ESSAYS ON THE MACROECONOMIC EFFECTS OF FISCAL POLICY: THE ROLE OF EXPECTATIONS, SECTORAL SHOCKS AND FINANCIAL BUBBLES

by

NIKOLAOS ALEXANDROS ANGELOPOULOS

A thesis submitted to the University of Birmingham for the degree of

DOCTOR OF PHILOSOPHY

Birmingham Business School

College of Social Sciences

University of Birmingham

December 2022

Abstract

What is the role of expectations in the business cycle? In the past decade, new theoretical and empirical tools were developed for tackling this question in a large body of literature. With this dissertation we contribute to this growing body of empirical research by presenting novel stylized facts and empirical methods pertaining to anticipated fiscal spending shocks and technological revolutions. In the first part, we introduce a novel approach to identify the economic effects of a shock in the components of public spending; public investment and public consumption. In the second part, we utilize an augmented version of this approach to isolate the exogenous variation in the components of public spending caused by two US wars and we proceed to ask the question of whether the effects of public investment and consumption are dependent on the monetary regime of the economy. In the final part of this dissertation we construct a novel indicator of anticipated technological booms and proceed to utilize it to establish novel stylized facts regarding the connections of technology shocks with financial boom and bust cycles.

Dedicated to all those who stood by me.

Acknowledgements

I would like to express my special thanks to my supervisors Prof. Christoph Gortz and Prof. Kaushik Mitra for giving me the opportunity for this project as well as for their support and guidance. I am also grateful to the Department of Economics of the University of Birmingham for their financial and administrative support throughout this project. Finally, I am deeply grateful to my family and friends for their constant support throughout the duration of this thesis.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Do anticipated shocks have explanatory power over the fluctuations of the business cycle? This is a question that has gained considerable traction in the past decade with the literature developing a substantial amount of novel empirical and theoretical tools to address it. With this thesis we contribute to the literature by presenting novel empirical methods and facts for a range of anticipated shocks.

The connecting thread of this thesis is the recognition that the nature of many economic shocks is anticipated by the agents and the applied econometrician should account for those effects. Our explorations however are not confined solely within the realm of expectational shocks. Rather, we use this observation as our starting point to undertake empirical exercises in a number of questions that are of substantial interest to the academic economists and the policymaker. Are the components of public spending, namely public investment and consumption, and their effects statistically distinct? If so, how does one go about identifying their economic impact while also minimizing the risk of confounding their shock processes? Are the effects and the propagation of those public spending components dependent on the monetary regime of the economy? Are technological shocks correlated and ultimately causally related with financial boom and busts? To approach these questions we have developed novel empirical methods and document new, to our knowledge, empirical facts on the fiscal multipliers and the macrofinancial impact of technology booms.

In the first part of this thesis presented in Chapter 2, we ask the question of whether the effects the public investment and consumption multipliers should be considered and analyzed separately. The convention in the existing literature is to analyze the components of the fiscal multiplier as a single object. Therefore, by asking whether the components of public spending are distinct we are also implicitly addressing the question of whether the estimates of the fiscal multipliers provided by the literature are suffering from external validity. The starting point in our approach is to utilize the information contained in the exogenous variation of the defence spending variables; defence investment and defence consumption. In doing this, we are approximating the exogenous components of public investment and consumption with the variation of their defence spending counterparts. Our empirical method lies in isolating the unexplained by the major economic aggregates variance of those defence spending variables. We implement this method without imposing any theoretical priors on whether the shock is anticipated or not; in that way we let the data speak for itself naturally lending weight to the components of the public spending shock that are anticipated or unanticipated. We find that the extracted shocks with this method have high explanatory power over the fluctuations of public investment and consumption. In our main finding, we document that the public investment and public consumption shocks have distinct effects on the economy.

In the third Chapter of this thesis we utilize a slight variation of our agnostic econometric framework developed in the second part to ask the following question: are the components of public spending liable to state dependence when the economy is at a Zero Lower Bound (ZLB) state. As in the second chapter we utilize the information contained in the defence spending variables to shed light onto this question. Our starting point is to bring additional information about the ZLB states by using an extended quarterly US dataset starting at 1939. In this manner, we can bring in the information from a long ZLB period that the US economy experienced in the '40s. To address potential issues of confounding the shock processes that might arise due to a dataset featuring the idiosyncratic historical circumstances of the 1940s we amend our econometric framework by identifying the shock as the shock that maximizes the unexplained variance contribution to the variable of interest while at the same time minimizing the contribution of the potential confounder. Our main results indicate that the public investment is the component of government spending more liable to state dependence.

In the final Chapter of this thesis we explore macrofinancial dynamics emerging in the aftermath of a technology boom. In other words, we ask the question: do technology booms lead to credit and financial expansions and ultimately busts? In other words, what is the link between technology booms and financial bubbles? The data underlying our exercise is a long panel dataset at an annual frequency of 15 industrialized countries spanning the entire history of finance capitalism from 1870-2008. We construct a novel indicator of technology booms using an extended patent dataset and identify the years during which each country experienced a technology boom. The usage of the patent dataset makes our indicator variable an inherently forward looking one and in this sense, perturbations to that variable can be broadly interpreted as an anticipated technology shock. With this indicator at hand, while we are not proceeding with making formal identification claims, we proceed to document novel stylized facts regarding technological revolutions and their macrofinancial implications. A technology boom is followed by a credit expansion and an increase in asset prices. It is also followed by asset price bubbles. However, we find that the relationship may necessarily be causal: while we find evidence that a technology boom directly leads to a credit market expansion we also document that it increases the probability of a stock market boom and eventual bust while also increasing the probability of a housing boom that doesn't ultimately result in a bust.

Chapter 2

Identifying Public Investment and Consumption Shocks

2.1 Introduction

On the occasion of the large fiscal packages that governments around the world are planning to distribute in order to protect and reform their economies, the public debate on the nature and the effects of fiscal spending is as vital as ever. What are the dynamic effects of an exogenous change to government spending? How large is the fiscal multiplier? What are its properties and its implications for the economy? These are questions right in the center of policy debates around the world.

A revitalization of the debate on government spending is not happening for the first time. Following the 2008 crisis and the large fiscal packages employed to deal with that recession, the research interest around fiscal multipliers, a topic that hadn't gathered much attention in the literature until then, gained considerable traction, both at the empirical and the theoretical front. Important contributions were made including new methods to identify exogenous changes in policy, a standardization of the measurement of fiscal multipliers and the incorporation of state dependence, as Ramey (2019) succinctly reviews.

Despite these leaps in applied research however, a lot of issues are yet to be resolved. For example, on whether the fiscal multiplier has non linear state dependent properties, the evidence is still fragile. A wide range of values has been suggested for the size of the multiplier, with the evidence pointing to it being less than or equal to one, albeit this issue is yet to be resolved. An additional unresolved prominent issue with important practical and theoretical implications is that of whether the composition of government spending matters. Consider the government spending variable that the literature uses to deduce the properties of the fiscal multiplier; it is an aggregate of different government spending activities, that is, it is an aggregate of government consumption, investment and transfers and as such it is treated in the literature. To be more precise it is an aggregate of all the different government purchase activities, and as such, unlike spending, it doesn't include interest payments. This aggregation takes place under the implicit assumption that the different purchase activities that a government engages in, all have the same effects in the economy. This treatment begs or rather leaves open the question of whether we can rightfully treat it as a single theoretical and empirical object or whether its components have distinct effects on the economy. This is an important question, essentially asking whether the multiplier estimates of the literature are externally valid and can thus be used to draw analogies and deduce general properties or whether they are dependent on the composition of the empirical object that each author uses. In this paper, we attempt to tackle this question.

In the public sphere, it appears to be part of the conventional wisdom that government investment is not only a distinct object from government consumption but also a more effective tool in the hands of the policy maker [1]. The 'Golden Rule of Investment', the notion that public investment pays for itself due to it acting as a positive externality on the productivity of private inputs (Baxter and King (1993)) and is thus superior to government consumption is both intuitive and theoretically sound and indeed this notion appears quite frequently in the public debate. Despite the attractiveness of the idea that public investment has more potent effects on the economy than public consumption however, little formal empirical research has been directed towards validating that claim. Moreover, the limited existing empirical results attempting to

[1]For example Krugman (2016) notes in the New York Times "... the simplest, most effective answer to a downturn would be fiscal stimulus — preferably government spending on much-needed infrastructure, but maybe also temporary tax cuts for lower- and middle-income households, who would spend the money. "

estimate separate public consumption and investment multipliers present conflicting findings. To begin with, a prominent early attempt to estimate separate multipliers was Perotti (2004) who presents no evidence of any superiority of government investment to government consumption, followed by Auerbach and Gorodnichenko (2012) who dissagregate the public spending multiplier to its consumption and investment parts and report different behaviors and multipliers of its constituent parts with public investment exhibiting stronger effects. Both studies follow along the identification lines of the seminal Blanchard and Perotti (2002) paper, that is they employ a Structural Vector Autoregression with a recursive identification based on assumptions about the timing of each shock at any given quarter, with Auerbach and Gorodnichenko (2012) adding a non linear qualification. More recently, there have been attempts to re-conciliate the limited empirical findings with a solid theoretical framework on why the two multipliers might or might not be different. It is in this spirit that Boehm (2020) presents a theoretical mechanism justifying potentially distinct effects of the government investment from the consumption multiplier while also empirically testing his mechanism in a Jordà (2005) Local Projections (LP) framework. In the same vein, Ramey et al. (2020) provide a review of the existing literature, both in the empirical and the theoretical aspect, with a focus on the infrastructure part of the public investment multiplier.

Estimating the properties and size of public investment and consumption multipliers is far from trivial; a battery of econometric challenges is associated with it. The identification problem of government spending is well known; is there a systematic way to discriminate between endogenous and exogenous discretionary changes in public spending? This issue becomes even more pressing when one wishes to isolate exogenous public consumption or investment changes, that is to dissagregate public spending. Various solutions have been proposed for the general public spending identification problem with the most prominent ones being imposing assumptions about the timing of government spending (Blanchard and Perotti (2002)) and utilizing military spending information (Ramey and Shapiro (1998)) under the assumption that defence purchases do not enter the production function. Extra care should be taken with the identification method then, to isolate an exogenous change in policy. An additional and well recognized issue is that government spending shocks are usually anticipated before they start materializing and thus, through rational expectations, their effects propagate into the economy before the actual spending starts. This is an issue connected with the very extensive literature on news shocks, as kickstarted by Beaudry and Portier (2004) and since Ramey (2011) seminal contribution showing that government spending is anticipated, it has been a constant theme of the fiscal multiplier literature to take these effects into account. In the context of government spending, the standard way to deal with this is to extend the econometric system with additional information, usually in the form of a forecast variable.

More specifically to issues related to our exercise, when trying to construct dissagregated estimates of government spending multipliers, one observes that government investment and consumption are very strongly comoving, with a correlation coefficient of almost one. If one is not careful with the identification , there is the danger of spuriously mixing the shock processes of the one component with the other. Some papers (Ilzetzki, Mendoza, and Végh (2013), Perotti (2004)) have tried to deal with this issue by constructing constrained Impulse Response Functions (IRFs) where the IRF of the non shocked component is artificially held to zero. In this way the IRF of the 'pure' government investment or government consumption shock is isolated. Further, there is the problem of endogeneity of public capital to the production function in the long run, which puts constraints on our ability to safely estimate long term properties of public investment.

With these observations in mind, in this paper we employ a Structural VAR (SVAR) setting to identify, estimate and assess the short and medium run impact and multipliers of government investment and government consumption. Taking the hint from the an extensive literature (most notably Ramey and Shapiro (1998), Barro (1987)), our starting point to tackle the endogeneity issue is to use military spending variables, with the underlying assumption that military related and consequently war time government spending is largely exogenous to the economy. A problem with government spending Structural Vector Autoregeressions (SVARs) is, as already mentioned, that government spending is often anticipated by the agents and thus the econometrician needs to be careful to utilize enough information to capture these anticipation effects. In more technical terms, for our shock to be truly structural the econometrician's information set needs to span that of the agent's. This is an insight derived from the news literature and related to well known issue of non fundamentaleness of the identified shocks (Forni and Gambetti (2014)). To tackle this we choose as our government investment and consumption shocks the combination of linear SVAR residuals that account for the maximal unexplained by the model volatility of the variable of interest for a pre set horizon. This approach derives from Ben Zeev and Pappa (2017) (BZP) contribution of identifying a government spending news shock as the shock that accounts for the maximum residual volatility of aggregate defence spending and has no contemporaneous effect on the economy. Note however, that we don't impose the constraint of no contemporaneous effect here. In this sense, the only assumption that guides us is that military spending is exogenous. In contrast, the Blanchard and Perotti (2002) identification operates under the assumption that agents are surprised by a change in government spending while the Ramey (2011) and BZP (2017) identifications operate under the assumption that the agents fully expect the spending shock and the shock has no contemporaneous impact. In this debate we do not take a stand with our identification approach being rather agnostic and with the identified shocks not directly corresponding to any theoretical object. However, this approach allows us to extract the maximum amount of information about the processes under examination and thus, in principle, informational deficiency or non invertibility of the system should not be a problem. One way to think about the extracted shocks are as objects giving us broad information about the size and properties of the multipliers of interest.

Our main results are the following; firstly and most importantly, in a United States (US) post 1954 sample, we find evidence that public consumption and public investment have distinct effects on the economy. A short to medium term public consumption shock has expansionary effects while a public investment shock does not have any significant effects on the economy possibly due to it crowding out private investment. Surprisingly, we also find that despite the fact that we impose no constraints or qualifications to disentangle the two perfectly comoving variables, our identification approach is able to sufficiently isolate the 'pure' defence investment and consumption components without exhibiting any significant confounding of the processes.

Additionally, our results indicate that there is a difference between the medium run government investment and consumption multipliers. This difference between the multipliers is statistically significant. This is in accordance with Boehm (2020) mechanism and Perotti (2004) empirical findings. When it comes to our identified shocks, we find that they can explain almost all of the variation of the variables of interest and a reasonable part of the variation of the rest of the SVAR variables in our econometric system. Further, we demonstrate that our retrieved shocks have high informational relevance over the variations of the generic public investment and consumption variables and are able to capture equally well anticipated and unanticipated shocks. Additionally, we find that our shock processes mainly convey information about the short run, that is 0-6 periods after the impact of the shock.

Our paper is structured as follows: in Section 2, we provide a snapshot of the state of the literature. Section 3 motivates our use of defence consumption and investment to approximate the public consumption and investment shocks and dwells on the informational content of these variables. Section 4 provides the estimation and the identification framework. Our main results for the defence consumption and investment shock are presented in Section 5 and 6. Section 7 dwells with the informational content of the extracted shocks. Finally, in section 8 we explore the effects over which horizon are driving our results and in section 9 we present an additional exercise to make sure there are no confounders in our shock processes.

2.2 A Brief History Of Identifying Government Spending

In the effort to resolve the issue of the endogeneity of government spending to the state of the economy, a seminal contribution using quarterly aggregate time series data was made by Blanchard and Perotti (2002). Under the assumption that the quarterly frequency of data is sufficiently high and informationally dense enough to address any potential endogeneity issues, Blanchard and Perotti use institutional information to isolate the automatic response of fiscal policy to economic conditions and by implication to identify the discretionary response and utilize a SVAR to construct impulse responses. Their leading assumption that government spending is predetermined can be readily transferred to a Structural VAR (SVAR) framework solved with government spending ordered first in a Cholesky decomposition. By implication, any contemporaneous associations of government spending with macroeconomic aggregates can be treated as causal influences. This identifying assumption has since then been widely used in the literature either in an SVAR or a large macro-econometric framework (see for example Mountford and Uhlig (2009); Perotti (2005); Caldara and Kamps (2008)).

Despite the widespread adoption of the Blanchard-Perotti method however, as Barro and Redlick (2011) note, it is hard to be optimistic about the ability of this approach to truly isolate causal channels of government spending to macroeconomic aggregates, at least when the data dealt with is in the form of aggregate time series for non defence purchases. Small variations of general government spending and the fact that it most likely responds to fluctuations in the economy are the main reasons behind this pessimism. Another widely adopted approach to isolate exogenous variation is to instead utilize the defence spending components of government spending. The underlying assumption here is that the drivers of defence spending are quite plausibly exogenous to the state of the economy and do not affect the household's Marginal Rate of Substitution (MRS), something that cannot be as easily assumed about government spending at the state and the local level which most probably responds to aggregate economic conditions. The downside with the defence spending component of government spending however, is that a small number of events might be accounting for a disproportionately large fraction of the variation of defence spending as we shall later see.

Aside from the general issue of isolating the exogenous component of government spending, there is also the issue that government spending events are most likely anticipated, as first recognized by Ramey and Shapiro (1998). This insight, connecting this branch of literature with the extensive literature on news shocks, initiated by Beaudry and Portier (2004), Beaudry and Portier (2006), has large implications for any identification attempt. The effects of a government spending shock start propagating in the economy, through rational expectations, before the shock materializes, and the additional issue the researcher has to face is how to account for those anticipated effects.

In a seminal contribution, Ramey (2011), tackles both of these identification issues by constructing a military news variable that captures the anticipated component of defence spending and bringing it in as a variable in the econometric system of interest. Since then, the literature on fiscal multipliers mainly deals with the anticipation problem by bringing extra information in the econometric system under scrutiny, either in the form of expectational variables (Forni and Gambetti (2016), Caggiano et al. (2015)) or in the form of variables constructed based on (mainly war related) narratives (Ramey (2011), Ramey and Zubairy (2018)). An important exception is Ben Zeev and Pappa (2017) who utilize Barsky and Sims (2011) method and identify a fiscal news shock as the linear combination of reduced form VAR residuals that maximizes the forecast error variance of defence spending and is contemporaneously orthogonal to it; as with Barsky and Sims (2011) the crucial assumption here is that defence spending is largely exogenous to the economy.

A common element in all of the previously mentioned papers is that government spending is treated as an aggregate object, with no attention paid on whether its constituent parts might be having statistically different effects. Is this aggregation hiding potentially useful information though? As mentioned in the introduction, apart from an intuition existing in the public sphere that government consumption cannot be aggregated with government investment, some papers have tried to deal with this issue. Albeit, this body of literature is still quite limited in size. Starting with an early seminal study, Aschauer (1989b) tried to isolate the long term effects of public infrastructure spending on the economy. An attempt to explicitly estimate separate multipliers for government spending and consumption has been made by Perotti (2004) who explicitly recognizes that government spending is an aggregate of elements which might be associated with distinct effects and provides his own estimates of 'pure' public investment and consumption shocks in a SVAR setting, finding no significant difference between the two. More recently, Auerbach and Gorodnichenko (2012) address the same issue in a non linear SVAR framework, this time providing evidence for differing effects of the two multipliers. Ilzetzki, Mendoza, and Végh (2013) following the same methodology as Perotti, give their own estimates of government investment in an SVAR setting for a number of countries, finding no robust evidence that government investment is more effective in high income countries. At the theoretical front there have been two recent important contributions from Boehm (2020) and Ramey et al. (2020) who present evidence that counter conventional wisdom, government investment might have adverse effects on the economy in the short and medium run due to the crowding out of private investment it induces. Both papers complement their theoretical finding with an empirical investigation. Ramey et al. (2020) also adopts a longer run perspective arguing that in the long run public investment can have significantly positive effects on output and productivity.

2.3 Fluctuations and Informational Content of Defence Spending

Central to our identification is the assumption that the episodes that trigger large variations in defence spending are exogenous to the economy. This assumption appears reasonable; there is no theoretical justification or intuition as to why political and war events that are responsible for the variation in defence spending should be entering the production function, as long as the infrastructure of the country remains intact and no mass conscription takes place.

Assuming exogeneity of defence spending, as a large part of the relevant literature does, one can then approximate an exogenous change in the components of fiscal policy through variations in military spending. One can think of surges in military spending in this context as a 'natural experiment' generating the necessary for the econometrician exogenous variation. A condition that needs to hold for this approximation to be accurate however, is that it should be the case that variation in defence spending accounts for a considerable part of the exogenous variation of the government spending variables.

The informational content of the aggregate defence spending variable has already been scrutinized extensively in the literature. Nevertheless, since we are focusing on its components and not the aggregate, we provide some additional motivation for our identification by presenting and analyzing the informational content of defence investment and consumption.

2.3.1 The US History of Defence Spending

To begin this preliminary analysis, we first want to gauge the variation of the defence spending variables. To analyze the extent of the exogenous variation of defence variables, Figure 2.1 plots the quarterly per capita real defence investment and consumption change expressed as a ratio to the previous year's real per capita GDP for the time period 1947-2019. Under the assumption of exogeneity of military spending, this graph should give us a good idea about the events that inform our identification.

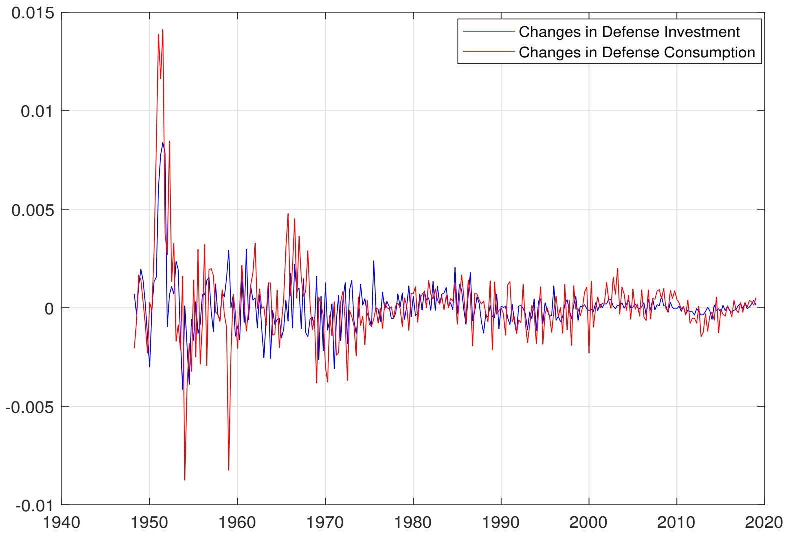

Figure 2.1: Changes in Defense Purchases, 1947–2019 (expressed as ratios to the previous year's GDP).

What kind of information can we extract from this figure? At low frequencies, the movements of the two variables closely coincide. At higher frequencies however, defence consumption appears to be much more volatile than defence investment, with quite sharp ups and downs. This sharper volatility renders, in principle at least, a government consumption shock more easily identifiable than a government investment shock. It could also suggest that defence consumption is more prone to manipulations of the policy maker, rendering the identification of its discretionary variations easier than the defence investment case. Further, we observe a clear decline in volatility from the '70s onwards, raising some suspicions that from that point on there might be a different stochastic process driving the variables.

A major spending surge in the early '50s stands out in the graph, both in defence investment and consumption. This is clearly associated with the beginning and the end of the US involvement in the Korean War. The variables are clearly volatile throughout the '50s, with defence consumption being more so. From the mid to late '60s there is another surge in the variation of defence consumption mainly, associated with the gradual escalation of the US involvement in the Vietnam War. From there afterwards, there is a clear fall in the volatility of the two variables, with the exception of a gradual spending surge in the '80s, apparently associated with President Reagan's arm buildup, the largest in peacetime history, in the context of increasing Cold War tensions of the period.

What have we learned from this graph? In the post war history of US defence spending, the Korean War clearly stands out as the most informative event. The sheer magnitude of this spending event immediately raises the question of whether any identified shock from this sample would be externally valid, or in other words, would the properties of the shock stem from the information contained in the war event and if so how generalizable would this shock be. This is a subtle and important question that can be tackled separately and could make its own paper. In this paper we wish to avoid it and to make our identification as robust to confounders as possible at the cost of potentially throwing away valuable exogenous variation. We thus exclude the Korean war from the sample, so that the bulk of the information in our identification scheme won't be coming from major war events. [2].

We therefore focus on the post Korean War period with our sample starting at 1954. Given this sample, if we wish to identify an exogenous spending shock, the main source of information is in the period before the '70s with two major sources of volatility associated with war events, one sudden and the second more gradual. Defence consumption appears more volatile than defence investment, maybe pointing to its more flexible nature, not requiring the extensive planning and time that an investment program is typically associated with. Another property that readily stands out from the graph is the subsequent to the Korean War decline in the values

[2]In doing this, we ignore Barro and Redlick (2011) warning that ' It seems unlikely that there is enough information in the variations in defense outlays after 1954 to get an accurate reading on the defense spending multiplier.'

of both variables in the early '50s.

We have seen that the major variations in defence spending are mainly associated with war related events by looking at the defence spending variables. A related but distinct question is whether these variations contain enough information to allow us to sufficiently approximate exogenous policy changes of government investment and consumption. In other words, for how much of the total variation of total government consumption and investment does the defence variable variation account for.

To engage this question at a preliminary level, in Figure 2.18 in Appendix C we plot the ratio of the defence spending component to its corresponding general government spending component. We see that each defence spending component consistently accounts for more than 20% of the respective government spending component for the whole of the sample, with that value being over 50% throughout the '50's.

As a first more precise parse to this question, we regress the government spending variable of interest on four lags of the defence spending variable of interest and obtain the R^2 of the regressions. This exercise should give us a first hint on how much of the total variation of government spending is due to defence spending. Table 2.1 reports the results:

Table 2.1: Variance Contribution Of Defence Spending To Total Gov Spending

Regression	R^2
Gov Spending On Def Spending (Total)	0.9746
Gov Investment On Def Investment	0.9337
Gov Consumption On Def Consumption	0.9809

We see that for all components of public spending, defence spending variation accounts for the overwhelming majority of its variation, with the R^2 ranging from 93% to 97 %.

The military variables accounting for most of the variation is one thing. More importantly though, they should account for most of the exogenous variation. To parse this question we run the following exercise; figure 2.2 plots the total public and defence investment and consumption in real log per capita levels for the post Korean War period. The lower panels depict the general government spending variable, while the upper panels depict the defence spending variables. The vertical lines correspond to important war related events in our sample, which under the literature's assumption are largely exogenous. Looking at the government variables we see that each major defence spending event in our dataset is associated with clear upswings in the corresponding government spending variable. Upon inspection therefore, military spending events seem to be responsible for a good portion of the variation in the components of government spending and this variation corresponds with exogenous events.

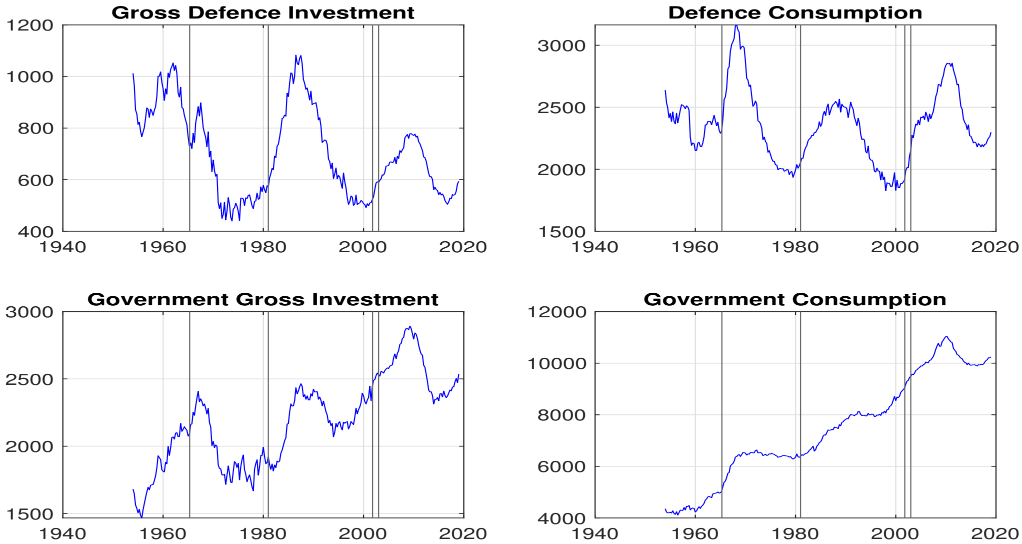

Figure 2.2: *Notes:* Variables expressed in real per capita terms. Upswings interpretation : (1)1964Q3 : Gulf of Tonkin Incident, Beginning of the US involvement in the Vietnam War; (2) 1980Q4 : Beginning of Reagan Presidency; (3) 2001Q3 : 9/11; (4) 2003Q1: US invasion of Iraq.

We can therefore safely conclude that the components of defence spending contain enough information and have enough variation that could allow us to reliably identify causal effects of changes in government consumption and investment on aggregate variables. As already mentioned one should keep in mind however the concern that the identification might be overly influenced by the two major war related event.

2.3.2 What Constitutes Defence Consumption and Defence Investment ?

We showed in the previous section that the defence spending variables contain enough exogenous variation that should in principle allow for a robust identification of changes in the components of government spending. Before we proceed with describing our identification strategy however, we must first answer the question of what kind of information those defence spending variables contain. What are they composed of? This is an important question related to the issue of whether the economy would respond to military spending in the same way it would respond to other types of government purchased. In this section therefore we take a close look into the public accounting methods behind the military spending figures and what kind of spending is hidden in the two defence variables.

According to the National Income and Product Accounts (NIPA) [3], defence consumption is made up of the following large categories of spending; for the first category more than half of defence consumption is spent for compensation of general government employees and consumption of general government fixed capital. The second big category of defence consumption is consumption of intermediate goods and services, wherein the intermediate goods are mostly made for weapon systems (aircrafts, ships, vehicles, electronics) and are highly durable final products. This composition corresponds to a significant degree to the composition of general government spending wherein the better part (meaning consistently more than half of the total spending) of it goes into government employees compensations and the rest into the purchase of intermediate goods. Given this accounting, there is no reason to believe that, if stimulated accordingly, they shouldn't have approximately the same demand inducing effects that a stimulation of general government consumption would have. Using defence consumption as an approximation of general government consumption shouldn't thus pose a problem. Thus, defence consumption should provide a good approximation for general government consumption.

Things get more involved when one delves into the composition of defence investment however. Defence investment can be divided into three categories; spending on military structures, spending on equipment, whereby the equipment consists of weapons and weapon delivery systems (air crafts, missiles, ships, vehicles) and spending on intellectual property products that include software and research and development. Some clarifications need to be made here regarding the accounting method; as Perotti (2004) reviews, according to the 1993 System of National Accounts weapon delivery systems should be classified as government consumption.

[3] U.S. Bureau of Economic Analysis, "Table 3.11.5. National Defense Consumption Expenditures and Gross Investment by Type"

The US public accounting system is unique in the sense that it classifies them as government investment. As mentioned earlier, the 'Golden Rule of Investment' states that public investment pays for itself since in the long run it adds to public capital. This proposition though implies that a big part of public investment goes into infrastructure and research and development, since these are the traditional channels through which Total Factor Productivity (TFP) is raised.

Looking at Figure 2.3 that plots the time series of the contribution of each component to total defence investment, we see that spending on structures makes up for around 10% of defence investment in most of the sample with the exception of the Korean War. Investment on equipment and on intellectual property products contribute about an equal portion of 40-50% for most of our sample.

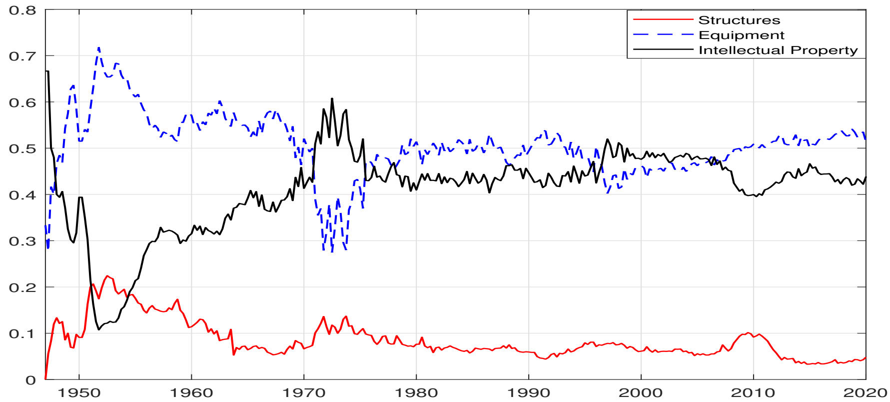

Figure 2.3: Composition of Defence Investment

This composition might raise some concerns as to whether the shock we are identifying is a valid approximation to a regular public investment shock. Indeed, when one speaks of a public investment shock what comes to mind is mainly infrastructure spending and how it adds to the productivity of the factors in an economy. Coming to the components of defence investment, spending on structures is the closest to what we mean by infrastructure spending, although much of the literature on the social returns to public capital concerns non defence public capital (Aschauer (1989)). Given the small spending of defence investment that goes into structures, we don't expect a diffusion of defence investment into TFP through this channel, as is a common assumption in the literature. Defence Investment effects could however diffuse into

TFP through the Research and Development channel, as military R&D makes up for most of total R&D throughout our sample. Indeed, Moretti, Steinwender, and Van Reenen (2019) show that military investment in R&D results in productivity gains.

Given these considerations and the fact that military equipment which makes the majority of defence investment has a life of more than a decade, it would be reasonable to assume that any defence investment shock would be close to a shock on the public consumption of durable goods. However, the possibility of an increase in TFP through defence investment is still entirely possible through the Research and Development channel.

2.4 Econometric Framework

2.4.1 Data

For our exercise we employ an SVAR with 11 variables. We make the choice to utilize a relatively big SVAR due to the potential fragility of small SVARs (Forni, Gambetti, and Sala (2019)). The variables we include in our system, in the following order, are: Defence Consumption, Defence Investment, GDP, Private Investment, Private Consumption, Unemployment, Labor Hours, an Average Tax Rate variable, inflation, the 3-month T-Bill and Fernald (2014) utilization adjusted TFP. All data are from the US. An advantage with using US defence spending data, is that the US never suffered a major destruction of infrastructure. These variations should thus be mainly associated with demand effects.

Following standard practice, our measure of private investment includes consumer expenditure on durables, while that of consumption consists of non durables and services. As the average tax variable we use the federal receipts to output ratio. The National Income and Product Account (NIPA) variables have been converted to real using a GDP deflator.

Contrary to standard practice in the SVAR literature, we are not normalizing the variables by log transforming them. Rather, the normalization we choose is to divide the NIPA variables by an 2nd degree polynomial estimate of trend GDP. [4] In this way, the estimated IRFs should

[4] The change of normalization does not affect our results qualitatively though. The main results with log

correspond to dollar to dollar changes instead of percentage elasticities as is usually the case. We do this following Owyang, Ramey, and Zubairy (2013) and Ramey and Zubairy (2018) suggestions, since the logarithmic transformation could possibly bias our multiplier estimates. We delve deeper into this matter in the fiscal multipliers section.

Our dataset spans the period 1954Q1-2019Q1. Details on how the data is constructed are provided in the Data Appendix. Due to the big size of our system, we employ Bayesian methods to estimate the SVAR. We use 4 lags and a standard Minnesota prior, controlling for stationary and non stationary variables. The posterior is drawn from a Gibbs sampling algorithm of 20000 draws with a burn in rate of 16000 draws.

2.4.2 Identification Method

We now turn to the description of the empirical method. Following the SVAR literature, our guiding principle is that the shock should be a linear combination of the residuals of the estimated SVAR.

Following Ben Zeev and Pappa, 2017 , who in turn build on Barsky and Sims' (2011) method, we base our identification strategy of the fiscal shocks on medium run restrictions. Fundamentally in this approach, a structural public investment or consumption news shock should be the linear combination of SVAR residuals that maximizes the forecast error variance of defence investment or consumption and has no contemporaneous effect on its own variable. Notice that critical in this assumption is that both of our defence spending variables should be exogenous to the economy, a reasonable assumption as we demonstrated earlier.

However, in contrast to Ben Zeev and Pappa, 2017 , we are not going to impose a zero contemporaneous impact constraint; we proceed this way because of the danger of throwing away valuable information by constraining the problem. Instead, we are going to follow a more agnostic approach, in the spirit of Angeletos, Collard, and Dellas (2020), and identify a public spending shock as simply the shock that accounts for the maximal volatility of the variable of interest. This should produce a generic public spending shock, without an exact theoretical

normalization are presented in Appendix D.

object to correspond with, but it should nonetheless inform us on the properties of each kind of public spending.

To formally describe our model, we start with the reduced form moving average representation of the SVAR :

$$Y_t = B(L)u_t \tag{2.1}$$

where Y_t is an $m \times 1$ vector of endogenous variables, $B(L) = I_m + B_1 L - ... + B_p L^P$ is the matrix polynomial in the lag operator L where the $B_i, i = 1, ..., p$ are $m \times m$ parameter matrices and u_t is a zero mean, m-dimensional error process of the reduced form with a variance covariance matrix $E[u_t u_t'] = \Sigma_u$. Identification then amounts in finding a linear mapping $u_t = A\epsilon_t$ between the prediction error and the orthogonal shocks contained in ϵ_t, where A is the impact matrix.

The structural Vector Moving Average (VMA) representation of the SVAR is then going to be:

$$Y_t = C(L)\epsilon_t \tag{2.2}$$

where $C(L) = B(L)A$.

The matrix A must satisfy $\Sigma_u = E[A\epsilon_t \epsilon_t' A] = AA'$. This restriction isn't however sufficient to identify A. To see this, consider that for all permissible matrices A, there exists an alternative matrix $\tilde{A}$ such that $\tilde{A}Q = A$, where Q is a $m \times m$ matrix with orthonormal columns, that also satisfies $\tilde{A}QQ'\tilde{A}' = \Sigma_u$.

Consider the h-step ahead forecast error of the i_{th} variable y_i in Y as

$$y_{i,t+h} - E_t[y_{i,t+h}] = \sum_{\tau=0}^{h} B_{i,\tau}\tilde{A}Q\epsilon_{t+h-\tau} \tag{2.3}$$

where $B_{i,\tau}$ represents the i row of the matrix of MA coefficients at horizon τ. The share of the forecast error variance of variable i attributable to shock j at horizon h is then

$$\Omega_{ij}(h) = \frac{e_i'(\sum_{l=0}^{h} B_l \tilde{A} q q' \tilde{A}') e_i}{e_i'(\sum_{l=0}^{h} B_l \Sigma_u B_l') e_i} \tag{2.4}$$

where q is the j column of the orthogonal matrix Q, such that $q'q = 1$ and e_i is a column vector with 1 in the i_{th} position and zeros elsewhere.

To identify the defence spending shock, assume that the variable from which we wish to extract the shock that maximizes the error variance is ordered in position p of the SVAR. Then, the impact vector of the public spending shock is the vector q solving the following problem :

$$q = \underset{q}{\operatorname{argmax}} \sum_{l=0}^{h} \Omega_{p,p}(h) \tag{2.5}$$

subject to the following constraint:

$$q'q = 1 \tag{2.6}$$

In words, we are searching for the orthonormal vector q that is explaining as much as possible of the unexplained forecast error variance of the defence spending variable ordered in position p in the SVAR. [5]. This method might raise concerns due to it being prone to spuriously mixing shock processes, especially since the two defence variables exhibit almost perfect comovement. We address these concerns later on in this chapter.

2.5 Results

In this section, we present the 20 period IRFs to the defence consumption and defence investment shocks. The responses are interpreted as deviations from the steady state values. To assess statistical significance 68% and 84% posterior density intervals are used. As a first check on the validity of our identification we are reporting the Forecast Error Variance Decompositions of the variables to the respective shock with a 68% and 84% posterior density interval.

[5] This is a method very reminiscent of Principal Component Analysis (PCA) in machine learning.

2.5.1 Defence Consumption Shock

We start with the defence consumption news shock; what are the main findings? Firstly, the IRFs follow the hump shaped news shock response, typical in the literature; this acts as a first sign that we are to a substantial degree capturing the anticipated effects of the shocks of interest. Notice also that most of the variable IRFs revert back to zero after 8 periods. This is a clear indication to the short run nature of the shock we are identifying.

The effects of the defence consumption shock are expansionary. Output and private consumption rise on impact and unemployment falls. Private investment also reacts by rising significantly, albeit only for 4 periods. On the other hand, the tax and the inflation rate as well as the nominal interest rate don't react significantly and neither do hours worked. Interesting is the short term rise in Total Factor Productivity which reverts back to zero after about 10 periods.

Overall, the defence consumption shock has some short term stimulative effects on the economy, raising output and lowering unemployment up to the 8th period of the horizon. This seems to be in line with the priors we had that public consumption has by its nature more short term effects on the economy.

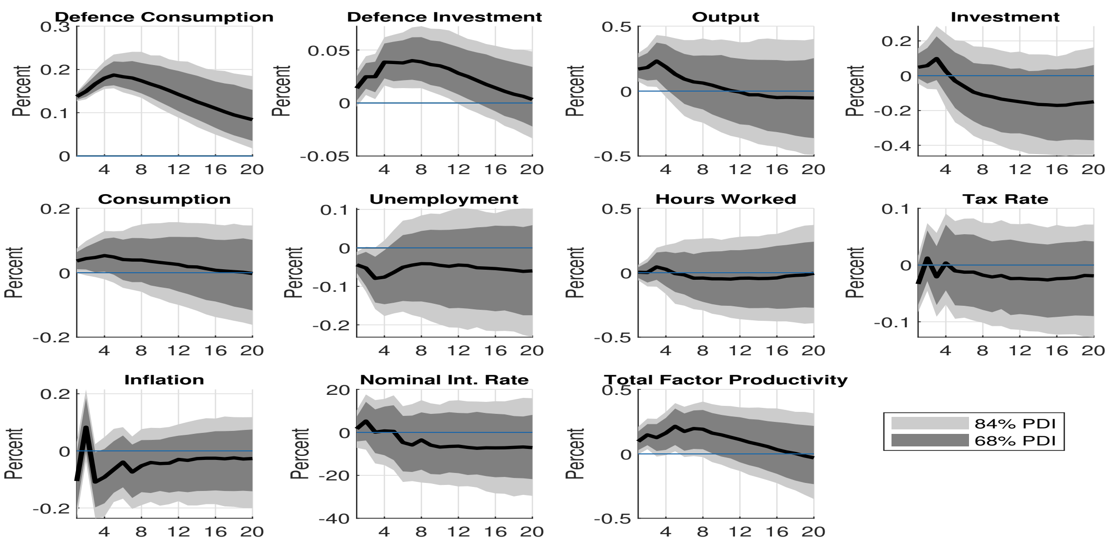

Figure 2.4: IRFs to Defence Consumption Shock. 84% and 68% Posterior Density Intervals.

In Figure 2.5 we plot the Forecast Error Variance Decompositions to gauge the statistical

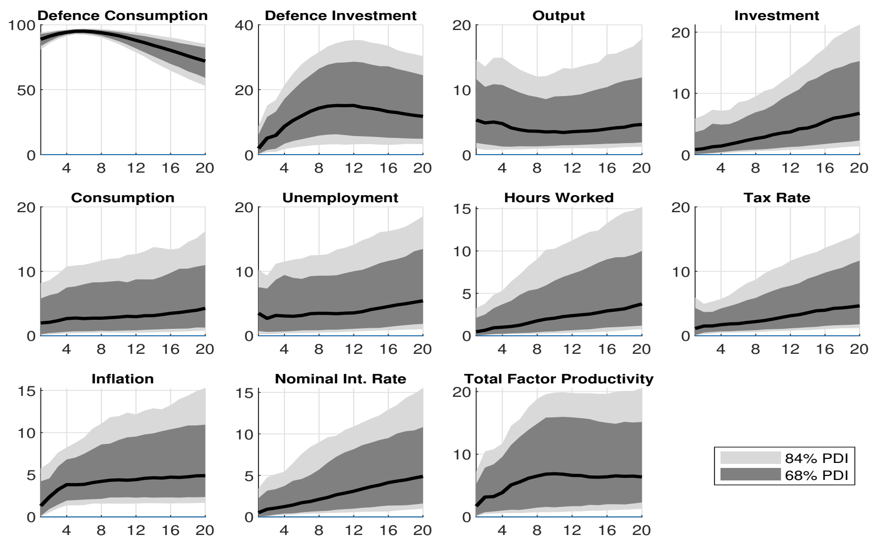

Figure 2.5: Variance explained by Defence Consumption Shock.

information our identified defence consumption news shock contains about the variation of the variables in the system. The shock significantly contributes to almost 100 % of the defence consumption variation and up to 10% of the defence investment variation. What stands out is that the shock doesn't significantly contribute to any of the aggregate variables.

What do these numbers tell us? Firstly, that the shock has little contribution to most of the business cycle variables in our SVAR apart from output and unemployment. This is to be expected; government spending shocks should not account for much of the business cycle variation and war related spending events in particular are scarce. This exercise also acts as an initial parse to the question of confounding, indicating that we are not mixing up processes with the other variables of the system. The fact that it explains up to 100 % of defence consumption is also encouraging, suggesting that we are capturing all of the exogenous process moving defence consumption. Upon first inspection, the fact that our shock contributes up to 10 % of the variation of defence investment appears worrying, possibly suggesting that we are also capturing part of the defence investment process in our identification. We address this concern in section 9 of the present chapter.

2.5.2 Defence Investment Shock

Next, we move on to the defence investment news shock IRFs. Figure 2.6 reports the IRFs to a one standard deviation defence investment news shock. Things are different from the defence consumption shock case. The defence investment variable doesn't exhibit the hump shaped pattern but rather a steep decrease, indicating a surprise shock. However, the other variables of the system are hump shaped thus indicating that the shock doesn't have an immediate effect but rather propagates gradually in the economy.

The defence investment shock moreover appears to have a recessionary effect on the economy. Output, hours worked and private investment fall significantly. Interestingly, private consumption rises significantly on impact as does the tax rate. What is behind this rise in consumption is unclear. The negative response of output is likely due to the steep fall of private investment. The negative response of private investment to a public investment shock is an indication that there are the private investment is being crowded out, likely due to the rise in the real interest rate as in Boehm (2020). In this context, the response of the monetary variables is interesting; both are falling albeit by the same magnitude thus indicating that the real interest rate remains unchanged. If this is true then it is a potential hint against the suggested mechanism of the real interest rate rise being behind the fall in private investment. Finally, interesting is the response of Total Factor Productivity that is not reacting significantly distinctly from zero, counter to the intuition and theoretical insight that public investment adds to productivity. Overall therefore, the defence investment shock doesn't have the stimulative short term effects on the economy that we found the defence consumption shock to have.

What kind of information does the defence investment news shock contain about the variation of the variables? Figure 2.7 reports the relevant FEVDs. The shock is contributing up to 90 % of the defence investment variance, which is encouraging, suggesting again that our process is capturing most of the process. It is additionally contributing to only 10% of the defence consumption variance. As is the case with the defence consumption shock there is not much contribution to the variation of the main business cycle variables. It is mostly contributing to investment, inflation and hours worked and these in the medium to long term.

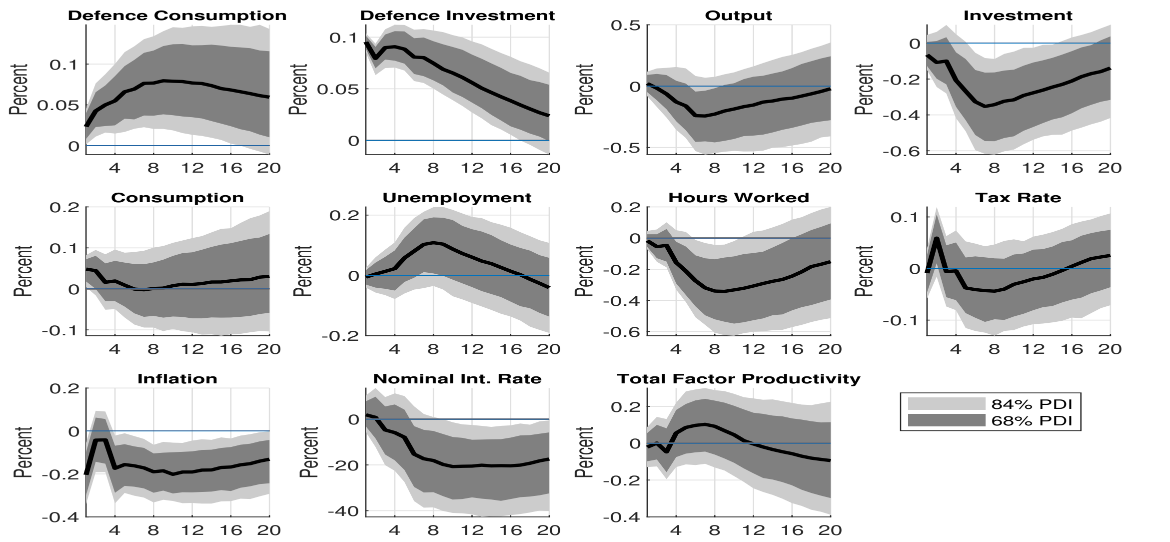

Figure 2.6: IRFs to Defence Investment Shock. 84% and 68% Posterior Density Intervals.

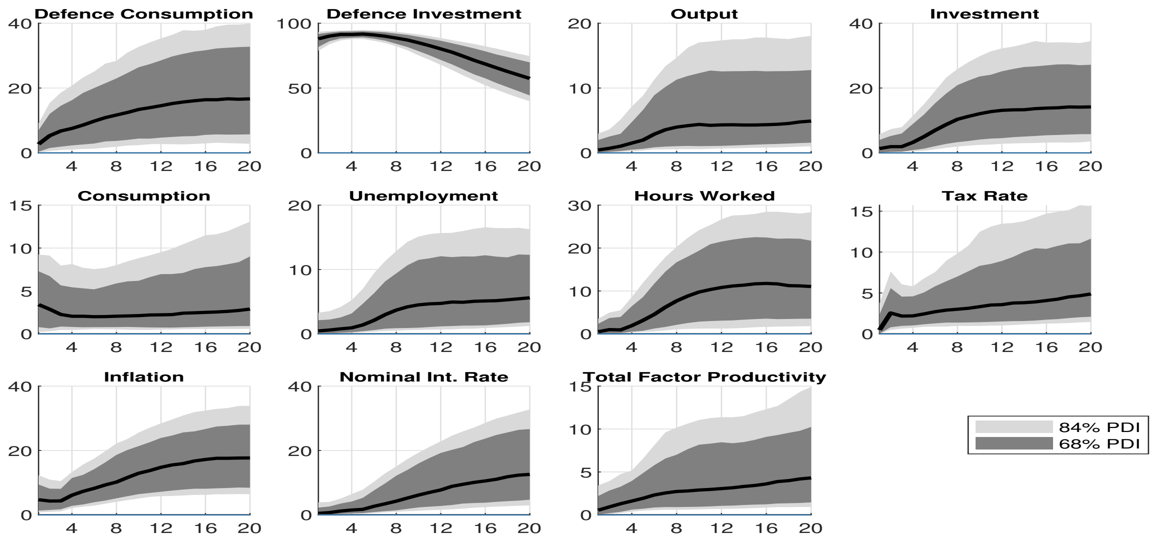

Figure 2.7: Variance explained by Defence Investment Shock.

Overall, in contrast to the IRFs, the FEVs of our two shocks show similar patterns; each shock is contributing almost 100% to its own variable and around 20% to the other defence spending variable. To other variables of the system neither of the shocks is contributing anything significantly with the contributions ranging from 10% to 20%. Thus the government spending shocks, as approximated with the defence investment variables, don't appear to have significant influence over the business cycle.

A Closer Look Into The Mechanics Of The Defence Investment Shock

An important question with regards to the defence investment shock is what causes this sharp fall of private investment. This is an issue with important implications for both the policy maker and the theoretician. Theory predicts that the crowding out of private investment is caused by the high elasticity of demand for private investment which is very sensitive to an increase in the real interest rate caused by a public spending shock. However, in our system both the nominal interest rate and the inflation rate fell in response to the shock, suggesting that the real interest rate was unchanged. So the mechanism driving the deep drop of private investment is potentially somewhere else to be found.

Perotti (2004) suggests a culprit; a plausible explanation for the strong negative reaction of private investment to the public spending shocks might be hidden in the national accounting practices of the US. Quoting from Perotti (2004):

> "Most goods purchased by the government and classified as government investment take more than one quarter (the accounting period) to be produced. National account systems record government purchases of goods with long production processes with two different methodologies. According to the first (called work put in place method under accrual accounting, and progress payment method under cash accounting), each quarter the addition to the value of the unfinished good by a private contractor is recorded directly as government investment in the government accounts. According to the second method (called delivery method under accrual accounting, and payment method under cash accounting), each quarter the addition to the value of the good by the private contractor is recorded as work in progress in the inventories of the private sector. When the good is completed and delivered to the government, the whole value of the final good is recorded at once as government investment, while private inventories are decreased by the same amount."

Since private inventories are a part of private investment, the delivery accounting method should generate a mechanical negative correlation between private and public investment. The

US uses the delivery accounting method for its defence machinery and equipment. Does the data support this hypothesis however? A way to gauge into this issue is to decompose private investment into its two major constituent parts, that is fixed private investment and change in private inventories, insert them in the system and plot their IRFs; figure 2.8 presents the results from this exercise;

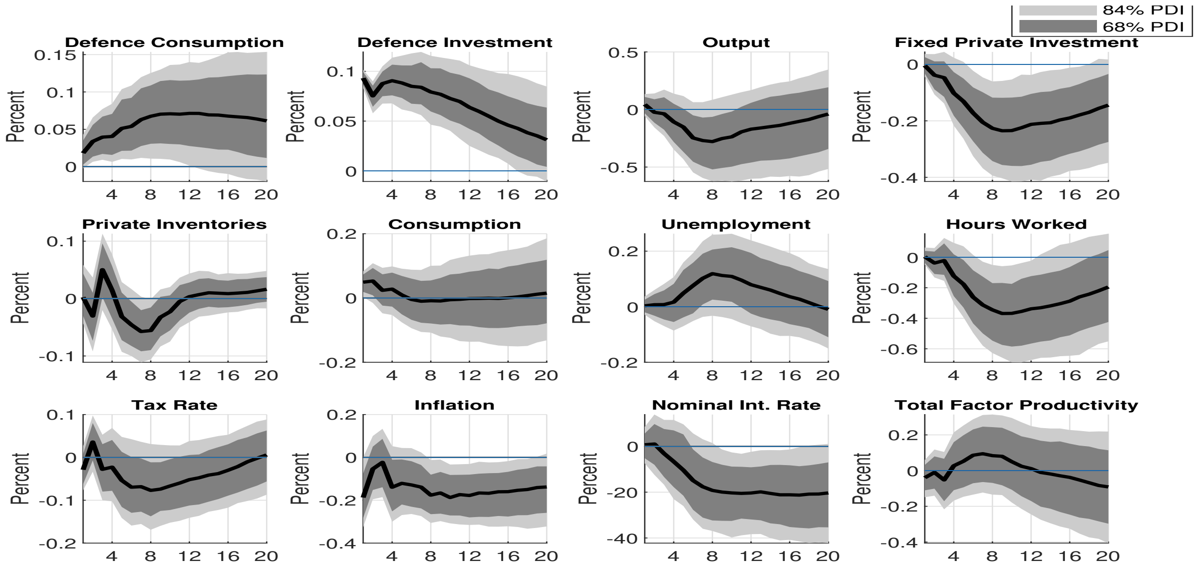

Figure 2.8: IRFs to Defence Investment Spending Shock with Private Investment Decomposed Into Fixed Private Investment And Change In Private Inventories.

What do we learn from this figure? It appears that the accounting hypothesis that the fall of private investment is due to the mechanical fall of private inventories doesn't hold. Private inventories rise and fall significantly but the fall in fixed private investment is much larger and much more reminiscent of the shape of the general private investment IRF, suggesting thus that the culprit is not to be found in Perotti's accounting hypothesis. A shock in public investment does indeed crowd out private investment. Given these considerations, the most probable mechanism appears to be the fact that through public investment, the government is absorbing the resources from the private sector rather than some mechanical reaction stemming from the construction of the variables.

2.6 Fiscal Multipliers

In this section, we discuss the fiscal multipliers associated with our shocks. The question on how a multiplier is to be defined has long been debated in the literature. Initial measures defined the multiplier as a ratio of the peak of output to the initial spending shock (Blanchard and Perotti, 2002) and variations of this measure have been widely utilized in the literature Auerbach and Gorodnichenko (2012), (Auerbach and Gorodnichenko (2013)). Mountford and Uhlig (2009) argued instead for an approach that takes the dynamics more explicitly into account. They propose instead the integral multiplier, that is the the multiplier is more appropriate to be measured as the integral of output response divided by the integral of government spending response. We therefore consider the integral multiplier in our analysis.

An additional issue arises when it comes to the construction of these multipliers. Typically, the multiplier is constructed through converting the IRF elasticities to dollar equivalents. In SVAR analyses, the IRF elasticities are expressed in percentage terms and do not directly reveal the government spending multipliers. The literature deals with this problem by using an ex post conversion factor based on the sample average of the ratio of GDP to government spending.

As Ramey and Zubairy (2018) note however, since the ratio of GDP to government spending tends to fluctuate a lot in samples, this method tends to bias the multiplier estimates upwards. This should be especially true for defence spending variables that exhibit large fluctuations at and around war related events. To deal with this issue, they propose using Gordon and Krenn (2010) transformation, wherein the variables entering the SVAR are not expressed in logarithms but rather they are divided by an estimate of trend GDP. This transformation directly puts all the NIPA variables in the same units, so that the multipliers can be estimated directly. For the purposes of our analysis, we employ a second degree polynomial estimate of trend GDP as the normalization variable.

Following these, the integral multipliers can be simply estimated as the simple cumulative ratio of output to the defence spending variable we are interested in:

$$M = \frac{\sum_{i=1}^{h} Y_i}{\sum_{i=1}^{h} G_i} \tag{2.7}$$

The estimates with our shocks yield the following results with the 84% Probability Density Intervals in parentheses:

Table 2.2: Output Multipliers

Horizon	Defence Investment	Defence Consumption	Total Defence Spending
4	-0.54	1.18	-0.09
	[-2.15,1.378]	[0.18,2.17]	[-0.37,0.16]
8	-1.67	0.78	-0.15
	[-3.89,0.94]	[-0.41,1.96]	[-0.48,0.19]
12	-2.11	0.59	-0.11
	[-4.6,1.25]	[-0.87,1.94]	[-0.5,0.28]
16	-2.17	0.41	-0.04
	[-5.34,1.84]	[-1.37,1.91]	[-0.49,0.4]
20	-2.05	0.29	0.04
	[-5.85,2.41]	[-1.77,1.93]	[-0.45,0.54]

The public investment multiplier is never significantly different from zero. Same is not true for the public consumption multiplier which is significantly positive and greater than one at least in the short term, up until period 4 and then is not significantly different from zero. This result hints towards a limited efficiency of the investment multiplier when compared with the consumption multiplier. For context we are also presenting the total defence spending multiplier in the third column. Notably, the multiplier of aggregate defence spending is never significantly different from 0, showcasing the lack of sufficient exogenous variation in US defence spending in the aftermath of the Korean War. Our estimates for the multipliers of the defence spending components are consistent with the closest paper to ours in the literature, Boehm's (2020) main conclusion [6], that when it comes to short term shocks, public consumption is superior to public investment.

Finally, and importantly to our initial question, we would like to more precisely test whether the defence investment and consumption multipliers are different. The Bayesian framework provides us with the flexibility to do this. For any given draw we calculate the difference

[6] The qualitative aspects of our estimates also match those in the literature. For reference, Boehm's (2020) estimated multipliers are presented in Appendix C.

between the multipliers and store it. We then compute the median of the response. In Figure 9 we then report the median of the differences between the defence consumption and the defence investment multiplier. A positive value indicates a higher defence consumption multiplier higher than the defence investment one.

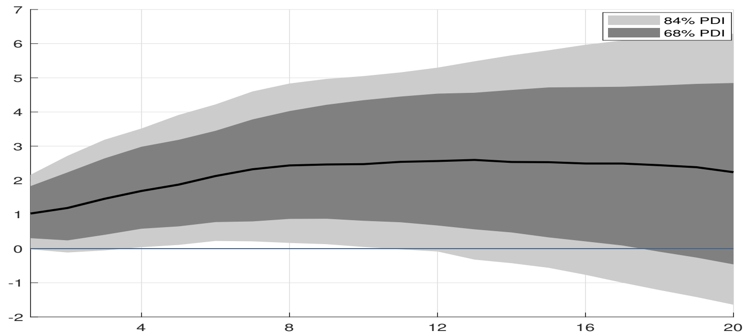

Figure 2.9: Difference Between Defence Consumption and Defence Investment Multiplier.

In the figure we see evidence that the two component of government spending are indeed statistically distinct objects. The defence consumption multiplier is larger than the defence investment one within the 68% posterior probability range, thus implying that fiscal variables should not be treated as aggregates. Note of course that this result comes with the caveat that the economy reacts the same to a defence spending shock as with a general spending shock. But even with this caveat, this is important evidence pointing to the fact that the components of fiscal spending having distinct effects on the economy.

2.7 Explorations with the identified shocks

What is the nature of the information contained in the identified shocks? Are we spuriously mixing the process of one defence variable with the process or the endogenous movement of the other? In this section, we perform exercises to gauge the informational content of the shocks; we first plot the time series of the identified shocks and check how they correspond to historical events associated with increased government spending. Secondly, by noting that the identified shocks can act as an instrument for the aggregate public spending variables, we can test their relevance by performing standard instrument tests. Thirdly, we employ a constrained

IRF method, as the literature in the past has suggested, to isolate 'pure' public consumption and investment shocks and see if this has any effect on our results.

2.7.1 Time series of identified shocks

Which historical events are our shocks capturing? Answering this question could provide us with some more context regarding the external validity of our shocks. Figure 2.10 plots the historical time series of the two identified shocks. We see a significant correspondence of the identified shock events in terms of timing. Despite this temporal correspondence, there are significant differences in the magnitude of the shocks. A dense cluster of public investment and consumption shocks, both negative and positive, is formed at mid '50s period, the period of the end Korean War and the subsequent escalation of the Cold War. The occurrence of shocks remains quite dense throughout the '50s and the early '60s, compared with the rest of the sample. This is to be expected since that period was associated with an escalation of the Cold War tensions.

Another cluster of shocks is formed from the mid '60s till the mid '70s, culminating in the early '70s. This corresponds to the period of the full military commitment of the US in the Vietnam War, beginning in 1965 with the Gulf of Tonkin incident and beginning to end in the early '70s until the fall of Saigon in 1975.

The '80s is another period of relative turbulence with the reescalation of the Cold War tensions and the military buildup program of president Reagan. This is mainly captured in the sharp positive upswing of defence investment. Finally, the '90s and the '00s are a period of relative stability in terms of military spending shocks with a few upswings in the '90s and one in the early 2000s probably associated with the beginning of the War on Terror.

The shocks we are identifying thus appear to be capturing information that can readily be associated with war episodes and/or political events. As we argued in Section 2 these events can be safely assumed to be exogenous to the economy. We can therefore conclude that with our process we are capturing the main drivers of exogenous variation of the defence spending components which implies that we are capturing the main sources of exogenous variation of the public spending components.

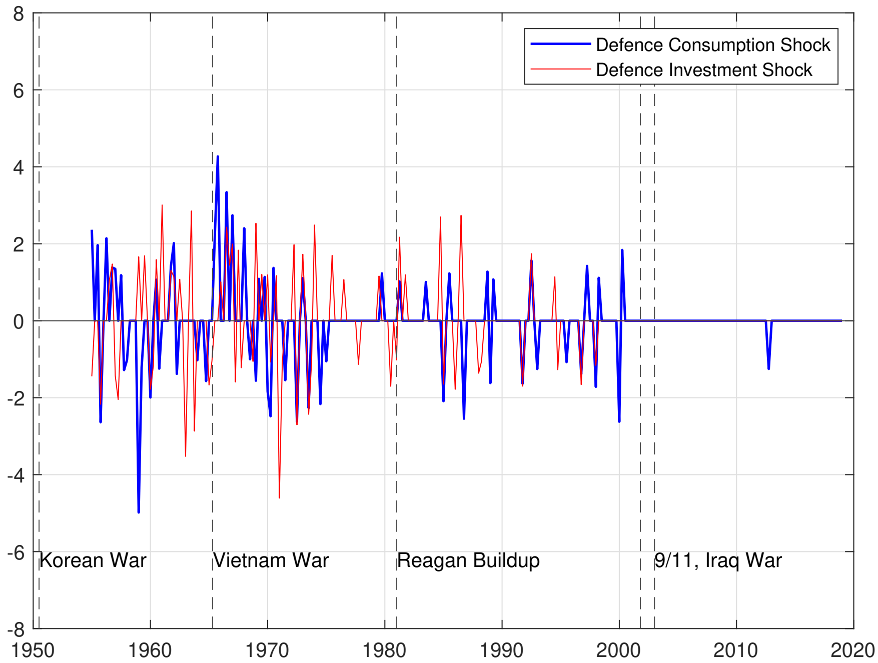

Figure 2.10: Time series of military consumption and investment news shock.
Notes: The value of the time series is equal to zero is it less than one standard deviation in absolute value and is equal to its actual value otherwise. Both series have been standardized.

2.7.2 Local Projections

Apart from the traditional SVAR method of producing impulse responses wherereby one iterates the reduced form coefficient matrix, on can also produce them using Jorda's (2005) local projection framework. In principle, local projections and the SVAR impulse responses are equivalent objects as Plagborg-Møller and Wolf (2021) recently showed so in this sense this exercise can also serve as a robustness test of our baseline estimates; if our VAR is not misspecified the SVAR IRFs and the local projections shouldn't be far apart.

Multipliers with Local Projections

Local projections is a framework bearing analogies to direct forecasting, that has been developed to guard against the dangers that a misspecified SVAR could create when constructing the impulse responses.

To see how local projections work assume that we have an identified shock ϵ_t. Then the impulse response to a shock at time t of our variable of choice y over horizon h is estimated by :

$$y_{t+h} = m_h \epsilon_t + controls + e_{t+h} \tag{2.8}$$

with the coefficient m_h being used as the estimate of our impulse response to the shock ϵ_t.

For our purposes local projections can be used to estimate the integral fiscal multiplier by reformulating the problem the following way :

$$\sum_{i=0}^{h} y_{t+i} = \beta_h + m_h \sum_{i=0}^{h} g_{t+i} + x_h(L)\gamma_{t-1} + trend + e_{t+h} \tag{2.9}$$

where the dependent variable y is the sum of real GDP and the independent variable g is the sum of government spending from t to $t + h$, and γ is a vector of control variables. The government spending variable g can be instrumented with our identified shock as Ramey and Zubairy (2018) note. Then, the estimated coefficient m is the integral multiplier for each horizon.

We proceed with this method using as control variables the 4 lags of output, the two defence spending variables and the two identified spending shocks. The multipliers this method yields presented in Appendix C are qualitatively similar to our baseline IRF estimates. The defence investment multipliers are statistically insignificant while the defence consumption multipliers in the initial periods are positive. This is a first hint that our estimates are robust to misspecification and we are efficiently estimating the two multipliers.

Instrument Relevance

By noticing that (2.9) is essentially an instrumental variable problem wherein we are instrumenting the relevant government spending variable with our identified shock, we can gauge the informational relevance of our shocks by constructing the first stage F-statistic for instrument relevance at each horizon. For the F-statistics significance thresholds, we use the Olea and Pflueger (2013) thresholds at the 5% and 10 % level. As a rule of thumb, F statistics below these thresholds indicate a possible problem with the relevance of our instruments. Figure 2.11 presents the results :

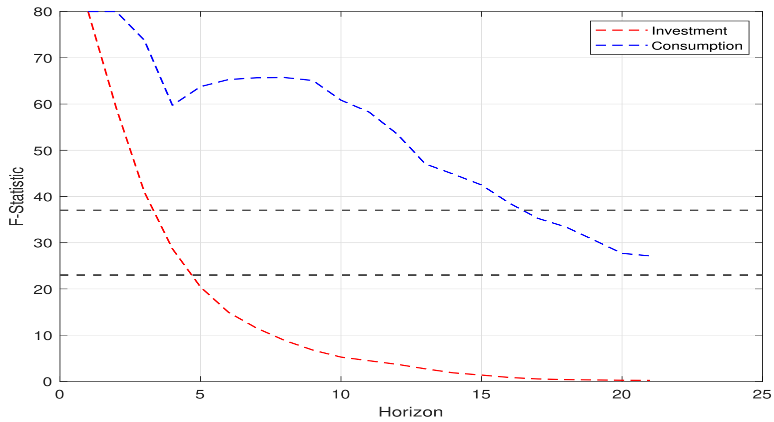

Figure 2.11: First-stage F-statistics for government spending shocks.
Notes: The F-statistics are based on the regression of the sum of government spending from t to t+h on the shock at t, plus the lagged control variables. Values above 80 have been capped at 80. The horizontal dashed lines are the Olea and Pflueger, 2013 5% worst case bias (upper line) and 10% worst case bias (lower line) thresholds.

How are our shocks faring in this exercise? Our defence consumption shock is showing high relevance throughout the sample. It surpasses the 10% threshold in most horizons, with the exception of the very last ones. The shock's relevance peaks at the initial horizons and progressively declines until it reaches its lowest point in the 20th horizon. This shows us that our identification is capturing both the surprise and the anticipated effects of the public consumption shock.

The defence investment shock also exhibits informational relevance. This relevance is however concentrated in the first 5 periods after the shock hits the economy. It is highly relevant, surpassing the 10% threshold from the impact until the 4th period in our horizon. Thereafter the relevance of the shock falls to statistical insignificance.

2.7.3 Constrained IRFs

The literature has suggested that to disentangle the effects of government consumption from government investment one could proceed in the following way; estimate an SVAR with the two government spending variables in it and construct the IRFs of the, say, government investment shock, by obtaining that shock from the SVAR and then constraint the IRF of government consumption. This way the 'pure' effect of the spending shock is captured (Perotti (2004), Ilzetzki et al (2013)).

This methodology of constrained IRFs however has its shortcomings; most importantly, it violates the Lucas critique in the sense that the resulting IRF should not be invariant to changes of government policy since it was obtained by constraining the effects government consumption. In this regard, it cannot be considered structural.

Nevertheless, this exercise would allow us to approximately gauge whether we are confusing the effects of the defence investment with the defence consumption shock and vice versa. We construct the constrained IRFs by calibrating a series of shocks to the variable we are interested in constraining that would keep its response fixed to zero throughout the horizon of the IRF. The technical details of this method, as laid out in Bachmann and Sims (2012), are presented in Appendix A. Figures 2.12 and 2.13 then present the results for the defence consumption shock with the defence investment response being constrained to zero.

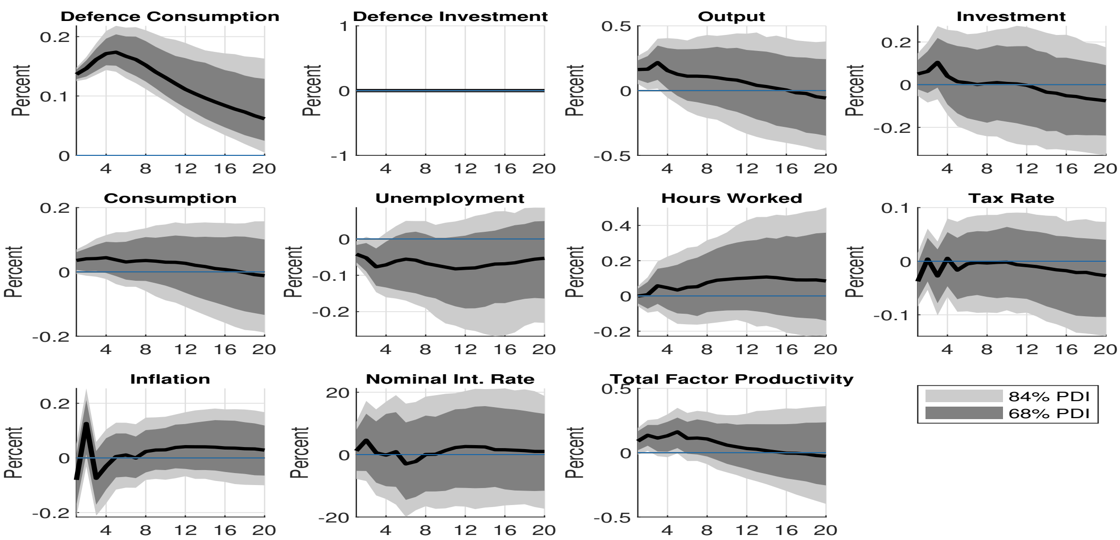

Figure 2.12: Counterfactual IRFs to a Defence Consumption Shock with Defence Investment constrained to zero.

For both of the 'pure' shocks considered, the reaction of the economy is virtually identical

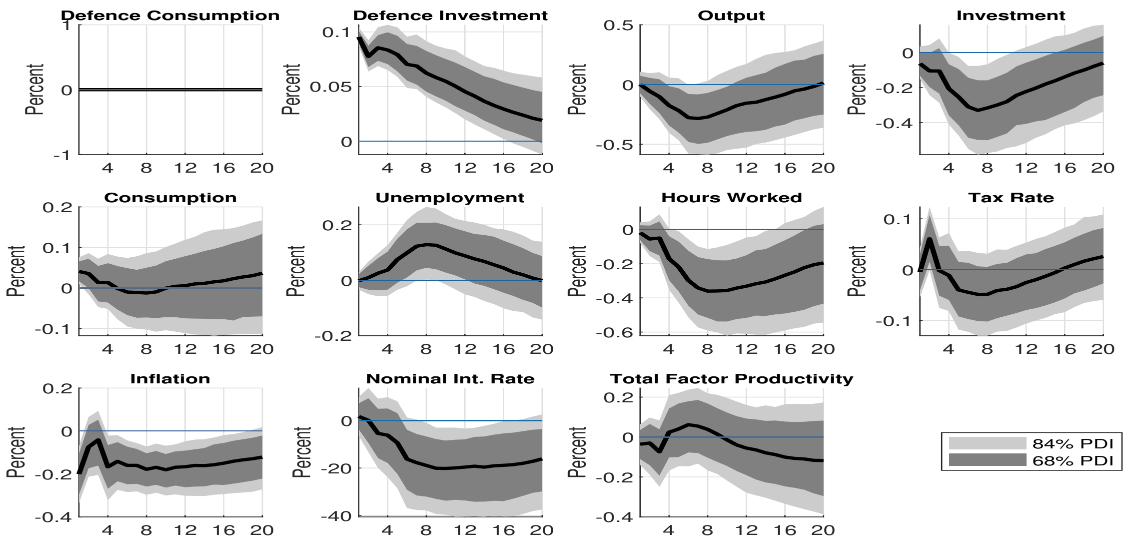

Figure 2.13: Counterfactual IRFs to a Defence Investment Shock with Defence Consumption constrained to zero.

with the benchmark case with all the qualitative and quantitative features being the same. This exercise thus suggests that our max share method without any constraints is enough to capture the pure public consumption and investment shocks without confusing the response of the one component with the response of the other.

2.8 Effects Driving Our Results

As already mentioned, theory suggests that a fiscal news shock process can be broken down into an anticipated and an unanticipated component. As Forni and Gambetti (2016) report both components can have sizable and significant effects. A significant constraint of our agnostic approach is that it is not providing us with a clear cut picture about which component is driving our results since the identification is not driven by any theoretical considerations. Although the IRFs and the instrument relevance test seem to attest to the fact that our shocks are mainly informative about the short and medium run horizons and don't convey much information on impact about the longer run, we would like to approach this a bit more systematically.

To this purpose, we perform the following exercise: we employ the agnostic maxshare approach, laid out in section 4.2, but this time extracting the shock that maximizes the unexplained variance over a subset of the initial 20 period horizon. Then, we construct the IRFs and for each subset of horizons we check which one matches better with our baseline. If the medium horizon IRFs match better with the baseline for example, this would imply that the effects of the shock start propagating over the medium term etc. In a sense, the philosophy of this exercise is very close to Dieppe, Francis, and Kindberg-Hanlon (2021) Non Accumulated Max Share (NAMS) approach. To get a more complete picture, along with the IRFs we also compare the FEVs from this exercise with the baseline FEVs.

We are targeting three horizon intervals; a short run (0-6 periods after the shock), a medium run (6-16 periods) and a long run (16-20). We are first considering the IRFs to a consumption shock. Figure 2.14 presents the results for the defence consumption shock;

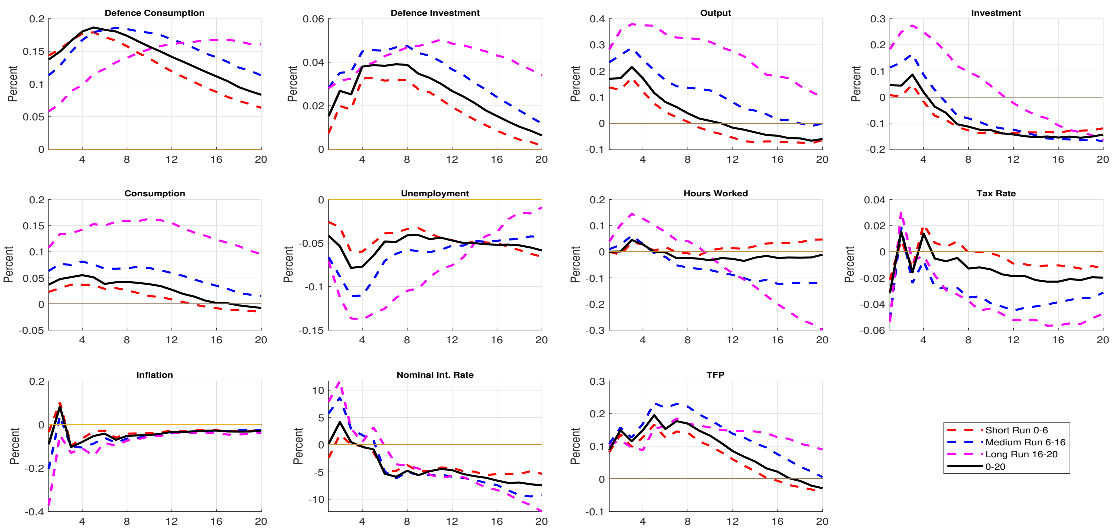

Figure 2.14: IRFs to a Defence Consumption Shock. Targeted at Different Horizon Intervals

The first thing that is noticeable from this figure is that the long run IRF is the one most distant from the baseline, suggesting thus that our method isn't capturing any process that propagates over the long run. On the other hand, the baseline IRF resembles very much the medium and mainly the short run IRF, suggesting thus that we are mainly capturing information about the short run horizon.

Turning to the defence investment shock, the same is true as with the Defence Consumption shock as show in in Figure 2.15. Namely, the short run IRF is the one closest to our baseline IRF, thus suggesting that we are mainly capturing information about the short run properties of the shock.

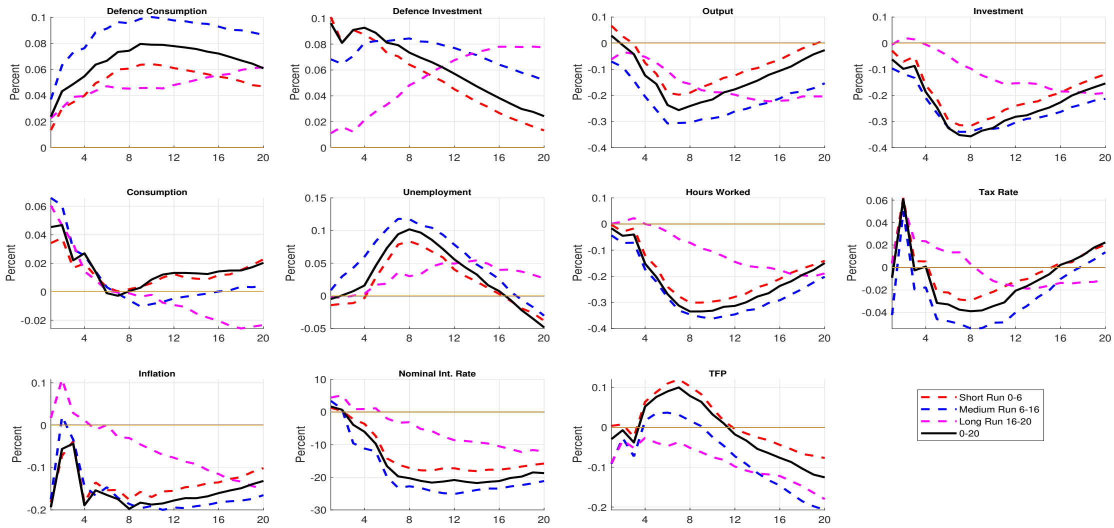

Figure 2.15: IRFs to a Defence Investment Shock. Targeted at Different Horizon Intervals

The FEVs of this exercise are presented in Appendix D.

In conclusion, the dominant process driving the two defence spending variables is a short run one, with a medium run one closely following, with the time ranges defined as above. In contrast, the shocks don't appear to be providing much information about the long run.

2.9 Guarding Against Confounding

A substantial danger with our identification is confounding different shock processes. This danger is inherent in our identification method. Since the two shocks we wish to identify stem from two almost perfectly comoving variables, whose exogenous variation is a result of a single choice of the policy maker, it could be expected that the shock process that contributes most to the unexplained variance of the variable of interest as solved by the max share algorithm could also be contributing to the variance of the perfectly comoving variable.

As a first parse to this question, we wish to plot the total FEV of the two shock processes under consideration; if the two processes jointly explain more than 100% of the variance of any of the variables of interest, then this is an indication of confounding and a misspecified model. To further guard against this concern we are going to perform the instrument relevance test of Section 7, this time to assess whether our identified shocks contain any relevant information for the other defence spending variable; that is, we are testing whether our consumption shock contains information about the public investment variable and vice versa.

Finally, to take the point home and guard against the possibility of the IV test results being misleading, we check the robustness of our qualitative results by considering an alternative identification scheme, devised to restrict any spurious mixing of shocks.

2.9.1 Total FEV

Figure 2.16 plots the FEVs of the combined defence consumption and defence investment shock. Taking into account inevitable measurement errors, our two shocks appear to be faring quite well; they both account for almost 100% of the variance of the two defence spending variables. This fact in combination with the single shock FEVs is a first indication that we are not confounding the processes.

Interesting in their own regard are the FEVs to the other variables. The total contribution of the two shocks to them remains quite low peaking at below 10 %. Exceptions to this trend are the inflation and the nominal interest rate where the response peaks at 20 %. These results are consistent with the literature (Ben Zeev and Pappa, 2017) which point to only a limited effect of the public spending shocks to the business cycle variation.

2.9.2 Instrument Test

If we are mixing up the defence consumption and defence investment processes, then the identified defence investment shock should contain information about defence consumption and vice versa. Therefore, to see if we are confounding the shock processes we can simply perform the IV relevance test outlined above, testing for the instrument relevance of the consumption shock to public investment and vice versa. A low relevance score in these tests should indicate

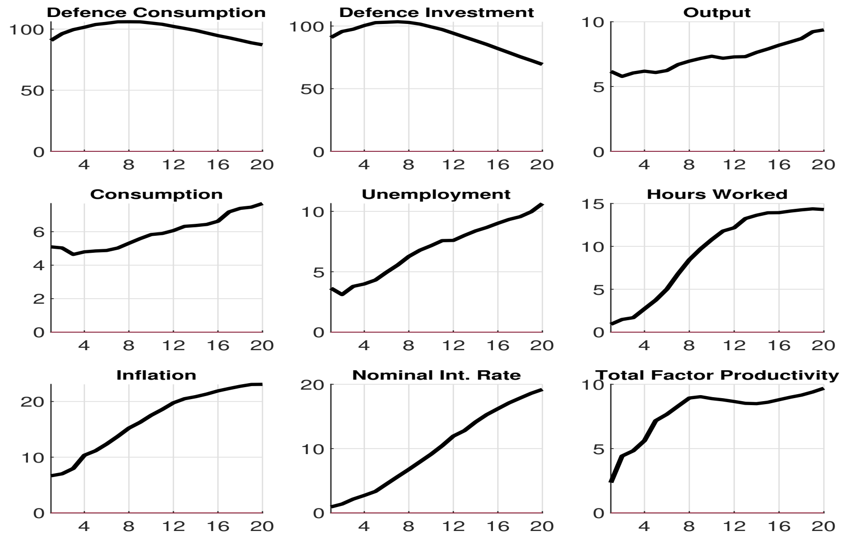

Figure 2.16: Variance explained by Both Defence Consumption and Defence Investment Shock. 90% Posterior Density Intervals.

that we are not confounding processes in our identification. Figure 2.17 presents the results of this exercise:

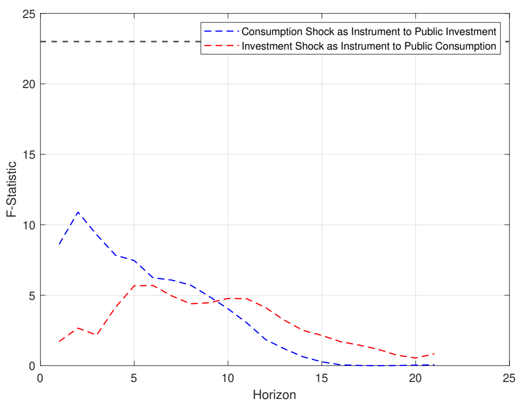

Figure 2.17: First-stage F-statistics for government spending shocks.
Notes: The F-statistics are based on the regression of the sum of government spending from t to t+h on the shock at t, plus the lagged control variables. V The horizontal dashed line is the Olea and Pflueger, 2013 10% worst case bias (lower line) threshold.

Both F-statistics are quite lower than the Olea and Pflueger (2013) threshold, thus indicating that the shocks don't contain any relevant information for the similar variables. We can therefore conclude with a margin of relative safety that we are not mixing up the shock processes.

2.10 Conclusions

In this paper, we propose an identification of public consumption and public investment shocks. In the spirit of Uhlig (2003), Ben Zeev and Pappa, 2017 and Angeletos et al (2020) we are employing an agnostic identification method, utilizing the unexplained variation of the variables of interest from a large SVAR containing all the main macroeconomic aggregates. The necessary exogeneity is derived from utilizing defence spending variables in the position of the general government spending variables, under the assumption that war events are exogenous to the economy.

Our results for the period 1954 to 2019, that is a period without large war events, suggest that the components of government spending are statistically distinct with each having each own idiosyncratic effect on the economy. In particular, we find that the defence consumption shock has stimulative effects on the economy, raising output and private consumption and lowering unemployment in the short run. On the other hand, the defence investment shock and its associated multiplier is never statistically different from zero and is associated with a substantial drop in private investment. Utilizing the dynamic multipliers, we find that the defence consumption multiplier is statistically higher than the defence investment multiplier. Additionally, the defence investment multiplier is never significantly different from zero, a finding that comes in agreement with key facts in the literature about the public investment multiplier.

Our shocks do a good job in explaining the variation in public investment and consumption in short and medium term horizons as shown with the relevant FEVs and with the instrument relevance tests we conducted. Importantly, we also demonstrate that we are not confounding shock processes in our identification, a serious concern given the nature of our identification process and the variables used.

Our results have implications for the fiscal multiplier literature both at the theoretical and the empirical front; if the components of public spending have distinct effects then this implies that the size of the general fiscal multiplier is dependent on its composition. In other words, many multiplier estimates might be suffering from external validity. On the theoretical side of things, this has implications on how researchers are modelling the fiscal spending; the modelling of public investment should be treated separately from the modelling of public consumption. On the empirical front, these results imply a lack of external validity to the until now majority of multiplier estimates with the composition of the multiplier playing a role in the effects it will have on the economy.

2.11 Technical Appendix*

2.11.1 A. Bayesian Estimation

Following along the lines of the literature on Bayesian estimation of VAR models, we estimate the posterior distribution of our model with a Gibbs sampling algorithm. We use a Minnesota prior for the prior coefficient distribution. In this appendix, we discuss the specifics of the estimation process.

The linear model is specified as follows;

$$Y_t = B(L)Y_{t-1} + u_t \tag{2.10}$$

To estimate it we specify a Minnesota prior for the coefficient matrix $B(L)$. The posteriors are then drawn from a Gibbs sampling algorithm. More formally, the process is the following;

1) We set the priors for the covariance matrix Σ and the coefficient matrix B with a standard Minnesota prior. We assume that the coefficient matrix follows a normal distribution $p(B) \sim N(\tilde{B}, H)$ and and the covariance matrix an inverted Wishart distribution $\Sigma \sim IW(\tilde{S}, V)$. Set a starting value for Σ. The hyper-parameters that control the priors take the following values; $\lambda_1 = 0.2$, $\lambda_2 = 0.5$, $\lambda_3 = 1$, $\lambda_4 = 10^5$.

2) We first sample the VAR coefficients from the conditional posterior distribution specified as follows; $F(b|\Sigma, Y_t) \sim N(B^*, V^*)$. The mean of the distribution B^* is defined as:

$$B^* = (H^{-1} + \Sigma^{-1} \otimes X_t' X_t)^{-1}(H^{-1}\tilde{B} + \Sigma^{-1} \otimes X_t' X_t \hat{B})$$

and V^* is defined as

$$V^* = (H^{-1} + \Sigma^{-1} \otimes X_t' X_t)^{-1}$$

Once $B*$ and $V*$ have been calculated and the distribution has been parametrized, the VAR coefficients are drawn from the posterior distribution:

$$b^1 = B^* + [b \times V^{*1/2}]$$

3) The next step is to draw the covariance matrix Σ from its conditional Inverted Wishart distribution, conditional on the draw of the coefficient matrix in step 2.

4) Repeat steps 2 and 3 20000 times. We keep the last 4000 draws from these iterations to form the empirical distributions of the parameters of interest. This is the so called burn in period. The last draws are used to calculate impulse responses and any other structural analysis the researcher might wish to perform.

2.11.2 Constructing The Counterfactual IRF

The process of constructing the counterfactual IRF is as follows; the main idea is to construct a second shock process that offsets the response of the variable of interest to the original identified shock. To describe the process more formally, suppose we have solved for the structural VAR and we have a structural representation of this form;

$$A_0^{-1} Y_t = B(L)\epsilon_t$$

where ϵ_t is the vector of structural shocks and A_0 is the impact matrix such as $A_0\epsilon_t = u_t$, with u_t being the vector of the reduced form VAR residuals.

The impulse response is constructed with the matrix A_0 that governs the instantaneous relationships among the model variables. Suppose we wish to hold the response of the i variable to a shock of the n variable at zero. To do this then, we will need to impose the restriction on the element $A_{0(k,n)} = 0$. To implement this restriction, we will calibrate the element ϵ_i to zero out the effect of the shock of the n variable on variable i.

Suppose we have the structural form of the system:

$$Y_t = \sum_{j=1}^{p} A_0 A_j Y_{t-j} + A_0\epsilon_t$$

which can be more compactly be written as a VAR(1) by writing it in companion matrix form where $Z_t = [Y_t Y_{t-1} ... Y_{t-p-1}]$ and Λ collects the coefficient matrices:

$$Z_t = \Lambda Z_{t-1} + A_0\epsilon_t$$

Now, let e_i be a selection row vector $1 \times n$ with a one in the i_{th} place and zeros elsewhere. Let $A_0(q)$ be the q^{th} column of A_0. The impulse response of variable i to structural shock q at horizon $h = 1, ..., H$ is:

$$\Phi_{i,q,h} = e_i \Lambda^{h-1} A_0^{-1}(q)$$

We hold the variable of interest i constant to a shock at each forecast horizon by setting $\Phi_{i,q,h} = 0$. We do this by creating a hypothetical sequence of shocks ϵ_k so as to force this relationship to hold in every relevant horizon. To do this the shock needs to be calibrated as such that it offsets the impact of the structural shock derived from the process of variable n. In more formal terms:

$$A_0(i,n) + A_0(n,n)\epsilon_{m,n} = 0 \equiv \epsilon_{m,n} = -\frac{A_0(i,n)}{A_0(n,n)}$$

for the first horizon shock. In the following horizons the subsequent shock can be calculated recursively as:

$$\epsilon_{m,n} = \frac{\Phi_{i,q,h} + \sum_{j=1}^{h-1} \epsilon_n \Lambda^{h-j} A_0(n)\epsilon_{n,j}}{\epsilon_n A_0(n)}, h = 2, ..., H$$

Given this sequence, we can then compute the modified impulse responses of the variables in the system to the structural shock of interest as :

$$\Phi_{i,q,h} = \Phi_{i,q,h} + \sum_{j=1}^{h-1} \epsilon_n \Lambda^{h-j} A_0(n)\epsilon_{n,j}$$

2.12 Data Appendix

The data is from the Federal Reserve Economic Database (FRED). TFP is the utilization adjusted TFP produced by Fernald (2012). Table 3 describes the original data and the transformations applied.

Table 2.3: Description of Data

Data	Mnemonic	Freq.	Transf.
Government consumption expenditures: Federal: National defense	A997RC1Q027SBEA	Q.	-
Federal Government: Real National Defense Gross Investment	DGI	Q.	-
Gross government investment	A782RC1Q027SBEA	Q.	-
Gross Domestic Product	GDP	Q.	-
Government Consumption Expenditures and Gross Investment	GCE	Q.	-
Personal Consumption Expenditures: Nondurable Goods	PCND	Q.	-
Personal Consumption Expenditures: Services	PCESV	Q.	-
Gross Private Domestic Investment	GPDI	Q.	-
Personal Consumption Expenditures: Durable Goods	PCDG	Q.	-
Population Level	CNP16OV	M.	EoP.
Employment Level	CE16OV	M.	EoP.
Unemployment Rate	UNRATE	M.	Avg.
Nonfarm Business Sector: Average Weekly Hours	PRS85006023	Q.	-
Gross Domestic Product: Implicit Price Deflator	GDPDEF	Q.	-
3-Month Treasury Bill: Secondary Market Rate	TB3MS	M.	Avg.
Total Factor Productivity (Growth Rate)	dtfputil	Q.	-
6th Degree Polynomial Real GDP Trend	-	-	-
Average Marginal Tax Rate	AMT	Q.	-
Nonfarm Business Sector: Hours Worked for All Employed Persons	HOANBS	Q.	-

Note: Q: Quarterly, M: Monthly, EoP: end of period, Avg: quarterly average.

Table 2.4: Variables in the VAR

Real Defence Consumption per Capita	`DC=100*log(A997RC1Q027SBEA)`
Real Defence Investment per Capita	`DI=100*log(DGI)`
Real Government Spending per Capita	`DC=100*log(GCE)`
Real GDP per Capita	`Y=100*log(GDP)`
Real Investment per Capita	`I=100*log(PCDG+GDPI)`
Real Consumption per Capita	`C=100*log(PCND+PCESV)`
Unemployment Rate	`UNRATE = UNRATE`
Labor Hours (LH)	`LH = 100*log(PRS85006023*CE16OV/CNP16OV)`
Average Marginal Tax Rate (AMTBR)	`AMTBR=100*AMT`
Inflation	`INF = 100*log(GDPDEF/GDPDEF(-1))`
Interest Rate	`TBILL=TB3MS/4`
Utilization Adjusted TFP	`TFPU=100*cumsum(dtfputil/400)`

Notes: All variables have first been converted from nominal (`N`) to Real Per Capita (`RP`) terms with `RP=(N/CNP16OV)*1000*(1/GDPDEF)`.

2.13 Appendix C

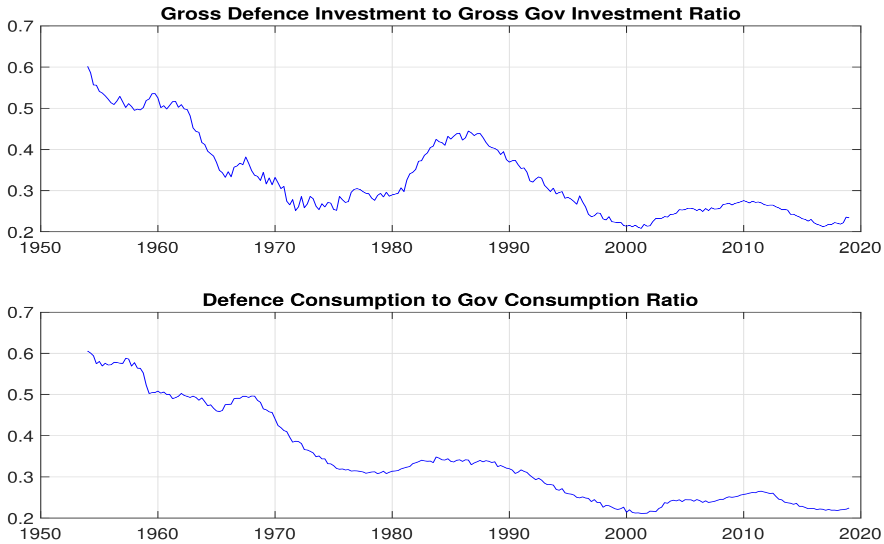

Figure 2.18: Military Spending To Government Spending Ratio

Table 5 presents Boehm's (2020) multiplier estimates for each component of government spending.

Horizon	Public Investment	Public Consumption
4	-0.08	0.76
	(0.22)	(0.35)
8	-0.06	0.85
	(0.36)	(0.35)

Table 2.5: Boehm's (2020) Multipliers with Local Projections

Table 6 presents the multipliers obtained from our exercise when the impulse responses are constructed with Jorda's (2005) Local Projections (LP) framework.

Horizon	Defence Investment	Defence Consumption
4	-2.3	1.38
	(3.6)	(0.63)
8	-5.9	1
	(4.69)	(0.68)
12	-7	0.98
	(7.4)	(0.73)
16	-9.5	0.77
	(20)	(0.85)
20	-8	0.4
	(30)	(0.95)

Table 2.6: Fiscal Multipliers with Local Projections
Notes: The values in parentheses under the multipliers give the HAC robust standard errors.

2.14 Appendix D

2.14.1 Impulse Responses With Log Normalization

In this section we present the responses to our shocks when the literature standard log normalization is considered. The results remain qualitatively unaltered, suggesting that the normalization is not affecting the qualitative impact of our shock on the variables in our econometric system.

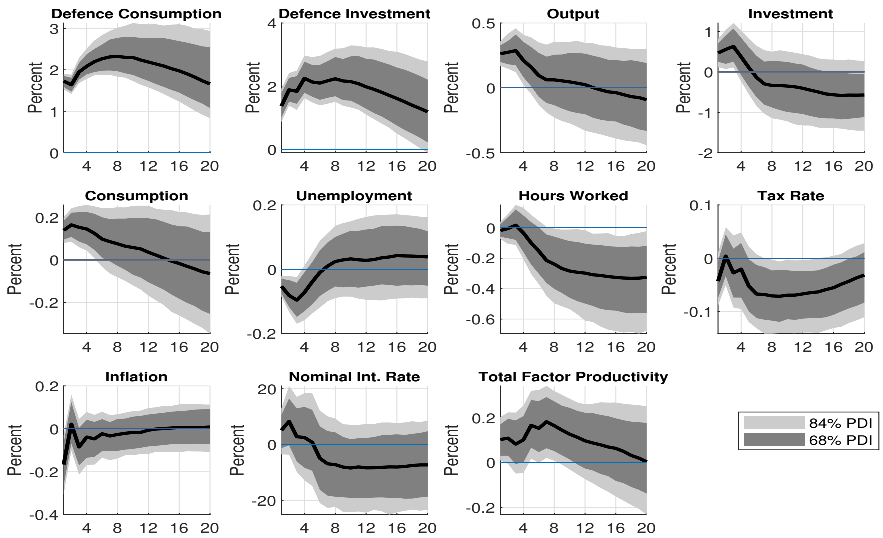

Figure 2.19: Impulse Response Functions to A Defence Consumption Shock; Variables Normalized with Logs.

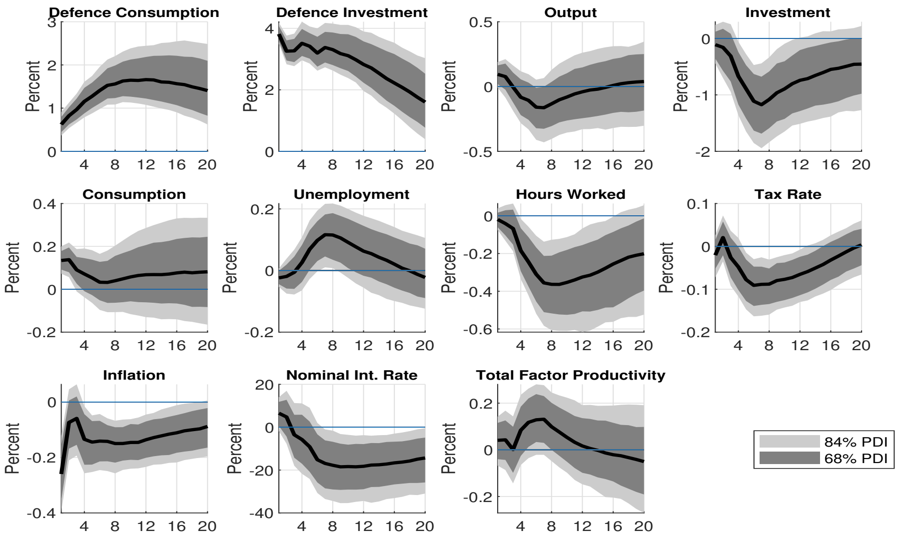

Figure 2.20: Impulse Response Functions to A Defence Investment Shock; Variables Normalized with Logs.

2.14.2 Forecast Error Variance Decompositionsfrom the Non Accumulated Max Share Exercise

In this section we present the Forecast Error Variance Dempositions for the NAMS exercise of section 8.

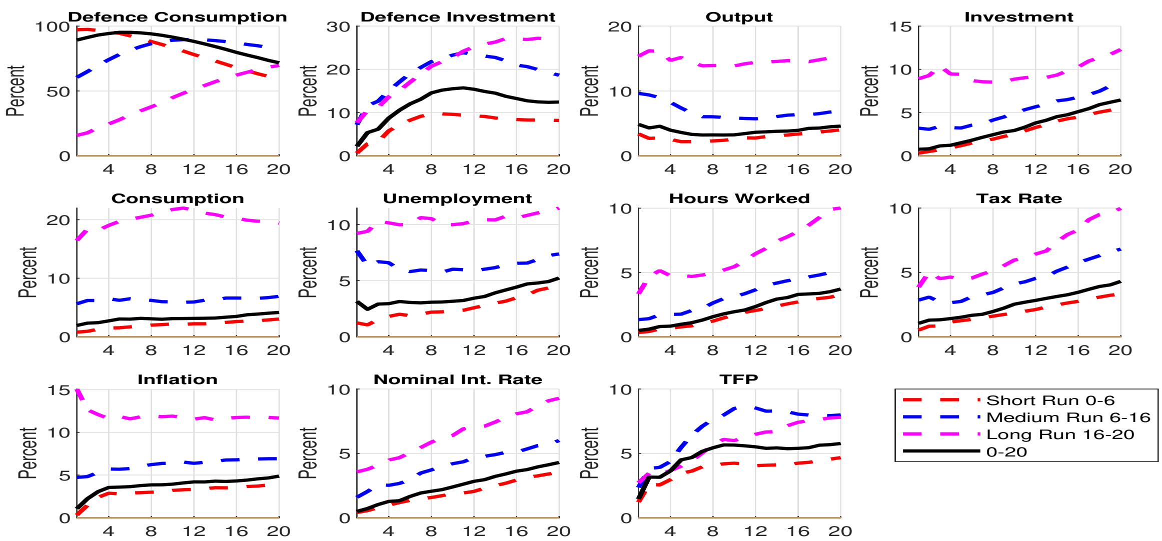

Figure 2.21: FEVs to a Defence Consumption Shock. Targeted at Different Horizon Intervals

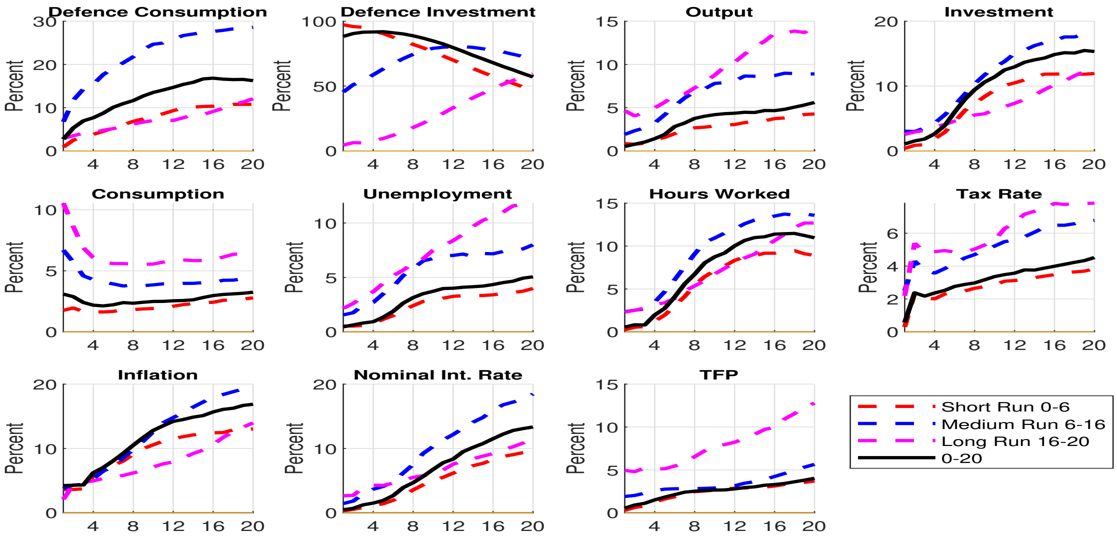

Figure 2.22: FEVs to a Defence Investment Shock. Targeted at Different Horizon Intervals

Chapter 3

Government Consumption and Investment Multipliers at the Zero Lower Bound; Evidence From Two Wars

3.1 Introduction

How large are the output multipliers in response to a shock in public investment or consumption when the nominal interest rates are at the Zero Lower Bound (ZLB)? Does the composition of government purchases matter more in ZLB times? How do shocks in the components of government purchases propagate and do the propagation channels vary across different states? Should one distinguish between the two types of public spending?

In this paper we attempt to engage the question of whether the components of government spending have different properties when the economy faces a ZLB constraint; does the composition of fiscal spending matter and does it exhibit a different behaviour in ZLB times.

A wide array of theoretical considerations can motivate our question; the issue of the size and properties of the fiscal multiplier has been extensively scrutinized in the literature. While in the canonical Real Business Cycles (RBC) framework (see Baxter and King (1993)) the size of the multiplier is independent of the constraints of the monetary policy, within the

New Keynesian framework there are lots of good reasons to believe that the multiplier would be larger in times when the monetary policy is constrained by the ZLB. The mechanism is relatively straightforward; in New Keynesian models the increase in output in response to a fiscal shock first goes through an increase in private consumption and investment. If the economy had hit a ZLB constraint, the increase in private consumption and investment would be more pronounced; the constrained response of the real interest rate would ameliorate any crowding out effects that the interest rate might have had on private consumption and investment. This would ultimately lead to a greater increase in output and a greater fiscal multiplier. The underlying assumption here is that a key feature of the reaction private spending is its sensitivity to movements of the real interest rate. Standard medium scale New Keynesian models such as Woodford (2011), Eggertsson (2011) and Christiano, Eichenbaum, and Rebelo (2011) all provide theoretical evidence of how the multiplier could in this way rise above one. In another sub-strand of the New Keynesian literature, wherein the ZLB period is caused by a non fundamental confidence shock (Mertens and Ravn (2014)), government spending is deflationary, increasing the real interest rates and thus reducing private consumption and investment and resulting in a multiplier lower than one.

Taking the issue of estimating fiscal multipliers one step further one can ask whether the composition of government spending matters for the size and properties of the fiscal multiplier. Or to be more specific, do the main components of fiscal spending, that is public investment and consumption, have distinct effects on the economy? This is a question that can be asked for both the linear and the non linear, for our purposes the ZLB, case. At the theoretical front, some earlier literature has cited the productivity of government capital as the main motivation to distinguish among public investment and consumption multipliers (Aschauer (1989b) and Aschauer (1989a)). In this line of reasoning, public investment is mostly associated with its infrastructure spending part. More recently, Ramey et al. (2020) provides some support to the hypothesis that infrastructure spending has beneficial long run effects that make it distinct from public consumption spending. At a more general level and considering public investment as a whole, Boehm (2020) provides a model as to why government investment might in fact have a lower output multiplier in the short to medium run than government consumption, with the main culprit being the high intertemporal substitution of private investment leading to strong crowding out effects induced by a rise in the real interest rate.

With our paper we wish to take these theoretical considerations onto the data. In doing so, we also seek to propose novel ways of addressing several challenges associated with estimating fiscal multipliers. Estimating fiscal multipliers at the ZLB is notoriously difficult. ZLB episodes are being exceedingly rare and not barely persistent enough in the post war advanced economies to allow the econometrician to draw any reliable conclusions. Another challenge is that those episodes generally coincide with recessions thus rendering it hard for the researcher to separate typical slack from the ZLB effects. On a high level, these can be regarded as informational deficiency problems. It follows that generally, the way to tackle the main problems associated with the ZLB is through a careful choice of sample. In this regard, a seminal contribution is by Ramey and Zubairy (2018) who utilize a large quarterly historical dataset with US data, exploiting the variation induced by three Wars (WW1, WW2 and the Korean War) to estimate the properties of the fiscal multiplier at the ZLB. In another approach, Miyamoto, Nguyen, and Sergeyev (2018) use a novel dataset from Japan, the only advanced economy that experienced a protracted ZLB episode in the post war period.

When it comes to the issue of empirically distinguishing among the different components of government spending, little empirical work has been conducted. Boehm (2020) provides some evidence that the public investment multiplier is smaller than the public consumption multiplier for short lived shocks. Perotti (2004) and Ilzetzki, Mendoza, and Végh (2013) also provide some evidence for the quantitative inferiority of public investment. An additional challenge, acknowledged in Ilzetzki, Mendoza, and Végh (2013) is the strong comovement between public investment and consumption which entails the danger of confounding the shock processes in the identification. With this paper we contribute to the empirical literature in two ways: firstly we utilize an extended historical dataset to tackle the data availability problems associated with the ZLB estimation. Secondly, we propose a novel agnostic identification that minimizes any potential confounding between the shock processes of the two highly correlated components of government spending in our extended sample.

Our dataset therefore includes two large war episodes that the US was engaged in; the World War 2 (WW2) and the Korean War. The large variation those episodes induced on the macroeconomic variables comprise the main source of information in our identification. Employing a Threshold VAR (TVAR) to capture the non linear data generating process wherein the econometric model changes according to a predetermined threshold, our starting point for the identification is the Structural Vector Autoregression (SVAR) framework whereby the identified shocks are linear combinations of the residuals of a Vector Autoregressive system. We utilize the agnostic identification approach that we employed in Chapter 2; our main assumption is that shocks extracted and constructed from the unaccounted variation of defence consumption and investment are exogenous and should serve as good approximations to a generic public consumption and public investment shock.

To minimize any potential confounding between the two processes, we employ an augmented version of Uhlig (2003) max share algorithm that explicitly takes into account potential confounders in the extracted shock processes. To our knowledge, a first version this algorithm was introduced by Ben Zeev and Pappa (2017). It addresses the dangers of confounding by explicitly looking for the vector that is maximizing the variance contribution to the variable of interest and at the same time is minimizing the contribution of the related variable.

Our main findings are as follows; public investment, as approximated by our defence investment shock, is the component of government spending more liable to state dependence in the form of the ZLB. In contrast, while we find a difference among public consumption multipliers across states, we argue that this difference is due to the idiosyncrasies of our sample and thus we find no concrete evidence of state dependence for defence consumption. In accordance with the theoretical and empirical literature on the topic, we present evidence that the mechanism behind the state dependence of public investment at the ZLB is the high elasticity of private investment to movements in the real interest rate. Additionally, we present evidence that public consumption and public investment shocks might be of inherently different nature in terms of how anticipated they are; a public consumption shock appears to not always be perfectly anticipated by the agents and has many elements that make it more akin to a surprise shock. In contrast, we find the public investment shock to be mostly anticipated, probably due to legislative and implementation lags associated with public investment projects.

Our paper is structured as follows; section 2 provides a review of the current state of research on the topic. In section 3 we delve into the details of our dataset and the intricacies associated with it. Section 4 lays out our econometric strategy and the augmented max share problem we solve for the identification to minimize concerns about confounding. Section 5 is the heart of our paper and presents the baseline structural results. In section 5 we also delve deeper into the propagation mechanisms of the shocks we observe. In section 6 we provide a discussion about the anticipated or unanticipated nature of the retrieved shocks. Finally, in section 7 we present our conclusions.

3.2 Literature

With our paper we contribute to the large body of work on estimating government spending multipliers. Seminal work in this research agenda is the paper of Blanchard and Perotti (2002) who utilized assumptions about the quarterly timing of government spending together with institutional information on the elasticities of the automatic stabilizers to isolate the automatic response to changing economic conditions and by implication isolate the discretionary component of public spending. In the past decade, the literature expanded significantly and a variety of different approaches have been utilized to identify a government spending shock ranging from excess stock returns (Fisher and Peters (2010)) to narrative war dates (Ramey (2011)) and utilizing the information of military variables (Ben Zeev and Pappa (2017)). A significant leap in the state of knowledge was the recognition of the influence fiscal foresight has on the transmission of the shocks [1] (see among others Ramey and Shapiro (1998), Fisher and Peters (2010), Forni and Gambetti (2016), Caggiano et al. (2015)) and by implication the recognition

[1]This is related to a more general argument about the importance of anticipated shocks in the business cycles. The archetypal anticipated shock is a technology one, as has been argued by Barsky and Sims (2011), Görtz, Tsoukalas, and Zanetti (2022), Gambetti et al. (2022) among others.

that models relying on current shocks are prone to suffering from non fundamentaleness (Forni and Gambetti (2014)), that is these shocks do not embed all the information processed by rational agents in their decision making.

A limited substrand of this literature concerns the importance of the composition of public spending packages and more specifically whether public consumption is statistically distinct from public investment. This is an important question, essentially asking whether the estimated multipliers suffer from external validity in the case they are conditional on the composition of the fiscal package. While in the public sphere the notion that, at least in the long run, the public investment multiplier is more potent than the public consumption one due to infrastructure spending, the idea being that infrastructure projects eventually result in productivity gains, this idea has been challenged by the literature in recent times. More specifically, in the short to medium term, there is the notion that public investment crowds out private investment, at a much higher degree than public consumption crowds out private consumption, and as a result a public investment shock might result into recessionary effects. Towards this direction, important empirical contributions have been made by Perotti (2004), who finds no difference between the two public spending multipliers, and Boehm (2020) who finds evidence pointing towards a limited effectiveness of public spending. In contrast, Auerbach and Gorodnichenko (2012), utilizing an SVAR framework and a Blanchard and Perotti (2002) identification scheme, report a higher public investment multiplier.

On another tangent relating to the estimation of fiscal multipliers, the fact that the size and the propagation of fiscal shocks might be highly conditional on the state of the economy, while highly intuitive, has only recently began getting formalized into a rigid empirical and theoretical framework. In the past decade a few papers have tackled the issue of estimating the multiplier across different states of the economy. The state dependence is predominantly modelled either as a recession/expansion or a ZLB/ non ZLB period.

Regardless of the researcher's choice of the model, estimating multipliers at the ZLB state is an arguably more challenging task than estimating them in recessionary times. The reason is that ZLB periods are generally rare and not barely persistent enough to contain enough information that would allow for a robust extraction of a structural shock. This data deficiency issue has been tackled either via utilizing long historical datasets (Ramey and Zubairy (2018)) or data from Japan (Miyamoto et al (2018)), a country with a long ZLB episode in the '90s.

The ZLB has been studied extensively in the theoretical literature. It plays an important role in medium scale DSGE New Keynesian models and its interaction with a government spending episode creates interesting dynamics (see for example Woodford (2011),Eggertsson (2011), and Christiano, Eichenbaum, and Rebelo (2011)). In this sense, our paper is related to this body of literature, testing whether the predictions of New Keynesian models hold to the data. At the ZLB the lack of reaction of the real interest rate to the government spending shock implies that there is no crowding out of private spending and thus the multiplier should be higher. Boehm (2020) makes this hypothesis more particular, arguing that this state dependence is especially true for the public investment component of fiscal spending, since private investment might be exhibiting a higher elasticity to interest rate movements than private consumption.

3.3 Data

3.3.1 War As The Main Source Of Exogenous Information

The empirical staple of modern macroeconometric analysis is the post World War 2 US data. The majority of empirical analyses engage with some subsample of the post war US aggregates. This data is indeed adequate for the inference of statistically strong conclusions for a wide array of macroeconomic questions. Of more relevance to our issue of estimating fiscal multipliers and as we showed in Chapter 2, the post Korean War defence spending variables, that is the aggregates from 1954 onwards, contain enough information to allow us to extract and create from their residuals strong instruments for public consumption and investment. This sample space however might not be enough for our present task of extracting conclusions about the properties of the components of government spending in the ZLB state. Any robust estimate would require that the state not only spans a sufficient portion of the sample but also that the exogenous changes, in our case the exogenous changes of defence spending, are sizeable enough to be identifiable.

The only US post Korean war ZLB episode taking place in the period 2007-2015 is of sufficient persistence but is not associated with big exogenous changes in the defence spending variables thus potentially rendering any extraction of structural spending shocks statistically fragile.

To meet this challenge a sample is needed that contains both a persistent ZLB episode and a sizeable enough exogenous variation in the defence spending variables. We therefore wish to add to our analysis the informationally rich 12 year period from 1939 to 1951. To gauge the informational content of this period, Figure 3.1 plots the defence consumption and defence investment starting from 1939 to 2019 at a quarterly frequency. The construction of these variables for the period 1939-1947 is described in more detail in section 3 and in the Data Appendix of the present chapter. From the graph we can derive some interesting information regarding the period 1939-1951. First of all, the defence expenditure shock experienced in that period is of immense magnitude, dwarfing anything that followed. This is an observation that also shows in the magnitudes of the impulse responses across states as we will see in the results sections. Indeed, the US in that period was operating a war economy, wherein all the productive activities of the society were concentrated around the war effort. Apart from a huge, very arguably exogenous, defence spending shock, an additional reason we chose to include this period is that during those 12 years the nominal interest rate stayed relatively constant. Indeed, Ramey and Zubairy (2018) classify this period as a ZLB period. This classification becomes more apparent if one sees the deviations of the nominal interest rate from the Taylor Rule prescriptions, showcased in the Data Appendix. Therefore, the whole decade of the '40s provides us with information on a huge government spending shock in a time of protracted monetary accommodation, thus rendering it appropriate for our purposes.

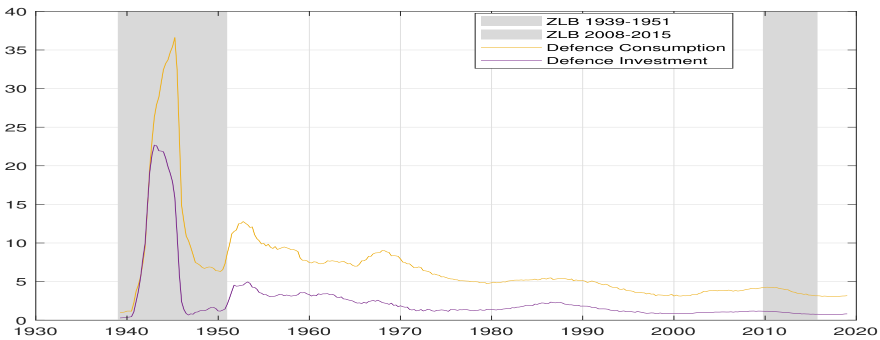

Figure 3.1: Log Real Defence Spending Variables and ZLB periods 1939-2019.

The richness of the information contained in a sample including the 1939-1951 period is undeniable, but there are also challenges associated with this choice that the researcher needs to take into account. This was a period of extraordinary circumstances. The US economy was exiting the Great Depression and transforming itself into a full blown war economy wherein almost all of its productive capabilities resources were being allocated to the war effort. Coupled with the large variations in defence spending, huge variations were also observed in all other major macroeconomic aggregates. This naturally begs the question of whether conclusions extracted from such a sample are generalizable and whether they can help us draw conclusions and analogies to the normal business cycle experience. In other words do results stemming from an identification overly relying on 'abnormal' war events have external validity?

There are additional challenges regarding the possible confounders in such a sample space. Ramey and Zubairy (2018) provide some elegant considerations to questions regarding the use of the '40s as an econometric source of information; consider that with the '40s, being the main source of information, apart from coinciding with the last phases of the Great Depression, one can think of the price rationing practices that were widely introduced during the war or the sense of patriotic duty that might have been pushing the labor productivity upwards. At a very high level however, despite the aforementioned issue, these periods are the closest a macro econometrician can have to natural experiments, that is events outside of the endogenous workings and feedback loops of the economy with effects sizeable enough on the economic aggregates to be informative about the nature of many economic processes. Thus, not utilizing this period would be equivalent with throwing valuable information away. More specifically to the US case, Barro and Redlick (2011) argue that its WW2 experience provides an excellent opportunity to study the effects of government spending shocks. A number of reasons are cited; the changes in defence spending in that period are exogenous and very sharp, including both positive and negative values, and more importantly the US was the only country that did not face a huge loss of human life and physical capital, thus rendering demand effects dominant in the study of fiscal spending. These considerations are equally, if not more, true for the Korean War, a war of much more limited destruction for the US but of also a great spending shock.

On a similar tangent and more specifically to our case, the '40s being the main period of exogenous variation, apart from the identification problem the state of a war economy creates, naturally relates our exercise to the question of what took the US economy out of the Great Depression. Accounting for such a state dependence is out of the scope of this paper, but the literature converges in the main culprits being the lax monetary policy and the explosion in defence spending. Instead of taking a stand on this question, one angle to look into our exercise is to think of it as a questions of whether it was the interaction between the monetary policy and the defence spending that mattered rather than them being autonomous components.

The choice of this dataset also introduces one more subtle issue related with the identification of shocks that we explicitly address with our novel identification method. The agnostic max share method to construct shocks, as introduced by Uhlig (2003) and utilized by Angeletos, Collard, and Dellas (2020) relies on the assumption that the variation of an exogenous variable not explained by a VAR system is structural. By construction it captures the unexplained variation of the exogenous variable and 'packs it up' into a shock. Our two variables of interest here, defence consumption and investment, under the very reasonable assumption that their induced variation is exogenous, both faced very sharp increases and drops in their value contemporaneously. Due to this almost perfect correlation that they have, any identification method is thus likely to confound the variation of the one process with the variation of the other. We contribute to this issue by tackling this issue head on and solving an augmented version of the unexplained variance problem as discussed further below.

3.3.2 Data Description

As we already hinted in the introduction, to identify the public consumption and investment shocks and to discriminate across states we stitch together a dataset spanning the period 1939Q2-2019Q1. Following Burnside, Eichenbaum, and Fisher (2004) we use a four variable VAR as a baseline to capture the main properties of the shock and to gauge the propagation of the shock in more detail we augment the VAR with a variable of interest in the 5th position. The data series we are using for our baseline exercises are; defence consumption, defence investment, GDP and a measure of the average tax rate, constructed as the ratio of Federal Government Current Receipts to GDP. For the rotating variables we use private consumption, private investment, manufacturing wage and a measure of the real interest rate. Details on how the data series are constructed can be found in the Data Appendix. The composition of the Defence Investment and Consumption variables are also described in the data appendix. [2].

The National Income and Public Accounts (NIPA) of the US provide us with publicly available quarterly data back to 1947Q1. Our main source for the 1939-1947 period is Ramey's (2011) and Ramey and Zubairy's (2018) dataset. The description of the dataset before 1947 is available in the Data Appendix. Regarding the main variables of our analysis, defence consumption and defence investment, we retrieve their pre 1947 values from the data publicly offered by the Bureau of Economic Analysis (BEA). However, this dataset is only at an annual frequency. To convert it into quarterly we dissagregated those variables using Denton (1971) simple interpolation method. As the interpolating variable we use Ramey's Defence Spending variable. More details on Denton's Method are provided in the Data Appendix. Figure 3.23 in the same appendix also presents the defence consumption and investment series resulting from this process. Additionally, we interpolated the government investment and consumption annual series using as an interpolation variable the government spending series again provided by Ramey (2011). This series is presented in the Appendix.

The variables are expressed in real per capita terms. To convert them into their per capita

[2]Of utmost importance for identifying government spending shock is that the data series should be at a quarterly frequency. Data series in lower frequencies than that should not be informationally rich enough to capture the full dynamics and subtleties of the prompt response of the economic agents to the spending shock.

equivalents we divide them by the Total Population instead of the Non Institutionalized Over 16 population as is the convention in the literature, due to data availability issues of the latter. Details on the transformation are presented again in the Data Appendix.

An additional issue regarding the transformation applied on the variables arises when one wishes to estimate the size of the fiscal multipliers from the Impulse Response Functions (IRFs). The literature utilizes log levels of the variables so that the IRFs have a natural interpretation as percentage changes. This is adequate for most cases, however if we are interested in estimating the multiplier as Gordon and Krenn (2010) and Ramey and Zubairy (2018) point out, the following issue arises; since the multiplier conveys information about the added value of one dollar of government spending on output, the percentage terms obtained by the IRFs need to be converted back to dollar terms. The literature relies on an ex post conversion wherein the percentages are multiplied by the sample average of the ratio of GDP to Government Spending Y/G. The problem is that this ratio might be having huge variations throughout the sample thus biasing the multiplier estimates. To tackle this we follow Gordon and Krenn's (2010) and Ramey and Zubairy's (2018) suggestion and transform all the variables into the same units before commencing with the estimation and identification of the model, so that there is no need to convert the percentage terms into their dollar equivalents by dividing them with a trend GDP variable. The trend variable we are using is a 2nd degree polynomial estimate of the trend of real per capita GDP.

One last issue regarding our data is the variable that determines whether the system is in a ZLB state. For this, we utilize Ramey and Zubairy's (2018) indicator variable, being 1 if the system is in the ZLB and 0 otherwise. The ZLB periods under this classification are presented on Figure 1.

3.4 Econometric Strategy

In this section, we lay out our econometric plan for the estimation and identification of the state dependent public consumption and investment multipliers. We begin by describing our VAR framework of choice to accommodate the non linear Data Generating Process. We continue by laying out a novel identification, with Uhlig's (2003) max share tailored in such a way that minimizes potential cofounding of the correlated shock processes.

3.4.1 Threshold VAR

Our primary aim is to capture the state dependent dynamics associated with shocks to the components of public spending. The underlying assumption here is that there is more than one Data Generating Process (DGP) driving our dataset. To capture this process then, the baseline linear VAR analysis is not sufficient.

To capture the non linearity in the DGP we choose to employ a Threshold VAR (TVAR), a generalization of the the threshold autoregressive models (TAR) of Tong (1978) and Tong (2012). Fundamentally, the idea of a TVAR is simple; there is a different set of coefficients driving the dynamics at each state and a simple threshold variable triggers the state, switching values between 0 and 1. To split the coefficients the TVAR utilizes information from a threshold variable that determines the prevailing regime and thus provides a natural framework for gauging the state dependent properties of the multipliers.

Notice that in contrast to other non linear VAR specifications (more prominently the Smooth Transition TVAR), the TVAR doesn't allow for a gradual transition from one state to another. The transition here is completely controlled by the dummy state variable switching between 1 and 0. In this regard we model the monetary state of the economy with a TVAR for reasons of tractability. In doing so, we follow the literature, most notably Galvao and Marcellinio (2014) and Ramey and Zubairy (2018) who also utilize a TVAR to study the economy under different conditions of the monetary regime.

To introduce some notation, let Y_t denote the vector of endogenous variables, and I_{t-1} an indicator function taking the value of 1 if the model is in the ZLB regime and 0 otherwise. Notice that the indicator function is lagged one period. Our econometric specification is then as follows :

$$Y_t = I_{t-1} \sum_{i=1}^{p} B_i^1 Y_{t-i} + (1 - I_{t-1}) \sum_{i=1}^{p} B_i^2 Y_{t-i} + \epsilon_t \tag{3.1}$$

where B^1 is the coefficient matrix when the model is in the ZLB state and B^2 when the model is in the normal state. Y_t is the vector containing the variables that describe the system and ϵ_t is a vector containing the reduced form residuals. Finally, I_t is an indicator variable that is 1 if the system is in a ZLB state and 0 otherwise. Implicit in this specification is that the variance process and consequently the variance covariance matrices are dependent on the regime. We therefore postulate that $\epsilon_t^l \sim N(0, \Sigma_\epsilon^l)$ for $i = 1, 2$ where l represents the state of the economy.

Expressed this way, the TVAR specification allows for non linearities through two channels; on impact of the shock through the regime specific variance covariance matrices and dynamically through the coefficient matrices. Impulse responses are then constructed from the regime specific coefficient matrices.

A problem when dealing with states that span only a limited amount of observations, as is the case with the ZLB, is that the parameters to be estimated are too many in relation with the data points, thus increasing the noise in our estimation and even leading to the coefficient matrix having explosive roots. Indeed, this is a problem noted by Burnside, Eichenbaum and Fisher (2004). To deal with this issue we are going to follow Burnside, Eichenbaum and Fisher (2004) and Ramey (2011) and estimate a simple baseline model of four variables to reduce the number of parameters that need to be simultaneously estimated. Then, to study the dynamics more extensively we are going to augment the variable of interest to the last position of the system.

To be specific, our benchmark model for the estimation of the multipliers will consist of 4 variables as contained in vector Y; defence consumption, defence investment, GDP and a measure of the average tax rate. The aggregate variables are expressed in real per capita terms and have been normalized by a measure of trend GDP, the fitted values of the second degree polynomial of per capita output. Using defence consumption and investment together in a model is necessary to avoid confounding issues of one process with the other that might arise due to their high correlation and similar characteristics that they share. Controlling for the tax rate is also deemed necessary and is a standard practice in the literature on fiscal multipliers. The baseline results of the model are robust to specifications with more variables, indicating that these 4 variables are enough to account for the effects we are interested in capturing. The underlying assumption here is that these 4 variables are enough to describe the properties of government spending and the information set is big enough to enable us to avoid any substantial endogeneity or omitted variable problems. This is not a far fetched assumption and many papers describe the empirics of government spending with tax and GDP as the main control (Barro and Redlick (2011), Ramey and Zubairy (2018)). To study the propagation of the shock in more detail we use in the reserved fifth position of the VAR system the private investment and private consumption variables in rotation. The dummy indicator that determined whether the economy is in a ZLB state or not is the one constructed by Ramey and Zubairy (2018).

The model is estimated with standard Bayesian methods that will provide a flexible framework for comparing multiplier estimates across states and components. We utilize a standard Minnesota prior for the coefficient matrix and an inverted Wishart prior for the covariance matrix. We sample from the prior distribution with a Gibbs sampling algorithm with 10000 draws and 9000 burns. We don't make any strong assumptions on the prior distribution and thus our results are not affected by this. The Bayesian framework however allows us to naturally check the statistical difference of our estimates across different states. We model the dynamics with 3 quarterly lags. Details on the estimation of the model are provided in the Technical Appendix.

3.4.2 Identification

To extract the structural shocks we follow the Stuctural VAR (SVAR) literature whereby the shock of interest is identified by solving for a linear combination of the reduced form model residuals of the VAR model. Our strategy, building on Chapter 1, derives from Uhlig's (2003) max share method and Ben Zeev and Pappa (2017) application of it to identify fiscal spending shocks. In its essence then, a structural public investment or consumption shock the linear combination of VAR residuals that maximizes the forecast error variance defence investment or consumption. Notice that we don't impose any zero contemporaneous impact constraints here as the empirical news shock literature usually does. We are thus maintaining an agnostic stance as to the theoretical object our identified shock corresponds with. In this sense one could interpret our shocks, as Angeletos et al (2020) point out, as a way of informing us about the generic properties and behavior of the components of public spending, without imposing any strict theoretical priors on it.

The simple max share method should not suffice for our purposes however. The reason is that, as mentioned earlier however, the inclusion of WW2 in ours sample entails concerns that the process as laid out above might be suffering from confounding shocks. The reason for this, is the contemporaneous upward surge in the variation of both variables during WW2. A process then that by construction captures variations and packs them up into shocks is bound to be 'confused' by this.

To tackle this concern, we seek to capture a vector that dynamically captures as much as possible of the exogenous variation of the defence variable of interest and as little as possible of the related defence variable. For this, we propose an augmented version of the baseline max share algorithm. In particular, and following Ben Zeev and Pappa (2017), we extend the baseline algorithm and identify the shock that maximizes the difference between the contribution to the five years variation in defence consumption and defence investment.

To formally describe our procedure we start with the reduced form moving average representation of the VAR. To save on notation we present the method for a linear case but the extension to a non linear framework is straightforward:

$$Y_t = B(L)u_t \tag{3.2}$$

where Y_t is an $m \times 1$ vector of endogenous variables, $B(L) = I_m + B_1 L - ... + B_p L^P$ is the matrix polynomial in the lag operator L where the $B_i, i = 1, ..., p$ are $m \times m$ parameter matrices and u_t is a zero mean, m-dimensional error process of the reduced form with a variance covariance matrix $E[u_t u_t'] = \Sigma_u$. Identification then amounts in finding a linear mapping $u_t = A\epsilon_t$ between the prediction error and the orthogonal shocks contained in ϵ_t, where A is the impact matrix.

The structural VMA representation of the VAR is then going to be:

$$Y_t = C(L)\epsilon_t \tag{3.3}$$

where $C(L) = B(L)A$.

The matrix A must satisfy $\Sigma_u = E[A\epsilon_t\epsilon_t' A] = AA'$. This restriction isn't however sufficient to identify A. To see this, consider that for all permissible matrices A, there exists an alternative matrix $\tilde{A}$ such that $\tilde{A}Q = A$, where Q is a $m \times m$ matrix with orthonormal columns, that also satisfies $\tilde{A}QQ'\tilde{A}' = \Sigma_u$.

Consider the h-step ahead forecast error of the i_{th} variable y_i in Y as

$$y_{i,t+h} - E_t[y_{i,t+h}] = \sum_{\tau=0}^{h} B_{i,\tau}\tilde{A}Q\epsilon_{t+h-\tau} \tag{3.4}$$

where $B_{i,\tau}$ represents the i row of the matrix of MA coefficients at horizon τ. The share of the forecast error variance of variable i attributable to shock j at horizon h is then

$$\Omega_{ij}(h) = \frac{e_i'(\sum_{l=0}^{h} B_l\tilde{A}qq'\tilde{A}')e_i}{e_i'(\sum_{l=0}^{h} B_l\Sigma_u B_l')e_i} \tag{3.5}$$

where q is the j column of the orthogonal matrix Q, such that $q'q = 1$ and e_i is a selection vector.

To identify the defence spending shock, assume that the variable of which the error variance we wish to maximize is ordered in position p of the VAR. Then, the impact vector of the public spending shock is the vector q solving the following problem :

$$q = \underset{q}{\operatorname{argmax}}\left(\sum_{l=0}^{h}\Omega_{1,1}(h) - \sum_{l=0}^{h}\Omega_{2,1}(h)\right)$$

subject to the following constraint:

$$q'q = 1 \tag{3.6}$$

.

or in words, we are looking for the orthonormal vector q that contains the shock process maximizing the variance contribution to defence consumption ordered first in the VAR and at the same time minimizing the variance contribution to the defence investment, ordered second. The same procedure is followed for the identification of the defence investment shock.

3.5 Multipliers in ZLB and Normal Times

In this section, we begin our explorations starting with the public consumption multipliers across ZLB and normal states. We present the Impulse Response Function (IRF) responses of the system to the shock of interest in each state and we proceed to quantify the multipliers and test whether any statistically significant state dependent effects are present. The IRFs for each state are constructed by the variance covariance matrix derived by the estimation of the TVAR described in (3.1) for each separate state as determined by the ZLB dummy.

Note that in all of what follows below, we are utilizing linear IRFs, meaning that the non linearity emerges only in the estimation and identification of shocks, not in the propagation. This is a choice made for reasons of tractability, as accommodating for non linearities in the propagation of shocks is far from computationally trivial [3]. Implicitly, there are two underlying assumptions behind linear IRFs; the first one is that the ZLB state lasts for more than 5 years (20 horizons in our IRFs) which corresponds to the empirics derived from the US experience, so that the IRFs accordingly don't switch states. The second assumption is that the magnitude of the shock is not enough to take the economy out of the ZLB state on impact. Both assumptions about the propagation of shocks in the ZLB state have been utilized and found to approximate the data well in Ramey and Zubairy (2018).

[3]A method the literature uses to accommodate for non linearities in the propagation of shocks is through the construction of Generalized Impulse Responses (GIRFs). Details on this framework and the computational challenges associated with it can be found in Koop, Pesaran, and Potter (1996)

3.5.1 Public Consumption Multipliers

Baseline

We begin our exercise with the baseline results for the public consumption shock. Figure 2 shows the impulse responses to a shock in defence consumption when monetary policy is constrained by the ZLB [4]. The shock has a clearly expansionary effect on the economy with the weight of the expansion falling exclusively on the initial periods after the impact. The variables peak in the second period, then slowly revert back to zero until they start rising again and stabilising at a higher level from period 16 onwards.

The path that the variables follow in the aftermath of the shock is not reminiscent of the typical bell shaped response to an anticipated shock. Rather, the sudden and sharp response of output and the defence variables suggests that the effects of the defence consumption shock had not been factored in by the agents very far ahead. Its effects, at least in the initial periods of the shock, clearly resemble a surprise shock. The very sharp decline of the variables after the 4th period points out to very temporary expansionary effects of the increased defence spending. Indeed for the rest of the horizons, the economy never reacts significantly differently from zero. Interesting in its own right is the tax rate that rises permanently and significantly with significant delay suggesting that the spending surge was financed through taxes imposed with a lag.

One question left unanswered by the inspection of the IRFs is whether these results are indeed driven by the ZLB state or by the extraordinary circumstances of the '40s that is the main source of information in this case. We will revisit this issue later in our explorations.

[4]In the appendix we present the Forecaste Error Variances from the baseline exercises.

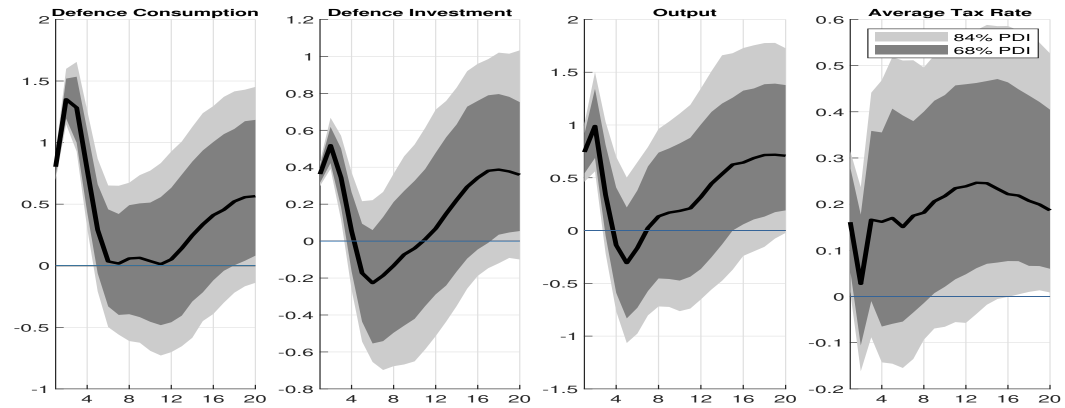

Figure 3.2: IRFs to Defence Consumption Shock in ZLB Regime. 68% and 84%Posterior Density Intervals.

Let's then move on to the propagation of the same shock in the normal state. Figure 3 presents the baseline IRFs from this exercise. In the normal state, the response to a defence consumption shock is clearly different, at least in a qualitative sense from the response in the ZLB state. The gradual unfolding of defence spending and the bell shaped response of output point out to clearly anticipated effects and that the agents had factored the effects of the shock in their calculations. After slowly rising the peak response of output occurs at period 9 and then starts reverting back to zero. Interesting again is the reaction of the tax rate, that rises significantly and permanently with a lag, suggesting again that the defence spending in the normal regime was again financed through taxes.

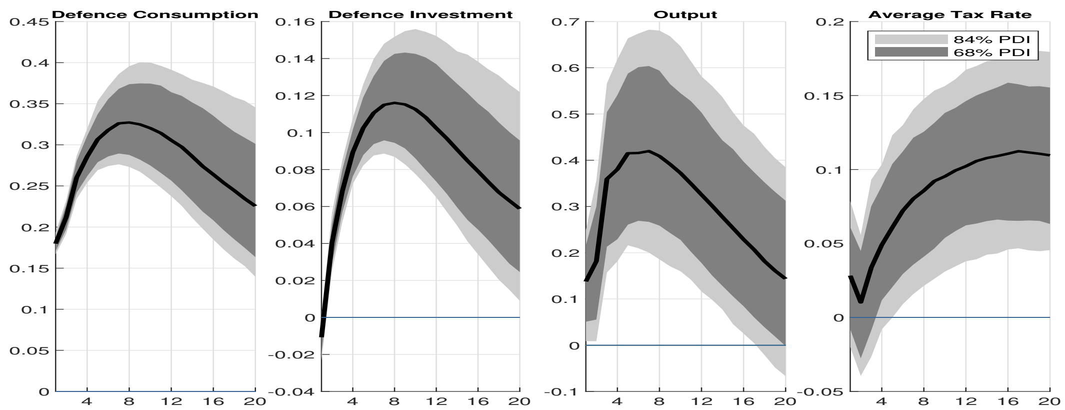

Figure 3.3: IRFs to Defence Consumption Shock in Normal Regime. 68% and 84%Posterior Density Intervals.

These baseline results establish that the reaction of the economy to a public consumption shock in the ZLB and the normal state is different both at a qualitative and a quantitative level. Next, we want to quantify these reactions and gauge their difference in a more precise way.

Quantifying the Multipliers

By just looking at the baseline graphs, a difference between the public consumption multipliers in the two different states is palpable. To be more precise as to whether there are any state dependent effects at hand, we need to gauge the multipliers associated with each state. Since all IRFs are expressed in the same units due to the normalization of choice, to quantify the multipliers one just needs to calculate the simple cumulative ratio of output to the defence consumption variable, the so called integral multiplier.

More formally, the integral multipliers can be simply estimated as the simple cumulative ratio of output to the defence spending variable we are interested in:

$$M = \frac{\sum_{i=1}^{h} Y_i}{\sum_{i=1}^{h} G_i} \tag{3.7}$$

The multipliers resulting from this exercise along with their 84% probability bands and the full sample linear multiplier as benchmark are presented on Table 3.1:

Table 3.1: Output Multipliers

Horizon	Normal	ZLB	Full Sample
4	1.10	0.49	0.66
	[0.42,1.77]	[-0.01,0.94]	[0.45,0.88]
8	1.22	0.39	0.52
	[0.59,1.85]	[-1.78,1.26]	[0.2,0.84]
12	1.21	0.99	0.5
	[0.61,1.8]	[-2.1,3.45]	[0.06,0.89]
16	1.14	1.21	0.49
	[0.57,1.74]	[-1.195,3.04]	[-0.07,0.94]
20	1.07	1.22	0.46
	[0.49 ,1.69]	[-0.58,3.2]	[-0.19,0.95]

The public consumption multiplier in normal times is positive with statistical significance and consistently greater than one, though not significantly so, while the ZLB public consumption multiplier is rising as time progresses though with widening density intervals. Interestingly, the multiplier at the ZLB is never significantly different from zero, a finding that seems to come in contrast with intuition and the suggestions of the literature that the fiscal multiplier is higher in the ZLB state. Noticeable are the widening probability intervals in the ZLB state. This points to the fact that at lower frequencies there is too much noise that our method is capturing and not enough information. It is clear that the bulk of the information and thus the effect of the public consumption shock at the ZLB unfolds on impact within the initial 4 periods.

Are these multiplier differences statistically significant however? Figure 3.4 plots the difference between the two multipliers at each horizon and the associated Posterior Density Intervals.

We see that the multiplier differences become statistically distinct in the medium term from the 5th to the 8th horizon. Within that time interval, the Normal times multiplier is significantly larger than the ZLB one. As we saw in the IRF section, this is due to the fact that the shock at normal times unfolds to its full effects over the medium term, pointing out again to its anticipated

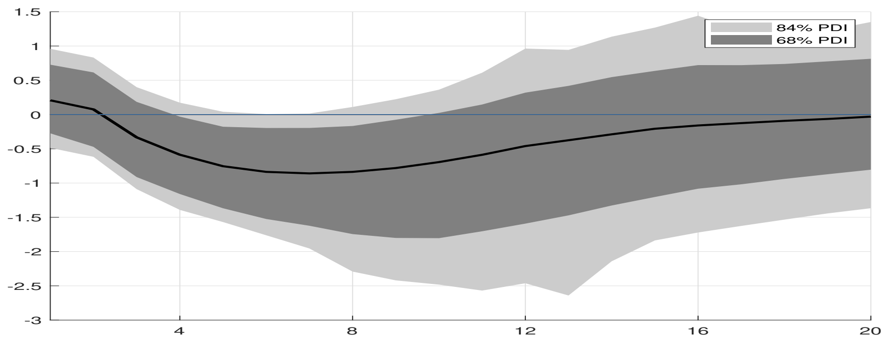

Figure 3.4: ZLB Consumption Multiplier minus 'Normal' Multipliers

nature. Having thus established the state dependence of the multipliers, the question that arises from this exercise is whether this difference is due to the distinct properties of the multiplier depending on the state or an idiosyncrasy arising from the sample choice with the ZLB period spanning for its most part war times.

A Look Into The Defence Consumption Shock Propagation Channels

While the above exercise is useful to establish the basic facts and serve as a basis for discussion, it doesn't shed any light as to what drives these results and the deeper workings of the shock in the economy. In this section we examine the effects of the public consumption shock on individual variables. As already mentioned, the procedure we follow for this exercise is to rotate the variables of interest in the 5th position, and when necessary the 6th position, of our VAR system and then consider their impulse responses to the shocks of interest, so in this case the defence consumption shock [5]. The variables we are going to consider for this exercise are private consumption, private investment, the real manufacturing wage and the real interest rate.

We begin with the responses of private consumption and private investment to our shock. Figure 5 reports the responses of private consumption to a public consumption shock in both states. Private consumption reacts significantly to the public consumption shock at the ZLB state only on impact while it rises with a delay in the medium term and significantly at the normal state. The path of private consumption in both states is interesting in its own regard. While the

[5]The baseline results do not change by the inclusion of this additional variable.

median of the response diverges very strongly across states in the initial 8 periods, after that the median of the response becomes identical.

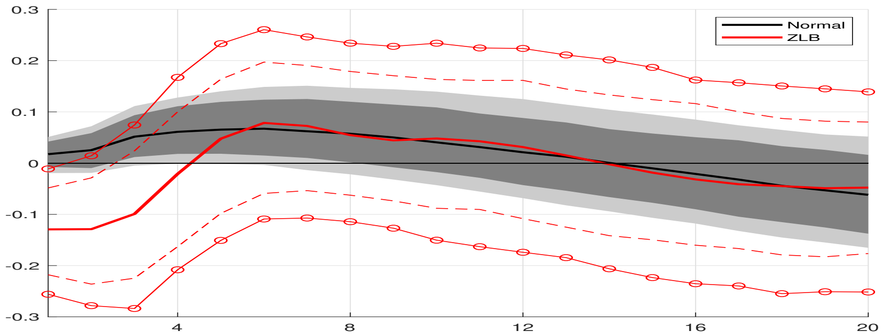

Figure 3.5: Response of Private Consumption to a Public Consumption Shock

This initial divergence implies that any state dependent effects of this shock fall within the first 8 periods. This on impact decline in consumption at the ZLB state is, at first sight, a finding that merits some additional investigation. We delve deeper into this issue by decomposing the response of consumption into the response of its major two constituent parts; consumption of services and consumption of non durable goods. Figures 3.6 and 3.7 report the IRFs of these components. We see that non durable consumption actually rises in response to the shock. Consumption in services, on the other hand, does not react significantly to the shock. We see that the path of the response of total private consumption resembles, in a qualitative sense at least, the path of non durable consumption. A reasonable hypothesis would be that the non significant reaction of services consumption is what drives the total consumption path to insignificance.

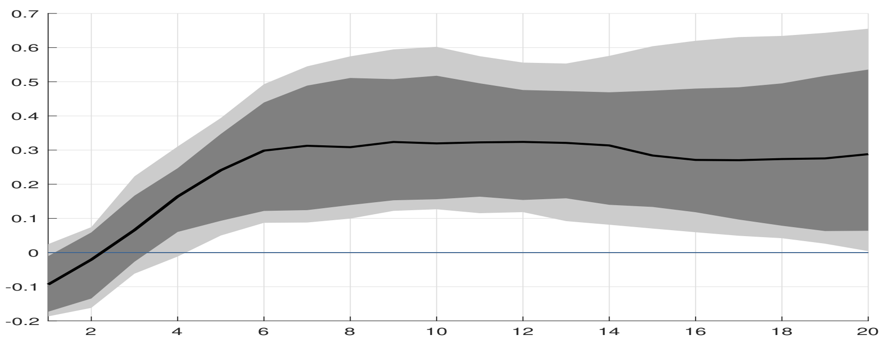

Figure 3.6: Response of Non Durables Consumption to a Public Consumption Shock in ZLB

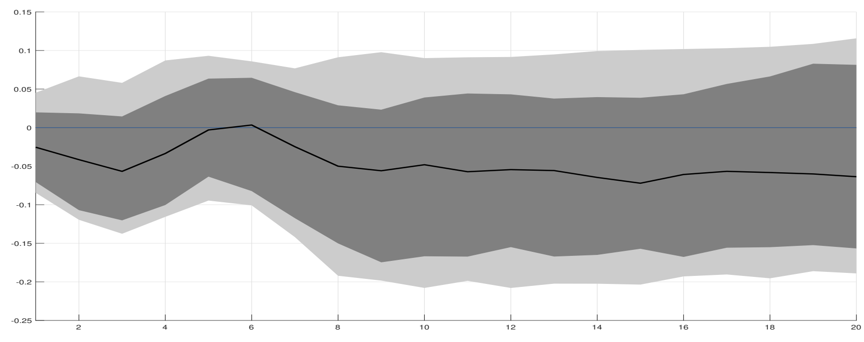

Figure 3.7: Response of Services Consumption to a Public Consumption Shock in ZLB

As we are informed from a consumer survey conducted in the period 1941-1944 (Henderson (2015)), the increase in income from that period due to the military spending programs was mainly getting absorbed by consumption on foods and other non durable goods such as clothing. In contrast services consumption within that time period appears to have declined. These observations closely mirror our findings in this section. It can therefore be assumed that our defence consumption shock extracted from the ZLB sample more closely captures the effects of the development of the war economy in the '40s rather than a generic public consumption shock. Hence, we can treat this result as broadly informative about the effects of public consumption as a more short term shock but it should be noted that it is not easily generalizable.

Next, we move on to the response of private investment to a public consumption shock. In Figure 8 we see that private investment never is significantly different from zero. This is a positive sign with theoretical implications; it might indicate that in DSGE models the economy is best described by a two sector model. It also appears that general equilibrium effects are not present.

The response of the wage to the public consumption shock which is also of theoretical and practical interest. Neoclassical models suggest that in the aftermath of public spending shocks, negative wealth effects come into play. In response to a shock the household optimally decreases its consumption and increases its labor supply, leading to a rise in output as a result. The real wage should thus fall in response to decreased labor supply before reverting back to its initial value. New Keynesian models on the other hand propose a number of mechanisms through

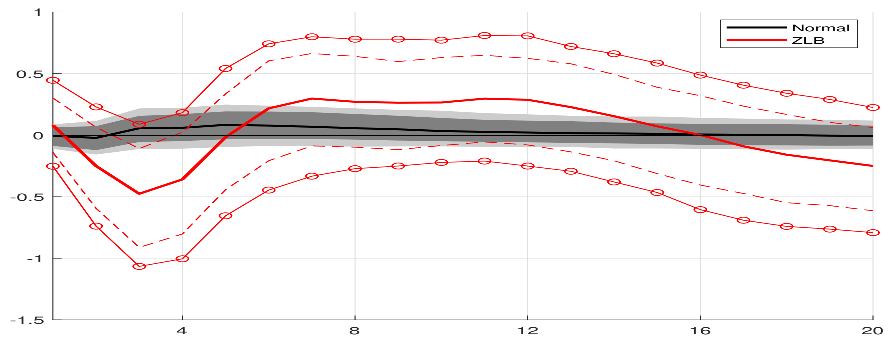

Figure 3.8: Response of Private Investment to a Public Consumption Shock

which wages rise with.

Given that the non linearity under examination is the ZLB, gauging the response of the real interest rate should be of some interest. Given the central place of the interest rate in the theoretical discussions about the fiscal multiplier, the impulse response is also of significant theoretical interest. Figure 9 presents the relevant result.

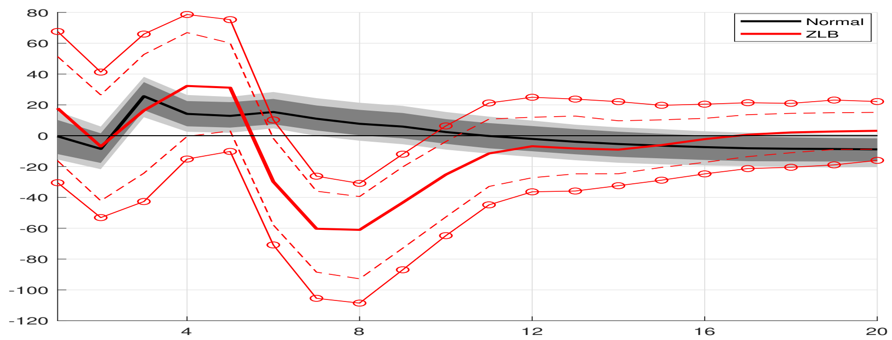

Figure 3.9: Response of Real Interest Rate to a Public Consumption Shock (Basis Points)

Reassuringly, the results are broadly as expected; the real interest rate significantly and with some persistency rises in normal times in the aftermath of a public consumption shock and its response is insignificant, with the exception of an interval between period 6 and 9, in ZLB times.

To conclude, the defence consumption shock in the normal regime has almost identical effects as those of a generic defence consumption shock analyzed in Chapter 2. The shock appears mostly anticipated with a significant expansionary effect and associated with a positive output multiplier. In the ZLB/War regime, the economy reacts in a very different manner. The shock is expansionary only in the very initial periods and after that it becomes insignificant.

We argued that this effect is not likely stemming exclusively from the general properties of the economy at the ZLB but rather from the state of war that spans most of the ZLB sample.

3.5.2 Public Investment Multipliers in Normal and ZLB Times

Next, we venture into the second part of our exercise which is the examination of the responses of the economy to a public investment shock in normal and ZLB times. In this section we broadly follow the structure presented above and commence our analysis with some benchmark results before moving on to the propagation mechanisms.

Baseline

As before, for our baseline we use real per capita defence consumption, defence investment, output and a measure of the average tax rate. Again, the normalization of choice allows us to interpret all the IRFs in the same units.

Figure 10 then shows the impulse responses to a defence investment shock when monetary policy is constrained by the ZLB. The response of the economy is again clearly expansionary. All variables follow the familiar from the news shock literature bell shaped path. The responses of the defence spending variables gradually rise indicating the gradual unfolding of a defence spending program. The output response to the defence investment shock follows the same pattern with its response being 0 on impact and peaking after 9 periods before reverting back to 0. As with the Defence Consumption case, the Average Tax Rate falls after 1 period and then rises significantly in the medium term before it starts reverting back to zero. This indicates that the shock has limited surprise elements and has been fully embedded in the expectations of the agents and its effects are priced in by them, in contrast to the public consumption shock at the ZLB.

Next, we move on to the responses of the economy to a defence investment shock in the normal regime. Figure 3.11 presents the IRFs of the benchmark model to a defence investment shock in normal times. Again the two defence spending variables follow a bell shaped path peaking in the medium term at horizon 7, indicating a gradual unfolding of the military spending

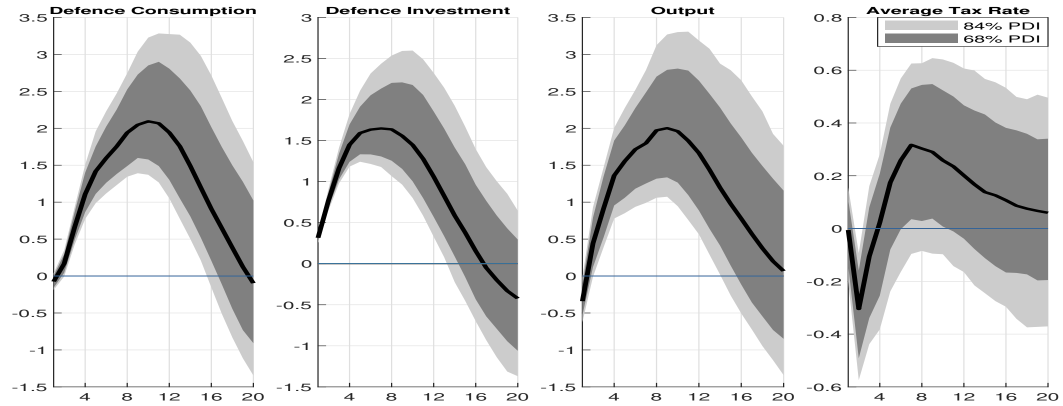

Figure 3.10: IRFs to Defence Investment Shock in ZLB Regime. 68% and 84%Posterior Density Intervals.

program. However, output this time responds positively only on impact and after that its response is indistinguishable from zero. The tax rate, following the pattern we have already seen, permanently rises indicating again that the military spending was financed through taxes. The differing response of output among states is suggestive of the fact that in the defence investment case there are potentially state dependent effects. Finally, taxes also rise permanently but this rise is only significant for the first 3 periods.

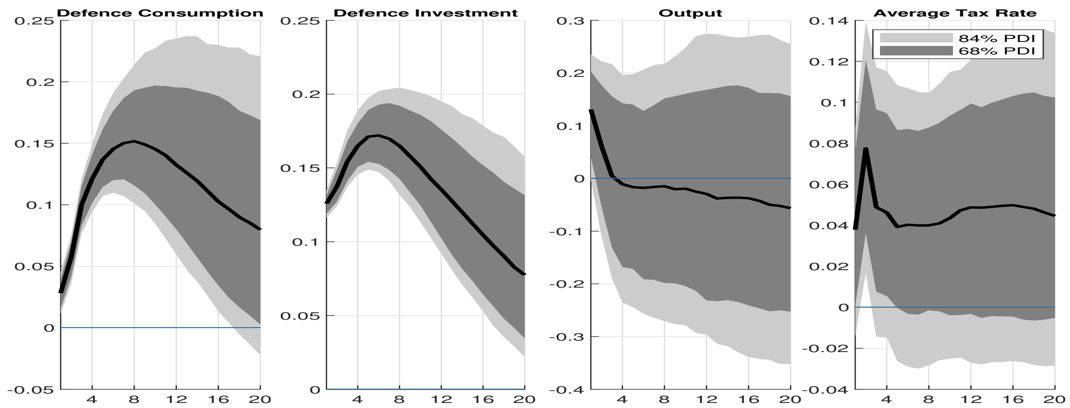

Figure 3.11: IRFs to Defence Investment Shock in Normal Regime. 68% and 84%Posterior Density Intervals.

As a more general note, the gradual propagation of the defence investment shock is an indication that its effects were actually anticipated. This result is not surprising given the nature of public investment. There are inherent delays associated with a public investment shock, even if its announcement came as a surprise. Time to build and time to implement features would allow the agents to price in the full effects of such a shock quite some periods ahead of its complete unfolding. This is in contrast to the defence consumption shock, which can potentially be a full surprise to the agents, as we saw in the defence consumption ZLB case.

Quantifying the Multipliers

The next step in our analysis is to gauge the multipliers associated with the defence investment shock. As with the defence consumption case, this should allow us to more precisely see if there are any state dependent effects. As with the defence consumption shock case, we utilize integral multipliers to construct the fiscal multipliers.

We begin by presenting the estimates of the simple integral defence investment multiplier in the Normal and the ZLB case. Table 3.2 presents the results with their 84% confidence bands:

Table 3.2: Defence Investment Output Multipliers

Horizon	Normal	ZLB	Full Sample Linear
4	0.33	0.68	0.68
	[-0.88,1.39]	[0.18,1.11]	[0.58,1.18]
8	0.11	0.97	0.68
	[-1.2,1.3]	[0.5,1.32]	[0.84,1.32]
12	0.05	1.13	0.68
	[-1.29,1.29]	[0.62,1.49]	[0.97,1.48]
16	0.02	1.28	0.68
	[-1.39,1.38]	[0.63,1.74]	[1.11,1.69]
20	-0.03	1.47	0.68
	[-1.6 ,1.5]	[0.58,2.2]	[1.26,1.99]

The qualitative and quantitative difference between the defence investment output multipliers is clear upon inspection of the table. The ZLB public investment multiplier is at all horizons significantly greater than zero while the Normal Times public investment multiplier is never statistically distinguishable from zero.

To be more precise with this claim however, we want to test whether there is a statistical significant difference between the public investment multipliers. Figure 12 then plots the time profile of the difference between the ZLB and the normal defence investment multiplier. While the ZLB normal times multiplier appears significantly larger on impact, as we move forward horizons, the ZLB multiplier gradually and significantly within the 68% posterior density interval takes over and becomes significantly higher. Hence, the greatest difference is to be traced in the medium to long term. This result reflects the fact that in normal times the surprise effect of the public investment shock is higher while in ZLB times the news effect is dominant.

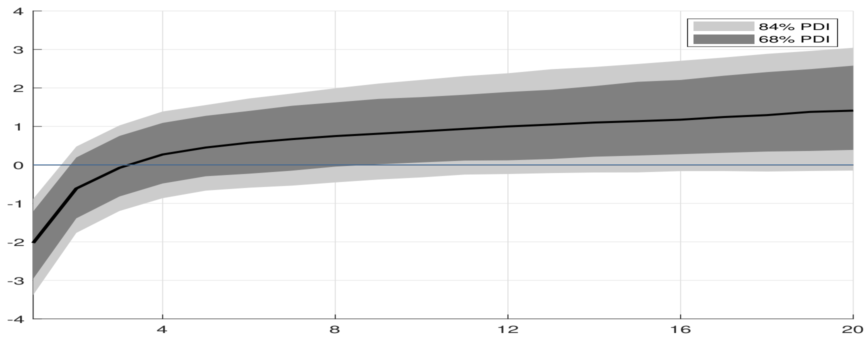

Figure 3.12: ZLB Public Investment Multiplier minus 'Normal' Multipliers

3.5.3 A Look Into The Propagation Channels Behind The Defence Investment Shock

The previous section presented the benchmark results for the public investment multiplier. To understand what drives these results we now delve into the propagation mechanism of the public investment shock. The process we go through is the same as in the public consumption case, namely we place the variable of interest at the 5th or 6th position of the VAR and obtain its impulse response.

Before we begin our analysis of the defence investment shock it would be useful to gauge the components of the defence investment variable. As per the US accounting system defence investment spending can be sub divided into three categories; spending on military structures, spending on equipment, whereby the equipment consists of weapons and weapon delivery systems (air crafts, missiles, ships, vehicles) and spending on intellectual property products that include software and research and development. A defence investment shock should therefore not only capture the effects of augmented infrastructure through spending on military structures but also the wider effects of the expansion of the military industry.

We begin our analysis by examining the response of private consumption to our defence investment shock. Figure 13 presents the IRF of private consumption in the ZLB and the Normal regime. What stands out immediately is the huge difference in variance across the two regimes. This variational difference is a pattern that will appear in all of our exercises involving defence investment and it's due to the fact that the ZLB state is associated with the huge variation in the defence spending variables due to WW2. The fact that private consumption reacts positively in the ZLB and negatively, albeit not significantly so in normal times, is a result that at first sight appears peculiar, especially given the insignificant response of private consumption in the defence consumption shock case. We can rationalize it if we consider that the defence investment shock captures the effects of the creation and/or expansion of the military industry.

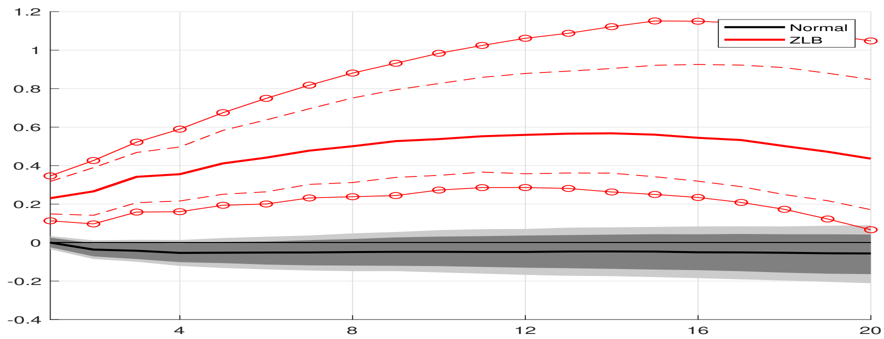

Figure 3.13: Response of Private Consumption to a Defence Investment Shock

Figure 14 shows the response of private investment. Private investment is falling in both cases as should be expected, a reaction that we hypothesize is due to its higher sensitivity to crowding out effects. In both cases the fall is culminating in between horizon 6 and 8.

The finding that private investment falls in response to the public investment shock in both

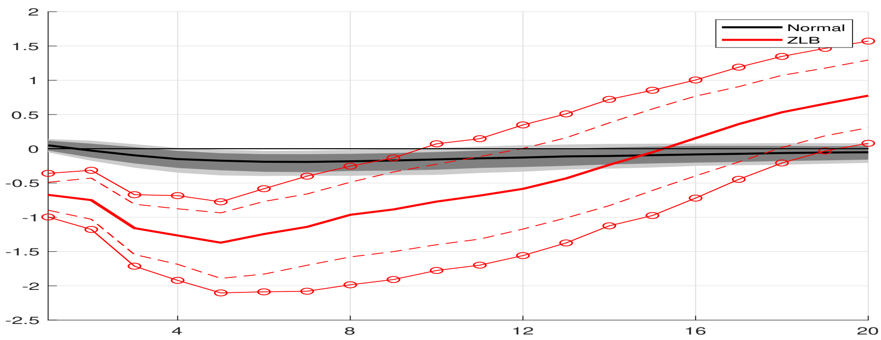

Figure 3.14: Response of Private Investment to a Public Investment Shock

regimes is of theoretical interest in itself. Boehm (2020) suggests that private investment is subject to strong wealth effects that are generated from the interest rate that crowds it out. Upon first inspection, this figure seems to be corroborating this view. Private investment is indeed crowded out in both states. Moreover, the fall in the normal regime should be larger than the fall in the ZLB regime due to the constrained response of the real interest rate. To check if our empirical framework supports this hypothesis we need to make the response of private investment in both regimes directly comparable. For that, we are going to construct the private investment multiplier to a public investment shock. Figure 3.15 presents the difference among these multipliers;

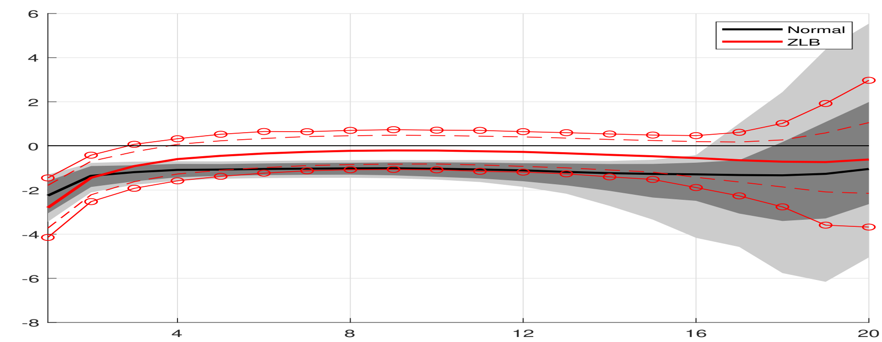

Figure 3.15: Private Investment Multipliers

Upon inspection of figure 3.15 it appears that indeed the fall of private investment is more pronounced in the normal regime than in the ZLB regime. More specifically in the ZLB regime, with the exception of the first 3 periods, the private investment multiplier is never significantly different from zero while in the normal regime the private investment multiplier is negative.

To complement this finding on the crowding out of private investment, we next check the response of the suspected cause of this phenomenon, namely the real interest rate. Figure 3.16 shows the response of the real interest rate to a public investment shock.

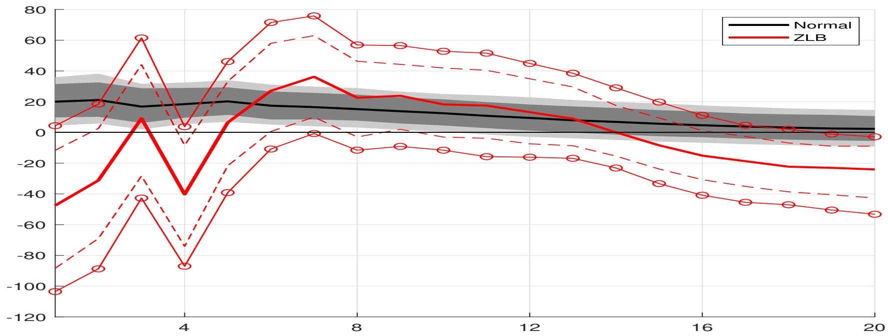

Figure 3.16: Response of Real Interest Rate to a Public Investment Shock (in Basis Points)

As it was expected, the real interest rate responds with a persistent rise in the normal case and a volatile but never significantly different from zero movement in the ZLB case. These IRFs hint strongly towards the existence of a causal relationship between the response of private investment and the response of the real interest rate. We delve deeper into the nature of this relationship deeper with an additional counterfactual exercise in the Appendix, wherein we demonstrate a direct relationship between this interest rate mechanism and the effect of the shock on output.

To conclude this section, the public investment shock in the ZLB has significant expansionary effects leading to a rise in output and private consumption. On the other hand, in the Normal regime the reaction of the economy is not as expansionary with output and private consumption not reacting significantly. Whether this difference is due to state dependent effects or the massive military spending program in the sample that informs our ZLB results is not resolved. A reasonable hypothesis would be that it is a combination of both effects. A hint towards state dependent effects is the response of private investment which falls in both states albeit more so in the normal state, as we saw from the private investment multipliers. Thus the results can be seen as generalizable. This is a finding that in juxtaposition with the real interest rate response in both states can be justified theoretically as a state dependent effect, namely the crowding out effects in the ZLB regime are less pronounced.

As a final comment, the central tension that looms over all of our findings, namely if it is the peculiarities of the sample or the state dependence the culprit driving the IRF differences hasn't been resolved in a precise way. The IRFs and some qualitative evidence allow us to puzzle together some evidence that point to state dependent effects. The extent to which these dominate remains however unclear and an open question to the literature.

3.6 Discussion: Anticipated or Unanticipated

Since Ramey (2011) seminal recognition that government spending shocks are in fact anticipated by the agents, it is a convention in the literature to use methods or bring in information that would allow the researcher to explicitly differentiate among the effects of surprise and anticipated shocks. As per convention, the Blanchard and Perotti (2002) identification of a government spending shock with its assumption of government spending being predetermined with respect to all the other variables in the system, serves as the unanticipated shock. To identify the anticipated shock there are a number of different approaches ranging from bringing in additional information in the form of forecasts to Ramey's narrative method.

Our identification approach doesn't take a stand on whether a government spending shock extracted from a particular sample is anticipated or unanticipated. Instead, it lets the data speak for itself, naturally lending weight to the forecast horizon where the unexplained variance is at its maximum.

In addition to the general considerations expressed in the previous sections regarding the nature of the shocks, we can also consider the historical idiosyncrasies of our sample. The explosion in federal expenditures in defence started in the second quarter of 1940, after the fall of France, at which time the war started looking more and more imminent for the US. Upward revisions of the imminent fiscal spending were being made public in popular publications. This interpretation seems even more plausible if we juxtapose it with narrative evidence from newspapers as cited for example in Gordon and Krenn (2010). We present some of these excepts in the Appendix. The heavy publicity that this explosion in public spending attracted made sure that the economic agents were informed about its forthcoming effects. But what was that defence spending about? If we go through the extracts from newspapers of the era cited by Gordon and Krenn (2010), as presented in the Appendix, we see that it was mainly about armaments to also meet the demand of the Allies at the other side of the Atlantic. Armaments are classified under the US National Accounts system as durable goods, and as per the US accounting convention they are classified under public investment. Thus, the agents were mainly informed about the forthcoming effects of defence/public investment in addition to the time to build lags inherent in public investment shocks.

How about defence consumption however? Why does it appears anticipated in the Normal sample but not so in the ZLB/WW2 sample? A priori, there are no theoretical explanations that can guide our reasoning as to why the anticipated effect of the shock is state dependent. A rough argument that can be made is by associating the rise in defence consumption with the conscription call of the early '40s. Consider that defence consumption is made up of compensation of general government employees and consumption of intermediate durable weaponry goods. In principle therefore, any increase in income induced by this shock is mainly through direct hiring into military service and increased wages rather than effects spilling over from the development of a war industry. It can be argued that the relatively sudden nature of a conscription call and of the increased wages associated with it make it much harder to capture its anticipated effects over frequencies as low as the quarterly frequency of our sample. After that, the war started and probably the additional income of the draftees must have failed to propagate in a war economy as it does in a normal economy generating a short term expansion.

So in conclusion, narrative evidence from that period seems to corroborate the fact that the main weight of the expansion fell on the defence investment program. Additionally, it appears that the special behavior of defence consumption in the ZLB state was not due to the ZLB per se but rather the special nature of defence consumption and the extraordinary circumstances of war.

3.7 Conclusions

In this paper we attempted to tackle the question of whether the components of public spending are state dependent when the state is the ZLB. This is a question of high practical interest for the policymaker. Our starting point is the utilization of the exogenous variation of the components of defence spending to construct the public investment and consumption shocks. In order to tackle the issue of informational scarcity for the ZLB state and maximize the amount of variation in the data, we employ an extended novel quarterly sample that reaches back to 1939. Our goal with that is to include the large ZLB episode of the '40s. While this sample choice ameliorates the informational deficiency problem, it created two new challenges. In the '40s the defence spending variables comove perfectly and are thus creating potential cofounding problems in the identification. An additional challenge is whether the empirical estimates derived with the '40s included are generalizable and broadly informative about the properties of the business cycle.

To tackle the first problem of the cofounding shock processes we utilize a novel identification method, in the form of an augmented max share algorithm as developed by Uhlig (2003). The philosophy of Uhlig's (2003) algorithm is to pack up the variance of the defence spending variable not explained by the model into a shock. Under the assumption that the variance if the variable is exogenous, which in the case of war spending variables is a reasonable approximation to reality, the retrieved shocks are structural. The problem is that this algorithm is liable to cofounding processes especially when there are two perfectly comoving variables as is the case with our dataset. To remedy this, we solve an augmented version of the algorithm by looking for the vector that maximizes the shock contribution to the variable of interest while minimizing the contribution to the comoving variable. Regarding the second problem of external validity, we utilize a series of qualitative and narrative data to resolve the question of whether our results are a product of the extraordinary circumstances of our sample or of the extracted structural shock processes.

Our results indicate that there is a statistically significant state dependence in the case of the defence investment multiplier. The defence investment multiplier is significantly positive in the ZLB while insignificant in the Normal state. We find some evidence that tthis result arises to some extent due to the muted response of the real interest rate to a spending shock in the ZLB regime and the consequent weaker crowding out of private investment from public investment. In contrast, we found no conclusive evidence of the public consumption shock exhibiting state dependent effects.

3.8 Technical Appendix

3.8.1 A. Bayesian Estimation

Following along the lines of the literature on Bayesian estimation of VAR models, we estimate the posterior distribution of our model with a Gibbs sampling algorithm. We use a Minnesota prior for the prior coefficient distribution. In this appendix, we discuss the specifics of the estimation process.

The linear model is specified as follows;

$$Y_t = B(L)Y_{t-1} + u_t \tag{3.8}$$

To estimate it we specify a Minnesota prior for the coefficient matrix $B(L)$. The posteriors are then drawn from a Gibbs sampling algorithm. More formally, the process is the following;

1) We set the priors for the covariance matrix Σ and the coefficient matrix B following the literature convention on the Minnesota prior. We assume that the coefficient matrix follows a normal distribution $p(B) \sim N(\tilde{B}, H)$ and and the covariance matrix an inverted Wishart distribution $\Sigma \sim IW(\tilde{S}, V)$. Set a starting value for Σ. The hyper-parameters that control the priors take the following values; $\lambda_1 = 0.2$, $\lambda_2 = 0.5$, $\lambda_3 = 1$, $\lambda_4 = 10^5$.

2) We first sample the VAR coefficients from the conditional posterior distribution specified as follows; $F(b|\Sigma, Y_t) \sim N(B^*, V^*)$. The mean of the distribution B^* is defined as:

$$B^* = (H^{-1} + \Sigma^{-1} \otimes X_t'X_t)^{-1}(H^{-1}\tilde{B} + \Sigma^{-1} \otimes X_t'X_t\hat{B})$$

and V^* is defined as

$$V^* = (H^{-1} + \Sigma^{-1} \otimes X_t'X_t)^{-1}$$

Once $B*$ and $V*$ have been calculated and the distribution has been parametrized, the VAR coefficients are drawn from the posterior distribution:

$$b^1 = B^* + [b \times V^{*1/2}]$$

3) The next step is to draw the covariance matrix Σ from its conditional Inverted Wishart distribution, conditional on the draw of the coefficient matrix in step 2.

4) Repeat steps 2 and 3 20000 times. We keep the last 2500 draws from these iterations to form the empirical distributions of the parameters of interest. This is the so called burn in period. The last draws are used to calculate impulse responses and any other structural analysis the researcher might wish to perform.

3.9 Appendix B: A Counterfactual Analysis

3.9.1 Counterfactual Analysis: Is The Real Interest Rate Crowding Out Private Investment?

A main attribute of the economy's reaction to the public investment shock as we saw was the significant fall of private investment in both states. This decrease was more pronounced in the normal state, with the private investment multiplier being negative while in the normal regime being indistinguishable from zero. This attribute of the public investment shock is quite constant throughout the empirical and the theoretical literature with the main culprit behind it being the private investment crowding out effects that the shock induces through the increase in the real interest rate.

If this interpretation is correct, then this could be the reason behind the larger multiplier in the ZLB state; the muted response of the real interest rate ameliorates the crowding out of private investment thus in turn leading to a larger output response. The main mechanism behind this is the high elasticity of private investment with respect to the interest rate. Indeed, this high elasticity could be the reason behind state dependence. Similarly, if private consumption is not sensitive to a shock in public consumption, then the ZLB doesn't play a role.

To test this we devise up a counterfactual exercise wherein we force the real interest rate to be constant in the IRF. Ideally, we would like to force the real interest rate at non zero values in the ZLB state. This would be a questionable practice however; how does one choose a non zero value and more importantly how does one choose a counterfactual non zero path ? Instead, we will take a more indirect route and force the response of the real interest rate to be zero in the normal state and gauge at the reaction of the system. We do this, following Bernanke et al., 1997 and Bachmann and Sims, 2012, by calibrating a shock to the real interest rate such that it will completely counteract the shock of the public investment on the real interest rate. If, as a result of this exercise, private investment for example falls less in the counterfactual scenario, then this would be an indication that a real interest rate channel exists and plays its role in the propagation.

The system we consider involves the 4 main variables, private consumption or private investment depending on the shock we are considering in the 5th place and the real interest rate in the 6th place. We mute the response of the real interest rate to the shock of interest and gauge at how the system reacts. The construction of the counterfactual is described in the relavant appendix of Chapter 2.

Figure 17 presents the results of this exercise for the public investment shock. The constrained response of the real interest rate leads to a lesser fall of private investment by about 0.3 and to a quite more sizeable response of output. There is little change observed in the paths of the other variables in the system however. This can be taken as evidence for the existence of the real interest rate channel in the crowding out effects which could be the culprit behind the distinct multipliers among states.

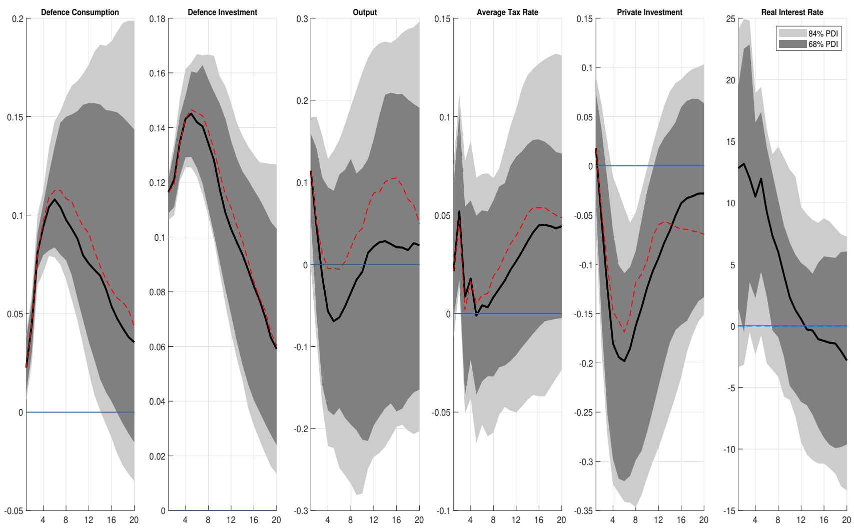

Figure 3.17: Counterfactual IRFs to Defence Investment Shock If Real Interest Rate Is Zero (in red). Black Line Is Baseline IRFs

Next, we want to conduct the same exercise but this time with defence and private consumption. This is of interest for two reasons; the first one is that the existence or non existence of a real interest rate channel would mean that the different responses to the defence consumption shock among states are due to the historical idiosyncrasies of the sample. Secondly, it would be of theoretical interest since it would act as evidence for or against the differing elasticities of private investment and private consumption to real interest rate movements. Figure 18 presents the results of the constrained IRFs for the defence consumption shock;

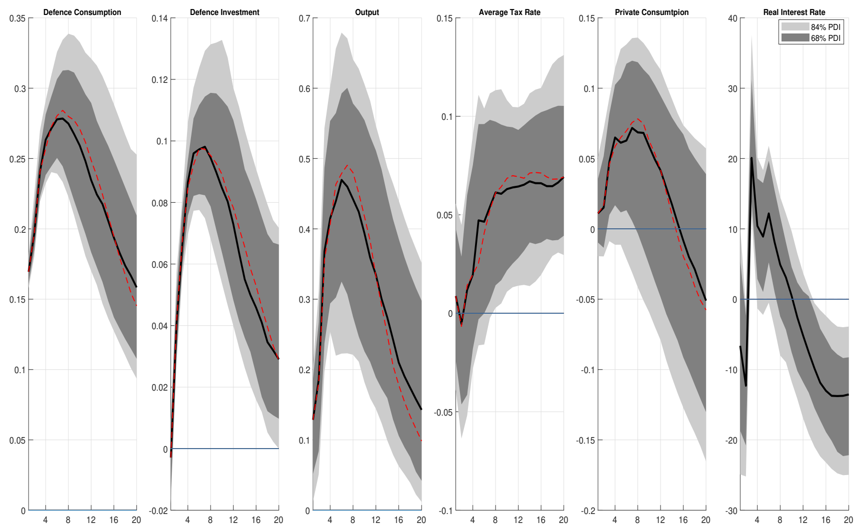

Figure 3.18: Counterfactual IRFs to Defence Consumption Shock If Real Interest Rate Is Zero (in red). Black Line Is Baseline IRFs

In this case, we see that the constrained IRF makes essentially no difference; the baseline and the counterfactual are essentially the same. The real interest rate has no effect on the response of private consumption, the agents don't seem to factor this in their maximization problems. Private consumption is rather getting crowded in than crowded out. Therefore, there is no reason to believe that the defence consumption shock would exhibit any state dependent effects and the state differences are due to the choice of our sample and the effects of the war.

3.10 Data Appendix

3.10.1 Data 1947-2019

The data is from the Federal Reserve Economic Database (FRED). TFP is the utilization adjusted TFP produced by Fernald (2014). Information on this data and the transformations applied can be found in the Data Appendix of Chapter 2 of this thesis.

3.10.2 Data from 1939-1947; Denton's Interpolation Method

For the years 1939-1947, all of the data points are taken from Ramey's (2011) dataset. Exceptions are the government consumption and investment and the defence consumption and investment variables that we constructed ourselves. To construct them we fetched their available annual data points from the Bureau of Economic Analysis' (BEA) website and dissagregated them using Denton's proportional interpolation method.

Denton's proportional method is described in detail in Bloem, Dippelsman, and Mæhle, 2001. To dissagregate annual variables one needs to have the annual series to be dissagregated and a quarterly series to be used as an indicator. Then, one needs to solve a minimization problems that keeps the annual benchmark series as close as possible to the indicator series. Mathematically, the problem to be solved is described as follows;

$$min_{X_1,..X_{4\beta},..,X_T} \sum_{t=2}^{T} (\frac{X_t}{I_t} - \frac{X_{t-1}}{I_{t-1}})^2 t \in [1, ..., 4\beta, ..., T]$$

under the restriction that the flow series' sum of quarters is equal to the quarters of the benchmark year:

$$\sum_{t=2}^{T} X_t = A_y, y \in [1, ..., \beta]$$

where X_t is the derived quarterly national account estimate for quarter t, I_t is the level of the indicator variable for quarter t, A_y is the annual data for year y, β^y is the last year for which an annual benchmark is available and T is the last quarter for which quarterly source data is available.

For our purposes, the indicator series is the defence spending series provide by Ramey(2011). The annual benchmark series are the defence consumption and defence investment series for the years 1939-1947 as provided by the BEA. The generated data points are presented in Figure 2 of the main paper. For the government consumption and government investment variables, the indicator variable we utilize is Ramey's (2011) government spending. Figure 19presents the interpolated data points for these:

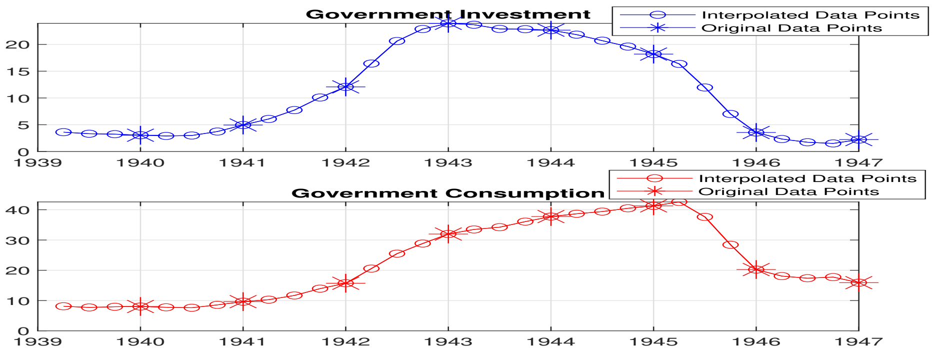

Figure 3.19: Actual and Interpolated Series For Government Consumption and Government Investment.

3.10.3 Composition of Defence Investment and Defence Consumption

Defence consumption is made up of the following large categories of spending; the first category is compensation of general government employees and consumption of general government fixed capital. The second category of defence consumption is consumption of intermediate goods and services, wherein the intermediate goods are mostly made for weapon systems (aircrafts, ships, vehicles, electronics), which are highly durable final products.

On the other hand, defence investment can be divided into three categories; spending on military structures, spending on equipment, whereby the equipment consists of weapons and weapon delivery systems (air crafts, missiles, ships, vehicles) and spending on intellectual property products that include software and research and development expenses.

3.10.4 Deviations From The Taylor Rule

Figure 20 presents the deviations of the nominal interest rate from a simple Taylor Rule's prescriptions for the time period 1939-2019. In accordance with Ramey and Zubairy, 2018 The Taylor Rule is constructed as follows;

$$nominalinterestrate = 1 + 1.5 \times yearlyinflation + 0.5 \times outputgap$$

wherein the output gap is constructed as the percentage difference between real GDP and potential GDP, measured as a fitted 6th degree polynomial and the inflation rate is on a year on year basis.

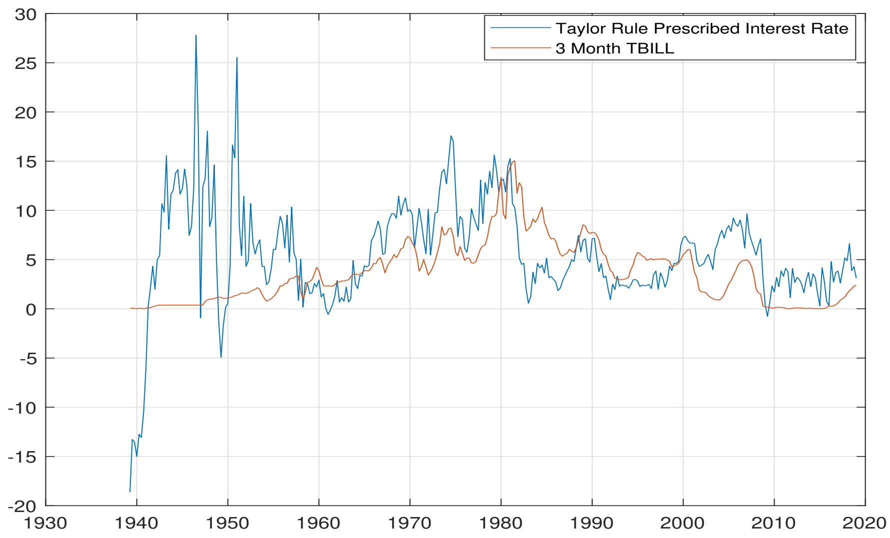

Figure 3.20: 3 Month T-Bill and Taylor Rule Prescribed Nominal Interest Rate

3.10.5 Newspaper Excerpts Anticipating The Increased Armaments Spending

In this section we present some of the popular press segments that indicate the wide anticipation of an increase in armament spending, as presented in **gordon2010end.**

For example on June 22, 1940 we read on Business Week:

> National Defense has become the dominant economic and social force in the United States today. It has created a new industry – armament – the ramifications of which will reach into every phase of our business life, and bring increased employment, higher payrolls, widening demands for machinery, and the construction of new factories."

Gordon and Krenn (2010) continue with their description of the economy prepping up its defence infrastructure expenses, pointing to clear anticipated news effects:

> The next big event that increased aggregate demand was the passage in September 1940 of the Selective Service Act that instituted the military draft and authorized an army of 1.2 million men. Military personnel on active duty increased from 458,000 on June 30, 1940, to 1,801,000 on June, 30, 1941 (Vatter, 1985, p. 8). On August 19, 1940, Congress passed another supplemental appropriations bill, raising authorized spending for the 1940-41 fiscal year from the 5 billion of June to 10 billion, including 1 billion (or one percent of GDP) for new factories to produce war material (New York Times, August 20, 1940). The Navy planned for a two-ocean navy "ready by 1944." FDR made what many regarded as an implausible promise to build 50,000 planes by the end of June 1942. The acceleration of economic activity is evident in the report "since June, Ford, General Motors, Chrysler, General Electric, Westinghouse, and practically all of our great mass-producing corporations have begun work on war orders, from radio equipment to twenty-five-ton tanks. But to transform the assembly lines means a terrific retooling."

3.11 Forecast Error Variance Decompositions Of The Baseline Models

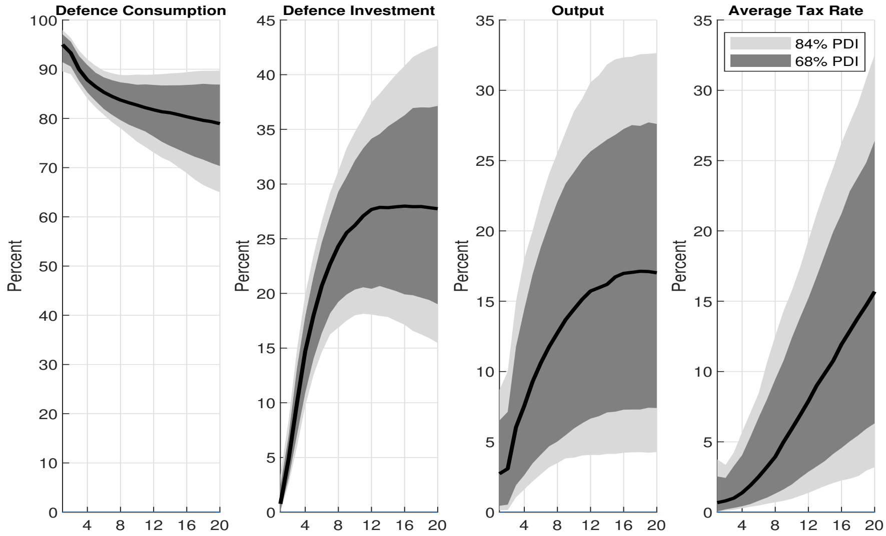

Figure 3.21: FEV To A Defence Consumption Shock At Normal Times

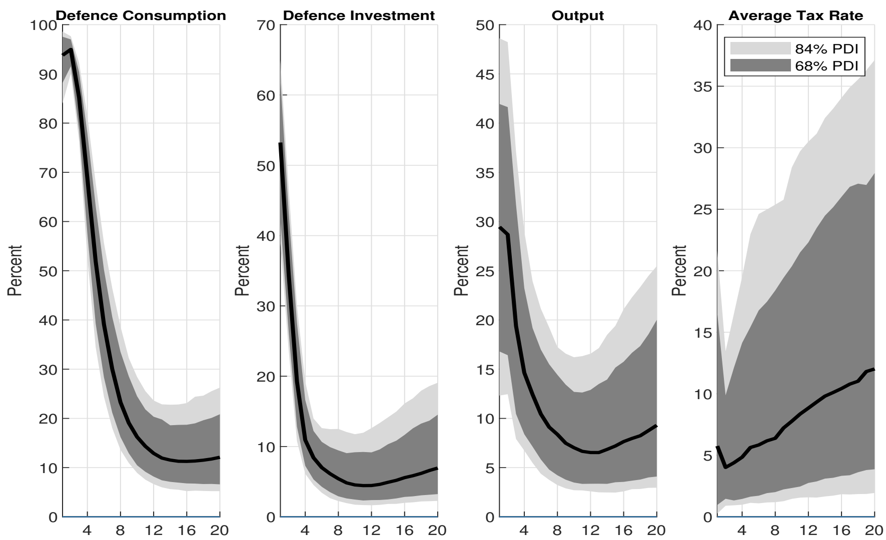

Figure 3.22: FEV To A Defence Consumption Shock At The ZLB

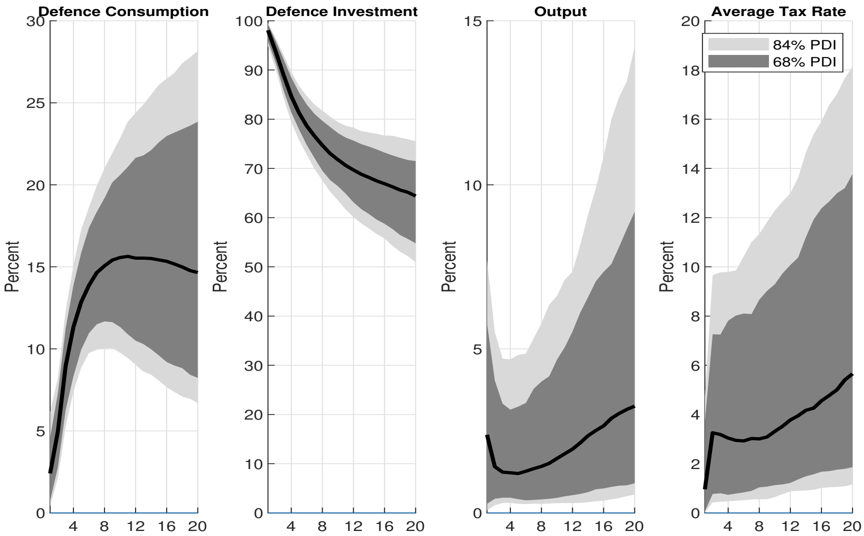

Figure 3.23: FEV To A Defence Investment Shock At Normal Times

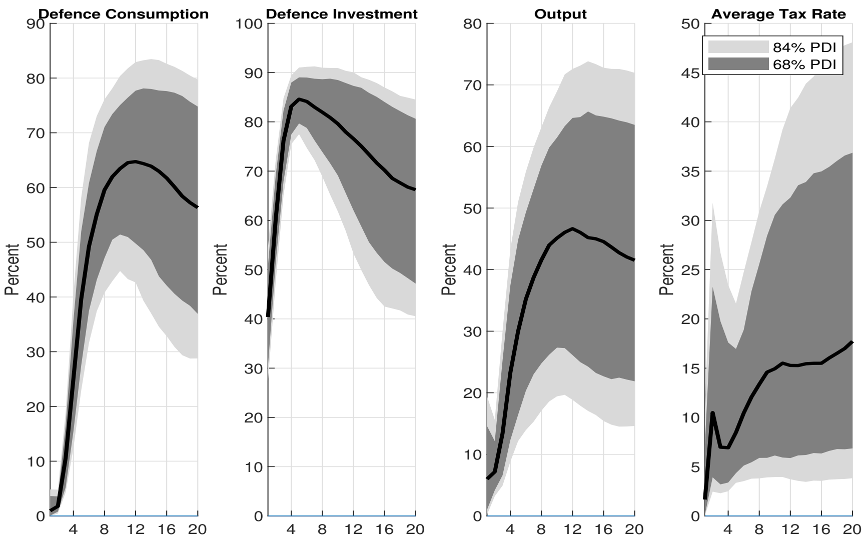

Figure 3.24: FEV To A Defence Investment Shock At The ZLB

Chapter 4

Technology Booms And Asset Price Bubbles

4.1 Introduction

Do technology booms lead to credit and financial expansions and ultimately busts? If so, what is the mechanism behind such a pattern? We contribute to this issue by developing novel measures of technology revolutions and financial bubbles and using them in a variety of econometric specifications to establish novel stylized facts and document deeper relations between the technology and the macrofinancial sector.

The idea that financial booms and busts are the results of predictable cycles rather than random events is by no means novel. Narrative evidence trying to uncover patterns has been abundant in the public sphere and is in a sense a widely adopted intuition of both pundits and commentators. Seminal in the formation of this view has been the work of Minsky (1977) and Kindleberger, Manias, and Crashes (1996) who provide a narrative account of how financial collapses unfold following a period of credit and stock market frenzy, or in their words 'the booms and the busts'. One can break down the hypothesis as follows; an exogenous event hits the economy, a 'displacement' in Kindelberger's terms, and if long and pervasive enough, it alters the economic outlook of the agents who start anticipating increased profit opportunities in the credit and stock market. This optimistic outlook leads to an expansion in the credit supply with businesses and individuals not wanting to miss out on the new financial opportunities laying ahead. Feedback loops are activated and the economy soon gets into a 'euphoria' state. Increased optimism about the growth prospects of the economy and financial speculation become the main characteristics of this phase. Soon, the agent's judgments becomes skewed, extrapolating their current profit margins to the future and expecting that they should continue. At this point their cognitive model diverges from rational expectations towards irrationality. As a result, assets are priced away from their fundamentals thus leading to the 'mania' state when asset bubbles are forming. This goes on until a negative event, the so called Minsky moment, bursts the bubble, panic ensues and cycle is led to its bust phase with asset prices collapsing, potentially taking down with them the whole economy. This is a pattern, according to Kindleberger, Manias, and Crashes (1996) entrenched in the structure of finance capitalism.

Despite the intuitive appeal of such a hypothesis, only recently has it started to come under the scrutiny of academic economists. In the past decade advances in the empirical literature have focused on establishing solid empirical evidence for the existence of the mania and bust phase of the above described cycle. Today, this is a well documented connection. In this regard, in a seminal paper, Jordà, Schularick, and Taylor (2013) utilizing a large historical panel dataset dating from 1870, provide evidence that financial crises are prceded by a credit market boom. López-Salido, Stein, and Zakrajšek (2017) provide evidence for a connection between high credit market sentiment and a future contraction in economic activity with a US dataset. More recently, Greenwood et al. (2020) provide evidence for the same pattern with a panel dataset spanning post war economies. In summary, these papers converge towards empirically establishing a common pattern, that an overheating credit market is a highly informative signal about the future detioration of the economic outlook. The implications of the existence of such patterns are substantial for the policymaker. If such patterns indeed exists, this implies that financial crises might very well be predicted and therefore also prevented.

While there is by now substantial empirical evidence that asset price bubbles and the ensuing financial crises might very well be a byproduct of credit expansions, little attention has been given to the first scale of the pattern suggested by Kindleberger and Minsky, namely what is the event that sets off the credit expansion and thus the beginning of the boom and bust cycle in the first phase. In looking for such an event, we are effectively looking to close down on the nature of the displacement event and whether there are any empirical regularities surrounding or whether it is a completely random event. In more formal econometric terms we are looking for a structural shock, or as Ramey (2016) puts it, for a primitive economic force that is uncorrelated with any other forces moving the economy and is economically meaningful. This displacement cannot be any shock however. As Kindleberger, Manias, and Crashes (1996) puts it 'a displacement consists of events that change the economic and financial environment, extend the horizon and alter the expectations'. So it is clear that we are looking for an event, forward looking in nature, that gets entangled in the forward looking expectations of the agents and fundamentally alters them.

Some culprits as to what that primitive shock that sets off the boom and bust cycle might be have been suggested; a rise in income inequality is a theoretically motivated such event (Mian and Sufi (2018) and Kumhof, Rancière, and Winant (2015)), with the underlying logic being that wealthy people are accumulating funds causing the non rich agents to borrow more. A more empirically motivated as well as supported by more casual considerations event is a financial liberalization shock as that in Japan, the Nordic countries in the '80s and Russia in the '90s. A body of research has focused on the implications of such a shock with emphasis in South America as Mian and Sufi (2018) concisely summarize. However, apart from being a low frequency event, or at least not high frequency enough to be the researcher to extract robust empirical patterns regarding its connection to financial crises, there are also questions regarding its fundamentaleness. As Sufi and Taylor (2021) point out it is very likely that liberalization 'is an opening gate enabling other fundamental forces in the local or global economy to play out'.

Another potential candidate as the trigger of a boom and bust cycle is a technology shock. This idea appears popular in the public sphere. Indeed, the archetypal case of a boom and bust cycle in the public debate is the dot com bubble of the '90s, where the anticipation for a widespread adoption of the Internet induced a speculative craze driving internet related stocks far away from their fundamentals valuation, until the expectations were revisited and a sell off followed. In the years following the dot com bubble, this chain of events was generalized into a hypothesis that technological revolutions and financial bubbles are going hand in hand is akin to fundamental law of capitalism. In a more systematic effort to empirically establish and model this pattern, Perez (2003) traces connections between technological developments and financial bubbles since the dawn of capitalism and the first Industrial Revolution up until the internet bubble, albeit in her model the unfolding of the dynamics plays out at exceedingly low frequencies that can last up to 50 years. The common baseline assumption in these explanations is that markets are irrational. In a Bayesian modelling approach, Pástor and Veronesi (2009) develop a DSGE model with Bayesian learning wherein the bubble like behavior is due to the stocks of a firm or a sector rising due to good news about future productivity and then falls as the technology becomes widely adopted and the risk falls. The empirical predictions of their model is backed up with evidence from the US rail road craze in 1830-1861 and the Internet dot com bubble of the '90s. Cao and L'Huillier (2018) similarly present evidence that the Great Recession, the Great Depression and the Japanese downturn of the '90s were preceded by a technological revolution.

Of central importance in these accounts, implicitly or explicitly, is the process by which agents form their expectations and engage with what they perceive as future economic opportunities. It is in this sense that the beginning of the financial cycle, the displacement event, can be regarded as a news shock, an information about future economic conditions that alters the agent's optimal behavior. A large literature has developed around identifying and scrutinizing the properties of such forward looking shocks, kickstarted by Beaudry and Portier (2004). The main shocks under scrutiny are technological news shocks. One of the early seminal papers in this literature strand that attempted to establish a connection between financial markets and technology news shocks was Beaudry and Portier (2006) wherein the information contained in stock prices was used to identify a technological news shock. More recently, Görtz, Tsoukalas, and Zanetti (2022) establish an explicit connection between the behavior of financial markets and technological news shocks by documenting that a news shock leads to laxer credit market conditions in a post war US dataset. Given these contributions, we wish to bring the insights developed by this literature strand onto the question of whether news about a future technological expansion is able to trigger the boom and bust financial cycle documented above.

The identification of such news shocks, especially the technology ones, is far from trivial. Since Barsky and Sims (2011) seminal paper, the standard way to go about it is by finding the linear combination of reduced form residuals of a Vector Autoregression (VAR) system that maximize the unexplained variance of a Total Factor Productivity (TFP) measure in the short to medium term but also has zero contemporaneous impact. Despite its widespread adoption this approach comes with a number of problems; TFP data is particularly prone to measurements errors. Additionally, there is a limited availability of this data, with the main TFP datasets spanning the post war period of the advanced economies. In response to these criticisms, a number of recent papers have started utilizing the information contained in the patent applications data to identify a technology news shock (Miranda-Agrippino, Hacioglu Hoke, and Bluwstein (2019), Cascaldi-Garcia and Vukotić (2022), Klein and Linnemann (2021)). The implicit assumption in the patent based approach is that patent applications act as a signal for the future prospects of the economy, conveying information to agents about future technological developments. The agents then update their information set and act accordingly with those signals.

In this paper, we seek to bring together these strands of literature in order to answer the question about the nature and properties of the displacement event that triggers the boom and bust financial cycles. More specifically, we wish to gauge whether a technological revolution is indeed associated with the financial cycles and in particular the formation of asset price bubbles that eventually burst. Our starting point for this is the insight from the news shock literature that patent data is highly informative about the effects of a coming technology shock. Using this insight, we construct a novel variable that identifies technology booms at an annual frequency.

With this indicator at hand we then proceed to ask a number of questions: are expectations about future technology a culprit behind credit booms? Are asset price bubbles and credit growth preceded by a technological boom? And if this is happening, to what extent are these relationships causal?

To obtain reliable answers to these questions, holding under scrutiny just one country for a limited time span wouldn't be optimal. In this sense we wish to diverge from the usual approach in macroeconometric exercises of taking the post war US data towards a more holistic approach. If as the narrative accounts suggest, this is a natural characteristic of the workings of a capitalist economy, more data needs to be brought in.

The gist of our strategy involves the utilization of a large historical dataset of 15 industrialized countries that dates back to 1870. Using this dataset we construct a technology boom indicator and pertube to gauge its effects on a number of macrofinancial variables. The aim of our exercise is not to establish causality but rather to document new stylized facts regarding the connection of technological shocks and asset price bubbles. For this reason, we take an agnostic approach regarding identification. We first seek to establish an empirical connection between credit and financial expansion and technology booms using simple unconditional specifications and we then proceed to test whether this connection can survive more complex specifications that take into account the dynamics and the feedback loops of the economy. Our reasoning is that if the technology booms can survive as independent drivers of the macrofinancial fluctuations in the more complex specifications, then the results of the unconditional are likely causal. If not then they serve as stylized facts broadly informing us about the sequence of events in the aftermath of the start of a technology revolution.

Our results show that technology booms are always followed by credit and financial expansions and a subsequent bust within a five year horizon. These relationships however may not always reflect deeper causal connections. With our conditional specification we show that technology booms directly lead to a limited credit expansion and increase the probability of a stock price bubble while also leading to a boom in house prices boom that is not followed by a bust. Furthermore, in our subsample analysis the technology boom even decreases the cumulated probability of a housing bubble.

The structure of our paper is as follows: in section 2 we present in detail the data underlying our analysis and their general trends. In section 3 we lay out our novel approach of identifying episodes of technology revolutions in the data using a long panel dataset. In section 4 we proceed with a simple econometric approach to establish stylized facts regarding the relationship between technology booms and the credit and asset price movements. In the same section we delve deeper into questions of causality, examining the deeper relations between technology booms and the asset and equity markets. In section 5 we delve into the question of whether technology booms are related with asset price bubbles and we lay out a parsimonious approach to identifying house and stock price bubbles. As in the previous sections, in Section 5 we first seek to establish some stylized facts before proceeding to more complex econometric exercises that help us answer questions about the causalities. Finally, in section 6 we present our results.

4.2 The Dataset

Underlying our exercise are two large historical panel datasets: the macro history dataset compiled and provided by Jordà, Schularick, and Taylor (2017) and Madsen's [Madsen (2007), Madsen (2010), Madsen, Islam, and Doucouliagos (2018)] long panel dataset collating data on patent applications. The combined data span the period 1870-2008 featuring macro financial data for 15 industrialized countries [1]. The data is at an annual frequency.

The main question we wish to ask in this paper is what are the macrofinancial dynamics emerging in the aftermath of a technology boom. It follows that our empirical emphasis will fall on the responses of the credit and asset price markets. Before delving into the econometric analysis, we first explore the data underlying our exercise.

4.2.1 Macrofinancial Variables

For the purposes of our exercise of studying the macrofinancial dynamics of an economy around a technology boom, we will focus on the response of the credit and the asset price markets. The main dependent variables in our analysis will be the following: for the credit market we will use a credit-to-GDP ratio and a real private loans per capita variable as provided in Jorda's

[1]The countries in our sample are the following: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, Norway, Sweden, Switzerland, UK and the US

dataset. For the asset prices we make use of the broad stock market index constructed by Jordà, Richter, et al. (2021) and the house price index as obtained from the same macro history dataset. The real private loans per capita variable has been constructed using total loans to non financial private sector deflated by the CPI and divided by the total population. In choosing this variable we follow Jordà, Schularick, and Taylor (2013) approach in approximating the lending activity of the economy. The credit-to-GDP variable represents the total bank credit to GDP ratio and approximates the activity of the financial sector; the data series for this variable is derived from Baron, Verner and Xiong (2021).

The choice of such a long dataset spanning 15 industrialized economies since the emergence of finance capitalism diverges from the usual practice of macroeconometric literature of using almost exclusively post World War 2 (WW2) data. Our choice can be motivated by a range of considerations: firstly, the boom and bust cycles are relatively low frequency events, occuring rarely. A dataset spanning exclusively the post war period could therefore be considered informationally deficient for our purposes and would severely limit the scope of our analysis. Secondly, as Jordà, Schularick, and Taylor (2013) observe, by employing a dataset that goes back to 1870 and features the majority of the advanced economies it can be argued that we don't have a sample; rather we have close to the entire population of modern economic experience under finance capitalism. This makes this dataset especially appropriate for the question we wish to study whereby the boom and bust finance cycles are treated as a deeply embedded characteristic of finance capitalism rather than a single occurrence event.

Next, we gauge the general trends that might be underlying our asset and credit variables. We start with the mean and the median of the panel dataset for the credit-to-GDP ratio in Figure 4.1. We see that there are two major upward long run trends; one in the period 1870-1913 and another in the post WW2 period. The ratio also exhibits some volatility episodes with sharp spikes notably in the 1920s and in the early '90s, coinciding with two periods of excess technological innovation as we shall see in the next section.

We then move on to the real private loans variable. Figure 4.2 displays the evolution of its mean and median. The volatility and the trends are less pronounced in this variable. Throughout

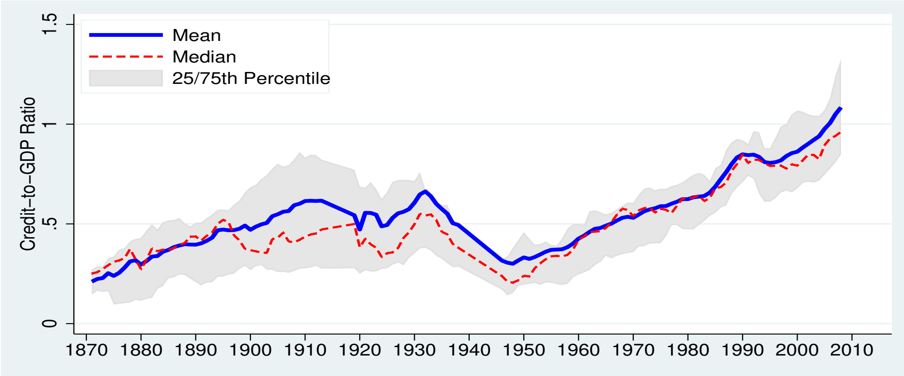

Figure 4.1: The solid line plots the mean of the credit-to-GDP ratio in the sample countries between 1870 and 2008. The red line is the median. The grey area is the interquantile range for the countries in the sample.

the sample there is a subtle upward trend indicating a persistent credit deepening. A sharp upward surge occurs again in the '1920s , as was the case with the credit-to-GDP variable. Noteworthy is also the deep drop and recovery in the '40s, no doubt a result of World War 2 (WW2).

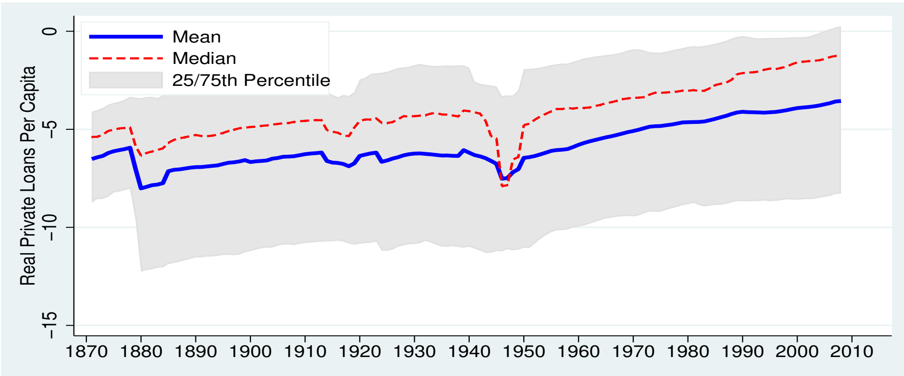

Figure 4.2: The solid line plots the mean of the log real private loans in the sample countries between 1870 and 2008. The red line is the median. The grey area is the interquantile range for the countries in the sample.

Next, we gauge the evolution of the asset price indices. Figure 4.3 displays the real stock market index (1990=100) and Figure 4.4 the real house price index underlying our analysis. On par with the general intuition, the stock price index is exhibiting notable more volatility with sharper movements, a well established empirical fact.

No clear underlying trend throughout the sample is visible for the stock price index on Figure 4.3. The 1870-1920 period saw a large cross country variation in both asset prices variables, as is evident from the wide interquantile ranges, with a relative stable mean and median. The

1920s decade stands out as we see a large rise followed by a large decrease. In the post war period the cross country volatility decreases significantly and from the 1980s onward there is a clear upsurge. This is followed by two very sharp drops in the late '90s and the late '00s, most likely capturing the aftermath of the busts of the dot com and the housing bubbles.

Coming to the house price index on Figure 4.4, we see that throughout the post WW2 there is a clear upward trend. House prices, as is the case with stock prices, remained largely stable in the 1870-1920 period. Notable again is that the cross country variation in that period is the largest in the sample as is evident from the ranges around the mean. As is the case for the other variables we looked into, in the 1920s there was a sharp upsurge in this case followed not by a collapse but by a stabilization in the house prices. We also observe an upsurge in the '00s and a brief one in the '90s, followed by a decrease.

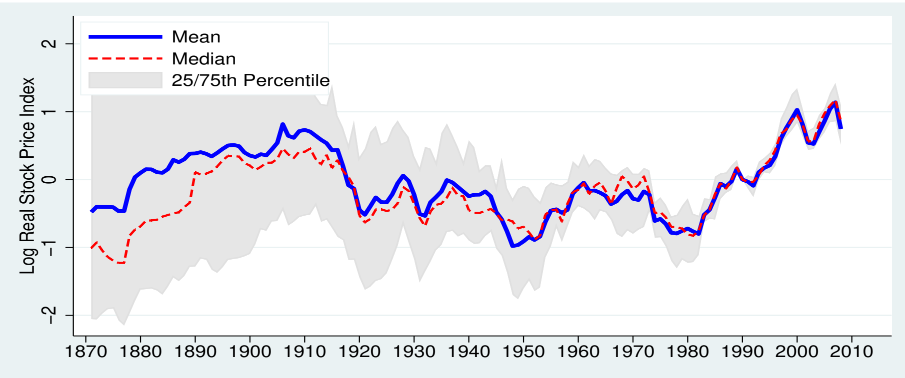

Figure 4.3: The solid line plots the mean of the log real stock price index variable in the sample countries between 1870 and 2008. The red line is the median. The grey area is the interquantile range for the countries in the sample.

What are the common patterns emerging from the movement of these variables? Firstly, it appears that in all of the variables the 1870-1920 period is distinct from the post war period; the trends, the variation and the movements of the variables are clearly different in those two periods. More specific observations can be made; the first one is that the 1920s is in all variables an idiosyncratic decade wherein a clear boom and bust occurred; a sharp increase in all variables in the early '20s followed by an equally sharp drop. The post war period is underlied by a general upward trend that reached its peak in the late '90s up until the Great Financial Crisis of 2008 (GFC). In contrast, there is no discernible common trend in the 1870-1920 period, which is also

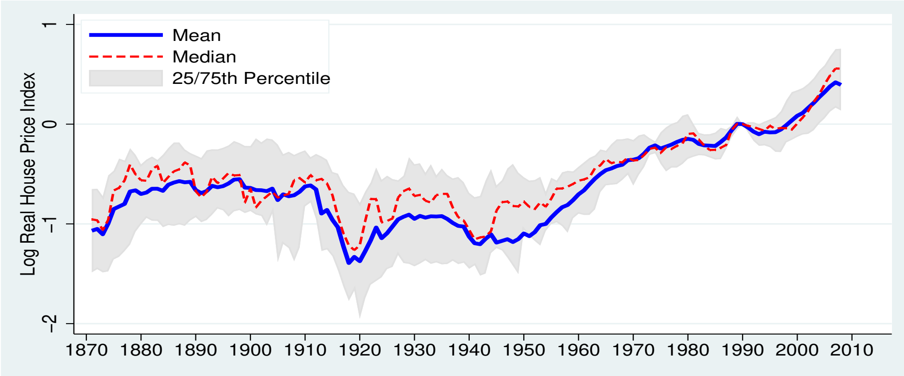

Figure 4.4: The solid line plots the mean of the log real house price index variable in the sample countries between 1870 and 2008. The red line is the median. The grey area is the interquantile range for the countries in the sample.

characterized by large cross country differences and considerable volatility. The preliminary analysis in this section therefore clearly points to discernible episodes of change in the data with the most obvious one being between the pre WW2 and the post WW2 period. The next question we wish to ask is whether any of these patterns are related to technology revolutions. For this, we need to first gauge the informational content of the patent dataset.

4.2.2 What Information Do The Patent Data Contain?

Of central importance in our approach is the utilization of the patent application dataset to identify the technology booms. For this reason we devote a separate section to the properties of the patent dataset. Before we delve into exploring the general properties of the dataset some general considerations regarding the information content of the patent application metric would be appropriate.

Utilizing the information of patent data to identify the effects of a technology shock has been a long standing practice in the literature. For example, Griliches (1998) notes that despite any weaknesses such datasets might pose, patent statistics are a unique resource for the analysis of technological developments due to their availability over such long timespans.

More recently, patent data has been utilized in the news shock literature to identify the effects of an anticipated technology shock (Cascaldi-Garcia and Vukotić (2022), Miranda-Agrippino, Hacioglu Hoke, and Bluwstein (2019)). But what is the nature of the information they contain? As Miranda-Agrippino, Hacioglu Hoke, and Bluwstein (2019) note, patent grants are an institutional framework created to standardize the introduction of technical innovations and provide to its owner the right to protect and to capitalize on her creation after a set amount of years. In this sense, a patent application is not same as the introduction of a new technology in the economy. From the time the application is submitted until the patent is granted and its effects are thus diffused in the economy, years might elapse. It would thus make sense to regard patent applications as a signal embedding information for technological changes that have not yet materialized, a measure of future technology. The agents observe the signal, form their expectations and adjust their behavior accordingly with what they expect the future economic outlook to be.

As a preliminary step in our analysis, we wish to gauge the properties of the data with a series of exercises in descriptive statistics, as in the previous section. The results here should hint to the existence of generic patterns and long term trends and should provide some guidance in the more formal econometric exercises that shall follow. This exercise should also help us discern any potential stochastic breaks and decide whether the dataset can be treated as the product of a single stochastic process or not. Figure 4.5 plots Madsen (2010) log of the patent applications pert capita of 15 industrialized country for the period 1870-2010. Upon first inspection, the first pattern that emerges is the rise in patent applications in most of the countries in the sample in the pre 1920 period. The period 1870-1920 is indeed the period associated with the Second Industrial Revolution, arguably the most cataclysmic of all technological revolutions. During that period the pace and size of changes was unmatched. Countries saw the introduction of a series of era defining inventions such as the introduction of electrical power, machine tools and the telephone.

The trend however clearly changes in the post WW2 era. The number of patent applications flattens out in most countries and there is a divergent behavior across the sample, in contrast with the pre WW2 period. While the evolution of the variable in most countries cannot match the pace that it had gained in the 1870-1920 period, there are some notable exceptions to this observation; Japan with a consistent upward trend and the US experiences a sharp increase in

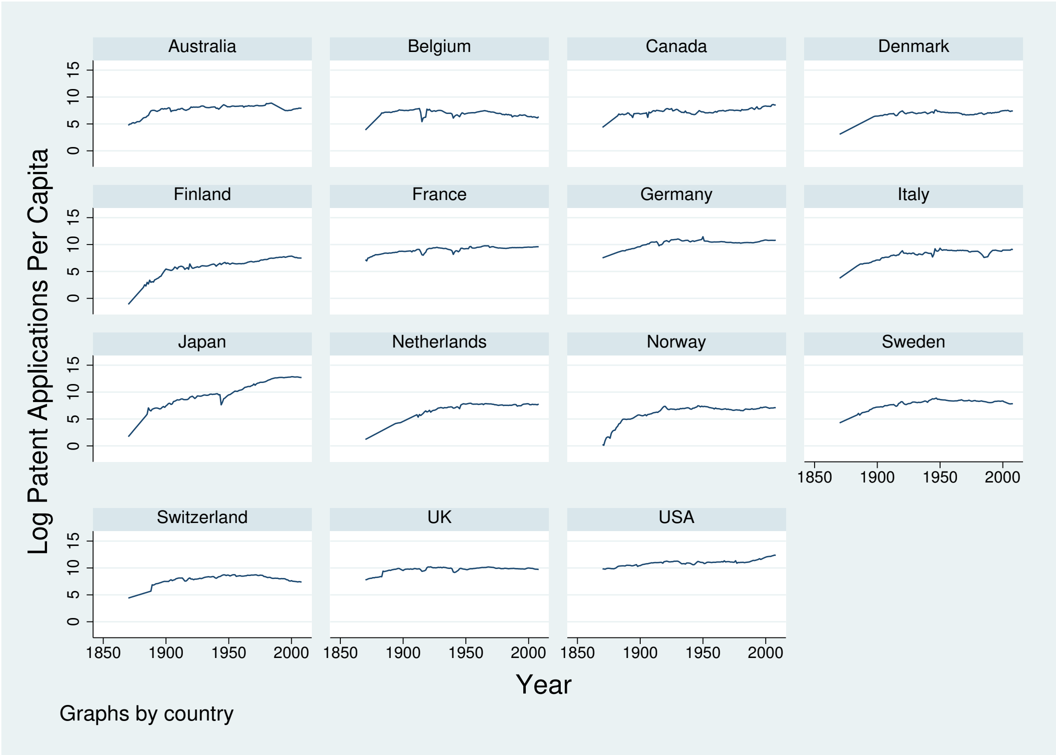

Figure 4.5: Log Patent Applications Per Capita.

patent applications in the '90s.

Those observations are further elucidated by the mean annual growth rate of patent applications within those historical subsamples as presented in Appendix . The pre-1920 period in all countries saw the sharpest increase in technological innovation.

Those observations and trends show much clearer if we resort to the mean and median plots as in the previous section. In Figure 4.6 we can observe the mean, the median and the interquantile range of the variable throughout the sample. The first thing that stands out is that there is a large divergence between the mean and the median indicating, as observed before, large cross country variation in patent applications. It is obvious that some countries are having much higher patent activity than others in the post war period and this gap keeps growing as we get closer to the present day. The second notable characteristic is the very steep slope of the mean, indicating that the countries skewing the sample experienced an explosion in patent applications while others remained stagnant.

This should intuitively make sense. In the post war period some countries, like the US and

Japan emerged as global technology hubs. A sharp increase in trade and knowledge inflows during that period means that the countries in the sample were able to import technological innovations more easily and were thus less reliant on producing heir own.

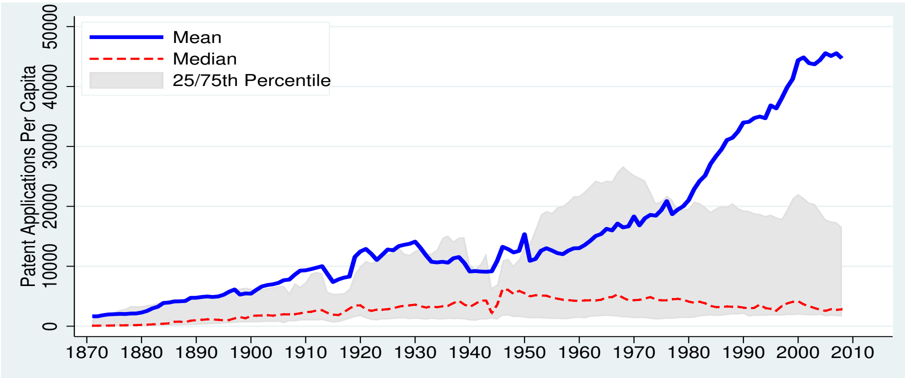

Figure 4.6: The solid line plots the mean of the log patent applications per capita in the sample countries between 1870 and 2008. The red line is the median. The grey area is the interquantile range for the countries in the sample.

The mean of the data therefore shows a clear structural break that should inform our approach. The fact that in the post WW2 period the fact that a handful of countries were leaders in the technological innovation and the main exporters of knowledge could potentially render a technology boom identification based on country specific patent variables more fragile.

It is for this reason and the observations that in the exercises that follow we are going to utilize a full sample benchmark specification and a post WW2 specification as a constant robustness check to our results. In addition, following Jordà, Richter, et al., 2021 in all of our empirical explorations below, we exclude the World War 1 and 2 years as well as 5 year windows around wars.

4.3 Empirical Identification of Technology Booms

In this section we propose and develop a novel way of identifying a technology boom using the information contained in the patent dataset. The key concept in our identification is a measure of 'excess' patent applications, where by excess we mean an upward deviation from a long run trend. The underlying assumption here is that patent applications constitute an accurate signal for the future technological developments in an economy given that they are observed by the agents. An upward surge in patent applications should therefore indicate that the agents are expecting a future technological boom and are adjusting their expectations and decisions accordingly. In defining a technology booms using the deviation of a variable from its long run trend, we take inspiration from a substantial body of literature that empirically identifies asset bubbles as a deviation of the variable above some specified threshold (Borio and Lowe (2002) and Detken and Smets (2004))

Given this, we propose the following criterion to trace a technology boom in the sample: we require that the log patent applications per capita variable becomes elevated from a country specific Hodrick Prescott (HP) patent applications filtered trend above a specific threshold. Simply put, our criterion traces a tech boom wherever there is an excessive positive difference between the HP filtered variable and the original variable.

To be more precise, we parameterize this criterion as follows: for the construction of the HP trend we employ a smoothing parameter of $\lambda = 100$ on our annual patent applications data. For the elevation threshold that determines whether a country is experiencing a technology boom or not we choose the standard deviation ϕ_i of the difference ζ_{it} between the trend and the log variable. Specifically and to account for the fact that there exist substantial spikes in the data the data that could potentially skew the variable, we require the difference ζ_{it} to be 0.5 times the standard deviation ϕ_i for the patent series of each country throughout the sample. Our technology boom is then an indicator variable that takes the following form:

$$TechBoomSignal_{it} = I(\zeta_{it} > 0.5\phi_i) \tag{4.1}$$

This criterion produces a discrete sequence of indicator variables that determine whether the country is experiencing a technology boom or not. Figure 4.7 plots the years that our algorithm captures as technology booms using the method specified above. Upon first inspection some clear trends emerge: we see that in the 20th century there are three decades characterized by a technological boom that are common across almost all of the countries in our sample. These periods are in the 1920s, the 1960s and the 1990s. These results make sense. These decades were indeed associated with waves of technological booms; the 1920s was the period of mass electrification in the advanced industrialized economies. In the 1960s the electronics boom took place as characterized by Malkiel (1999). And the 1990s was the period of the IT revolution, the most recent technological revolution that has been closely associated in the public sphere with the subsequent crash. [2]

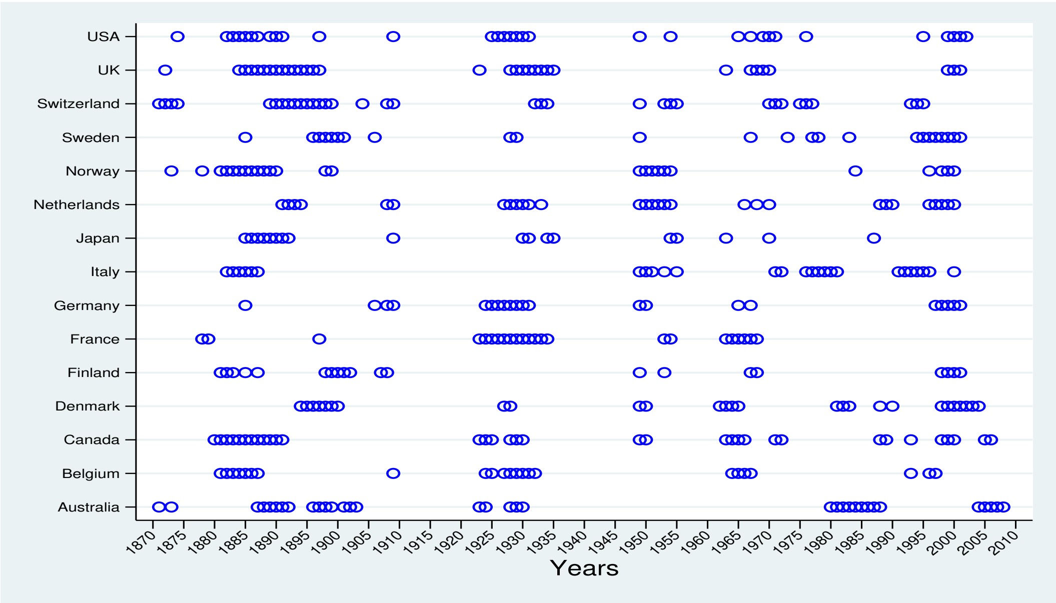

Figure 4.7: Years wherein there was a Tech Boom as detected by the signal of Eqn. 1.

Apart from these well established episodes of technological booms in the 20th century, there are also clusters of technological shocks forming in the latter quarter of the 19th century. This pattern again broadly corresponds with the historical experience; that was the period when the effects of the Second Industrial revolution had started rippling through the industrialized societies and catalyzing immense changes in economic life. We can therefore see that our indicator is able to capture technological booms in a relatively simple and precise manner.

To what extent however can these episodes, or at least their beginnings, be considered structural shocks? Can we use them as a valid proxy for an exogenous technological shock? As a general comment on this, a reasonable assumption would be that patent applications are

[2]The criterion we use is of course inherently ad hoc. We calibrated the parameters and the standard deviation such that they match the historical experience rather than estimated them. In doing so, we closely followed the approach developed by Jordà, Schularick, and Taylor, 2015. An estimation of those parameters is beyond the scope of our exercise.

to a certain degree predetermined in an econometric system, in the sense that when agents are submitting a patent application considerations regarding the state of the economy in the near future should not play a pivotal role. We don't wish to stretch this reasoning however and assume that a perturbation in patent applications can act as an exogenous shock so that we can maintain a flexibility in our approach.

A related valid concern might be that the indicator variable might by construction be capturing regulatory changes in the patent applications framework that makes it easier or harder for agents to submit patents, rather than a structural surge in patent applications. However, as Miranda and Agripino (2020) argues this shouldn't be a big concern for our purposes; it is highly unlikely that these regulatory changes were made with current or future economic conditions in mind. If anything therefore, instead of posing a problem they should provide some useful exogenous variation for our identification purposes.

4.4 Macrofinancial Impact of Technology Booms

With our technology boom indicator at hand our next task is to explore the macrofinancial consequences of such a boom. Is the coming of a technology boom, or the prediction of its arrival, associated with the boom and bust dynamics of the financial cycle? In other words, are the technology booms the culprit, the primitive shocks, that alter the expectations of the agents and lead to an overheating of the credit market and the subsequent asset price bubbles? And if so, what is the exact mechanism behind these processes? To approach these questions we utilize our indicator variable as the treatment in a variety of econometric specifications and gauge at the Impulse Response Functions (IRFs) of the dependent variables of interest.

Unpacking the causal relationships of such a question is by no means an easy task. The booms and bust cycle contains complex dynamics and causal chains unfolding in complicated sequences that might be challenging to untangle. A technology boom causes a credit boom which in turn causes an asset price boom which in turn leads to a bubble. As a starting point therefore we first wish to establish the existence of such a sequence of events as stylized facts.

In other words, in this section we ask the question: do technology booms indeed precede credit and asset price booms? To answer this, we are going to use an unconditional econometric specification, more akin to forecasting rather than a structural exercise.

This is the approach that will inform our process for the rest of the paper. Namely, we ask two separate questions, one concerning the stylized facts of processes in the aftermath of a technology boom and one wishing to look further into any deeper relationships that the data might be hiding.

4.4.1 Unconditional Paths: Establishing A Sequence Of Events

We follow Jorda's et al (2013) approach and start by considering a stripped econometric specification with minimal information. Here, we are firstly interested in characterizing a simple unconditional path of the cumulated response of a variable y to a treatment x at a time t. This is a very flexible approach that allows us to relax the assumptions about the exogeneity of the technology boom indicator. In what follows therefore, the treatment might or might not be random, as mentioned above, making causal claims isn't the purpose of this exercise. Econometrically, we proceed with the following specification:

$$CR(\Delta_h y_{it+h,\delta}) = E_{it}(\Delta_h y_{it+h} | x_{it} = x_{it} + \delta) - E_{it}(\Delta_h y_{it+h} | x_{it} = x_{it}) \tag{4.2}$$

where $CR(\Delta_h y_{it+h,\delta})$ denotes the average cumulated response of the variable across countries h periods ahead in the future and δ is the perturbation to the treatment variable x. The treatment in this case is the first occurrence in a discrete sequence of technology booms. If there is no sequence of such booms, a single occurrence still counts as a treatment. Notice that we are not conditioning on any information regarding the current or the future state of the economy in this specification. Our aim is simply to establish an association of events. More structural exercises establishing more formal claims will follow in the next section.

We start with the response of the credit market to the technology boom. The variables we are going to consider here are the credit-to-GDP ratio and the real private loans per capita variable. For the treatment variable x we are going to utilize the first year in a positive sequence of the technology boom indicators. Figure 4.8 plots the unconditional cumulated responses to a perturbation in the technology boom indicator with their 90% confidence intervals in the full sample and the post WW2 era. Both variables rise significantly in both samples in the aftermath of the technology boom. It appears at first sight therefore that a credit market overheating follows a tech boom. [3]

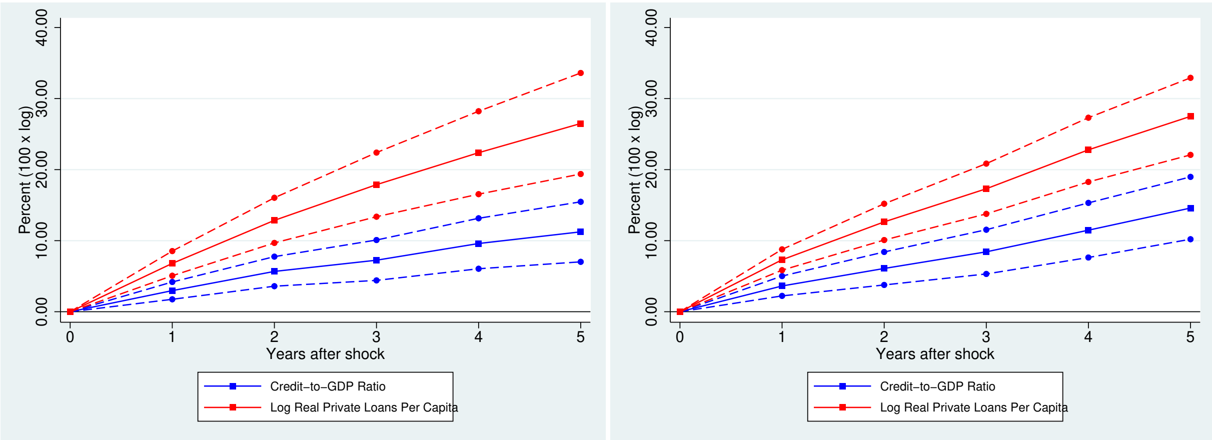

(a) Unconditional Response Of Credit Variables To A Technology Boom-Post WW2 Sample

(b) Unconditional Response Of Credit Variables To A Technology Boom-Full Sample

Figure 4.8: Unconditional Paths of Credit Variable: Post WW2 and Full Sample. 90% Confidence Bands

In the next iteration of our exercise, we are interested in the response of the asset prices in the aftermath of a technology boom. In doing this, we are interested in the behavior of the housing market and the equity market separately. In this we follow Jorda, Schularick and Taylor's (2015) findings that the unfolding and effects of a housing market boom and bust can differ substantially from an equity market one. In Figure 4.9 we gauge the unconditional responses of the real stock market index in the post WW2 and the full sample. The results demonstrate a clear upward surge in the stock market in the aftermath of the beginning of a technology boom in both samples. In both cases the increase is steep in the initial two years and thereafter it stabilizes.

On a similar note, we move on to the response of the housing market on Figure 4.10. The

[3]In the Appendix, figures 4.21 and 4.22, we present the IRFs for horizons up to 15 years. The shocks follow a non stationary path over such longer horizons. It should be kept in mind that the shocks in this case are not structural shocks but perturbations and don't have a clear causal interpretation. We are thus merely establishing a sequence of events rather than a deep structural relationship

picture is similar with the stock market. In both samples there is a significant rise in the value of the house price index in the aftermath of a technology boom. Unlike the stock price however, this rise doesn't flatten out after an initial upward surge but rather it is maintained throughout the five year horizon of our exercise. That's the case for both samples.

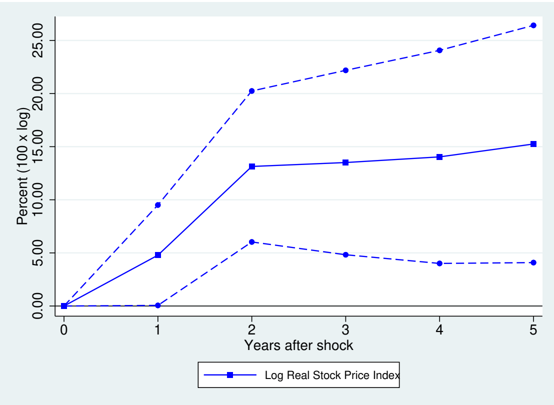

(a) Unconditional Response Of Real Stock Market Index (1990=100) To A Technology Boom-Post WW2 Sample

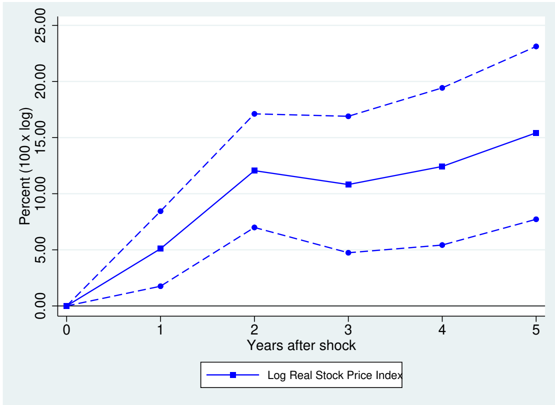

(b) Unconditional Response Of Real Stock Market Index (1990=100) To A Technology Boom-Full Sample

Figure 4.9: Unconditional Paths of Real Stock Price Index: Post WW2 and Full Sample. 90% Confidence Bands

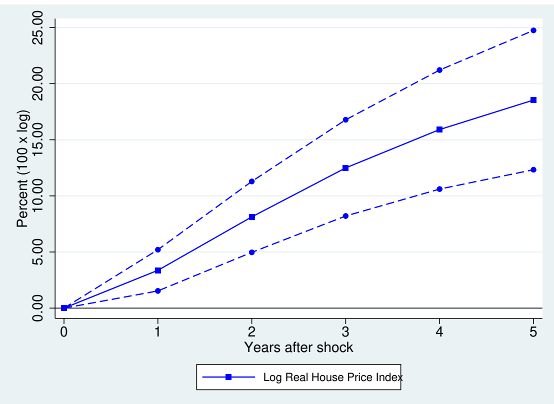

(a) Unconditional Response Of Real House Price Index To A Technology Boom-Post WW2 Sample

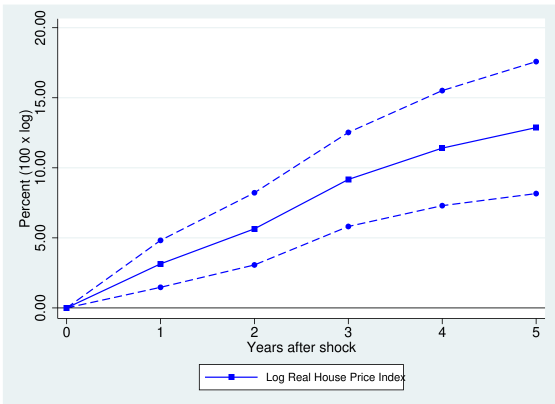

(b) Unconditional Response Of Real House Price Index To A Technology Boom-Full Sample

Figure 4.10: Unconditional Paths of Real House Price Index: Post WW2 and Full Sample. 90% Confidence Bands

As a final iteration in this exercise, we wish to gauge the reaction of the economy in the aftermath of the technology boom. Is a technology boom followed by an expansion? Figure 4.11

presents the response of the real per capita GDP variable to a perturbation to our indicator. GDP in both samples increases sharply in a significant manner, thus confirming that an expansion indeed follows.

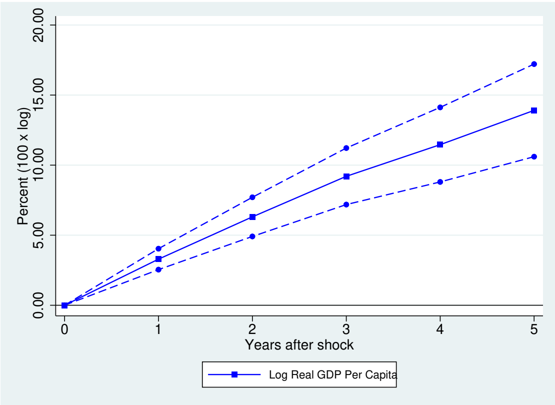

(a) Unconditional Response Of Log Real GDP Per Capita To A Technology Boom-Post WW2 Sample

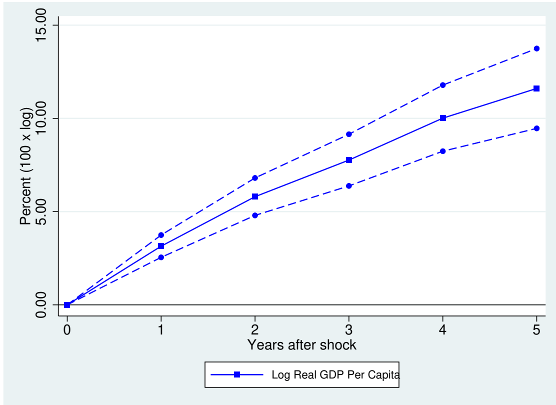

(b) Unconditional Response Of Log Real GDP Per Capita To A Technology Boom-Full Sample

Figure 4.11: Unconditional Paths of Log Real GDP Per Capita: Post WW2 and Full Sample. 90% Confidence Bands

What have we learned from this exercise? We view this exercise as a first step in establishing some stylized facts regarding the unfolding of technology revolutions. To our knowledge, we are the first to document in a large historical dataset that an overall expansion takes place in the aftermath of a technology shock as approximated by our patent indicator. As we have flagged many times already, we view these results as stylized facts rather than a product of econometric identification. This is a finding that meshes well with general intuition and it connects with the findings of large literature on the effects of anticipated technology shocks. What remains to be seen is if our treatment variable can survive as a driver of this expansion in more complex specifications or in other words whether it is the result of endogenous dynamics or rather a deeper economic relationship.

4.4.2 A Conditional Specification: Establishing Causality

In the previous section we established the existence patterns regarding the evolution of the economy in the aftermath of a technology boom. The boom is clearly followed by an expansionary period, with credit and asset variables following an upward path with no sign of downward movements within a five year horizon. While no causal claims were made, this is some first evidence in support of the narrative account of the boom and bust cycles, wherein a fundamental 'displacement', in our case the technology boom, alters the expectations of the agents and eventually leads to a bubble. Whether deeper relationships in the data hide behind these results, we study in this section.

One concern with the unconditional results is that we ignored the complex dynamics of economic processes, with the numerous feedback loops getting activated during them. Would the reaction of the credit and asset markets to a perturbation of the technology boom dummy be the same if we enrich the dynamics of our specification? By utilizing richer specifications we make it far less likely that the technology booms survive as a driver of the overall expansion we documented. If they survive however, it is much more probable that a causal relationship between technology booms and asset and credit market booms might be in place.

To unveil whether the data is hiding any deeper relationships our specification of choice is the Local Projections (LP) framework as developed by Jordà (2005). This is a simple semi parametric approach of constructing IRFs to shocks of choice. To gauge the responses of the credit and asset market variables we are then going to utilize the following specification:

$$\Delta_h y_{it+h} = a_i + \beta_h x_{it} + \sum_{j=0}^{p} \Gamma_j X_{it-j} + u_{it}, h = 1, ..., H \tag{4.3}$$

where a_i are country fixed effects, x_{it} is again the first occurrence in a sequence of our technology boom indicator variable and the vector X contains a history of $p = 1$ lags of the control variables at time t with coefficients Γ and u is the error term. The vector of controls includes the following variables: GDP, inflation rate, the short term interest rate, the private investment per capita rate and their first lags as well as the first lag of the patent applications per capita variable. Using country fixed effects is a simple and easy way to account for the cross country variation of the shock.

The notation $\Delta_h y_{it+h}$ denotes the percentage change of the dependent variable of interest h periods ahead for each of the countries $i = 1, .., N$ in our sample. The percentage change is given by 100 times the logarithmic difference of the variable. It thus has a natural interpretation as a growth rate.

The coefficient on the tech boom indicator β is the chief coefficient of interest here. We retrieve this coefficient after running forecasts with specification (3) up to $h = 5$ horizons ahead and use it to construct the Impulse Response Functions (IRF) of the variables of interest to the perturbation. The resulting IRF should then represent the conditional path for the cumulated response of each variable y controlling for a history X.

As before, we commence with the response of the credit market. Figure 4.12 presents the conditional paths of the credit variables to a perturbation in the technology boom indicator. The results differentiate themselves from the unconditional specification. For the credit-to-GDP variable we see that the paths are not uninterruptedly upward. Rather, they resemble the hump shaped response characteristic in the news shock literature in both samples. In the post WW2 sample credit-to-GDP moves upward until period 3 and then reverts back to insignificance. In the full sample the increase is even weaker, with the variable rising only in the 1st year of the aftermath of the technology boom.

Moving on to the loans per capita variable, we observe that it is insignificant in all horizons, with a notable exception of the 1st year in the aftermath of our shock in the post WW2 sample, wherein we observe marginal significance within the 90% bounds.

Into the specifics, we see that the technology boom survives as an independent driver in the increase in private loans per capita in both the full and the post war sample, at least for the first two years after the shock. However, this credit boom lasts only for a couple of years before the variable reverts back to statistical insignificance.

Next, we gauge the response of the asset variables to a technology boom. We start with the stock market index in Figure 4.13. Again, the results are different from the unconditional specification. The index never attains statistical significance in neither of the samples.

Turning to the real house price index in Figure 4.14 the picture we get is different. In contrast to stocks, the house prices rise significantly in both samples. In the post WW2 sample the rise is steep and upward until the 3rd year after the tech episode and from there after it stabilizes at

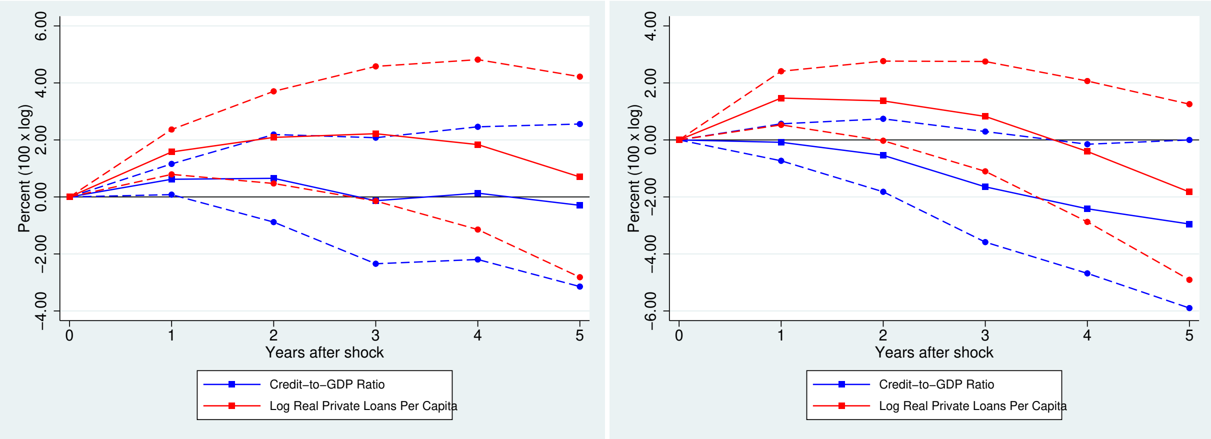

(a) Conditional Response Of Credit Variables To A Technology Boom-Post WW2 Sample

(b) Conditional Response Of Credit Variables To A Technology Boom-Full Sample

Figure 4.12: Conditional Paths Of Credit Market Variables: Post WW2 and Full Sample. 90% Confidence Bands

a level of about 5%. In the full sample the rise is significant in all horizons flattening out at the 3rd year after the episode. [4]

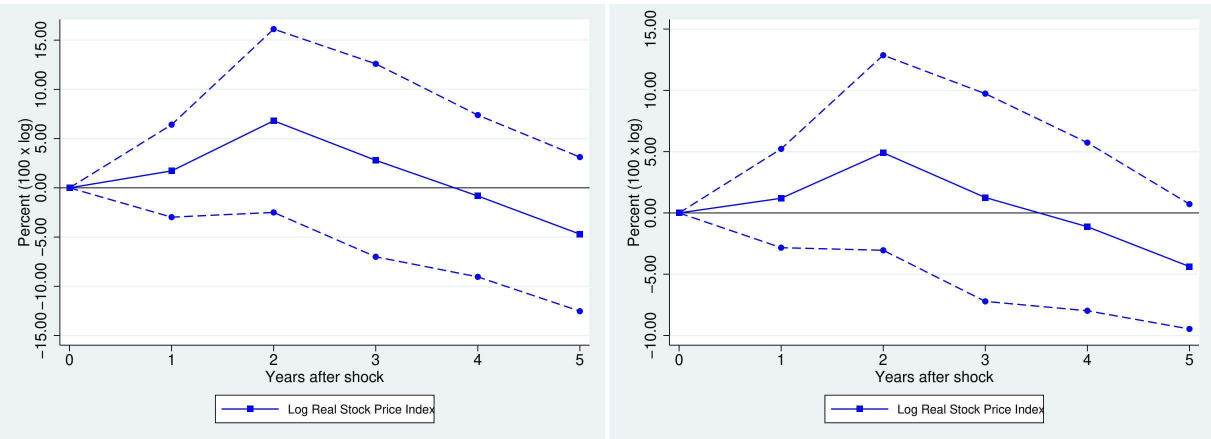

(a) Conditional Response Of Stock Market Index To A Technology Boom-Post WW2 Sample

(b) Conditional Response Of Stock Market Index To A Technology Boom-Full Sample

Figure 4.13: Conditional Paths of Stock Market Index: Post WW2 and Full Sample. 90% Confidence Bands

Finally, to examine the conditional path of the economy and to provide some context for the credit and asset market results presented above, we gauge the GDP per capita variable. Figure 15 presents the results. The path of the variable is upward in a bell shaped manner. There is an increase in the aftermath of the tech boom, this increase stabilizes and then starts decreasing

[4] In the Appendix, figure 4.23, we present the 15 year IRFs to the shock.

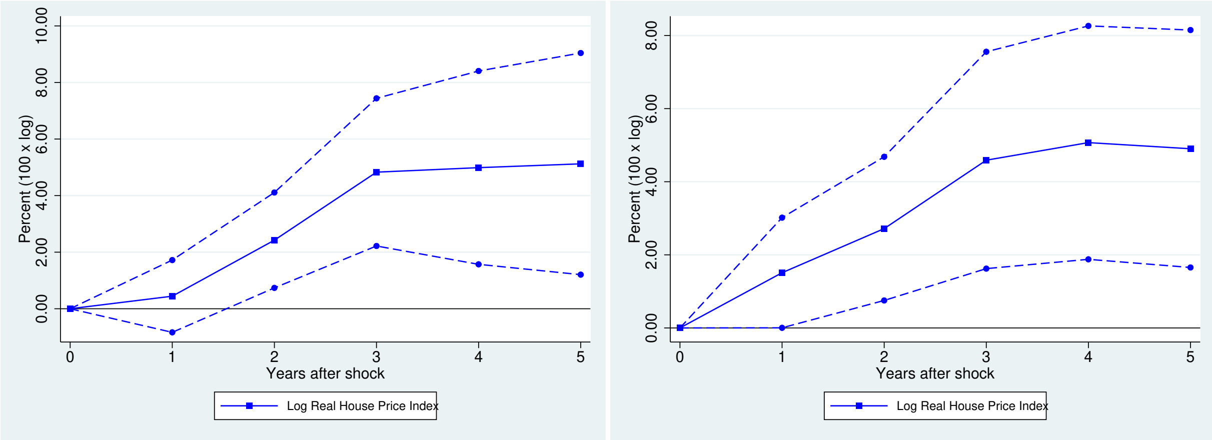

(a) Conditional Response of House Price Index To A Technology Boom-Post WW2 Sample

(b) Conditional Response of House Price Index To A Technology Boom-Full Sample

Figure 4.14: Conditional Paths of House Price Index: Post WW2 and Full Sample. 90% Confidence Bands

again. The shock propagates gradually until it the 3rd year whence its effect reaches its peak and after that its effects commence subsiding. The gradual propagation of the shock reaching its peak after some years is encouraging; it shows that we are indeed capturing an anticipated technological change as approximated with a TFP shock in the news shock literature .

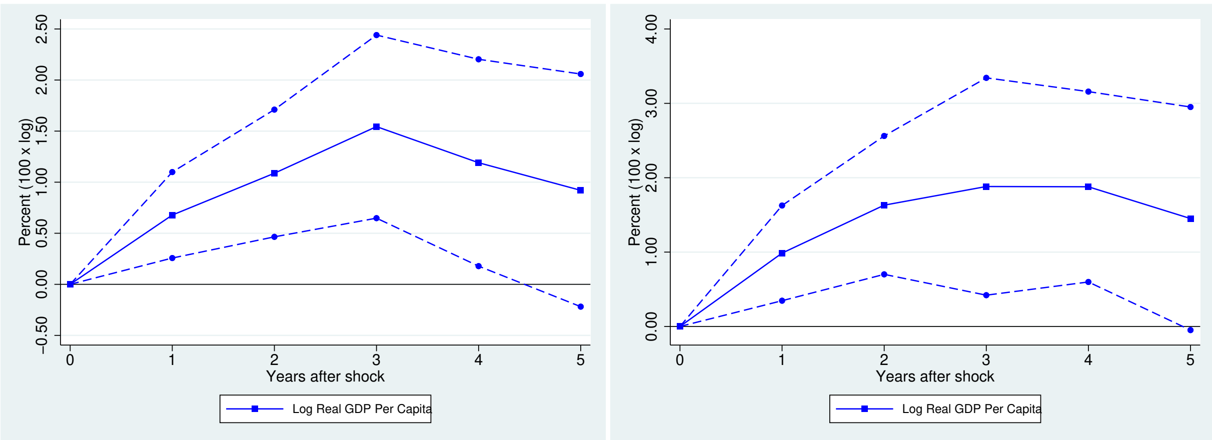

(a) Conditional Response of Log Real GDP Per Capita To A Technology Boom-Post WW2 Sample

(b) Conditional Response of Log Real GDP Per Capita To A Technology Boom-Full Sample

Figure 4.15: Conditional Paths of Log Real GDP Per Capita: Post WW2 and Full Sample. 90% Confidence Bands

So what are the main takeaways from this exercise? The fact that the house price and stock price response are different is in line with the findings of a large body of literature that separates treat these asset classes separately (Borio and Lowe, 2002; Detken and Smets, 2004; Jordà, Schularick, and Taylor, 2015). The conditional expansion of the house price could be hinting towards the fact that the technology shock leads to a permanent increase in house prices that is not followed by a sharp decrease. The reaction of the stock market index is more in line with the latest shock literature that finds no effect of a tech news shock on the stock prices. The credit market is led to an expansion with the private loans increasing significantly and peaking two years after the shock. However, overall, the credit expansion in the economy is not taking place at the same pace with the rise in GDP as we can see from the credit-to-GDP variable.

4.5 Do Technology Booms Lead To Asset Price Bubbles

We have established that a technology boom leads to a credit expansion and to an increase in house prices. On the other hand, while a stock price boom follows it, the technology episode doesn't appear to be directly causing it. What is the probability that those booms eventually turn into busts however? It is a regular empirical finding in the literature that an overheating of the credit market eventually leads to a higher probability of a an economic downturn (Jordà, Schularick, and Taylor (2013)). Is this the case with our exercise? In this section, we take our reasoning ones step further and move to the main question of our exercise; namely whether technology booms are associated and ultimately whether they are causally related with asset price bubbles.

For this exercise, employing the levels or the differences of the variables as we did in the previous section wouldn't be sufficient. We need a variable containing enough information that should capture the boom and bust dynamics of the markets. To this purpose, we construct and employ a measure of asset price bubbles.

With the term 'bubble' here we mean a deviation of an asset price variable above its fundamental value that is followed by a crash. While challenging, the literature has provided some guidance on how to empirically identify such episodes. For our purposes, we take inspiration from Jordà, Schularick, and Taylor (2013) approach and propose an identification of a bubble as the conjunction of two empirical criteria; one intended to capture the upwards deviation and one intended to capture the the subsequent decrease. In what follows, we lay out the identification criteria and the calibration of those.

Having identified the bubbles we will then proceed in the same way that we proceeded in the previous section; namely, we will first establish a sequence of association with an unconditional specification and then we will utilize a more complex specification to delve deeper into more causal questions.

4.5.1 Defining Bubbles

Our starting point in these exercises is the construction of an indicator variable capturing house and stock price bubbles. The underlying data for our exercise here are the house and stock price indices. We begin this task with an approach very similar with the construction of our tech boom indicator; we construct a variable using an upward deviation criterion from the variable's long term value. We then utilize an additional criterion to capture any potential drop in the variable's value.

Into the specifics of our approach, we require that the difference ζ between the log variable and its HP filtered trend ($\lambda = 100$) is higher than the standard deviation of that difference ϕ times a scaling parameter k. As is the case with the tech boom this produces an indicator variable that takes the value of 1 when the price is elevated with these criteria. To this we are also going to impose an additional criterion to capture the bubble's bust phase, a potential crash. We wish here to incorporate information on a subsequent crash that might follow the elevation episode. To capture this, we require that the real asset prices should fall by more than μ percentage points over a 3-year window looking forward from any year in the episode. This is in accordance with Jordà, Schularick, and Taylor (2013) approach of identifying bubbles. The bubble episode then should be the intersection of those two signals.

To introduce some notation, for the identification of the bubble we are going to use the following price elevation and a price correction signal:

$$PriceElevationSignal_{it} = I(\zeta_{it} > k\phi_i)$$

$$PriceCorrectionSignal_{it} = I(p_{it+3} - p_{it} > \mu)$$

$$BubbleSignal_{it} = (PriceCorrectionSignal_{it} = 1) \cap (PriceElevationSignal_{it} = 1)$$

We calibrate the signals as follows: for the house price bubbles the price elevation scaling parameter is $k = 0.5$. For the stock price bubble we require the same parameter of the price elevation signal to be $k = 1$. For the percentage drop in the price correction signal we require it to be $\mu = 10\%$ for the housing bubble and $\mu = 10\%$ for the stock price bubble. On figure 16 we present the years when a bubble episode occurred according to our criterion:

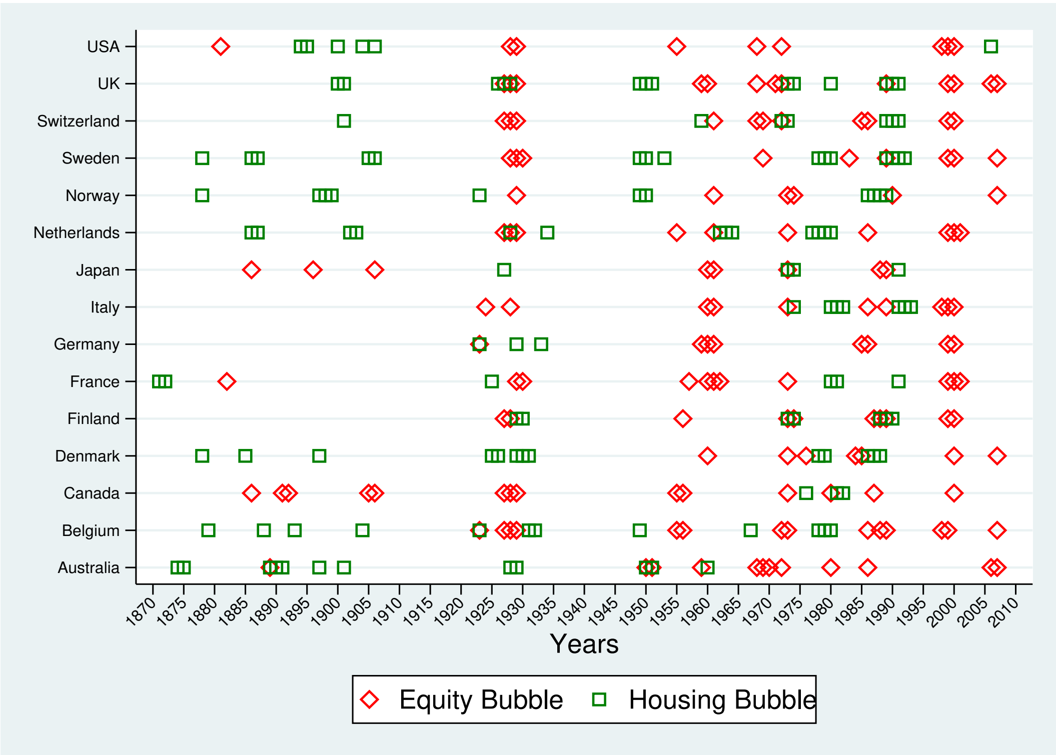

Figure 4.16: Stock And House Price Bubbles

At a first glance, the algorithm seems to be producing intuitive results with the bubble episodes broadly corresponding to the historical experience. When it comes to the stock market bubble we observe two main cross country clusters; one in the late '90s and one in the late

'20s. We know that both of those periods were indeed associated with stock market bubbles and a subsequent crisis. There are also two smaller clusters of stock market bubbles taking place in the late '50s-early '60s and in the mid '80s. Those indeed were periods of stock market surges related to tech booms; the early '60s was characterized by the electronics boom and the early to mid '80s was characterized from the biotech revolution (Pastor and Veronesi (2009)). Apart from those well documented episodes, our indicator appears to be capturing episodes characterizing only individual countries; we have for example the late '80s stock market bubble in Japan and the Australian stock price bubble of the late '60s.

Regarding the housing bubbles, again the results of our algorithm appear to make sense; in particular regarding the 2008 housing bubble, our indicator variable switches on in 2006 in the US. The Scandinavian housing bubbles of the late '80s are also represented in our sample as well as that of Japan from the same period. Some more obscure episodes are also captured like a real estate boom in Denmark in the early 1900s and the Australian real estate boom of the 1880s that crashed in the early 1890s.

Overall therefore, we can say that the results appear largely intuitive. They are also in accordance with the house and stock price bubbles presented by Jordà, Schularick, and Taylor (2015) in their own exploration of the housing and asset price bubbles.

4.5.2 Do Tech Booms Lead To Bubbles?

In this section we delve into the question at the heart of our exercise: what is the relationship between technology booms and asset price bubbles? Do technology booms indeed kick start a financial boom and bust cycle?

To approach this question we will proceed as we did in the previous section; we will first need to see if asset price bubbles indeed follow technology booms. We do this again using the unconditional specification and as a shock we consider a perturbation to the tech boom indicator. The dependent variable in this case is the bubble indicator variable. Following this, we will utilize a more complete specification to take into account the feedback loops and complex dynamics that emerge after a tech boom. If the probability of a bubble still rises after this conditioning a causal relation is likely to exist.

Are Tech Booms Related With Asset Price Bubbles? An Unconditional Exploration

As in the previous section we utilize the following unconditional specification to establish a chain of association:

$$Pr(Bubble_{it+h,\delta}) = E_{it}(Bubble_{it+h}|x_{it} = x + \delta) - E_{it}(Bubble_{it+h}|x_{it} = x) \quad (4.4)$$

where $Bubble_{it}$ is the bubble indicator variable and $Pr(Bubble_{it+h,\delta})$ denotes the probability of a bubble occurring across countries given a change δ in the treatment technology boom variable x. Since both the dependent and the independent variables are indicators, the coefficient on the independent variable can be interpreted as the increment in the probability of a bubble occurring given a technology boom, essentially reducing to the linear probability model. We begin therefore with the probability of a housing bubble and its 90% confidence intervals. Figure 17 plots the unconditional response of the housing bubble indicator in the post WW2 and a full sample:

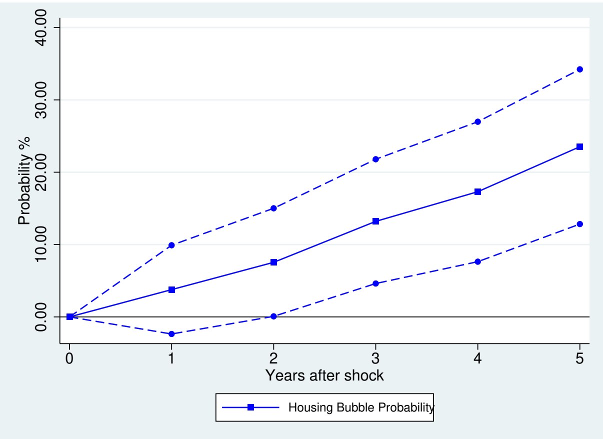

(a) Unconditional Response Of House Bubble Indicator To A Technology Boom-Post WW2 Sample

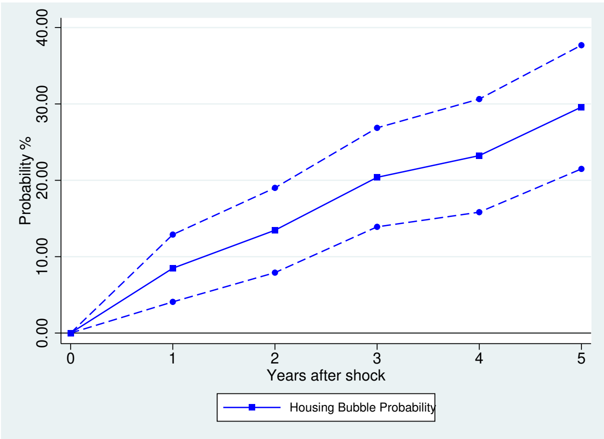

(b) Unconditional Response Of House Bubble Indicator To A Technology Boom-Full Sample

Figure 4.17: Unconditional Paths Of House Bubble Indicator: Post WW2 and Full Sample. 90% Confidence Bands

We see that the probability of a housing bubble in the aftermath of a technology boom rises significantly within the 90% confidence levels in both samples. In the post WW2 sample it becomes significant after the 2nd year and rises up to 20% within the 5 year horizon that we have set. In the full sample, the confidence intervals are considerably tighter with the probability being significant in each horizon rising up to 30% after 5 years. We therefore conclude that a housing bubble is more likely to occur after a technological revolution.

We next move on to the unconditional IRFs of the stock market indicator. Figure 18 plots this response for the post WW2 and the full sample.

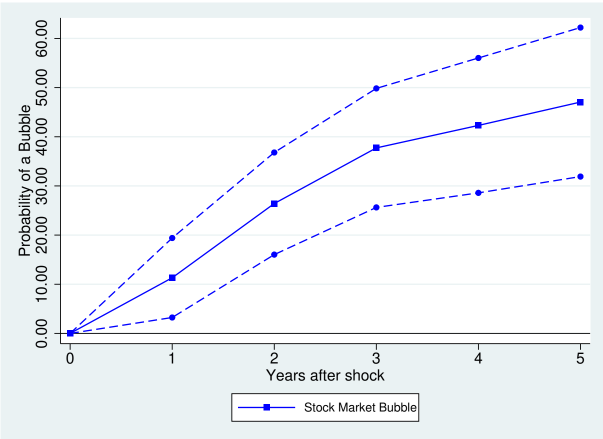

(a) Unconditional Response Of Equity Bubble Indicators To A Technology Boom-Post WW2 Sample

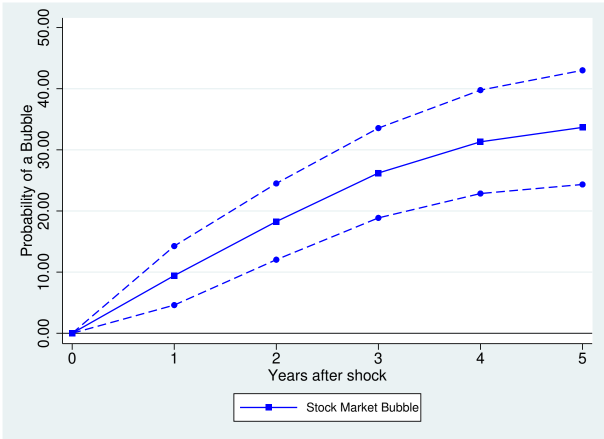

(b) Unconditional Response Of Equity Bubble Indicators To A Technology Boom-Full Sample

Figure 4.18: Unconditional Paths Of Equity Bubble Indicators: Post WW2 and Full Sample. 90% Confidence Bands

As is the case with the house bubble indicator, the incremental probability of a stock market bubble occurring after a tech boom rises significantly. In the post WW2 sample the probability is always positive and significant peaking at a 50% within the 5 year horizon. In the full sample the picture is the same. The increase in probability is always significant this time peaking to a 30% rise five years after the shock. As with housing bubbles therefore, we see that equity bubbles are also more likely to occur in the aftermath a technological revolution.

Do Tech Booms Cause Asset Price Bubbles? A Conditional Specification

In this section, we employ a richer specification to take into account the complex dynamics and feedback loops that are developing in the economy. To do this we are again going to employ the Local Projections framework to generate out IRFs and condition on a rich set of covariates. We again utilize the simple Local Projections framework to retrieve the IRFs. Our specification to gauge whether a technology boom is causally related to asset price bubbles is then the following:

$$Bubble_{it+h} = a_i + \beta_h x_t + \sum_{j=0}^{p} \Gamma_j X_{it-j} + u_{it}; h = 1, ..., H \tag{4.5}$$

where $Bubble_{it+h}$ is the bubble indicator variable, a_i are country fixed effects, x_{it} is the again the first occurrence in a sequence of our technology boom indicator variable and the vector X contains a history of p lags of the control variables at time t with coefficients Γ and u is the error term. The controls included are: GDP, inflation rate, the short term interest rate, the private investment per capita rate and their first lags as well as the first lag of the patent applications per capita variable.

Figure 19 plots the probability of a housing bubble in the aftermath of a technology boom after conditioning on the state of the economy. We see that in both samples the tech boom doesn't make the bubble more likely. Rather, in the post War sample the technology boom is apparently causing a decrease in the probability of a housing bubble.

This result in combination with the fact that the technology boom leads to an increase in the house price index from the previous section leads us to the following conclusion: while a housing bubble follows after the technology boom, the tech boom by itself not only is not the cause of that. Rather, it leads to a house boom with no bust in sight.

As the final iteration in our exercise, we move on to the response of the stock price index to the conditional specification of the technology boom. Figure 20 presents the relevant IRFs. The results here are again different amongst our two samples; in the post WW2 the tech boom is not significantly raising the probability of a bubble. In the full sample however, the stock bubble probability rises significantly after one year peaking at 15% in the 3rd year.

What are we to make of the results in this section? How can we reconcile the fact that a tech boom seemingly causes an increase in the probability of a stock market bubble in one of our samples and it doesn't affect the probability of a housing bubble with the results from section 4.2 wherein we saw that the stock prices aren't affected by it while the house prices rise? One way to approach the issue is to consider that the bubble indicator contains more information

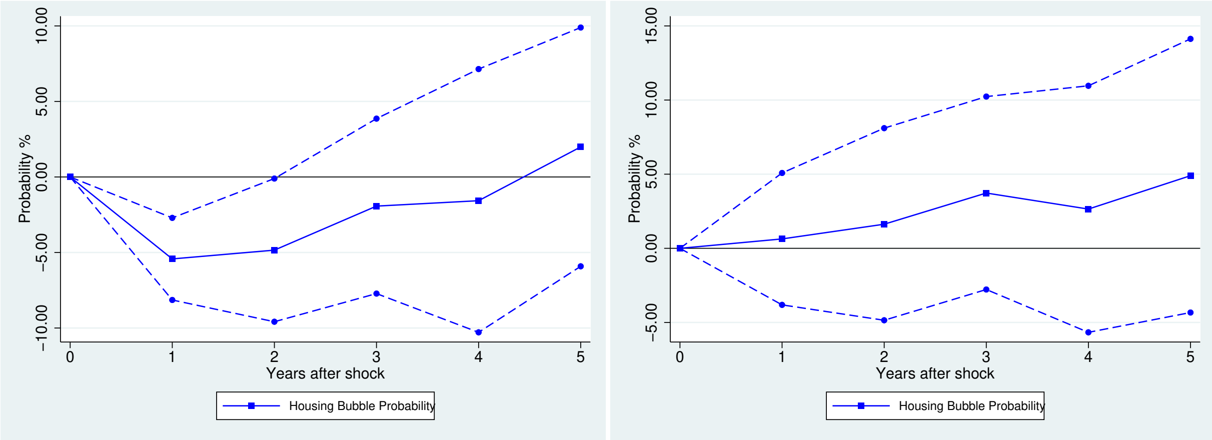

(a) Conditional Response Of House Bubble Indicator To A Technology Boom-Post WW2 Sample

(b) Conditional Response Of House Bubble Indicator To A Technology Boom-Full Sample

Figure 4.19: Conditional Paths House Bubble Indicator: Post WW2 and Full Sample. 90% Confidence Bands

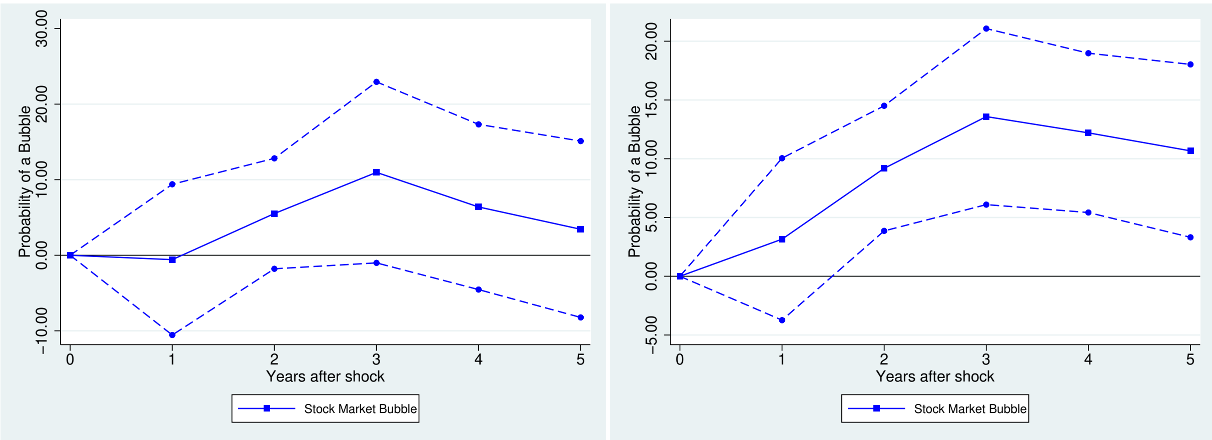

(a) Conditional Response Of Stock Market Bubble Indicator To A Technology Boom-Post WW2 Sample

(b) Conditional Response Of Stock Market Bubble Indicator To A Technology Boom-Full Sample

Figure 4.20: Conditional Paths Of Stock Market Bubble Indicator: Post WW2 and Full Sample. 90% Confidence Bands

than the raw variables. The fact that it is direclty causing a rise in the house prices connects well with the fact that the probability of housing bubble decreases. The technology boom in the case of the housing market is purely expansionary. In the case of the stock market similarly the fact that stock prices don't react significantly connects with the fact that in the full sample there is an increase in the probability of a bubble. The bubble indicator is capturing the effects of a potential crash which the raw variable isn't.

We can overall summarize the results from this section as follows; we have established that asset price bubbles are much more likely to happen after a technological boom within a five year horizon. However, the technological boom is not always the driver of these increased likelihoods. As a direct consequence of it the probability of a housing bubble is mostly unaffected and even decreases in a post WW2 sample. The probability of a stock price bubble on the other hand increases within a 3 year horizon as a consequence of it and is unaffected in the post WW2 sample. The deeper reasons of why a tech boom has seemingly different effects on the probability of different asset classes is beyond the scope of this paper.

4.6 Conclusions

Do technology booms indeed set in a boom and bust dynamic in the economy? Do they lead to an over expansion of credit and consequently to asset price bubbles that ultimately bust?

In this paper using a long panel dataset of 15 industrialized economies spanning the entire period of finance capitalism, we have attempted to shed some light on those questions. Using this dataset we constructed a novel technology boom indicator variable that is capturing information on anticipated technology revolutions. We consider this is indicator variable to be of enough generality so that it could potentially be utilized in various other econometric exercises that involve explorations with a technology shock.

With this indicator variable at hand, we proceed to construct a technological shock. We choose not to take a stance on the exogeneity of this shock. Rather, instead of tackling head on questions of identification and causality, we proceed to document stylized facts and deeper relations hiding in the data in connection with our technology boom indicator.

The main stylized fact we are documenting is that an overall credit expansion follows the technology revolution alongside with a boom in asset prices and an increase in real GDP per capita. Proceeding to gauge whether this stylized fact constitutes a causal relationship, we find that technology shocks do indeed cause a credit expansion and a house price boom but they have no causal effect on the stock prices.

In the final step of our exercise we examined whether those expansions ultimately lead to a bust; in other words are the technology booms the culprit behind asset price bubbles? To proceed with this question we first contribute to the literature by constructing an indicator variable for house and stock price bubbles in the spirit of Jordà, Schularick, and Taylor (2015). We then establish the stylized fact that those bubbles have an increased probability of occurring after a technology shock. However, this relationship is not causal. The technology boom only marginally increases the probability of a stock price bubble in our full sample and it has zero to negative effect on the probability of a house bubble.

These results are open to a multitude of interpretations from the theoretical literature. Potential explanations and detailed examination of the inner mechanics of the dynamics we documented are outside the scope of the present paper but offer themselves to future research. Our technology boom and bubble indicators are quite general and may be utilized to ask and answer a variety of questions within and outside the field of macroeconomics.

4.7 Appendix

4.7.1 A. Patent Growth Subsample Means

Table 4.1: Average Rate of Growth Rate of Patent Applications, Subsamples (Part 1)

	Australia	Belgium	Canada	Switzerland	Germany
1870-1920	6.416	7.919	6.250	7.639	6.195
1920-1945	1.146	-4.906	-2.545	1.514	0.398
1945-2008	-0.348	-0.277	2.724	-1.806	0.227

Table 4.2: Average Rate of Growth Rate of Patent Applications, Subsamples (Part 2)

	Denmark	Finland	France	UK	Italy
1870-1920	8.664	15.335	3.857	4.603	9.968
1920-1945	-0.269	-0.488	-1.732	-1.932	-3.773
1945-2008	0.223	1.892	1.486	0.212	2.138

Table 4.3: Average Rate of Growth Rate of Patent Applications, Subsamples (Part 3)

	Japan	Netherlands	Norway	Sweden	USA
1870-1920	14.944	10.422	14.458	8.035	2.780
1920-1945	-5.633	3.434	-0.727	2.195	-1.440
1945-2008	7.936	0.980	-0.083	-1.439	2.413

4.7.2 B. Long Run IRFs

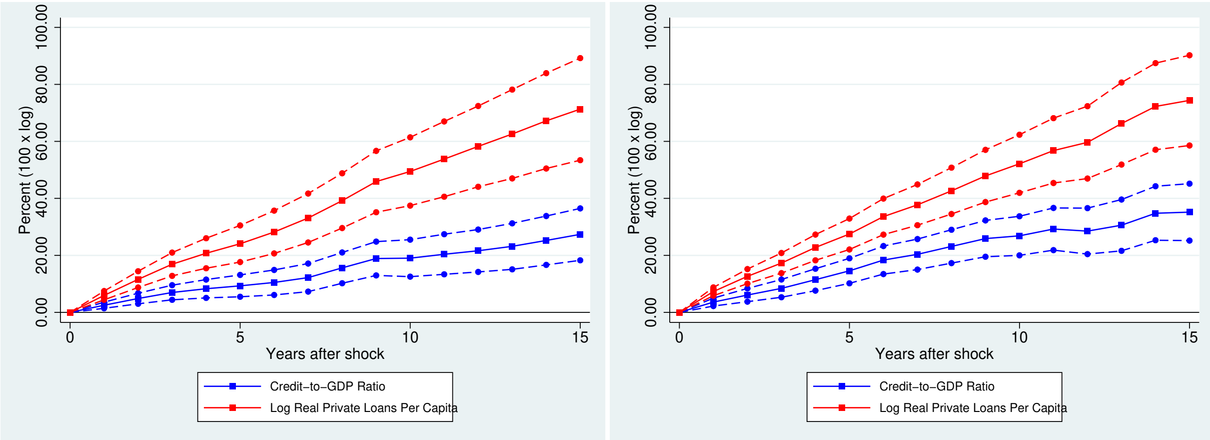

(a) Unconditional Response Of Credit Variables To A Technology Boom-Post WW2 Sample (b) Unconditional Response Of Credit Variables To A Technology Boom-Full Sample

Figure 4.21: Unconditional Paths of Credit Variable: Post WW2 and Full Sample. 90% Confidence Bands

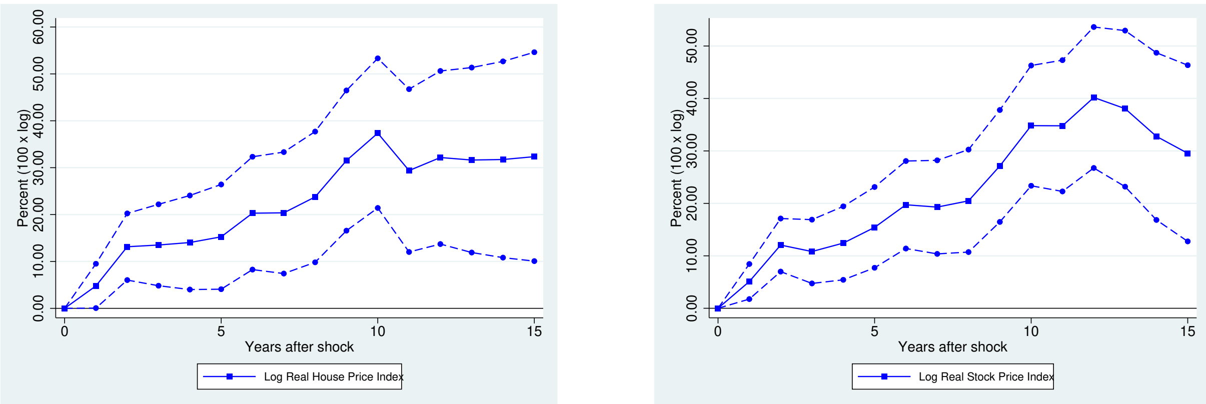

(a) Unconditional Response Of Real Stock Market Index (1990=100) To A Technology Boom-Post WW2 Sample (b) Unconditional Response Of Real Stock Market Index (1990=100) To A Technology Boom-Full Sample

Figure 4.22: Unconditional Paths of Real Stock Price Index: Post WW2 and Full Sample. 90% Confidence Bands

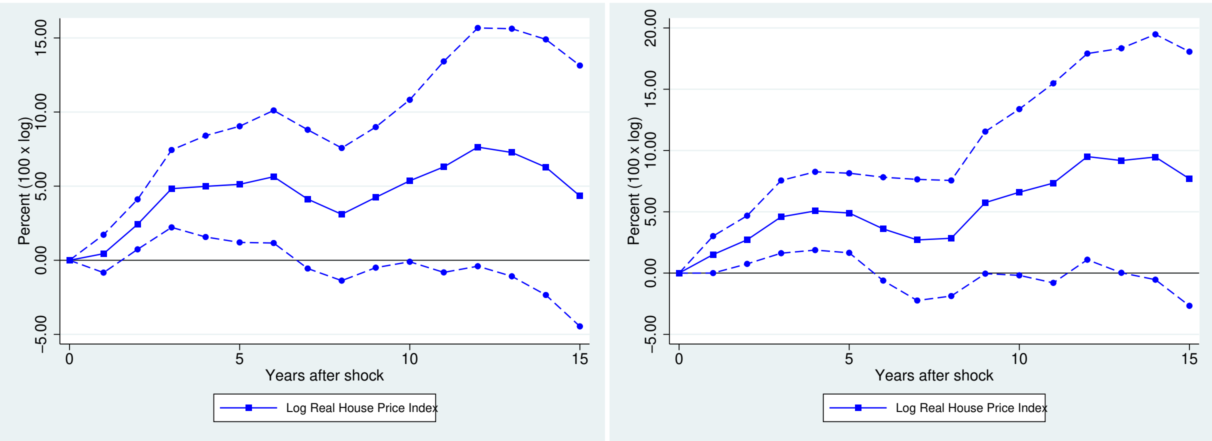

(a) Conditional Response of House Price Index To A Technology Boom-Post WW2 Sample

(b) Conditional Response of House Price Index To A Technology Boom-Full Sample

Figure 4.23: Conditional Paths of House Price Index: Post WW2 and Full Sample. 90% Confidence Bands

References

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2020). “Business-cycle anatomy”. In: *American Economic Review* 110.10, pp. 3030–70.

Aschauer, David Alan (1989a). “Does public capital crowd out private capital?” In: *Journal of Monetary Economics* 24.2, pp. 171–188.

— (1989b). “Is public expenditure productive?” In: *Journal of Monetary Economics* 23.2, pp. 177–200.

Auerbach, Alan J and Yuriy Gorodnichenko (2012). “Measuring the output responses to fiscal policy”. In: *American Economic Journal: Economic Policy* 4.2, pp. 1–27.

— (2013). “Output spillovers from fiscal policy”. In: *American Economic Review* 103.3, pp. 141–46.

Bachmann, Rüdiger and Eric R Sims (2012). “Confidence and the transmission of government spending shocks”. In: *Journal of Monetary Economics* 59.3, pp. 235–249.

Barro, Robert J (1987). “Government spending, interest rates, prices, and budget deficits in the United Kingdom, 1701–1918”. In: *Journal of Monetary Economics* 20.2, pp. 221–247.

Barro, Robert J and Charles J Redlick (2011). “Macroeconomic effects from government purchases and taxes”. In: *The Quarterly Journal of Economics* 126.1, pp. 51–102.

Barsky, Robert B and Eric R Sims (2011). “News shocks and business cycles”. In: *Journal of Monetary Economics* 58.3, pp. 273–289.

Baxter, Marianne and Robert G King (1993). “Fiscal policy in general equilibrium”. In: *The American Economic Review*, pp. 315–334.

Beaudry, Paul and Franck Portier (2004). "An exploration into Pigou's theory of cycles". In: *Journal of Monetary Economics* 51.6, pp. 1183–1216.

— (2006). "Stock prices, news, and economic fluctuations". In: *American Economic Review* 96.4, pp. 1293–1307.

Ben Zeev, Nadav and Evi Pappa (2017). "Chronicle of a war foretold: The macroeconomic effects of anticipated defence spending shocks". In: *The Economic Journal* 127.603, pp. 1568–1597.

Bernanke, Ben S et al. (1997). "Systematic monetary policy and the effects of oil price shocks". In: *Brookings Papers on Economic Activity* 1997.1, pp. 91–157.

Blanchard, Olivier and Roberto Perotti (2002). "An empirical characterization of the dynamic effects of changes in government spending and taxes on output". In: *the Quarterly Journal of economics* 117.4, pp. 1329–1368.

Bloem, Mr Adriaan M, Mr Robert Dippelsman, and Mr Nils Øyvind Mæhle (2001). *Quarterly national accounts manual: concepts, data sources, and compilation*. International Monetary Fund.

Boehm, Christoph E (2020). "Government consumption and investment: Does the composition of purchases affect the multiplier?" In: *Journal of Monetary Economics* 115, pp. 80–93.

Borio, Claudio EV and Philip William Lowe (2002). "Asset prices, financial and monetary stability: exploring the nexus". In.

Burnside, Craig, Martin Eichenbaum, and Jonas DM Fisher (2004). "Fiscal shocks and their consequences". In: *Journal of Economic theory* 115.1, pp. 89–117.

Caggiano, Giovanni et al. (2015). "Estimating fiscal multipliers: News from a non-linear world". In: *The Economic Journal* 125.584, pp. 746–776.

Caldara, Dario and Christophe Kamps (2008). "What are the effects of fiscal policy shocks? A VAR-based comparative analysis". In.

Cao, Dan and Jean-Paul L'Huillier (2018). "Technological revolutions and the three great slumps: A medium-run analysis". In: *Journal of Monetary Economics* 96, pp. 93–108.

Cascaldi-Garcia, Danilo and Marija Vukotić (2022). "Patent-based news shocks". In: *Review of Economics and Statistics* 104.1, pp. 51–66.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo (2011). "When is the government spending multiplier large?" In: *Journal of Political Economy* 119.1, pp. 78–121.

Denton, Frank T (1971). "Adjustment of monthly or quarterly series to annual totals: an approach based on quadratic minimization". In: *Journal Of The American Statistical Association* 66.333, pp. 99–102.

Detken, Carsten and Frank Smets (2004). "Asset price booms and monetary policy". In: *Macroeconomic Policies In The World Economy* 329, p. 189.

Dieppe, Alistair, Neville Francis, and Gene Kindberg-Hanlon (2021). "The identification of dominant macroeconomic drivers: coping with confounding shocks". In.

Eggertsson, Gauti B (2011). "What fiscal policy is effective at zero interest rates?" In: *NBER Macroeconomics Annual* 25.1, pp. 59–112.

Fernald, John (2014). "A quarterly, utilization-adjusted series on total factor productivity". In: Citeseer.

Fisher, Jonas DM and Ryan Peters (2010). "Using stock returns to identify government spending shocks". In: *The Economic Journal* 120.544, pp. 414–436.

Forni, Mario and Luca Gambetti (2014). "Sufficient information in structural VARs". In: *Journal of Monetary Economics* 66, pp. 124–136.

— (2016). "Government spending shocks in open economy VARs". In: *Journal of International Economics* 99, pp. 68–84.

Forni, Mario, Luca Gambetti, and Luca Sala (2019). "Structural VARs and noninvertible macroeconomic models". In: *Journal of Applied Econometrics* 34.2, pp. 221–246.

Gambetti, Luca et al. (2022). "The effect of news shocks and monetary policy". In: *Essays in Honour of Fabio Canova*. Vol. 44. Emerald Publishing Limited, pp. 139–164.

Gordon, Robert J and Robert Krenn (2010). *The end of the great depression 1939-41: Policy contributions and fiscal multipliers*. Tech. rep. National Bureau of Economic Research.

Görtz, Christoph, John D Tsoukalas, and Francesco Zanetti (2022). "News shocks under financial frictions". In: *American Economic Journal: Macroeconomics* 14.4, pp. 210–43.

Greenwood, Robin et al. (2020). *Predictable financial crises*. Tech. rep. National Bureau of Economic Research.

Griliches, Zvi (1998). "Patent statistics as economic indicators: a survey". In: *R&D and productivity: the econometric evidence*. University of Chicago Press, pp. 287–343.

Henderson, Steven W (2015). "Consumer spending in World War II: the forgotten consumer expenditure surveys". In: *Monthly Lab. Rev.* 138, p. 1.

Ilzetzki, Ethan, Enrique G Mendoza, and Carlos A Végh (2013). "How big (small?) are fiscal multipliers?" In: *Journal of Monetary Economics* 60.2, pp. 239–254.

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". In: *American Economic Review* 95.1, pp. 161–182.

Jordà, Òscar, Björn Richter, et al. (2021). "Bank capital redux: solvency, liquidity, and crisis". In: *The Review of Economic Studies* 88.1, pp. 260–286.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2013). "When credit bites back". In: *Journal of Money, Credit and Banking* 45.s2, pp. 3–28.

— (2015). "Leveraged bubbles". In: *Journal of Monetary Economics* 76, S1–S20.

— (2017). "Macrofinancial history and the new business cycle facts". In: *NBER Macroeconomics Annual* 31.1, pp. 213–263.

Kindleberger, Charles P, Panics Manias, and A Crashes (1996). *History of Financial Crises*.

Klein, Mathias and Ludger Linnemann (2021). "Real exchange rate and international spillover effects of US technology shocks". In: *Journal of International Economics* 129, p. 103414.

Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996). "Impulse response analysis in nonlinear multivariate models". In: *Journal of Econometrics* 74.1, pp. 119–147.

Krugman, Paul (2016). "A pause that distresses". In: *The New York Times URL https://www.nytimes. com/2016/06/06/opinion/a-pause-that-distresses. html*.

Kumhof, Michael, Romain Rancière, and Pablo Winant (2015). "Inequality, leverage, and crises". In: *American Economic Review* 105.3, pp. 1217–45.

López-Salido, David, Jeremy C Stein, and Egon Zakrajšek (2017). "Credit-market sentiment and the business cycle". In: *The Quarterly Journal of Economics* 132.3, pp. 1373–1426.

Madsen, Jakob B (2007). “Technology spillover through trade and TFP convergence: 135 years of evidence for the OECD countries”. In: *Journal of International Economics* 72.2, pp. 464–480.

— (2010). “The anatomy of growth in the OECD since 1870”. In: *Journal of Monetary Economics* 57.6, pp. 753–767.

Madsen, Jakob B, Md Rabiul Islam, and Hristos Doucouliagos (2018). “Inequality, financial development and economic growth in the OECD, 1870–2011”. In: *European Economic Review* 101, pp. 605–624.

Malkiel, Burton Gordon (1999). *A random walk down Wall Street: including a life-cycle guide to personal investing*. WW Norton & Company.

Mertens, Karel RSM and Morten O Ravn (2014). “Fiscal policy in an expectations-driven liquidity trap”. In: *The Review of Economic Studies* 81.4, pp. 1637–1667.

Mian, Atif and Amir Sufi (2018). “Finance and business cycles: The credit-driven household demand channel”. In: *Journal of Economic Perspectives* 32.3, pp. 31–58.

Minsky, Hyman P (1977). “The financial instability hypothesis: An interpretation of Keynes and an alternative to “standard” theory”. In: *Challenge* 20.1, pp. 20–27.

Miranda-Agrippino, Silvia, Sinem Hacioglu Hoke, and Kristina Bluwstein (2019). “When creativity strikes: news shocks and business cycle fluctuations”. In.

Miyamoto, Wataru, Thuy Lan Nguyen, and Dmitriy Sergeyev (2018). “Government spending multipliers under the zero lower bound: Evidence from Japan”. In: *American Economic Journal: Macroeconomics* 10.3, pp. 247–77.

Moretti, Enrico, Claudia Steinwender, and John Van Reenen (2019). *The intellectual spoils of war? defense r&d, productivity and international spillovers*. Tech. rep. National Bureau of Economic Research.

Mountford, Andrew and Harald Uhlig (2009). “What are the effects of fiscal policy shocks?” In: *Journal of Applied Econometrics* 24.6, pp. 960–992.

Olea, José Luis Montiel and Carolin Pflueger (2013). “A robust test for weak instruments”. In: *Journal of Business & Economic Statistics* 31.3, pp. 358–369.

Owyang, Michael T, Valerie A Ramey, and Sarah Zubairy (2013). "Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data". In: *American Economic Review* 103.3, pp. 129–34.

Pástor, L'uboš and Pietro Veronesi (2009). "Technological revolutions and stock prices". In: *American Economic Review* 99.4, pp. 1451–83.

Perez, Carlota (2003). *Technological revolutions and financial capital*. Edward Elgar Publishing.

Perotti, Roberto (2004). "Public investment: another (different) look". In.

— (2005). "Estimating the effects of fiscal policy in OECD countries". In: *Available at SSRN 717561*.

Plagborg-Møller, Mikkel and Christian K Wolf (2021). "Local projections and VARs estimate the same impulse responses". In: *Econometrica* 89.2, pp. 955–980.

Ramey, Valerie A (2011). "Identifying government spending shocks: It's all in the timing". In: *The Quarterly Journal of Economics* 126.1, pp. 1–50.

— (2016). "Macroeconomic shocks and their propagation". In: *Handbook of Macroeconomics* 2, pp. 71–162.

— (2019). "Ten years after the financial crisis: What have we learned from the renaissance in fiscal research?" In: *Journal of Economic Perspectives* 33.2, pp. 89–114.

Ramey, Valerie A et al. (2020). *The macroeconomic consequences of infrastructure investment*. Tech. rep. National Bureau of Economic Research.

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

Ramey, Valerie A and Sarah Zubairy (2018). "Government spending multipliers in good times and in bad: evidence from US historical data". In: *Journal of Political Economy* 126.2, pp. 850–901.

Sufi, Amir and Alan M Taylor (2021). "Financial crises: A survey". In.

Tong, Howell (1978). "On a threshold model". In.

— (2012). *Threshold models in non-linear time series analysis*. Vol. 21. Springer Science & Business Media.

Uhlig, Harald (2003). “What moves real GNP?” In.

Woodford, Michael (2011). “Simple analytics of the government expenditure multiplier”. In: *American Economic Journal: Macroeconomics* 3.1, pp. 1–35.