectors with a low frequency of price adjustment. We set the $\alpha_1 = 0.78$ and $\alpha_2 = 0.9$, in line with the evidence in Table 2. These parameters imply an average duration of price spells of about 4.5 months in sector 1 and of 10 months in sector 2. In other words, the average duration of prices is more than twice as high in sector 2.

We display the impulse responses of selected variables to government spending shocks in Figure 12. In the figure, we measure the percentage deviation from steady state due to the shock along the vertical axis. The horizontal axis measures time in months. The solid lines represents the scenario in which government spending increases in sector 1, while the dashed lines refers to the responses to a spending shock in sector 2. In both cases, we normalize the size of the shock such that aggregate government spending ("Big G"), upper-left panel, increases by one percent of GDP. Since government spending is exogenous and the persistence parameter ρ_k are identical across sectors, the dynamics of G are the same for both shock scenarios. The output responses, upper-right panel, differ considerably. Output increases only mildly in response to a shock in sector 1. Instead, the impact response is more than 4 times as large in case the shock originates in sector 2. We measure output and government spending in the same units and hence the output response provides a direct measure of the (impact) multiplier.

The crowding out of private expenditure determines the size of the multiplier. We show the response of consumption in the lower-left panel of the figure. Strong crowding out occurs in response to the spending shock in sector 1. Instead, crowding out is substantially smaller if the spending increase takes place in sector 2. The necessary monetary policy response to keep inflation in check rationalizes the difference in the crowding out across cases. In the figure we report the response of the interest rate in terms of annualized percentage points. Intuitively, the interest rate responds strongly because the shock is very short-lived. But the response of the interest rate is particularly strong when the shock originates in sector 1. The higher degree of price flexibility results in a larger inflationary response which requires a larger monetary response to stabilize inflation.

The degree of price stickiness shapes the extent of crowding out, and hence the multiplier. We contrast the results for our baseline scenario to a counterfactual in which we assume homogeneous pricing frictions across sectors to investigate this issue in more detail. Specifically, we set $\alpha_1 = \alpha_2 = 0.9$, that is, we increase the overall amount of price rigidity in the economy. We report results in Figure 13. The lines with circles refer to the counterfactual economy with homogeneous pricing fictions and the thin lines reproduce the results for our baseline calibration.

The output response to both shocks is similar—but, perhaps surprisingly, now the response is almost as weak as in response to a sector 1 shock in the baseline calibration. The reduced impact of the fiscal impulse obtains even though price stickiness is now higher than in the baseline. What matters for the transmission of sectoral shocks, however, is precisely in which sector prices are sticky. In fact, the multiplier of a shock arising in sector 2 increases if prices in sector 1 become more flexible.

Higher price stickiness in sector 1, instead, implies monetary policy must generate a larger reduction of private consumption in order to stabilize inflation, given the inflationary impact of higher spending in sector 2. This effect, in turn, means more crowding out and a smaller multiplier. Hence, it is not the degree of overall stickiness that determines the size of the fiscal multiplier but the relative price stickiness in the relevant sectors and the incidence of shocks.

Figure A.11 in the Online Appendix illustrates this mechanism in a more systematic way. We plot the output multiplier on impact in case the shock originates in sector 1 (red solid line) and in case the shock originates in sector 2 (blue dashed line). Throughout, we assume the pricing friction is unchanged in sector 2 ($\alpha_2 = 0.9$), but we vary the Calvo parameter in Sector 1 along the horizonal axis—all the way from zero (left) to one (right). Intuitively, raising the pricing friction in sector 1, also raises the multiplier in response to a sector 1 shock. However, it lowers the multiplier in response to a sector 2 shock. When we raise α_1 above the level of our baseline calibration (0.78), we increase the overall price stickiness in the economy and yet the multiplier in response to a sector 2 shock declines. The mechanism which underlies this result is straightforward. As sector 1 becomes more sticky, monetary policy has to compress consumption by more in order to stabilize inflation.

4.4.1 The Role of Monetary Policy and the Zero Lower Bound

Until now we have maintained the assumption monetary policy follows a strict inflation target. Section A.4.3 of the Online Appendix discusses results for a Taylor rule. Overall, monetary policy under the Taylor rule is more accommodating than under the targeting rule but results are qualitatively similar in both cases.

Nevertheless, monetary policy and, in particular, its interaction with the degree of price stickiness, is key for the fiscal transmission mechanism, especially when the zero lower bound on interest rates binds. In this case, higher government spending does not trigger an increase in interest rates. As a result, private consumption is crowded in and the fiscal multiplier larger is than in normal times—a result well known from one-sector models (Eggertsson, 2011; Christiano et al., 2011).

We now investigate the role of monetary policy and the ZLB in our multi-sector model. We assume a shock to the time-discount factor increases households' desire to save and, as a result, pushes the economy into the ZLB. We then contrast the effect of government spending shocks in sector 1 and 2.[14] We show results in Figure 14 for the case of a Taylor rule (result look very similar in case of inflation targeting). The blue solid lines show the adjustment to a shock originating in Sector 1, whereas the red dashed lines plot the responses to a shock originating in Sector 2. The lower-right panel shows the response of the policy rate: it is not responding

[14]We solve the model while allowing for an occasionally binding ZLB-constraint using the OccBin toolkit of Guerrieri and Iacoviello (2015).

during the first 3 month because the ZLB binds.[15] As a result, private consumption is now increasing on impact in response to both shocks (lower-left panel) and the multiplier slightly exceeds unity just as in the one-sector model.

The novel result is the ranking of the multipliers across sectors flips at the ZLB. The spending shock in sector 1 has now a larger effect on output than a shock in sector 2: the impact multiplier is about 1.2 in case of a sector 1 shock, the sector with more flexible prices, and just slightly above one in case of a sector 2 shock. The multiplier is larger in sector 1, because higher government spending in this sector has a stronger inflationary impact. In normal times, the inflationary pressure would trigger a stronger monetary contraction, which does not take place at the ZLB. The stronger inflationary impulse therefore translates into a stronger drop in the real interest rate and, eventually, a stronger crowding-in of private consumption.

4.4.2 The Role of Heterogeneity

In order to further illustrate to what extent sectoral heterogeneity matters for the fiscal transmission mechanism, we contrast the dynamic adjustment to an aggregate government spending shock across two model specifications. In each case, we raise government spending by the same amount, namely by one percent of steady state output and proportionally to the steady-state levels of spending in each sector. In the first scenario, we mimic a one-sector model as we let sector 1 "take over" the economy: $\gamma = \omega = n = 1$. In this case, we set $\alpha_1 = 0.9$, equal to α_2 in the two-sector baseline economy. In the second scenario, we study the adjustment of the baseline economy to an aggregate spending shock.

Figure 15 shows the results. The solid lines correspond to the case of the one-sector economy, the dashed lines correspond to our baseline economy. The dynamics differ markedly across scenarios. In particular, the output effect of the aggregate shock is stronger in our baseline model and the crowding out is weaker. To understand this result, recall our baseline calibration implies that sector 2 accounts for a larger share of government spending ($1 - \gamma = 0.69$). For this reason, it also accounts for a larger fraction of any additional government spending (assuming, as we do, that spending is raised in equal proportions across sectors). At the same time, sector 2 is the sector in which prices are more sticky. Put differently, an aggregate shock is similar to a shock that originates in sector 2. The output multiplier is still a bit smaller than in case of a genuine sector 2 shock (Figure 12) simply because in our baseline model only 69 (rather than 100) percent of the additional expenditure is spent on sector 2 goods. The multiplier of an aggregate shock in the two-sector model is, however, about 3 times as large as in the one-sector

[15] The fiscal impulse is small enough and does not change the duration for which the ZLB binds even though, in principle, the exit from the ZLB is endogenous and may be quicker if the fiscal stimulus is large (Erceg and Lindé, 2014).

model, even though the overall extent of pricing frictions is lower. Again, what matters is in which sector pricing frictions are large, not the overall extent: the multiplier tends to be large if the additional spending occurs in sectors in which the pricing frictions are high relative to the other sectors, rather than merely in absolute terms.

5 Conclusion

In this paper, we dissect the anatomy of Big G. A systematic analysis of the entire universe of procurement contracts of the US federal government allows us to establish five basic facts regarding the nature of government spending. To summarize, the five facts are:

1. **Government spending is granular;**
2. **Government spending has a sectoral bias;**
3. **Contracts are characterized by short duration;**
4. **Idiosyncratic shocks at the firm/ sectoral level drive aggregate variation; and**
5. **Government spending is concentrated in sectors with relatively sticky prices**.

We believe accounting for these facts is important and will improve our understanding of how fiscal policy works. As a first step in this direction, we calibrate a simple two-sector New Keynesian model that captures the five facts in a stylized fashion.

The fiscal transmission mechanism in the micro-founded two-sector model differs considerably depending on which sector the shock originates in. Importantly, while private expenditure is crowded out independently of where the shock originates, the crowding out can become arbitrarily large if the shock hits the sector in which private expenditure is concentrated and prices are flexible. Crowding out, instead, is limited in the case in which the shock hits the sticky sector. In this case, the output multiplier is also considerably larger, by a factor of four. We leave a more systematic quantitative exploration for future work.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Auclert, A., M. Rognlie, and L. Straub (2018). The intertemporal keynesian cross. mimeo.

Auerbach, A. J. and Y. Gorodnichenko (2012a). Fiscal multipliers in recession and expansion. In *Fiscal policy after the financial crisis*, pp. 63–98. University of Chicago Press.

Auerbach, A. J. and Y. Gorodnichenko (2012b). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy 4*(2), 1–27.

Auerbach, A. J., Y. Gorodnichenko, and D. Murphy (2019, January). Local Fiscal Multipliers and Fiscal Spillovers in the United States. NBER Working Papers 25457, National Bureau of Economic Research, Inc.

Baqaee, D. and E. Farhi (2019). Productivity and misallocation in general equilibrium. Technical report, forthcoming, Quarterly Journal of Economics.

Barsky, R., C. L. House, C. Boehm, and M. S. Kimball (2016). Monetary policy and durable goods.

Barsky, R. B., C. L. House, and M. S. Kimball (2007, June). Sticky-price models and durable goods. *American Economic Review 97*(3), 984–998.

Boehm, C. E. (2019). Government consumption and investment: Does the composition of purchases affect the multiplier? *Journal of Monetary Economics*.

Bouakez, H., O. Rachedi, and E. Santoro (2018). Sectoral heterogeneity, production networks, andthe effects of government spending. mimeo.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–75.

Chodorow-Reich, G. (2019, May). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy 11*(2), 1–34.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*(1), 78 – 121.

Cohen, L., J. Coval, and C. Malloy (2011). Do powerful politicians cause corporate downsizing? *Journal of Political Economy 119*(6), 1015–1060.

Corsetti, G., A. Meier, and G. J. Müller (2012, November). Fiscal Stimulus with Spending Reversals. *The Review of Economics and Statistics 94*(4), 878–895.

Demyanyk, Y., E. Loutskina, and D. Murphy (2019). Fiscal stimulus and consumer debt. *The Review of Economics and Statistics 101*(4), 728–741.

Eggertsson, G. B. (2011, June). What Fiscal Policy is Effective at Zero Interest Rates? In *NBER Macroeconomics Annual 2010, Volume 25*, NBER Chapters, pp. 59–112. National Bureau of Economic Research, Inc.

Erceg, C. and J. Lindé (2014). Is there a fiscal free lunch in a liquidity trap? *Journal of the European Economic Association 12*(1), 73–107.

Farhi, E. and I. Werning (2016). Fiscal multipliers. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, Chapter 31, pp. 2417–2492. Elsevier.

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral vs aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy 119*(1), 1–38.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Galí, J. (2015, March). *Monetary Policy, Inflation, and the Business Cycle: An Introduction* *to the New Keynesian Framework and Its Applications Second edition*. Number 10495 in Economics Books. Princeton University Press.

Galí, J., J. D. López-Salido, and J. Vallés (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association 5*(1), 227–270.

Galí, J. and T. Monacelli (2008, September). Optimal monetary and fiscal policy in a currency union. *Journal of International Economics 76*(1), 116–132.

Guerrieri, L. and M. Iacoviello (2015). Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics 70*, 22 – 38.

Hagedorn, M., I. Manovskii, and K. Mitman (2019). The fiscal multiplier. mimeo.

Hettig, T. and G. J. Müller (2018). Fiscal policy coordination in currency unions at the effective lower bound. *Journal of International Economics 115*(C), 80–98.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Liebman, J. B. and N. Mahoney (2017). Do expiring budgets lead to wasteful year-end spending? evidence from federal procurement. *American Economic Review 107*(11), 3510–49.

McKay, A., E. Nakamura, and J. Steinsson (2016, October). The power of forward guidance revisited. *American Economic Review 106*(10), 3133–58.

McKay, A. and R. Reis (2016). The role of automatic stabilizers in the u.s. business cycle. *Econometrica 84*(1), 141–194.

Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks? *Journal of Applied Econometrics 24*(6), 960–992.

Nakamura, E. and J. Steinsson (2014, March). Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review 104*(3), 753–92.

Ozdagli, A. and M. Weber (2017). Monetary policy through production networks: Evidence from the stock market. Technical report, National Bureau of Economic Research.

Pasten, E., R. Schoenle, and M. Weber (2019a). Price rigidities and the origins of aggregate fluctuations. *Unpublished manuscript, University of Chicago*.

Pasten, E., R. Schoenle, and M. Weber (2019b). Production networks and the propagation of monetary policy shocks. *Journal of Monetary Economics (forthcoming)*.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. NBER Working Papers 21978, National Bureau of Economic Research, Inc.

Ramey, V. A. (2019). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? *Journal of Economic Perspectives 33*(2), 89–114.

Ramey, V. A. and M. D. Shapiro (1998). Costly capital reallocation and the effects of government spending. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 48, pp. 145–194. Elsevier.

Ramey, V. A. and S. Zubairy (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of Political Economy 126*(2), 850–901.

Svensson, L. E. O. (2003, June). What is wrong with taylor rules? using judgment in monetary policy through targeting rules. *Journal of Economic Literature 41*(2), 426–477.

Woodford, M. (2003). *Interest and Prices*. Princeton University Press.

Woodford, M. (2011). Simple analytics of the government expenditure multiplier. *American Economic Journal: Macroeconomics 3*(1), 1–35.

6 Figures and Tables

Figure 2: Comparison of USASpending Data with General Government Expenditures

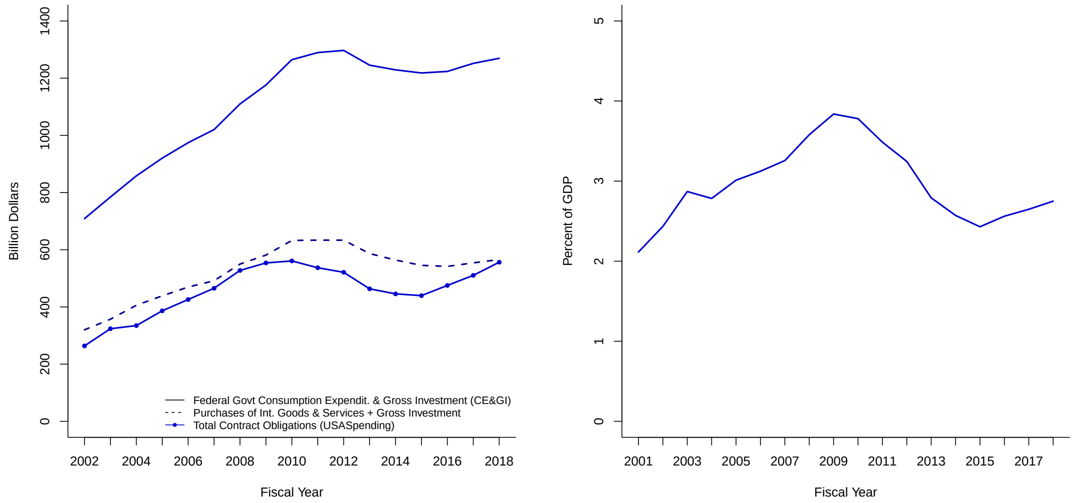

(a) USASpending vs NIPA Accounts

(b) Contracts as a Share of GDP

Note. This figure shows how aggregate contract obligations compare to Government spending as defined in the National Income and Product Accounts (NIPAs). The left panel shows that total contract obligations are roughly equivalent to federal government purchases of intermediate goods and services plus gross investment. The right panel shows that contract obligations account for about 2 to 4 percent of GDP.

Figure 3: Share of Obligations by Top Firms and Sectors

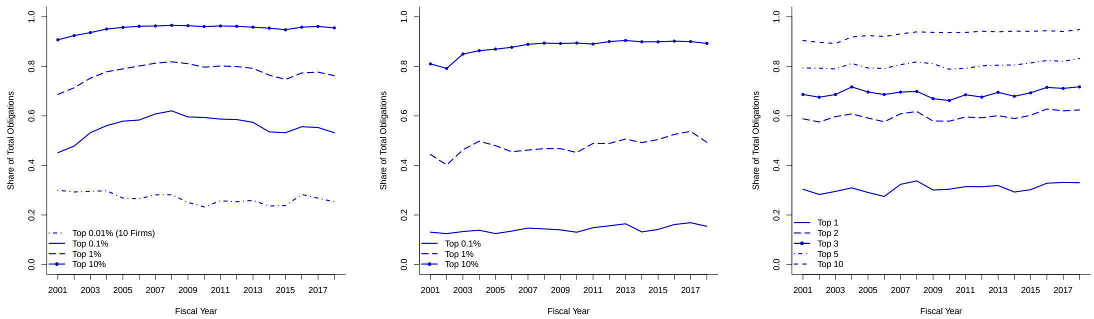

(a) Firms

(b) NAICS 6 Sectors

(c) NAICS 2 Sectors

Note. This figure shows the share of contract obligations awarded to the top shares of firms (the left panel) six-digit NAICS sectors (the middle panel) and two-digit NAICS sectors (the right panel).

Figure 4: Variance Decomposition: Within and Across Firms

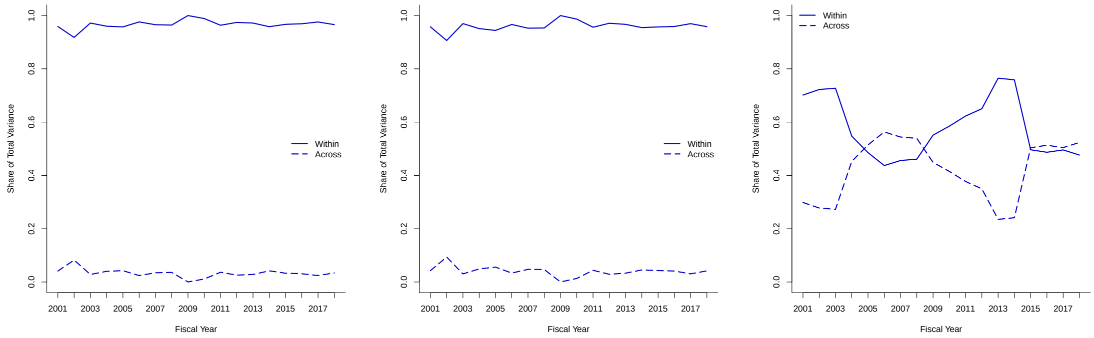

(a) Full Dataset (b) Top 20 Percent (c) Bottom 80 Percent

Note. This figure shows a decomposition of the variance of government spending into "within-firm" and "across-firm" variation. Specifically, total variance is given by:

$$\sum_f \sum_{i \in f} (g_{if,t} - \bar{g}_t)^2 = \underbrace{\sum_f \sum_{i \in f} (g_{if,t} - \bar{g}_{f,t})^2}_{\text{(a) Within Firm}} + \underbrace{\sum_f \sum_{i \in f} (\bar{g}_{f,t} - \bar{g}_t)^2}_{\text{(b) Across Firm}},$$

where i is an individual contract transaction and f is a firm. We plot each of the two RHS components as a share of the LHS. Panel (a) shows this decomposition for the full dataset, panel (b) restricts the sample to the top 20 percent of transactions, and panel (c) shows only the bottom 80 percent of transactions.

Figure 5: Density of Variance Decomposition Components

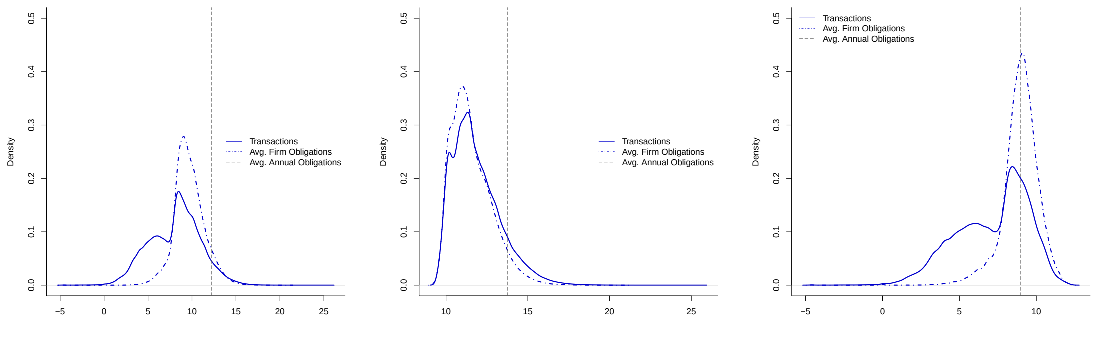

(a) Full Dataset (b) Top 20 Percent (c) Bottom 80 Percent

Note. This figure shows the density of each of the three components that underly the variance decomposition in Figure 4. The solid-blue line shows the density of the individual contract transactions—$g_{if,t}$, the dash-dotted line shows the density of average firm obligations—$\bar{g}_{f,t}$, and the dashed-black horizontal line shows the average annual obligations—$\bar{g}_t$. Panel (a) shows these densities for the full dataset, panel (b) restricts the sample to the top 20 percent of contract transactions, and panel (c) shows only the bottom 80 percent of transactions.

Figure 6: Q-Q Plot: Actual vs. Log-Normal

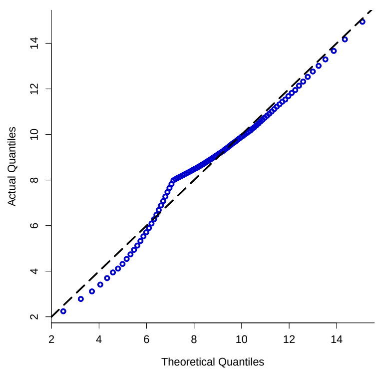

Note. This figure is a Q-Q plot with actual quantiles of log transactions on the y-axis and theoretical quantiles from a log-normal distribution with the same mean and standard deviation plotted on the x-axis.

Figure 7: Histogram of Log Transaction Value

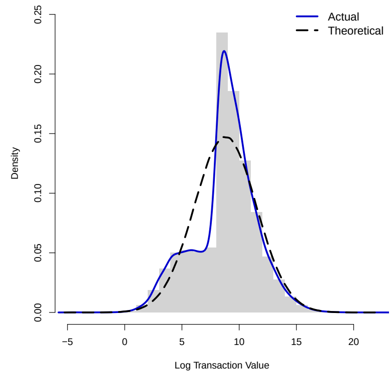

Note. This figure shows a histogram of log transaction obligations and the density of those log obligations. We also plot the density of a theoretical log-normal distribution with the same mean and variance.

Figure 8: Empirical CDF of Contract Durations

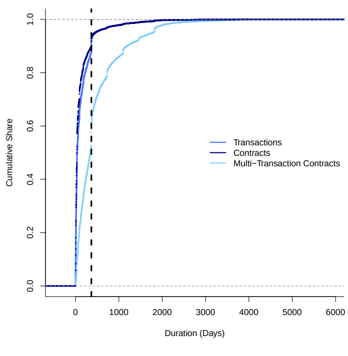

Note. This figure shows the empirical cumulative distribution function of the duration—the number of days between the start- and end-date—of transactions and contracts. The dashed black line marks 365 days. Contracts with negative durations or durations more than 5500 days (15 years) are excluded.

Figure 9: Initial and Modification Spending

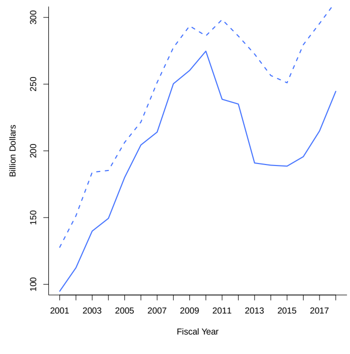

Note. This figure shows the levels of initial spending (any transaction that is *not* delineated a modification) and modification spending (transactions that are classified as modifications.

Figure 10: Decomposition of Sectoral Spending Growth

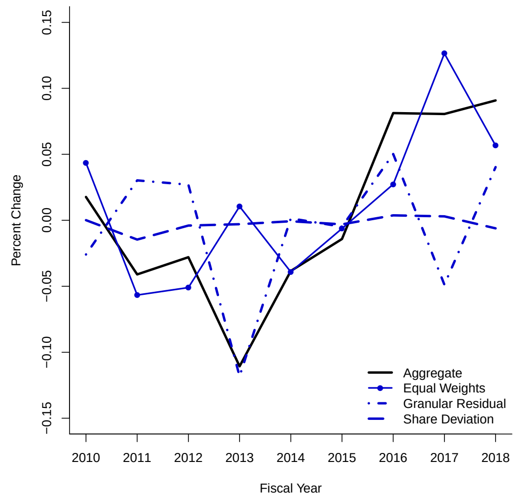

Note. This figure plots the individual components of government consumption growth, decomposed as in Foerster et al. (2011) as follows:

$$Z_t = \underbrace{\sum_{i=1}^{N} \omega_{i,t} z_{i,t}}_{(1)\ \text{Actual}} = \underbrace{\frac{1}{N}\sum_{i=1}^{N} z_{i,t}}_{(2)\ \text{Equal Weights}} + \underbrace{\sum_{i=1}^{N}\left(\bar{\omega}_i - \frac{1}{N}\right) z_{i,t}}_{(3)\ \text{Granular Residual}} + \underbrace{\sum_{i=1}^{N}(\omega_{i,t} - \bar{\omega}_i) z_{i,t}}_{(4)\ \text{Share Deviation}}$$

Figure 11: Sectoral Spending and Price Rigidity

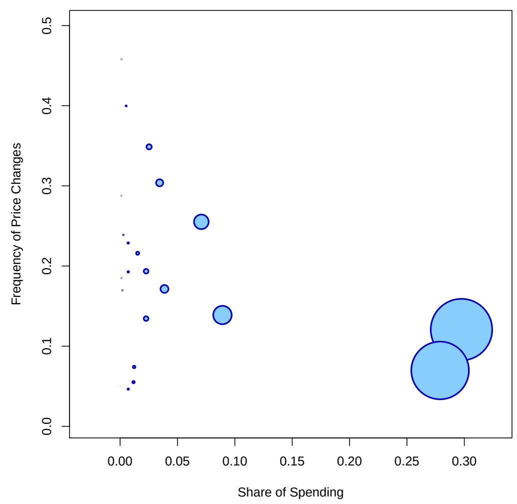

Note. This figure shows the average annual share of government spending in each two-digit sector (x-axis) plotted against the frequency of price changes in those sectors, based on BLS data. The size of the bubble corresponds with the average sectoral share of annual aggregate spending.

Figure 12: Dynamic Effect of Sectoral Shocks

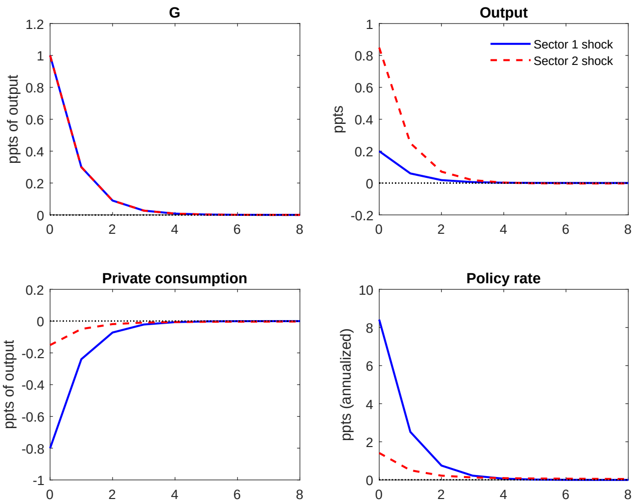

Note. Impulse responses to government spending shocks in two-sector model: sector 1 (solid line) vs sector 2 (dashed line). The shock is equal to one percent of output. The horizontal axis measures time in months. The vertical axis measures deviation from steady state in percentage points.

Figure 13: Dynamic Effect of Sectoral Shocks w/ homogeneous Pricing Friction

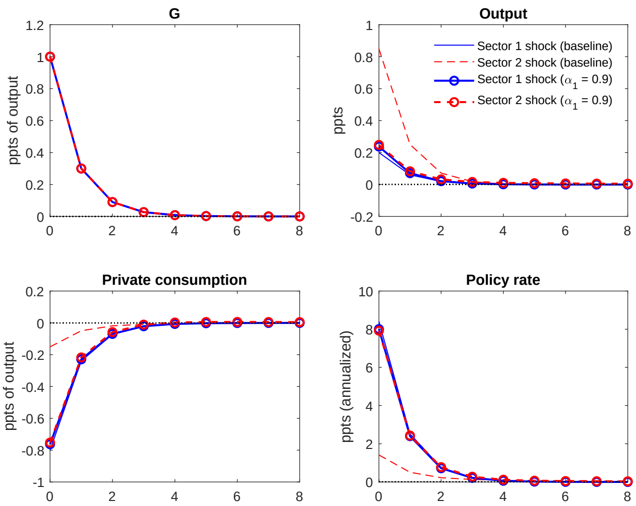

Note. Impulse responses to government spending shocks in two-sector model: sector 1 (solid line) vs sector 2 (dashed line). The shock is equal to one percent of output. The horizontal axis measures time in months. The vertical axis measures deviation from steady state in percentage points.

Figure 14: Dynamic Effect of Sectoral Shocks at Zero Lower Bound

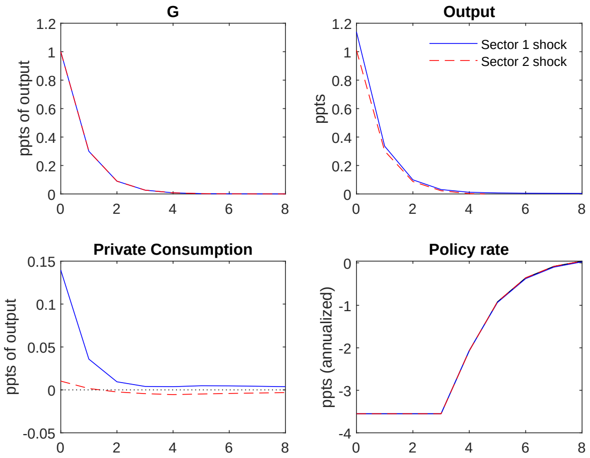

Note. Impulse responses to government spending shocks in two-sector model at zero lower bound: sector 1 (solid line) vs sector 2 (dashed line). The shock is equal to one percent of output. The horizontal axis measures time in months. The vertical axis measures deviation from steady state in percentage points. The zero lower bound binds because of shock to the time discount factor. The model is solved while allowing for occasionally binding constraints, as in Guerrieri and Iacoviello (2015).

Figure 15: Dynamic Effect of Aggregate Shock

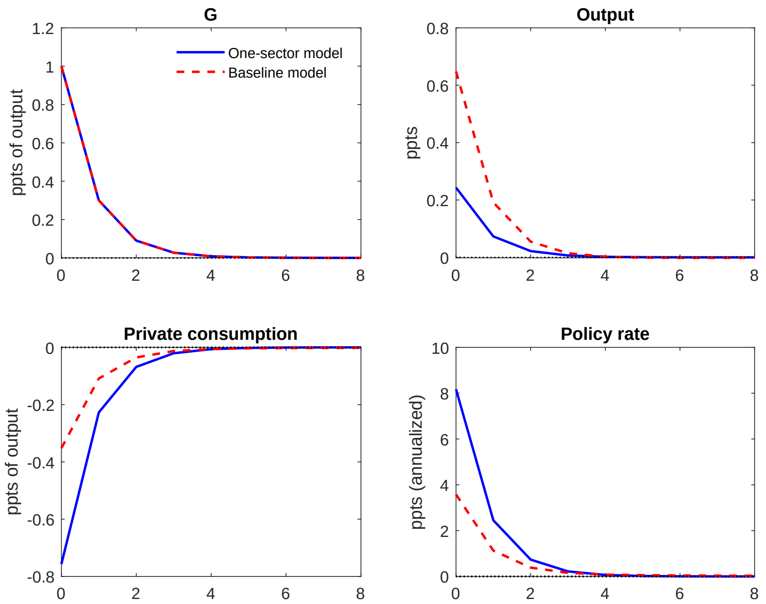

Note. Impulse responses to aggregate shock: one-sector model with $\alpha = 0.9$ (solid line) vs two-sector model (dashed line). The shock is equal to one percent of output. The horizontal axis measures time in months. The vertical axis measures deviation from steady state in percentage points.

Table 1: Explanatory Power of the Granular Residual

	Q=1000		Q=100	
	(1)	(2)	(3)	(4)
$\Gamma_t^{Q=1000}$	0.282***	0.280***		
	(0.080)	(0.079)		
$\Gamma_{t-1}^{Q=1000}$	−0.043	−0.046		
	(0.080)	(0.079)		
$\Gamma_{t-2}^{Q=1000}$	0.089			
	(0.080)			
$\Gamma_t^{Q=100}$			0.289***	0.278***
			(0.084)	(0.082)
$\Gamma_{t-1}^{Q=100}$			−0.030	−0.038
			(0.083)	(0.083)
$\Gamma_{t-2}^{Q=100}$			0.105	
			(0.084)	
Observations	15	16	15	16
R^2	0.558	0.506	0.539	0.478

*p<0.1; **p<0.05; ***p<0.01

Note. We run a regression of aggregate growth—$Z_t = \ln(G_t) - \ln(G_{t-1})$—on the granular residual and its lags: $Z_t = \beta_0 + \beta_1\Gamma_t + \beta_2\Gamma_{t-1} + \beta_3\Gamma_{t-2}$, where the granular residual is then given by $\Gamma_t = \sum_{i=1}^{K} \frac{g_{i,t-1}}{G_{t-1}}(z_{i,t} - \bar{z}_t)$. G_t is aggregate government consumption in year t, and $\bar{z}_t = Q^{-1}\sum_{i=1}^{Q} z_{i,t}$ is the average growth rate over the top Q firms.

Table 2: Spending Shares and Frequency of Price Changes

# of Sectors	**% of G**	% Value Added	Avg. Freq. In	Avg. Freq. Out	Avg. ρ In	Avg. ρ Out
1	30.93	6.29	0.12	0.14	0.3	0.28
2	59.89	13.3	0.09	0.2	0.27	0.31
3	69.14	16.22	0.1	0.22	0.29	0.28
4	76.48	20.39	0.13	0.2	0.24	0.36
5	80.51	24.72	0.14	0.21	0.23	0.39
6	84.13	30.66	0.17	0.19	0.3	0.35
8	89.12	40.5	0.19	0.15	0.32	0.29
12	95.53	55.3	0.17	0.18	0.34	0.34

Note. Avg. Freq. In refers to the average frequency of price changes in sectors within the given % of government spending. Avg. Freq. Out is the average frequency of price changes for all other sectors. The same interpretation is true for ρ, the persistence parameter of estimates AR1 processes.