Large and state-dependent effects of quasi-random monetary experiments*

Òscar Jordà† Moritz Schularick‡ Alan M. Taylor§

October 2016

Abstract

When open economies fix their exchange rate, they constrain monetary policy. The trilemma implies that arbitrage, not the central bank, determines how interest rates fluctuate. The annals of international finance thus provide quasi-natural experiments with which to measure how macroeconomic outcomes respond to policy rates. Using historical data since 1870, we estimate the local average treatment effect (L*ATE*) of monetary policy interventions with a trilemma instrument and discuss the connection with the population *ATE*. Using local projection instrumental variable methods we find that the effects of monetary policy are much larger than previously estimated, and that these effects are state-dependent. Using a novel control function approach we allow for possible spillovers via other channels and our results prove to be robust. Monetary policy has weak effects when output is depressed and when the economy has "lowflation." Our findings have profound implications for monetary economics.

JEL classification codes: E01, E30, E32, E44, E47, E51, F33, F42, F44

Keywords: interest rates, monetary experiments, trilemma, fixed exchange rates, instrumental variables, local projections

*We thank James Cloyne and Patrick Hürtgen for sharing their data with us. The paper has been improved by comments and suggestions from James Cloyne, Julian di Giovanni, Gernot Müller, and Ricardo Reis, and by feedback from presentations at the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of Cleveland, the Fourth CEPR Economic History Symposium, De Nederlandsche Bank, and the Bank of England. We are grateful to Helen Irvin for outstanding research assistance. Generous grants from the Institute for New Economic Thinking and the Volkswagen Foundation supported different parts of the data collection and analysis effort. We are grateful for their support. The views expressed in this paper are the sole responsibility of the authors and to not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

†Federal Reserve Bank of San Francisco; and Department of Economics, University of California, Davis (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

‡Department of Economics, University of Bonn; and CEPR (moritz.schularick@uni-bonn.de).

§Department of Economics and Graduate School of Management, University of California, Davis; NBER; and CEPR (amtaylor@ucdavis.edu).

1. Introduction

Few questions in economics have been studied and debated more intensively than the effects of monetary policy on output and inflation. But what do we really know? In the pre-crisis era, a fragile consensus existed that, at least in the short run, central banks can guide output and inflation by changing interest rates, with few doubts about the expected sign of the effects of policy changes: raising rates would be contractionary and deflationary, cutting rates would do the opposite (Romer and Romer 1989, 2004; Christiano, Eichenbaum, and Evans 1999).

In the wake of the Global Financial Crisis of 2008, central banks have struggled to achieve their targets causing some to question basic tenets by arguing that central banks could raise inflation by increasing interest rates (see, e.g., Cochrane 2016; Williamson 2016). As Bertrand Russell once said, every opinion now accepted was once eccentric. Time and further research will tell on what side of the ledger these new theories belong.

While few still disagree with the notion that money is neutral in the long-run, its short-run effects depend critically on various market frictions. As a result, determining whether monetary policy is a useful stabilizer is largely an empirical matter. Monetary policy reflects the central bank's response to the economic outlook. In the parlance of the policy-evaluation literature (see, e.g., Rubin, 2005), any measure of the policy effect is contaminated by confounders simultaneously correlated with the assignment mechanism and the outcome.

Empirical measures of the effect of interest rates on macroeconomic outcomes are fraught. Macroeconomic aggregates and interest rates are jointly determined. Agents form expectations about future outcomes. Economic mechanisms are dynamic. The traditional approach then goes something like this: control, as much as possible, for current and past information on macroeconomic variables. Assume that any variation left in interest rates is as good as random. Then calculate the effect of fluctuations in the unpredictable component of interest rates on outcomes.

Over time, a great deal of consensus has emerged almost exclusively built on the Post-WW2 U.S. experience. This consensus grew from a variety of empirical routes that share a common thread: the desire to control for the information used by the central bank in setting interest rates. Whether through a saturated model of observables (see, e.g., Christiano, Eichenbaum, and Evans 1999, and the long tradition of structural vector autoregressions of VARs); or through indirect measures based on financial data (e.g., Kuttner 2001; Faust, Wright, and Swanson 2004; Nakamura and Steinsson 2013; Gertler and Karadi 2015) the objective is the almost always the same: to achieve identification via selection-on-observables arguments. This paper tries a different approach.

Randomized control trials in macroeconomics are very rare. However, when countries peg their exchange rate to a "base" country, but leave capital to flow freely across borders, local and international interest rates also link up. Interest rates no longer reflect local conditions alone, and arbitrage under a peg constrains what the local central bank can do. This loss of monetary control in open pegs, an implication of the trilemma of international finance, is as close to a natural experiment as one can find in macroeconomics.

The general empirical validity of the trilemma, in recent times and in distant historical epochs, has been recognized for more than a decade: exogenous base country interest rate movements spill into local interest rates for open pegs (Obstfeld, Shambaugh, and Taylor 2004; Shambaugh 2004). Some important corollaries directly follow for empirical macroeconomics. A key contribution by di Giovanni, McCrary, and von Wachter (2009) exploits the resulting identified local monetary policy shocks to estimate impulse response functions for other macroeconomic outcome variables of interest using standard VAR methods with instrumental variables.[1]

In this paper we take the idea further: first, we use an instrumental variables local projections (LP-IV) approach which—unlike a VAR—is flexible, parsimonious, and adaptable to nonlinerities (which we find to be important);[2] second, we develop econometric methods to formalize the link between the average treatment effect of monetary policy in all economies (ATE) and the local average treatment effect (LATE) that can be estimated from open pegs using the trilemma instrument; third, we introduce a novel approach using control functions to place bounds on the bias that may arise in these types of estimates if there are confounding spillovers from base interest rates to the local economy via other channels (i.e., violations of the exclusion restriction); and fourth, we expand the sample from post-1970s to all of advanced economy macroeconomic history since 1870.

The logic of our approach can be summarized as follows. Fluctuations in base country interest rates (not explained by base country controls) are a source of exogenous variation when an economy pegs to this base country and keeps capital flows unfettered. This observation serves to construct our instrumental variable. Movements in domestic interest rates reflect several sources of endogenous and exogenous variation. However, movements correlated with our instrumental variable are as close to randomly assigned as it is possible to obtain given other domestic controls. Our identification relies on eras in which countries peg their exchange rates. Thus, instead of relying on the predominantly floating post-1970s era, we experiment in the laboratory of long-run macroeconomic history. This paper relies on annual macroeconomic data for 17 advanced economies from 1870 to 2013. Such a long and broad historical analysis has become possible for the first time as the result of an extensive multi-year data collection effort.

We are careful to state the assumptions required for identification since one contribution of our paper is to measure the attenuation bias generated when the selection-on-observables assumption, implicit in much of the literature, fails. In a first step, the results we obtain directly from the IV estimator are an estimate of the local average treatment effect (L*ATE*) for countries that peg the exchange rate and allow capital to move across borders. What about the other countries?

Of course, the extent to which the LATE results can be extrapolated to other economies depends on additional assumptions. We can show that countries exposed to the instrument have similar OLS responses to the interest rate treatment as those that are not. Finally, using the instrument in Romer

[1] In related work, di Giovanni and Shambaugh (2008) used the trilemma to investigate post-WW2 output volatility in fixed and floating regimes. Ilzetzki, Mendoza, and Végh (2013) partition countries by exchange rate regime to study the impact of a fiscal policy shock. In previous work (Jordà, Schularick, and Taylor 2015), we studied the link between financial conditions, mortgage credit, and house prices.

[2] See Jordà (2005); Owyang, Ramey, and Zubairy (2013); Jordà, Schularick, and Taylor (2015).

and Romer (2004) and Cloyne and Hürtgen (2014), we find that the LATE for the U.S. and the U.K., both of which historically have usually been base countries, is similar to the LATE for the pegging countries in our sample. This finding suggests that the initial LATE estimates are not driven by factors that also determine the decision to peg. For these reasons, we feel reasonably confident that our estimates provide a good approximation to the true, but unmeasurable ATE of interest rates on macroeconomic outcomes.[3]

Further reassurance follows in a section where we develop new methods to address potential spillovers from base interest rates into the local economy via channels other than the trilemma. To wit, how can one cope with the problem of a potential violation of the exclusion restriction? Drawing on the microeconometric literature on control functions we show a couple of interesting and useful results. First, although it might appear obvious which way this bias goes, careful calculation shows the direction is ambiguous. Second, even so, guided by economic reasoning on the relative magnitudes of local and domestic responses we can place plausible bounds on this bias; we can thus discipline the range of ATE responses that one can reasonably infer from the LATE estimates, and in fact the deviations appear to be small compared to our baseline LP-IV results.

Another innovation of our paper is to estimate the state-dependent effect of interest rates on macroeconomic outcomes. Stabilization policy is unlikely to be stimulative in a boom and contractionary in a bust. Recent studies on the fiscal multiplier (e.g., Auerbach and Gorodnichenko 2013ab; Ramey and Zubairy 2014; and Jordà and Taylor 2016) have found fiscal policy to have different effects depending on the state of the economy. Framed against a similar desire to understand monetary stabilization, our analysis sheds light on the results reported in Angrist, Jordà and Kuersteiner (forthcoming) and Tenreyro and Thwaites (2016). These papers show that interest rates have asymmetric effects, something we find as well. However, in our context, the source of the asymmetry comes from the difference in the economic environment in which each type of policy (accommodative versus contractionary) is implemented.

Advanced economies have recently been grappling with a low growth, low inflation environment dubbed by commentators as "lowflation." A burning question in central bank circles is about the effectiveness of monetary policy when faced with this predicament. To answer such a question, we again investigate the effects of monetary perturbations in lowfation environments. The historical data show that policy is rather ineffective in such cases, in a manner that reinforces our earlier analysis on state-dependent monetary effects. Thus our results have important implications for the current policy debate and more generally, for models of monetary economies.

[3]Although one might worry that a fix-float decision creates a potential endogeneity problem, we find in auxiliary regressions for our sample (not shown) that the decision to peg is largely unpredictable and one that is very persistent (in our sample, the average peg is 21 years). In addition, we guard against the risk that very short opportunistic pegs could bias our results by requiring a country to be pegged for two successive years when we construct our trilemma instrument. In other words, we think any endogeneity embedded in the decision to peg or leave a peg appears to be well contained at business cycle frequencies in our application. In addition, the monotonicity assumption requiring instrument and treatment to be positively associated appears to hold in the data.

2. The trilemma of international finance: a quasi-natural experiment

Throughout the history of modern finance, countries have managed international goods and capital flows using a variety of policies. The interplay of these policies had, at times, important consequences for domestic monetary conditions. We exploit situations when external conditions bleed into domestic policy as a way to identify exogenous movements in monetary conditions. Given the important role that this mechanism plays, we begin by presenting the main statistical properties of our instrumental variable before we discuss our overall empirical approach. The specific construction of this instrumental variable is one of the important contributions of our paper.

We suppose that any home country i at time t can either have a flexible (floating) or fixed (pegged) exchange rate with respect to some other *base* country. Let Δr_{it} denote the change in short-term nominal interest rates in the home country, and let $\Delta r^*_{b(i)t}$ denote the short-term nominal interest in country i's base country, which can differ across i, hence the notation $b(i)$. In the international historical dataset that we use (discussed in Section 3), these interest rates refer to three-month government bond rates, the closest measure to a policy rate that we were able to obtain (see, e.g., Swanson and Williams 2014 for an argument favoring medium term government rates to measure monetary policy effects).

Denote $\widehat{\Delta r^*}_{b(i)t}$ movements in base country $b(i)$ rates explained by observable controls for that base country, $X_{b(i),t}$. Thus, $\Delta r^*_{b(i)t} - \widehat{\Delta r^*}_{b(i)t}$ denotes unpredictable movements in base country interest rates. The idea is to take, not just the base country rate as exogenous with respect to the home country, but to control for any predictable movements in base interest rates that might have been expected by country i. Robustness checks based on directly using $\Delta r^*_{b(i)t}$ produced similar results to those reported below and are available upon request. Similarly, but for use later, we denote with X_{it} the set of domestic macroeconomic controls used in determining the home interest rate in the absence of a peg.

Going forward, we will be careful to separate country-year pairs into the corresponding subpopulations of *pegs*, *floats* and *bases*. For now, the variable $PEG_{it} = 0, 1$ simply denotes if country i is a *float* or a *peg* respectively. Sometimes exchange rates are managed over a small band around the peg. This was the case for several European economies in the lead up to the euro. This poses no difficulty for our instrument construction, however. We do not require fluctuations in home country rates to be perfectly explained by base country rates. All that is needed is for these two rates to be correlated. Fluctuations inside a corridor limit exchange rate variation such that wide interest rate differentials cannot persist in practice.

Over the span of nearly 150 years, countries have come in and out of fixed exchange regimes. However, once in such a regime, countries tend to stick to the arrangement. In contrast to Obstfeld and Rogoff (1995), who found the average duration of fixed exchange rates since 1973 to be about 5 years, we find in our longer-run sample that average to be about 21 years. Part of the reason is that

our sample includes advanced economies only, whereas theirs includes emerging market economies as well. Another is that our sample includes longer-lived peg episodes in the gold standard and Bretton Woods eras. Nevertheless, in defining out instrument, we require countries to have been in a peg for at least one year by interacting the variable PEG_{it} with PEG_{it-1} as in expression (1) below. This feature eliminates potential distortions of country that just joined a peg.

Finally, the variable $KOPEN_{it} \in [0, 1]$ denotes whether the home country is open to international markets or not. This capital mobility indicator is based on the index (from 0 to 100) in Quinn, Schindler, and Toyoda (2011). We use a continuous version of their index rescaled to the unit interval, with 0 meaning fully closed and 1 fully open.

Using this notation, our *trilemma instrument* is defined as

$$z_{it} \equiv (\Delta r^*_{b(i)t} - \widehat{\Delta r^*}_{b(i)t}) \times PEG_{it} \times PEG_{i,t-1} \times KOPEN_{it}. \tag{1}$$

In order to get a quick sense of the relationship between the endogenous variable and its instrument, Figure 1 shows scatter plots and fitted values from a regression of the change in short rates in the home country, Δr_{it}, on the instrument z_{it} defined in expression (1). These values correspond to those reported in columns 2 and 3 of Table 1. Table 1 reports first-stage regression results of the endogenous variable, Δr_{it} on the instrument z_{it}, without controls in columns 1–3 to provide the output for Figure 1, and then more formally with controls, in columns 4–6 to evaluate the strength of the instrument.

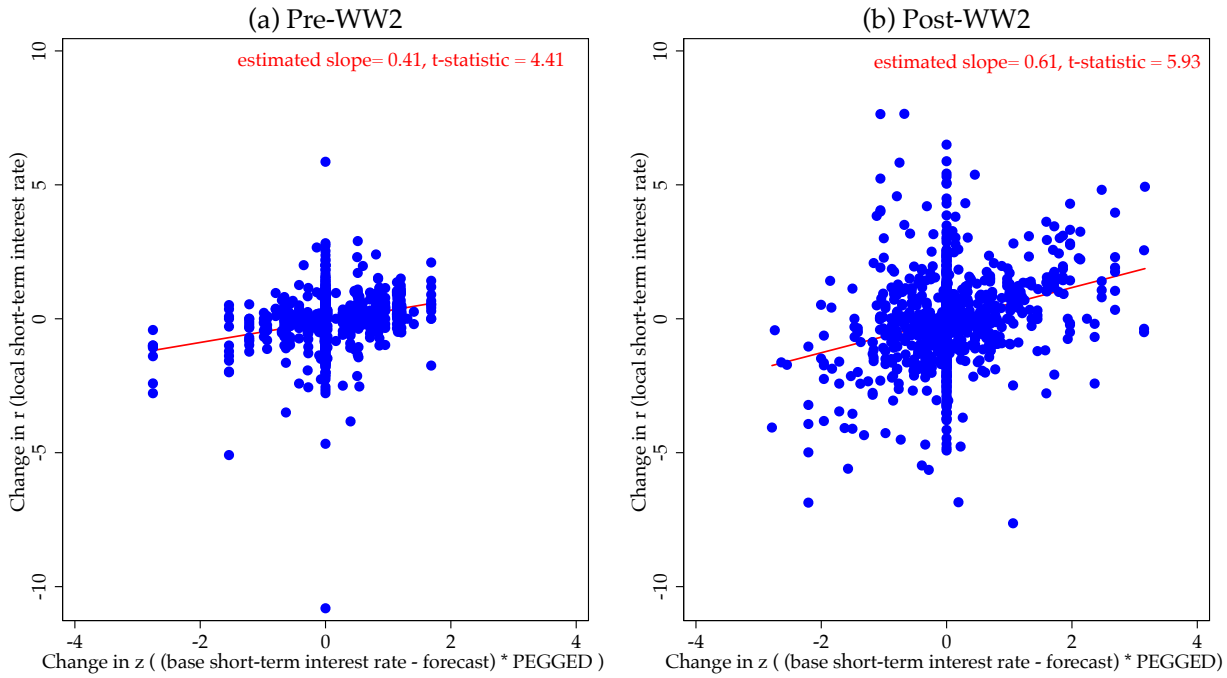
Table 1 reaffirms the evidence displayed in Figure 1 formally: Clearly z_{it} is not a weak instrument. Columns 4–6 refer to the formal first-stage regression with controls, country fixed effects and robust (clustered) standard errors (the regression also allows the coefficients of the controls to differ for the period 1973 to 1980 to account for the two oil crises as we also discuss below). The t-statistic on z_{it} in all cases is well above 3.5. Moreover, notice that the slope estimates in columns 4–6 are similar to those in columns 1–3, which suggests that the instrument contains information that is quite different from that contained in the controls.[4] We will pause further discussion of our instrument to introduce the statistical framework used for causal identification. First, we briefly discuss another of the novelties in this paper, the data.

3. The data

This paper relies on a large macro-historical dataset that we have assembled with the help of a small army of research assistants and generous support from many colleagues, and kept up to date over the past years. The database covers a broad range of real and financial variables for 17 countries at annual frequency for the modern era from 1870 to 2013. In Jordà, Schularick, and Taylor (2017), we

[4]The control list will be discussed in more detail below but it basically includes up to two lags of the first difference in log real GDP, log real consumption, investment to GDP ratio, credit to GDP, short and long-term government rates, log real house prices, log real stock prices, and CPI inflation.

Figure 1: *Relationship between change in short-rates in the home country and the trilemma instrument*



Notes: Pre-WW2 sample: 1870–1938 excluding WW1: 1914–1919. Post-WW2 sample: 1948–2013.

Table 1: *Relationship between change in short-rates in the home country and the trilemma instrument*

	No controls			With controls		
	(1) All years	(2) Pre-WW2	(3) Post-WW2	(4) All years	(5) Pre-WW2	(6) Post-WW2
Constant	-0.10***	-0.15**	-0.05***	-0.25**	-0.29***	-0.31
	(0.03)	(0.07)	(0.01)	(0.09)	(0.06)	(0.20)
z_{it}	0.54***	0.41***	0.61***	0.40***	0.49***	0.36***
	(0.08)	(0.09)	(0.10)	(0.08)	(0.14)	(0.10)
t-statistic	[6.93]	[4.41]	[5.93]	[5.16]	[3.49]	[3.50]
Observations	1922	866	1056	1223	312	911

Notes: *** p <0.01, ** p <0.05, * p <0.1. Standard errors in parentheses. Full sample: 1870–2013 excluding 1914–1919 and 1939–1947. Pre-WW2 sample: 1870–1938 (excluding 1914–1919). Post-WW2 sample: 1948–2013. Country fixed effects included in the regressions for columns 4–6. These regressions also include up to two lags of the first difference in log real GDP, log real consumption, investment to GDP ratio, credit to GDP, short and long-term government rates, log real house prices, log real stock prices, and CPI inflation. In addition we include world GDP growth to capture global cycles. See text.

describe the underlying data and its sources in great detail. The data come from a broad range of sources, including series constructed by economic historians and statistical offices, economic and financial history papers, yearbooks of central banks, as well as sources from central and private bank archives. The database contains the near-universe of available modern macroeconomic data for the following advanced economies: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, UK, and USA. Unlike the vast majority of empirical papers investigating the effects of monetary policy, our analysis rests both on a long time dimension and a large group of countries.

On the national accounts side, we have near complete coverage of real GDP data, interrupted only by a few war-time gaps. However, as we generally exclude war years from the empirical analysis this does not impact the results. Other than GDP, we also study the effects of changes in monetary conditions on consumption and investment decisions. For consumption and investment, we typically rely on country-level historical national account reconstructions as well as the data assembled by Barro and Ursúa (2010). On the financial side, we can work with a wide range of both quantity and price data. The credit data cover total bank lending to the non-financial private sector. We also have broad coverage of asset prices, again relying on numerous sources. The dataset comprises stock price indices as well as house price data from Knoll, Schularick, and Steger (forthcoming), the first major attempt to construct long-run house price indices for the advanced economies. A central variable for our analysis is the short-term interest rate. We typically use a risk-free short-term treasury bill rate with a maturity between 1 and 3 months. In some cases, we relied on central bank discount rates or short-term deposit rates of large financial institutions. Long-term interest rates are yields on government bonds with a maturity between 5 and 10 years.

3.1. Defining the base country in fixed exchange rate regimes

Part of the data construction effort consists of defining the base countries to which home countries pegged their exchange rates across different eras. This is done in Table 2. The possible base country interest rates used at different times in the history of exchange rate regimes correspond to the four rows in the table. The four major eras correspond to the four columns in the table. The table cells indicate which home countries correspond to each base in each era.

Prior to WW2, peg codings are taken from Obstfeld, Shambaugh, and Taylor (2004, 2005). After WW2 they are gleaned from Ilzetzki, Reinhart, and Rogoff (2008) and updates thereto. One exception with respect to this literature is that we do not code Germany as being pegged from 1999 onwards to emphasize the outsize role that Germany plays within the euro zone as a continuation of its central position with the EMS/ERM fixed exchange rate system. Prior to 1914 we treat the U.K. as the base for everyone, and after 1945 we treat the U.S. as the base for everyone, with the exception of EMS/ERM/Eurozone countries for which Germany is the base after 1973. In the interwar period, the choice of a suitable base country is more challenging and subjective given the instability of the interwar gold standard period; we follow Obstfeld, Shambaugh, and Taylor (2004) in using a hybrid

Table 2: *Selection of base country short-term interest rate for pegged exchange rates by era*

	Home country interest rates linked with the base			
Base country interest rate	Pre-WW1	Interwar	Bretton Woods	Post-BW
UK (Gold standard/BW base)	All countries		Sterling bloc: AUS*	
UK/US/France composite (Gold standard base)		All countries		
USA (BW/Post-BW base)			All other countries	Dollar bloc: AUS, CAN, CHE, JPN, NOR
Germany (EMS/ERM/Eurozone base)				All other countries

*We treat Australia as moving to a U.S. dollar peg in 1967.
Notes: See text, Jordà, Schularick, and Taylor (2015), and Obstfeld, Shambaugh, and Taylor (2004, 2005). Pre-WW1: 1870–1914; Interwar: 1920–1938; Bretton-Woods: 1948–1971; Post-BW: 1972–2013.

"gold center" short-term interest rate, which is an average of U.S., U.K., and French short term rates depending on which of the three countries was pegged to gold in a particular year; our results are not sensitive to this choice and we replicate our findings using any one of these three countries as the sole interwar base as in Obstfeld, Shambaugh, and Taylor (2004).

Base countries, such as the U.S. in the Bretton Woods era, are conventionally understood to pay no attention whatsoever to economic conditions in partner countries when making policy choices. Such behavior finds ample support in the historical record, as discussed in Jordà, Schularick, and Taylor (2015). Thus to peg is to sacrifice monetary policy autonomy, at least to some degree. It is therefore natural for us to treat the trilemma instrument, which we will denote as z_{it}, as an exogenous shifter of local monetary conditions in the home economy. The next section elaborates on the statistical design based on this idea.

4. Statistical discussion

Our empirical approach integrates elements of the potential outcomes paradigm of the Rubin Causal Model (Rubin 1974) into vector time series analysis using instrumental variable methods on local projections. In certain settings, fixed exchange rate regimes generate quasi-random variation in domestic interest rates, as we have seen. The extent to which base country interest rates are independent of home country economic conditions is key to satisfying the assumption of random assignment of the instrumental variable. The assumption that base country rate fluctuations affect home country outcomes via domestic rates only (the exclusion restriction) is natural when capital is allowed to flow freely.

The positive association between our instrument and domestic rates discussed in Section 2 fulfills the monotonicity assumption discussed in Imbens and Angrist (1994). Moreover, the relative size disparity between pegging and base economies suggests that a basic technical requirement, the stable unit treatment value assumption discussed below, is a reasonable starting point for the analysis. Finally, to allow for state-dependent policy effects, flexible functional forms, a large vector of controls (9 variables and their lags) and the panel structure of the data, we use instrumental variable local projections (LP-IV) with country fixed effects, a world business-cycle control, and cluster robust standard errors.

In a nutshell, our methods in this paper are as follows. First, the LP-IV setup allows us to estimate Local Average Treatment Effects (LATE) of interest rates on macroeconomic outcomes for the subpopulation of countries that peg their exchange rate and allow capital to move freely, the *pegs*. We find that, compared to the LP-IV estimates obtained with our set up, typical estimates based on regression control (such as VARs) display considerable attenuation bias of the average effect of policy interventions. Using a different instrumental variable based on Romer and Romer (2004) and Cloyne and Hürtgen (2014), we also compute LATE for the U.S. and the U.K. in the post-WW2 period, the best we can do to estimate the LATE for the subpopulation of *bases*. Using these LATE estimates limited to two bases we find that they are comparable to our LATE estimates for the wider sample of all pegs.

Second, we find similarities across OLS estimates for countries that peg versus float and similarities between the LATE obtained with two different instruments across two different subpopulations. We then argue that these findings, together with the persistence of exchange rate regimes (which seem to be mostly independent from the macroeconomic outlook), help us to build the case that our LATE estimates provide as good an approximation to the Average Treatment Effect (ATE) of monetary interventions as existing data currently allow. We now proceed to a more detailed and formal presentation of the methods.

4.1. Preliminaries and notation

Let $\boldsymbol{w} = (w_1, ..., w_k, ..., w_{K-1}, r)$ be a $K_w \times 1$ vector of random variables containing observations on $K-1$ macroeconomic outcomes (suitably transformed) plus a final additional variable r, a short-term interest rate variable. As a notational convention, we use boldface to denote random vectors, and normal type for random variables. Suppose we have a sample of panel data $\{\boldsymbol{w}_{i,t}\}_{i=1,t=1}^{N,T}$ where, e.g., for our application, $i = 1, ..., N$ denotes the 17 countries in our sample and $t = 1, .., T$ refers to the annual historical sample of nearly 150 years.

We will be interested in estimating the average effect of an experimental perturbation in r on to one of the outcome variables (element k) in $\boldsymbol{w}$ measured over time. Although in linear models such a concept is largely equivalent to the well-worn impulse response, we think that the definitions that we provide below are closer to the economic and policy relevant statistics one would want to characterize in practice.

Using the notation $y(h)$ to denote any random variable y observed h periods from the present, define the random variable $y(h) \equiv w_k(h) - w_k(-1)$ for any $k \in \{1, ..., K_w\}$, that is, the difference in the macroeconomic outcome variable w_k (that is, we also include r as a possible outcome) h periods in the future relative to last period. To make the notation specific to observations in the sample, $y(h)$ would simply be $y_{i,t+h} = w^k_{i,t+h} - w^k_{i,t-1}$. In order to keep the notation simple, we will omit the index k in the notation for $y(h)$ since, in the discussion that follows, we will only discuss one outcome variable at a time. On occasion, we will expand this notation to $y(r, h)$ which will refer to the potential outcomes one would observe for different values of the treatment (policy) variable r, h periods from today.

We will find it useful to construct the random vector $\boldsymbol{y} = (y(0), ..., y(h), ..., y(H-1))$ of dimension $1 \times H$ to collect the average response to r over time. Similar to the notation $y(r, h)$, we expand the notation $\boldsymbol{y}$ to $\boldsymbol{y}(r)$ to refer to the potential response paths one could observe for a given value of the treatment (policy) variable r. Next, we will define the vector of observable controls given by $\boldsymbol{x} = (w_1, ..., w_{k-1}, w_{k+1}, ..., w_{K-1}, r, \boldsymbol{w}(-1), ..., \boldsymbol{w}(-p))$. This vector is of dimension $1 \times K_x$, where $K_x = K_w(p+1) - 1$. Importantly, this vector includes contemporaneous observations for all the macro-outcomes (except the k^{th}, which appears as a left-hand side variable) and the interest rate r, as well as up to p lags of all the variables. Observations of the vector of controls from a sample over time and across countries are denoted $\{\boldsymbol{x}_{i,t}\}_{i=1,t=1}^{N,T}$.

Finally, it is important to note that we distinguish between population conditions involving generic random variables over which we make assumptions (which will be referenced without indices), and random variables observed in a finite sample (which will be referenced by subindices).

4.2. Selection on observables

This section discusses ideal conditions under which covariate control strategies (such as typical structural VAR settings) would allow identification of the average treatment effect (ATE) of monetary policy. In particular, suppose that, conditional on $\boldsymbol{x}$ the value of r is determined independently of the potential outcome $\boldsymbol{y}(r) = (y(0, r), y(1, r), ..., y(H-1, r))$. In other words, if r reflects the choice of interest rates made by a policymaker, the assumption would imply that $\boldsymbol{x}$ contains all the information used by the policymaker in determining r up to random variation, and therefore the value of r conditional on $\boldsymbol{x}$ is as if it had been randomly assigned. If r were a discrete variable that took on only a small number of values, such as $r \in \{-1, 0, 1\}$, then we would be in the familiar territory of the policy evaluation literature (see, e.g., Angrist 1995) and the full panoply of matching estimators would be at our disposal (for examples based on inverse propensity score weighting see Angrist, Jordà and Kuersteiner (forthcoming) and Jordà and Taylor (2016)). Here we assume that r takes on continuous values and therefore we seek identification through a different avenue.

Traditional identification of causal effects in monetary economics relies implicitly on assumptions that can be helpfully compared with the potential outcomes paradigm. The first assumption that we state is common in cross-sectional studies. In such settings it is assumed that there is no interference

between units (here country-year pairs). That is, the potential outcomes of one unit are independent of the treatment assignment to another unit. This assumption is commonly referred to as *Stable Unit Treatment Value Assumption* or SUTVA (Cox 1958; Rubin 1978) and is generally untestable.

In an international panel setting such as ours, this assumption deserves further discussion. SUTVA is a strong assumption to make regarding *treatment* assignment. The choice of policy interest rates in one country could affect the potential outcomes of another (specially those of a close trading partner). Violations of SUTVA would require explicit modeling of these interactions. The approach that we take here is more pragmatic. In addition to allowing for heterogeneity, we try to control for global factors that would tend to make policy choices correlated. Moreover, we correct standard errors for possible clustering as further insurance. That said, SUTVA is less problematic when evaluating the *instrument* assignment.

Economies that peg their exchange rate are usually small relative to the base economy. In such cases, a reasonable working assumption is that *pegs* have little effect on the policy choices made by *bases*. Moreover, despite having common base economies, pegging countries display considerable heterogeneity in their interest rate choices, as shown in Figure 1. Inference on violations of SUTVA and its consequences undoubtedly deserve further exploration, but that exploration is beyond the scope of this paper. In fact, there is relatively little research available that we can cite on this issue. For all these reasons, we will not explore violations of SUTVA further.

We assume that the process determining r can be written as $r(\boldsymbol{x}, \boldsymbol{\psi}, v)$, where the parameter vector $\boldsymbol{\psi} \in \Psi$ is assumed to be fixed (in order to avoid variation in policy due to structural breaks in the policy reaction function rather than through random chance via v). A more detailed discussion of this assumption is provided in Angrist, Jordà, and Kuersteiner (forthcoming). In other words, here we focus on policy interventions that do not come from changes in the policy reaction function.

Here, we state these ideas somewhat formally with the following assumption, which in one form or another, has characterized most empirical work on monetary policy to date. Based on the definitions for $\boldsymbol{y}(r), \boldsymbol{x}$, and r, the selection-on-observables condition (sometimes also called conditional ignorability, or conditional independence) can be stated as:

Assumption 1 (Selection on observables, or conditional ignorability)**.**

$$\boldsymbol{y}(r) \perp r \mid \boldsymbol{x} \qquad \textit{for } \boldsymbol{\psi} \in \Psi \textit{ fixed},$$

where $\perp$ indicates conditional statistical independence and $y(r)$ refers to the set of potential outcomes y could take for a given value of r. We say that perturbations in r given $\boldsymbol{x}$ are independent of the **potential** *outcome $\boldsymbol{y}(r)$ for $\boldsymbol{\psi} \in \Psi$ fixed.*

At first glance, the statement of Assumption 1 may seem peculiar since it appears to indicate that the outcome is orthogonal to the treatment given controls. However, notice that $\boldsymbol{y}(r)$ refers to the *potential* outcomes $\boldsymbol{y}$ can take, not to its realization, $\boldsymbol{y}$. That is, in ideal random assignment of treatment, potential outcomes should not affect the treatment chosen. We are not assuming that the treatment effect is null.

Selection-on-observables requires any observable information relevant in determining the treatment to be included in the conditioning set. This ensures that any remaining variation in treatment assignment is as good as random. Responses estimated by local projections can be specified and estimated as single equations with rich conditioning sets (in terms of explanatory variables), but in a parametrically conservative manner from the point of view of the number of lags required for consistency (see Jordà 2005 for an explanation of this point). In a panel context, where VAR specifications can be parametrically taxing, single equation panel estimation is particularly convenient.

Notice that Assumption 1 implies that any available information, specially that contained in contemporaneous variables, should be included in the conditioning set. To put it in the parlance of orderings in Cholesky decompositions from VARs, selection on observables would suggest that to calculate the responses to a given shock, the variable being shocked should always be ordered last. And if one were interested in calculating the responses to another shock, it would be preferable to reorder the variables again so that the corresponding variable is ordered last as well.

Instead, Assumption 1 can be more naturally accommodated using local projections. Rather than trying to zero out some variables, one can simply include them in the conditioning set. Since each variable can be analyzed separately, there is no need to exclude any contemporaneous information in any of the equations (other than the left hand side variable, naturally). Assumption 1 emphasizes the desirability to condition on as much relevant information as possible to ensure that any remaining variation in the treatment variable is as close to random as possible.

Under Assumption 1, the average dynamic response of the $1 \times H$ vector of outcomes $\boldsymbol{y}$ to a perturbation in r of size δ relative to the counterfactual of no perturbation, $\delta = 0$, all conditional on $\boldsymbol{x}$, would be a natural statistic of interest. For that, we define

$$I(\boldsymbol{y}, r, \delta, \boldsymbol{x}) \equiv \mathbb{E}(\boldsymbol{y}(r+\delta) \mid \boldsymbol{x}) - \mathbb{E}(\boldsymbol{y}(r) \mid \boldsymbol{x}) = \mathbb{E}(\boldsymbol{y}(r+\delta) - \boldsymbol{y}(r) \mid \boldsymbol{x}). \tag{2}$$

The notation $\mathbb{E}(.)$ refers to the ideal population expectation one would compute over the distribution of potential outcomes for all values of $r + \delta$ and its counterfactual r. Each of these random variables, $\boldsymbol{y}(r+\delta)$ and $\boldsymbol{y}(r)$, cannot be simultaneously observed. However, the expectation can be estimated by imposing further assumptions.

In principle, for a given δ, the function $I(\boldsymbol{y}, r, \delta, \boldsymbol{x})$ would depend on the level of the policy rate r, and the values of the conditioning set $\boldsymbol{x}$. A natural starting point is to assume a linear relationship,

$$\mathbb{E}(\boldsymbol{y}(r) \mid \boldsymbol{x}) = \boldsymbol{A}(r, \boldsymbol{x}) + r\boldsymbol{\gamma}(r, \boldsymbol{x}) + \boldsymbol{x}\boldsymbol{B}(r, \boldsymbol{x}). \tag{3}$$

Notice that all coefficients are written as a function of r and $\boldsymbol{x}$. When treatment is binary, it is customary to allow for the conditional mean parameters to differ between treatment and control. Generalizing this principle, one could imagine estimating expectations that differ across the different values of r. Since this is parametrically impractical, we assume for now that the coefficients do not depend on r. We relax this assumption later on. Similarly, the effect of treatment could vary

depending on $\boldsymbol{x}$. An increase in interest rates may have different effects when the economy is in recession than when it is in expansion. Again, we will relax this assumption as well. Therefore, we simplify expression (3) further and write instead,

$$E[\boldsymbol{y}(r) \mid \boldsymbol{x}] = \boldsymbol{A} + r\boldsymbol{\gamma} + \boldsymbol{x}\boldsymbol{B}, \tag{4}$$

where we replace the notation $\mathbb{E}$ to E to denote that we pursue an approximation to the ideal, and where $\boldsymbol{A}$ is a $N \times H$ matrix that collects fixed-effects coefficients for all N countries across all H horizons with $\boldsymbol{A} = (\boldsymbol{\alpha}^0, ..., \boldsymbol{\alpha}^{H-1})$. The effect of the treatment on the outcome is measured by $\boldsymbol{\gamma}$, a $1 \times H$ vector of average experimental responses, with $\boldsymbol{\gamma} = (\gamma_0, ..., \gamma_{H-1})$. Finally, $\boldsymbol{B}$ is a $K_x \times H$ matrix of coefficients such that $\boldsymbol{B} = (\boldsymbol{\beta}_0, ..., \boldsymbol{\beta}_{H-1})$ with $\boldsymbol{\beta}_h$ a $K_x \times 1$ vector for $h = 0, 1, ..., H-1$.

Given a finite sample for $\{(\boldsymbol{y}_{i,t}, r_{i,t}, \boldsymbol{x}_{i,t})\}_{i=1,t=1}^{N,T}$, a natural estimate of γ_h could be obtained by OLS using the following H panel regressions:

$$y_{i,t+h} = \alpha_i^h + r_{i,t}\,\gamma_h + \boldsymbol{x}_{i,t}\boldsymbol{\beta}_h + u_{i,t+h} \quad \text{for } h = 0, 1, ..., H-1. \tag{5}$$

Notice that the right-hand side regressors in (5) are common to $h = 0, 1, ..., H-1$ and, therefore, this can be estimated as a system of seemingly underrated regressions (SUR) or individually. This is nothing but the usual local projection estimator (see Jordà, 2005). Such an estimator is especially convenient for a panel context as it can be generalized relatively easily for state-dependence, as we shall see later. Moreover, it includes a Panel VAR assumption of the data generating process (DGP) as a special case.

The linear conditional mean assumption in expression (4) imposes constraints on the responses estimated that are not often appreciated. As noted earlier, $I(y, h, \delta)$ are independent of the specific values that r and $\boldsymbol{x}$ take. The average response to a perturbation δ in r is therefore the same whether the economy is expansion or recession, above or below potential, in a credit boom or a credit bust, and other important states of the economy. Similarly, the response is symmetric, that is $I(y, h, \delta) = -I(y, h, -\delta)$. However, a more general setup can easily allow for asymmetry and this will be important later in this paper. As Angrist, Jordà, and Kuersteiner (forthcoming) and Tenreyro and Thwaites (2016) have shown, there appear to be considerable differences in the effectiveness of monetary tightening relative to loosening. We will investigate a few of these extensions below. Next, we introduce appropriate machinery to exploit our instrumental variable.

4.3. Instrumental variables and local average treatment effects (LATE)

Moving away from these ideal conditions, note that Assumption 1 relies on having a rich set of controls that properly characterize all movements in r up to random variation. A more realistic scenario is that fluctuations in r could reflect unobserved factors that also affect the macroeconomic outcomes. This violation of Assumption 1 would result in inconsistent estimates of $I(y, h, \delta)$. We will see that this assumption is in fact violated in the data in one of our subpopulations.

We introduce additional assumptions to deal with such situations. Suppose that an instrument, z is available, as in our application, although in the more general case it is natural to think of a vector of instrumental variables $z_1, ..., z_m$. We now define $\boldsymbol{x}^*$ to be the same as $\boldsymbol{x}$ but with the element r omitted, that is, $\boldsymbol{x}^* \equiv (w_1, ..., w_{k-1}, w_{k+1}, ..., w_{K-1}, \boldsymbol{w}(-1), ..., \boldsymbol{w}(-p))$. This is appropriate, since r is the treatment or policy variable for which we have an instrument. Finally, we can then define $\boldsymbol{z} = (z_1, ..., z_m, \boldsymbol{x}^*)$. Using these definitions we now discuss a set of alternative assumptions to Assumption 1. These assumptions can be thought of as the time series version of some of the conditions in Angrist, Imbens, and Rubin (1996) but for continuous data as in Imbens (2014). Note that the specific timing conventions are motivated by the structure of our data. Moreover, because later on we will analyze state-dependent policy effects, we are mindful to specify these assumptions conditional on the state s used to stratify the analysis. For now the state is left unspecified.

In what follows we stick with the single instrument notation (i.e., $m = 1$) to facilitate the exposition without loss of generality. Note that by definition, for country i, $\boldsymbol{x}^*_{it}$ refers to controls observed at time t and its lags.

Assumption 2 (Unconfoundedness or random assignment of the instrument)**.**

$$z \perp \boldsymbol{y}(r), r(z) \mid \boldsymbol{x}^*, s. \tag{6}$$

This assumption is basically the instrument validity assumption of traditional econometric studies. We require the instrument to be independent of potential outcomes and of potential policy choices given controls. Naturally, one expects the instrumental variable to have predictive power for the policy variable r. However, the assumption states that fluctuations in the instrument are not determined by the potential policy choices that are likely to be made and denoted $r(z)$. This is one of the features of the potential outcomes notation that is perhaps less common in traditional time series studies. Similarly, the independence of z is with respect to the potential macroeconomic outcomes one could observe given the policy choice $r(z)$, which in turn potentially depends on the value of z. Again, the assumption does not mean that z and $\boldsymbol{y}$ are independent of one another. As explained in Imbens (2014), there are several ways in which this assumption can be stated. For our purposes and under the assumption of linearity, we could relax this assumption and invoke instead conditional mean assumptions. However, we think it is preferable to maintain the more general statement of the assumption for now.

Next, we state the exclusion restriction.

Assumption 3 (Exclusion restriction)**.**

$$\boldsymbol{y}(r, z, \boldsymbol{x}^*, s) = \boldsymbol{y}(r, z', \boldsymbol{x}^*, s) \quad \textit{for all } z, z'\textit{; given } \boldsymbol{x}^* \textit{ and } r\textit{, and state } s. \tag{7}$$

This assumption indicates that any effects of the instrument on the outcome come directly through the policy variable. Since these may be different depending on the state and on the conditioning history up to the point of intervention, we are careful in specifying each element.

The next assumption that we state is essentially the usual relevance assumption in instrumental variables analysis, augmented to account for potential differences given the state s.

Assumption 4 (Relevance)**.**

$$z \not\perp r \mid \boldsymbol{x}^*, s. \tag{8}$$

The final assumption we make is that of *monotonicity*.

Assumption 5 (Monotonicity)**.**

$$\textit{If} \quad z > z' \quad \textit{then} \quad E[r \mid z, \boldsymbol{x}^*, s] \geq E[r \mid z', \boldsymbol{x}^*, s] \quad \textit{for all} \;\; \boldsymbol{x}^*, s. \tag{9}$$

The monotonicity assumption plays an important role in the identification of the LATE although formal tests of the assumption are usually not available. Several forms of stating this assumption have appeared in the literature, although there are few results that provide guidance in our empirical setting. The stricter forms are based on higher-order stochastic dominance arguments. Here we prefer to make a statement in terms of mean dominance since it has a more natural connection with the results presented in Section 2—when interest rates increase in the base country, interest rates are predicted to increase in the pegging country.

Linearity assumptions further facilitate estimation in practical settings. The equivalent to expression (4) requires two further assumptions. First, we assume linearity of the conditional mean of r given $\boldsymbol{z}$, that is,

$$E[r \mid \boldsymbol{z}; s] = \boldsymbol{\mu}^s + \boldsymbol{z}\boldsymbol{\theta}^s \equiv \tilde{r}^s, \tag{10}$$

where $\boldsymbol{\mu}^s$ collects country fixed effects, and $\boldsymbol{\theta}^s$ is a $(m + K_x - 1) \times 1$ vector of coefficients. In principle, we allow the first stage to depend on the state variable s, although in our application we assume the coefficients are constant across states. In addition to economizing on estimated regressors, the assumption is not too restrictive. It simply means that there is a tight connection between the instrument and the policy variable, regardless of the state, as the trilemma mechanism would suggest.

Second, we assume linearity of the exclusion restriction:

$$E[\boldsymbol{y}(\tilde{r}) \mid \boldsymbol{z}; s] = \boldsymbol{A}_s^* + \tilde{r}\boldsymbol{\gamma}_s + \boldsymbol{x}^*\boldsymbol{B}_s^*, \tag{11}$$

where $\boldsymbol{A}_s^*$ is a $N \times H$ matrix that for a given state s, collects fixed effects coefficients for all N countries across all H horizons, that is $\boldsymbol{A}_s^* = (\boldsymbol{\alpha}_s^{0,*}, ..., \boldsymbol{\alpha}_s^{H-1,*})$; $\boldsymbol{\gamma}_s$ is the same $1 \times H$ vector of average experimental responses $\boldsymbol{\gamma}_s = (\gamma_{0,s}, ..., \gamma_{H-1,s})$ we were calculating under expression (4); and $\boldsymbol{B}_s^*$ is a $K_x - 1 \times H$ matrix of coefficients such that $\boldsymbol{B}_s^* = (\boldsymbol{\beta}_{0,s}^*, ..., \boldsymbol{\beta}_{H-1,s}^*)$, with $\boldsymbol{\beta}_{h,s}^*$ a $K_x - 1\times$ vector for $h = 0, 1, ..., H - 1$.

The equivalent to expression (5) would consist of the two-step estimator based on expressions

(10) and (11), that is:

$$
\begin{aligned}
r_{i,t} &= \mu_i^s + \boldsymbol{z}_{i,t}\boldsymbol{\theta}^s + \epsilon_{i,t} \qquad \Longrightarrow \qquad \hat{r}_{i,t}^s = \hat{\mu}_i^s + \boldsymbol{z}_{it}\hat{\boldsymbol{\theta}}^s, \\
y_{i,t+h} &= \alpha_{i,s}^h + \hat{r}_{i,t}^s \gamma_{h,s} + \boldsymbol{x}_{i,t}^* \boldsymbol{\beta}_{h,s}^* + \eta_{i,t+h}.
\end{aligned} \tag{12}
$$

Thus the notation allows for situations in which policy responses may be state dependent. With this machinery now in place, an estimate of the state-dependent counterpart to expression (2) is easily seen to be:

$$
I(y, h, \delta, s) = \hat{\gamma}_{h,s}\delta\,; \qquad \text{for } h = 0, 1, ..., H-1\,;\, s. \tag{13}
$$

5. Subpopulation experiments, population effects

In Section 2 we showed that the trilemma instrument introduced there meets Assumptions 4 and 5. The strategy in the next few sections consists of the following steps. First we compare estimates of monetary interventions under the assumption of conditional ignorability for the three subpopulations defined earlier: country-year pairs in which an economy pegged its exchange rate, denoted as *pegs*; country-year pairs in which an economy allowed its exchange rate to float, denoted as *floats*; and country-year pairs in which a given economy served as the base economy to which others pegged their exchange rate to, denoted as *bases*. The objective is to compare estimates of γ in expression (4) across subpopulations.

On a first pass, the assumption of conditional ignorability is rejected for pegs. Estimates of LATE based on LP-IV differ substantially from estimates based on LP-OLS. There is considerable attenuation bias. Next, using a different instrument, we are able to calculate LATE for two economies that serve as base economies over the sample for which the instrument is observed: the U.K. and the U.S. during most of post-WW2. This part of the analysis reveals the significant extent of attenuation bias due to failure of conditional ignorability in two representative base country-year samples. It also allows us to compare the LATE estimates against those obtained with our IV for the subpopulation of pegs and thus to serve as a cross-check. These causal estimates of treatment effects will set the stage for the last part of the analysis in which we investigate state-dependence.

Is our estimated LATE a good approximation to the unobserved ATE in all bins under different monetary and exchange-rate regimes? We argue that it is.

On paper, pegs approach systematic monetary policy very differently than floats. In the latter, systematic policy prescribes the interest rate to target explicit inflation and economic activity goals. In the former, a different systematic policy prescribes the interest rate by just mirroring the base country to maintain the peg, while seeking other stability goals implicitly. In practice, pegs seek the same price and output stability objectives as bases. They do it by taking advantage of the credibility earned by the central bank in the base country.

Consider a highly stylized setting, with many details omitted for clarity. Suppose we are in a symmetric world with one base economy and one other economy that can either peg or float. If

business cycles in the two economies were perfectly synchronized, and if the central bank in the base economy behaved optimally given the same policy goals as the other economy, the decision to peg or float would be indeterminate from the point of view of systematic monetary policy.

The world is much more complicated, of course. But for us the relevant question is this: what would allow us to extrapolate the LATE measured for pegs to the LATE we are unable to estimate for floats? Don't they have very different approaches to systematic monetary policy? There is a two pronged response to this question. First, we have succinctly articulated a case for why differences between pegs and floats may be smaller than feared at first glance—indeed, our stylized example is a case where the systematic policies are equivalent. Second and perhaps more importantly, the causal effect identified for pegs is based on non-systematic movements of base country interest rates. These are as close to an exogenous source of variation as we can get. Importantly, they are, by construction, silent about the policy choices a country makes.

In this sense, our approach is no different than the approach followed by a vast literature on empirical monetary policy. Unless the fundamental mechanics of the macroeconomy differ between pegs and floats, there would be no reason to expect this causal effect to be too different between the two subpopulations.

5.1. Subpopulation (local) average treatment effects under conditional ignorability

Consider the three subpopulations defined earlier: *pegs*, *floats*, and *bases*. A country can fall into any of these three bins throughout its history. For example, during Bretton Woods, Germany was in a *peg*, to the dollar. With the end of Bretton Woods, and later the introduction of the European Monetary System, we consider Germany to become a *base*. And there are other periods where we classify Germany as a *float*, as was the case for much of the interwar period. Other than bases, all other countries are either floats or pegs, depending on the period.

When presenting results, we will always measure and display the outcome variable in deviations relative to its initial value in year 0, with units shown in percent of the initial year value (computed as log change times 100), except in the case of interest rates where the response will be measured in units of percentage points. The treatment variable or impulse will be defined as the one-year change in the short-term interest rate in year 0, and normalized in all cases to a 1 percentage point, or 100 basis points (bps) increase.

The vector of control variables $\boldsymbol{x}^*$ introduced earlier includes a rich set of macroeconomic controls consisting of the first-difference of the contemporaneous values of all variables (excluding the response or outcome variable), and up to $p = 2$ lags of the first-difference of all variables, including the response variable. The list of macroeconomic controls is: log real GDP per capita; log real real consumption per capita; log real real investment per capita; log consumer price index; short-term interest rate (usually a 3-month government security); long-term interest rate (usually a 5-year government security); log real house prices; log real stock prices; and the credit to GDP

ratio. The sources for all of the variables used, covering 17 countries and all years since 1870, are described in more detail in Jordà, Schularick, and Taylor (2017), and its online appendix.

In almost all respects, this estimation setup produced stable outcomes. However, in line with the well-known "price puzzle" literature (e.g., Eichenbaum 1992; Sims 1992; and Hanson 2004), we found substantial instability in the coefficients of these control variables driven by the postwar high-inflation period of the 1970s. The traditional resolution of this puzzle has been to include commodity prices as a way to control for oil shocks. Given the constraints of our data, we choose to address this issue by allowing the controls to take on a potentially different coefficient for the subsample period of years from 1973 to 1980 inclusive, thus bracketing the volatile period of the two oil crises.

We begin by estimating the effect of an interest rate intervention on output (measured by real GDP per capita) and prices (measured by the CPI) only in the interest of saving space. Central Bank mandates usually relate to economic activity and price stability so these two variables are good reference points. These results are provided in Table 3. These estimates are based on a panel regression that allows the relevant coefficient estimates to vary for each of the three subpopulations that we consider: pegs, floats and bases. The table evaluates whether estimates across subpopulations are statistically different from one another. In addition, we also provide a joint test that, over the 5 horizons considered, the effect of interest rates on output and prices is zero. The analysis is conducted over the full and the post-WW2 samples.

Consider the output responses first. These are reported in columns 1 to 3. Full sample results indicate some minor differences across subpopulations. The p-values of the null that the coefficients are equal is reported in column 4. The differences are qualitatively minor, however. Turning to the post-WW2 results, p-values from column 4 suggest that if anything, the differences are even less important over this sample. Generally speaking, the coefficient estimates have the expected signs. An increase in interest rates causes output to decline. Note that in all cases the effect is statistically different from zero as reported in the rows labeled $H_0 : ATE = 0$.

Turning to the price responses, reported in columns 5 to 7, the picture turns murky. The overall effect of an interest rate increase on prices in the full sample is essentially null for pegs and floats (columns 5 and 6), with a p-value of 0.30 and 0.93 respectively. Bases display an overall negative response, as would be expected. Perhaps for this reason the null that the coefficient estimates across subpopulations are equal (column 8) has p-values that are all below 0.10. The picture changes somewhat for the post-WW2 subsample. Some negative signs appear for $h = 4$ but the responses are generally not very different from zero in the statistical sense.

The main takeaways from Table 3 can be summarized as follows. Viewed from the lens of the output response, there would be little evidence that anything is amiss. The responses across subpopulations are similar, they have the expected signs, and their effects are statistically significant. Looking at the post-WW2 results, the year 4 estimates indicate that a one percent increase in interest rates would make output (on average across subpopulations) about 0.5 percentage points smaller than it would otherwise have been, about a 0.1 percent in annual rate of decline. Of course, the price

Table 3: *LP-OLS. Real GDP per capita and CPI price responses to interest rates*

Responses at years 0 to 4 (100× log change from year 0 baseline).

(a) Full sample	*Output response*			*P=F=B*	*Price response*			*P=F=B*
Year	Pegs (1)	Floats (2)	Bases (3)	*p*-value (4)	Pegs (5)	Floats (6)	Bases (7)	*p*-value (8)
$h = 0$	0.07**	-0.13	0.22***	0.01	0.15**	0.35***	-0.06	0.01
	(0.03)	(0.10)	(0.06)		(0.06)	(0.11)	(0.08)	
$h = 1$	-0.13	-0.46***	-0.17*	0.17	0.25**	0.63**	-0.13	0.02
	(0.10)	(0.15)	(0.09)		(0.11)	(0.24)	(0.14)	
$h = 2$	-0.15	-0.67***	-0.47***	0.05	0.27*	0.53	-0.27	0.07
	(0.12)	(0.20)	(0.14)		(0.14)	(0.39)	(0.18)	
$h = 3$	-0.19	-0.66***	-0.33***	0.20	0.22	0.25	-0.57***	0.08
	(0.14)	(0.20)	(0.11)		(0.22)	(0.49)	(0.25)	
$h = 4$	-0.10	-0.53***	-0.19	0.26	0.17	0.16	-0.83**	0.09
	(0.17)	(0.21)	(0.12)		(0.31)	(0.57)	(0.35)	
$H_0 : ATE = 0$	0.03	0.03	0.02		0.30	0.01	0.93	
Observations		1247				1247		
(b) Post-WW2	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$h = 0$	0.03	-0.01	0.13**	0.19	0.16***	0.21**	0.05	0.22
	(0.02)	(0.08)	(0.05)		(0.05)	(0.09)	(0.06)	
$h = 1$	-0.18	-0.29*	-0.28**	0.72	0.29***	0.45**	0.13	0.23
	(0.12)	(0.15)	(0.10)		(0.09)	(0.20)	(0.11)	
$h = 2$	-0.27*	-0.50**	-0.64***	0.12	0.23	0.35	0.00	0.46
	(0.15)	(0.20)	(0.13)		(0.14)	(0.31)	(0.16)	
$h = 3$	-0.36**	-0.46**	-0.63***	0.24	0.03	-0.02	-0.35*	0.40
	(0.15)	(0.20)	(0.10)		(0.21)	(0.40)	(0.19)	
$h = 4$	-0.32	-0.42	-0.68***	0.24	-0.09	-0.21	-0.59**	0.32
	(0.19)	(0.25)	(0.11)		(0.29)	(0.51)	(0.27)	
$H_0 : ATE = 0$	0.00	0.02	0.00		0.00	0.00	0.19	
Observations		891				891		

Notes: *** p <0.01, ** p <0.05, * p <0.1. P denotes pegs, F floats, B bases. Cluster robust standard errors in parentheses. Full sample: 1870–2013 excluding WW1: 1914–1919 and WW2: 1939–1947. PostWW2 sample: 1948–2013. The column *P=F=B* displays the p-value of the null that for a given horizon h, estimates of the corresponding elasticity are equal across subpopulations. $H_0 : ATE = 0$ refers to the null that the coefficients for $h = 0, \ldots, 4$ are jointly zero for a given subpopulation. See text.

responses would offer a less reassuring picture, in great part, because the responses are generally insignificant and often have the "wrong" sign. The next step is therefore to examine the estimates of the local average treatment effect for the pegs and for bases based on two different instruments.

5.2. LP-IV LATE estimates for pegs and bases

Using our trilemma instrument, we compare LP-OLS with LP-IV estimates for the subpopulation of pegs making sure to match the samples used. LP-OLS estimates in Table 4 correspond to those in columns 1 and 5 in Table 3. Slight differences in the parameters reflect differences in the available samples for each subpopulation once the instrument is factored in. Under the assumptions of Section 4, we calculate the LATE of an interest rate intervention and evaluate any attenuation bias from violations of conditional ignorability using a Hausman test. Table 4 summarizes the main results.

The table is organized as follows. The output and price responses in columns 1 and 4 are LP-OLS estimates over the same sample as the LP-IV estimates reported in columns 2 and 5. Column 3 reports the p-value of the Hausman test of the null that the entry in column 1 is equivalent to the entry in column 2. Column 6 reports the Hausman test between columns 4 and 5 instead. We check whether the trilemma instrument is weak with Kleibergen-Paap tests. Finally, we test the null that all LATE coefficients are jointly zero and report the p-value of the test in the row labeled $H_0 : LATE = 0$.

The first check is to compare the LP-OLS responses reported in columns 1 and 4 here, with those reported in Table 3 in columns 1 and 5. Recall that in Table 3 we estimate the model using all observations, but allow coefficients to vary by subpopulation. The differences are relatively minor, owing to slight differences in the sample used given the availability of the instrument.

The important result of Table 4 is the size of the attenuation bias between LP-OLS and LP-IV. The differences are economically sizable and statistically significant as indicated by the Hausman tests of columns 3 and 6. Conditional ignorability clearly fails. Using LP-OLS estimates (column 1) and the full sample, output would be estimated to by about 0.06% lower four years after an increase in interest rates of 1%. In contrast, the LP-IV effect is measured to be nearly a 3% decline, or about a 0.6% annualized rate of lower growth. A similar pattern is observable for the price response. Full sample LP-OLS estimates are largely insignificant and often have the wrong sign. LP-IV estimates are sizable, significant, and have the right sign.

Comparing the full sample results with the post-WW2 results we find differences in the output response to be relatively minor. The price response, however, becomes somewhat delayed. The LP-IV response suggests that on impact and the year after, the price response is essentially zero although by year 4, prices are expected to be about 2% lower than they were four years earlier. Tests for weak instruments suggest the trilemma instrument is relevant and tests of the null that the LATE estimated with LP-IV is statistically different from zero. Interest rates have a strong effect on output and prices for the subpopulation of pegs.

Table 4: *LP-OLS vs. LP-IV. Attenuation bias of real GDP per capita and CPI price responses to interest rates. Trilemma instrument. Matched samples*

Responses at years 0 to 4 (100× log change from year 0 baseline).

(a) Full sample	*Output response*		*OLS=IV*	*Price response*		*OLS=IV*
Year	LP-OLS	LP-IV	p-value	LP-OLS	LP-IV	p-value
	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.10*	-0.22*	0.01	0.09	-0.22	0.11
	(0.04)	(0.13)		(0.05)	(0.20)	
$h = 1$	-0.16	-1.05***	0.00	0.22**	-0.70**	0.01
	(0.10)	(0.23)		(0.10)	(0.33)	
$h = 2$	-0.19	-2.00***	0.00	0.11	-1.61***	0.00
	(0.15)	(0.35)		(0.14)	(0.44)	
$h = 3$	-0.21	-2.31***	0.00	-0.08	-2.91***	0.00
	(0.19)	(0.44)		(0.22)	(0.70)	
$h = 4$	-0.06	-2.97***	0.00	-0.17	-3.88***	0.00
	(0.22)	(0.63)		(0.32)	(0.92)	
KP weak IV		48.14			42.76	
$H_0 : LATE = 0$		0.00			0.01	
Observations	667	667		667	667	
(b) Post-WW2	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.06*	-0.03	0.31	0.07	0.19	0.45
	(0.03)	(0.08)		(0.05)	(0.16)	
$h = 1$	-0.13	-0.90***	0.00	0.18**	0.10	0.78
	(0.10)	(0.24)		(0.08)	(0.29)	
$h = 2$	-0.20	-1.89***	0.00	0.09	-0.50	0.11
	(0.14)	(0.37)		(0.13)	(0.37)	
$h = 3$	-0.23	-2.03***	0.00	-0.13	-1.35***	0.01
	(0.17)	(0.42)		(0.22)	(0.45)	
$h = 4$	-0.15	-2.62***	0.00	-0.30	-1.96***	0.00
	(0.21)	(0.63)		(0.33)	(0.57)	
KP weak IV		37.03			33.86	
$H_0 : LATE = 0$		0.00			0.01	
Observations	522	522		522	522	

Notes: *** p <0.01, ** p <0.05, * p <0.1. Cluster robust standard errors in parentheses. Full sample: 1870–2013 excluding WW1: 1914–1919 and WW2: 1939–1947. PostWW2 sample: 1948–2013. *Matched sample* indicates LP-OLS sample matches the sample used to obtain LP-IV estimates. *KP weak IV* refers to the Kleibergen-Paap test for weak instruments. $H_0 : LATE = 0$ refers to the p-value of the test of the null hypothesis that the coefficients for $h = 0, \ldots, 4$ are jointly zero for a given subpopulation. *OLS=IV* shows the p-value for the Hausman test of the null that OLS estimates equal IV estimates. See text.

The next exercise is reported in Table 5 and involves comparing these results with estimates based on a different instrumental variable and subpopulation. We turn to the Romer and Romer (2004) instrument updated by Cloyne and Hürtgen (2014) for the U.S. and the U.K. We denote the instrument as *RRCH*. Both the U.S. and the U.K. can be thought of as belonging to the subpopulation of bases and thus provide the best approximation of the LATE results for this group. The results in this table are organized in a similar manner as Table 4. The top panel reports the output and price responses estimated by LP-OLS and LP-IV using the RRCH instrument. We note that the RRCH instrument is available only from 1969 to 2007 for the U.S. and 1976 to 2007 for the U.K. Thus, to make the results more directly comparable, the bottom block of Table 5 provides similar results to those in Table 4 but limited to the same years (1969–2007).

Consider this comparison first. Despite the more abbreviated sample reported in the bottom half of Table 5, the results based on the trilemma instrument are essentially the same as those reported in the bottom half of Table 4. Thus, any differences between pegs are bases is not explained by a different historical setting. The more interesting results appear in the upper block of the table. Note that the results for the U.S. and the U.K. are based on a relatively small sample (71 observations). The conclusions based on this part of the analysis are colored by this limitation.

Concerning the output response, a Hausman test confirms that differences between LP-OLS and LP-IV estimates are minor. The responses typically have the correct sign and are of a similar magnitude and shape. They are only slightly attenuated with respect to the corresponding responses calculated for the pegs in the bottom half of the table. The differences are more apparent for the price response. The LP-OLS response of prices to an interest rate intervention (column 4) is close to zero economically and statistically, with coefficients that often have the wrong sign. The LP-IV response has a similarly muted response initially but it becomes increasingly negative. By year 4, prices are expected to be about 2.5% lower than they would otherwise would be, a response that is very similar to that reported in the bottom half of the table and to that in Table 4.

The main findings so far can be summarized as follows. First, we report strong evidence of attenuation bias when comparing LP-OLS to LP-IV estimates broken down by subpopulation. We take this evidence to suggest that conditional ignorability, despite the extensive set of controls that we include, is insufficient to achieve identification of the causal effect of interest rates on output and prices. More broadly, this result casts doubts on the impulse responses reported using standard VARs predicated on common identification assumptions not based on external instruments.

Second, the responses calculated for pegs using our trilemma instrument and those calculated for the U.S. and the U.K. using the RRCH instrument are reasonably close. We conclude from these results that the LATE estimates for pegs and those for bases are somewhat similar. Therefore going forward, we use both instruments in tandem to maximize statistical power while recognizing that the results that we report will be driven primarily by the subpopulation of pegs (given the differences in sample size). Although we can only speculate about the causal effects of interest rates for floats, the similarities of LP-OLS estimates reported in Table 3, and the results in Table 5 provide some reassurance that interest rate responses are similar for this subpopulation as well.

Table 5: *LP-OLS vs. LP-IV. Attenuation bias of real GDP per capita and CPI price responses to interest rates. U.S. vs U.K. Trilemma versus RRCH instrument. Matched samples*

Responses at years 0 to 4 (100× log change from year 0 baseline).

(a) RRCH IV	*Output response*		*OLS=IV*	*Price response*		*OLS=IV*
Year	LP-OLS	LP-IV	p-value	LP-OLS	LP-IV	p-value
	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.11	0.39***	0.07	0.12	0.43*	0.19
	(0.03)	(0.16)		(0.13)	(0.23)	
$h = 1$	-0.25	-0.23	0.90	0.47	0.83**	0.27
	(0.20)	(0.23)		(0.13)	(0.33)	
$h = 2$	-0.74	-0.57	0.75	0.65**	0.79	0.82
	(0.14)	(0.53)		(0.02)	(0.62)	
$h = 3$	-1.19*	-0.69	0.55	0.08	-0.59	0.52
	(0.10)	(0.82)		(0.39)	(1.04)	
$h = 4$	-0.97*	0.14	0.21	-0.51	-2.52*	0.16
	(0.11)	(0.89)		(0.69)	(1.42)	
KP weak IV		13.12			12.85	
$H_0 : LATE = 0$	0.00	0.00		0.00	0.00	
Observations	71	71		71	71	
(b) Trilemma IV	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.04	0.00	0.71	0.07	0.16	0.48
	(0.02)	(0.09)		(0.05)	(0.13)	
$h = 1$	-0.12	-0.85***	0.00	0.18	0.04	0.60
	(0.13)	(0.22)		(0.10)	(0.26)	
$h = 2$	-0.16	-1.61***	0.00	0.10	-0.69*	0.05
	(0.18)	(0.32)		(0.14)	(0.41)	
$h = 3$	-0.15	-1.57***	0.00	-0.08	-2.17***	0.00
	(0.21)	(0.37)		(0.21)	(0.60)	
$h = 4$	-0.08	-1.49***	0.00	-0.17	-3.49***	0.00
	(0.25)	(0.37)		(0.34)	(0.81)	
KP weak IV		16.63			15.35	
$H_0 : LATE = 0$	0.05	0.00		0.01	0.00	
Observations	372	372		372	372	

Notes: *** p <0.01, ** p <0.05, * p <0.1. Cluster robust standard errors in parentheses. RRCH refers to the Romer and Romer (2004) and Cloyne and Hürtgen (2014) IV. U.S. sample: 1969–2007. U.K. sample: 1976–2007. *Matched sample* indicates that the LP-OLS sample matches the sample used to obtain LP-IV estimates. *KP weak IV* refers to the Kleibergen-Paap test for weak instruments. $H_0 : LATE = 0$ refers to the p-value of the test of the null hypothesis that the coefficients for $h = 0, \ldots, 4$ are jointly zero for a given subpopulation. *OLS=IV* shows the p-value for the Hausman test of the null that OLS estimates equal IV estimates. See text.

6. A novel LP-IV LATE spillover correction

Simply stated, the exclusion restriction in Assumption 3 says that the instrument affects the local outcome only through its correlation with the intervention variable. There is no connection between the local outcome and the instrument otherwise. Economic arguments supporting this assumption are strong. Our instrument is the *unpredictable* component of base-country interest rates given observable base-country controls (scaled by the degree of local openness). Surprise movements in base-country interest rates are then assumed to explain, at least partially, the exogenous variation in pegging-country interest rates that cannot be explained by that country's controls.

But for the deeply skeptical, a priori arguments for exclusion restrictions often fall on deaf ears. However convoluted the argument, for them the possibility always exists for spillover effects going from the instrument to the outcome through channels other than the endogenous variable. As an example of how that economic argument may go, consider an oil shock. Such a shock could affect global markets. Although we control for global business cycle effects directly (to avoid such contamination), such a shock may induce unpredictable movements in base-country interest rates that are correlated with movements in pegging-country outcomes, notably inflation. Such movements in inflation could, to some extent, be unrelated to the domestic interest rate response but be incorrectly attributed to the interest rate channel under the exclusion restriction. In this example, the IV estimator would be biased relative to its true value.

We now show how our setup offers a novel way to assess potential contamination from instrument spillover effects by taking advantage of the subpopulation of floats. The main ideas can be explained with a simple example. Consider the scalar case where, as before, y is the outcome variable, r is the intervention, z is the instrument. We abstract from state dependence and from controls to keep things simple.

In general, the standard IV setup can be summarized as follows,

$$y = r\,\gamma + z\,\delta + u, \tag{14}$$

$$r = z\,\theta + \epsilon, \tag{15}$$

where $E(r\,u) \neq 0$, but $E(z\,u) = 0$, and the second line in the expression is the familiar first-stage regression in IV estimation.

The exclusion restriction takes the form of an assumption that $\delta = 0$. If this restriction does not hold, it is easy to see that:

$$\hat{\gamma}_{IV} \xrightarrow{p} \gamma + \frac{\delta}{\theta}.$$

In other words, the bias induced by the failure of the exclusion restriction depends on both the size of the failure (δ) *and* the strength of the instrument (θ). Weaker instruments will tend to make the bias worse. This point was made in, for example, Conley, Hansen and Rossi (2011).

Our identification of interest rate effects is thus far based on the subpopulation of pegs. Now we will show how we can use an auxiliary regression from the subpopulation of floats to manage potential violations of the exclusion restriction. In economies whose exchange rate moves freely, central banks retain full control of interest rate policy. For these reasons, our LP-IV LATE has focused on the peg subpopulation. However, even in floats, central banks may still chose, for example, to smooth fluctuations in the exchange rate in response to interest rate policy in other large (base) economies. In addition, there may also be spillover effects in floats from large economies.

We argue that floats contain useful information that we now exploit. In our application, the subpopulation of floats is such that $z = 0$ by construction since $PEG_{it} \times PEG_{i,t-1} = 0$ in expression (1). Consider redefining the instrument for the float subpopulation as $z = KOPEN \times (\Delta \hat{r}_b^* - \Delta r_b^*)$ (omitting indices). That is, ignore the trilemma for a moment. We will still maintain the assumption that $E(z\,u) = 0$ since large economies (those in the base subpopulation) are unlikely to use information on the macroeconomic outlook of float economies when setting interest rates.

Next, consider estimating (14) using OLS (the appendix derives the more general results with controls). Estimates of the intervention effect γ and the spillover effect δ will be biased as long as $E(r\,u) \neq 0$ and $\theta \neq 0$. However, it is easy to show (under standard regularity conditions) that the OLS estimates of expression (14) are such that

$$\left.\begin{array}{lll} \hat{\gamma}_{OLS} & \xrightarrow{p} & \gamma - \lambda \\ \hat{\delta}_{OLS} & \xrightarrow{p} & \delta + \theta\,\lambda \end{array}\right\} \quad \text{with bias} \quad \lambda = \frac{E(r\,u)E(z^2)}{E(z\,r)^2 - E(z^2)E(r^2)}. \tag{16}$$

Expression (16) provides a very intuitive way to understand the OLS bias in the outcome equation for floats. If $E(r\,u) > 0$, then $\lambda > 0$ and the effect of domestic interest rates on outcomes will be attenuated by the bias term λ. Similarly, the spillover effect will be amplified by an amount $\theta\,\lambda$. This amplification will be larger the stronger the correlation between r and z as measured by the pseudo first-stage coefficient θ. The results in Section 5 suggest that there is considerable attenuation bias in $\hat{\gamma}$ when comparing OLS versus IV estimates. The implication is that simple OLS will tend to make the spillover effect seem larger than it really is, and the interest rate response smaller than it really is. But of course, if $E(r\,u) < 0$, then $\lambda < 0$, and the sign of the biases is reversed, and all of the above statements would be inverted. Put simply: the direction of this bias is ambiguous, and cannot be assumed ex ante to go one way or the other.

In order to quantify these biases, Table 6 reports estimates based on expression (5) for real GDP per capita and the price level, augmented with the presently-redefined z_{it}. We define the instrument z_{it} using the UK as the base country for all float economies before 1939 (we drop war years associated with WW1 and WW2), the U.S. for all economies in the Bretton Woods period (1946-1971) and thereafter for Australia, Canada, Japan and the U.K., and Germany after Bretton Woods (from 1972 onwards) for any remaining country (all of which are European). This is the same specification as that used in Tables 3 and 4 extended with the instrument z, but this time using the subpopulation of floats. Table 6 also reports the coefficient associated with the pseudo-first stage

regression of the domestic interest rate on the instrument defined for a floating economy. This is the parameter θ in expression (16).

Table 6 makes clear the intuition behind expression (16). The interest rate responses of real GDP per capita reported in column 1 are economically small. They are statistically insignificant for the full sample estimates reported in panel (a) of the table, and only significant in years 3 and 4 in the Post-WW2 sample reported in panel (b). In contrast, the response to the instrument (think of it as a shock to the base country interest rate) is almost three times larger and significant. Price responses follow a different pattern, with responses to the own interest rate shock of the wrong sign, but responses to the base country interest rate (measured by the instrument) of the correct sign. This is a feature we will return to in the results reported below.

Finally, we note that the regression of the domestic interest rate on the instrument and the control set is generally non-zero, but about half the magnitude of the coefficient estimated in the first stage regression for the peg subpopulation and reported in Table 1. Compare 0.20 for the float subpopulation with 0.40 for the peg subpopulation using full sample estimates. These results are consistent with those reported in Obstfeld, Shambaugh and Taylor (2005).

Expression (16) does not allow us to proceed any further, in general. Absent further information or assumptions, we know nothing about γ and δ. However, in our particular problem we argue that economic logic can supply a constraint on this relationship that will allow us to achieve identification. We take as a reasonable assumption that the direct effect from r, the domestic interest rate, is probably stronger than that coming from the spillover, z. Note that, because both variables relate to interest rates, their scales are comparable. Without loss of generality, suppose that $\gamma = \alpha\delta$, that is, the true domestic interest rate effect on outcomes is a scaled version of the spillover effect from the foreign interest rate. We will hold that the economically plausible prior is that $\alpha \geq 1$.

Now, it is easy to see that

$$\hat{\delta}(\alpha) = \frac{(\hat{\delta}_{OLS} + \hat{\theta}\hat{\gamma}_{OLS})}{1 + \alpha\hat{\theta}} \quad \xrightarrow{p} \quad \delta(\alpha). \tag{17}$$

To make further progress, taking α as given, we can use a control function approach to correct our LP-IV LATE estimates for biases due to potential spillover effects. Without loss of generality, expression (14) can be rewritten as

$$(y - z\hat{\delta}(\alpha)) = r\gamma + u + z(\hat{\delta}(\alpha) - \delta(\alpha)).$$

Moreover, the usual moment conditions imply that

$$E(z(y - z\hat{\delta}(\alpha))) = E(z\,r)\gamma + E(z\,u) + E(z^2(\hat{\delta}(\alpha) - \delta(\alpha))),$$
$$\text{with } E(z\,u) = 0\,, \text{ and } (\hat{\delta}(\alpha) - \delta(\alpha))\frac{1}{N_p}\sum_{j}^{N_p} z_j^2 \quad \xrightarrow{p} \quad 0,$$

Table 6: *LP-OLS. Real GDP per capita and CPI price responses to domestic and base-country interest rates. Full and post-WW2 samples for subpopulation of exchange rate float economies*

Responses at years 0 to 4 (100× log change from year 0 baseline).

(a) Full sample	*Output response*		*float=z*	*Price response*		*float=z*
	float rate	z	*p*-value	float rate	z	*p*-value
	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	-0.10	0.06	0.29	0.55**	-0.14	0.08
	(0.14)	(0.09)		(0.21)	(0.23)	
$h = 1$	-0.27	-0.17	0.80	1.36***	-0.35	0.05
	(0.23)	(0.21)		(0.44)	(0.47)	
$h = 2$	-0.26	-0.61**	0.46	1.75***	-0.72	0.03
	(0.28)	(0.21)		(0.56)	(0.66)	
$h = 3$	-0.25	-1.10***	0.11	1.70**	-1.23	0.04
	(0.28)	(0.32)		(0.67)	(0.85)	
$h = 4$	-0.11	-1.42***	0.04	1.84**	-1.80*	0.03
	(0.29)	(0.44)		(0.75)	(0.98)	
float on z	0.20***			0.22***		
first stage estimate	(0.06)			(0.06)		
Observations	266			266		
(b) Post-WW2	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	-0.06	-0.05	0.93	0.43*	-0.31	0.04
	(0.10)	(0.10)		(0.21)	(0.18)	
$h = 1$	-0.25	-0.30	0.86	0.81*	-0.60	0.05
	(0.15)	(0.25)		(0.40)	(0.42)	
$h = 2$	-0.25	-0.84**	0.31	1.08*	-0.84	0.07
	(0.25)	(0.36)		(0.57)	(0.60)	
$h = 3$	-0.41**	-1.46***	0.03	0.99	-1.23	0.09
	(0.18)	(0.33)		(0.65)	(0.74)	
$h = 4$	-0.46**	-1.84***	0.01	1.13	-1.72*	0.06
	(0.16)	(0.47)		(0.75)	(0.88)	
float on z	0.18*			0.16**		
first stage estimate	(0.09)			(0.07)		
Observations	207			207		

Notes: *** p <0.01, ** p <0.05, * p <0.1. Cluster robust standard errors in parentheses. *float rate* is the rate on short-term (usually 3-months) government bonds of the domestic economy considered and indexed by i; z is defined as $z_{it} = KOPEN_{it} \times (\Delta \hat{r}^*_{b(i),t} - \Delta r^*_{b(i),t})$. $r^*_{b(i),t}$ refers to the base country rate on short-term (usually 3-months) government bonds. We use the U.K. as the base before 1939; we use the U.S. for the Bretton-Woods era (1946-1971) for all countries, and for Australia, Canada, Japan, and the U.K. after Bretton-Woods (1972 onward); and we use Germany for all remaining economies (all in Europe) in the post Bretton-Woods era (1972 onward). *float* = *z* refers to the test of the null that the coefficient for the float rate is equal to that for the instrument and report the p-value of the test. *float on z* refers to the coefficient of the regression of the float rate on the instrument (called θ in the text in expression (14)). See text.

as long as

$$\frac{1}{N_p}\sum_{j}^{N_p} z_j^2 \xrightarrow{p} Q_z < \infty, \text{ and } N_f \longrightarrow \infty \text{ as } N_p \longrightarrow \infty,$$

with N_f and N_p the size of the subpopulation of floats and pegs respectively.

Our IV estimator *corrected for potential spillover effects* is then constructed by subtracting the spillover term from the outcome variable in the standard IV coefficient estimator,

$$\hat{\gamma}(\alpha) = \frac{\frac{1}{N_p}\sum z_j(y_j - z_j\hat{\delta}(\alpha))}{\frac{1}{N_p}\sum z_j r_j}, \text{ with } \sqrt{N_p}(\hat{\gamma}(\alpha) - \gamma(\alpha)) \xrightarrow{d} N(0, \ddot{\sigma}_\gamma^2), \text{ and } \ddot{\sigma}_\gamma^2 = F(\sigma_\gamma^2; \sigma_\delta^2; \alpha; k),$$

where $\ddot{\sigma}_\gamma^2$ denotes the population variance of $\hat{\gamma}(\alpha)$ using our modified IV estimator, σ_γ^2 denotes the variance one would typically estimate by ignoring uncertainty from the auxiliary regression on the subpopulation of floats, $k = \sqrt{N_f/N_p}$, that is the square root of the ratio of the relative sample sizes for the subpopulation of floats and pegs, and σ_δ^2 is the variance of the spillover effect estimated from the auxiliary regression on the subpopulation of floats. Naturally, we assume the sample sizes of both subpopulations go to infinity at an appropriate rate so that $k < \infty$. Notice that the variance would also be a function of α.

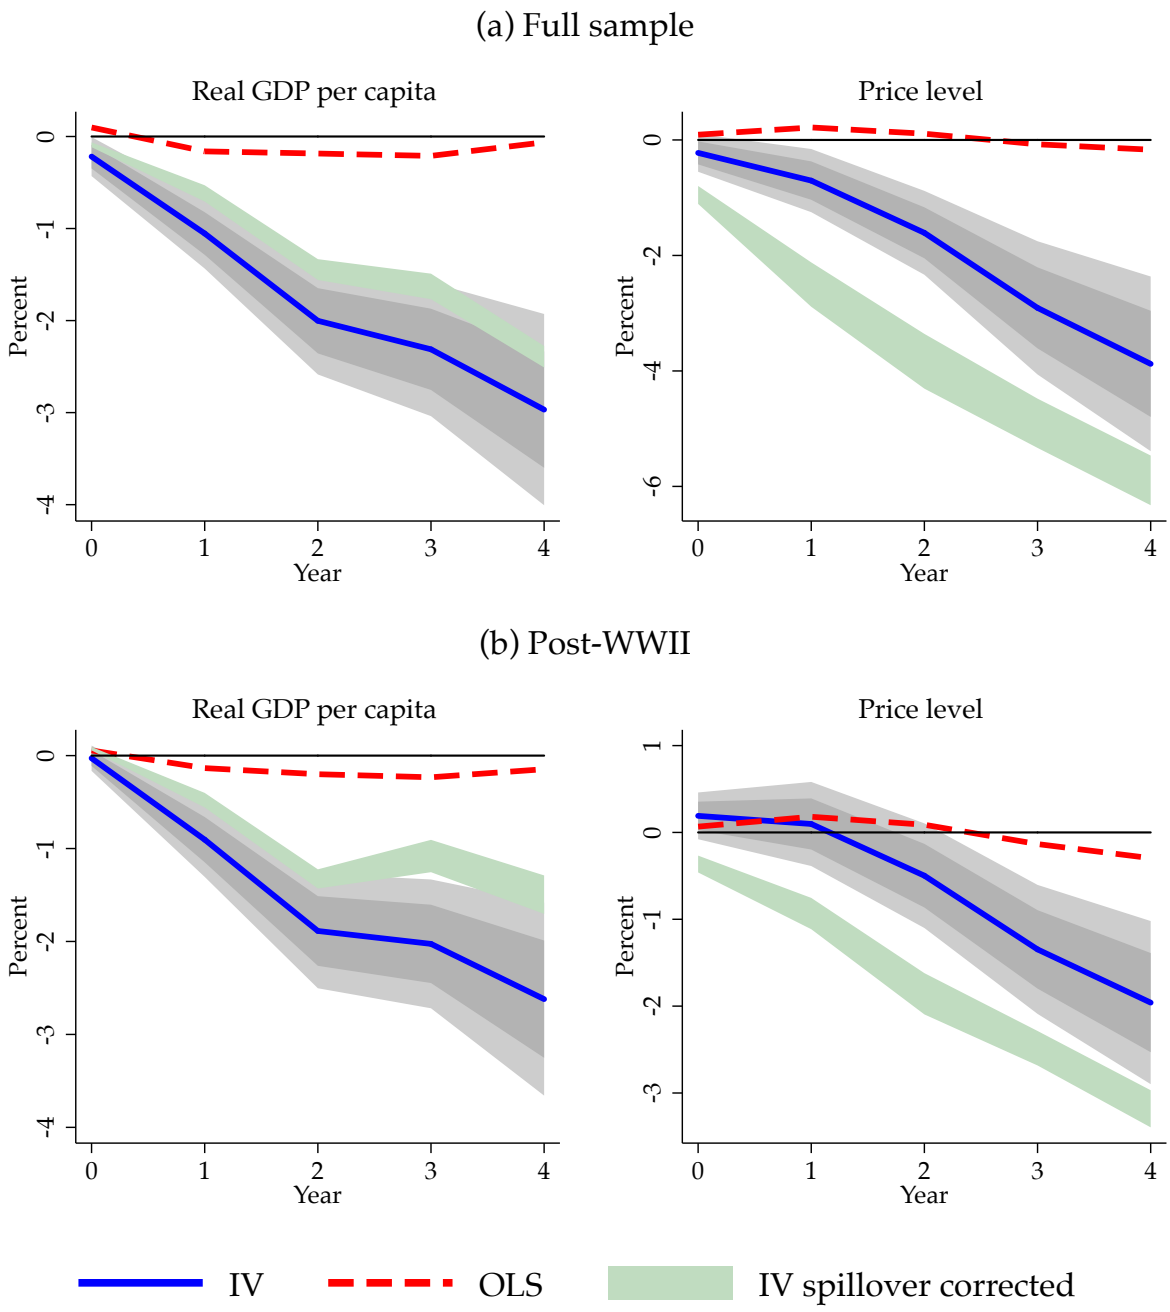
The appendix provides a more formal derivation for a panel setting such as ours and to account for controls. But the intuition is clear. One can use the subpopulation of floats to estimate the spillover effect channel (the δ in our example). Using this estimate, a control function approach to correct the usual IV estimator consists of simply netting out the spillover channel from the outcome variable. The corrected estimate of the treatment effect γ in (16) will be consistent and asymptotically Gaussian with variance inflated by the auxiliary regression uncertainty.

However, note that all of these results above depend on α, which is as yet an unknown parameter. Given our application, we have argued that it is natural to assume that $\alpha \geq 1$; that is, we think it is reasonable to assume that local interest rates have stronger effects on the local economy, under floating exchange rates, than do base country interest rates. In order to make progress, the corrected results that we report below use a range of values of α from 1 to 4. Economically speaking, this means that the local interest rate effect is 1 to 4 times larger than the effect coming from interest rates driven by, say, a large external central bank.

Using this control function method to correct the IV estimator, Figure 2 shows results based on the subpopulation of pegs but using the subpopulation of floats to adjust the results. The figure displays responses to a 1 percent increase in local interest rates for real GDP per capita and the price level using alternative estimators. The red-dashed lines show the responses reported in Table 4 in columns 1 and 4 based on LP-OLS using the rich control set described in the previous section. As we remarked earlier, the responses for GDP generally have the correct sign but are economically and statistically small. This is in line with the attenuation bias described in expression (16).

Next, the solid blue line with associated point-wise error bands in grey show the LP-IV estimates reported in columns 2 and 5 of Table 4. As we noted then, the responses are considerably larger,

Figure 2: *Real GDP per capital and CPI price responses to a 1 percent increase in interest rates. LP-OLS, LP-IV, and spillover corrected LP-IV*



Notes: Full sample: 1870–2013 excluding WW1: 1914–1919 and WW2: 1939-1947. LP-OLS estimates displayed with dashed-red line, LP-IV estimates displayed with a solid blue line and 1 S.D. and 90% confidence bands, LP-IV spillover corrected estimates displayed as a light green area using $\alpha \in [1, 4]$. See text.

both statistically (the null $H_0 : LATE = 0$ is rejected at the 1 percent level) and economically. Finally, the shaded region in light green displays the range of impulse responses that would result from our spillover adjustment for $\alpha \in [1, 4]$.

Several results deserve comment. First, notice that the correction for spillover effects tends to attenuate the responses as compared to our preferred LP-IV estimates. For the response of real GDP per capita, the correction suggests that the output response is about 0.5 to 1 percentage points less negative by year 4 than that reported using LP-IV alone. In the post-WW2 era this means that the cumulative change in output due to a 1 percent increase in rates is probably closer to -1.5% (about -0.3% per annum) than to -2.5% (about -0.5% per annum).

Interestingly, the response of prices is amplified. The reason is easy to see in Table 6. The LP-OLS estimate of the effect of domestic interest rates on prices is positive rather than negative. Meanwhile, the effect of base country interest rates is strongly negative. Thus, the correction makes the LP-IV price response more sensitive to interest rates, specially on impact. This feature has often been an achilles heel of the VAR literature, with either a price response that has the wrong sign or a price response that remains largely muted for a prolonged period of time even as the response of output shows a more immediate response.

7. The causal effects of interest rates on the macroeconomy

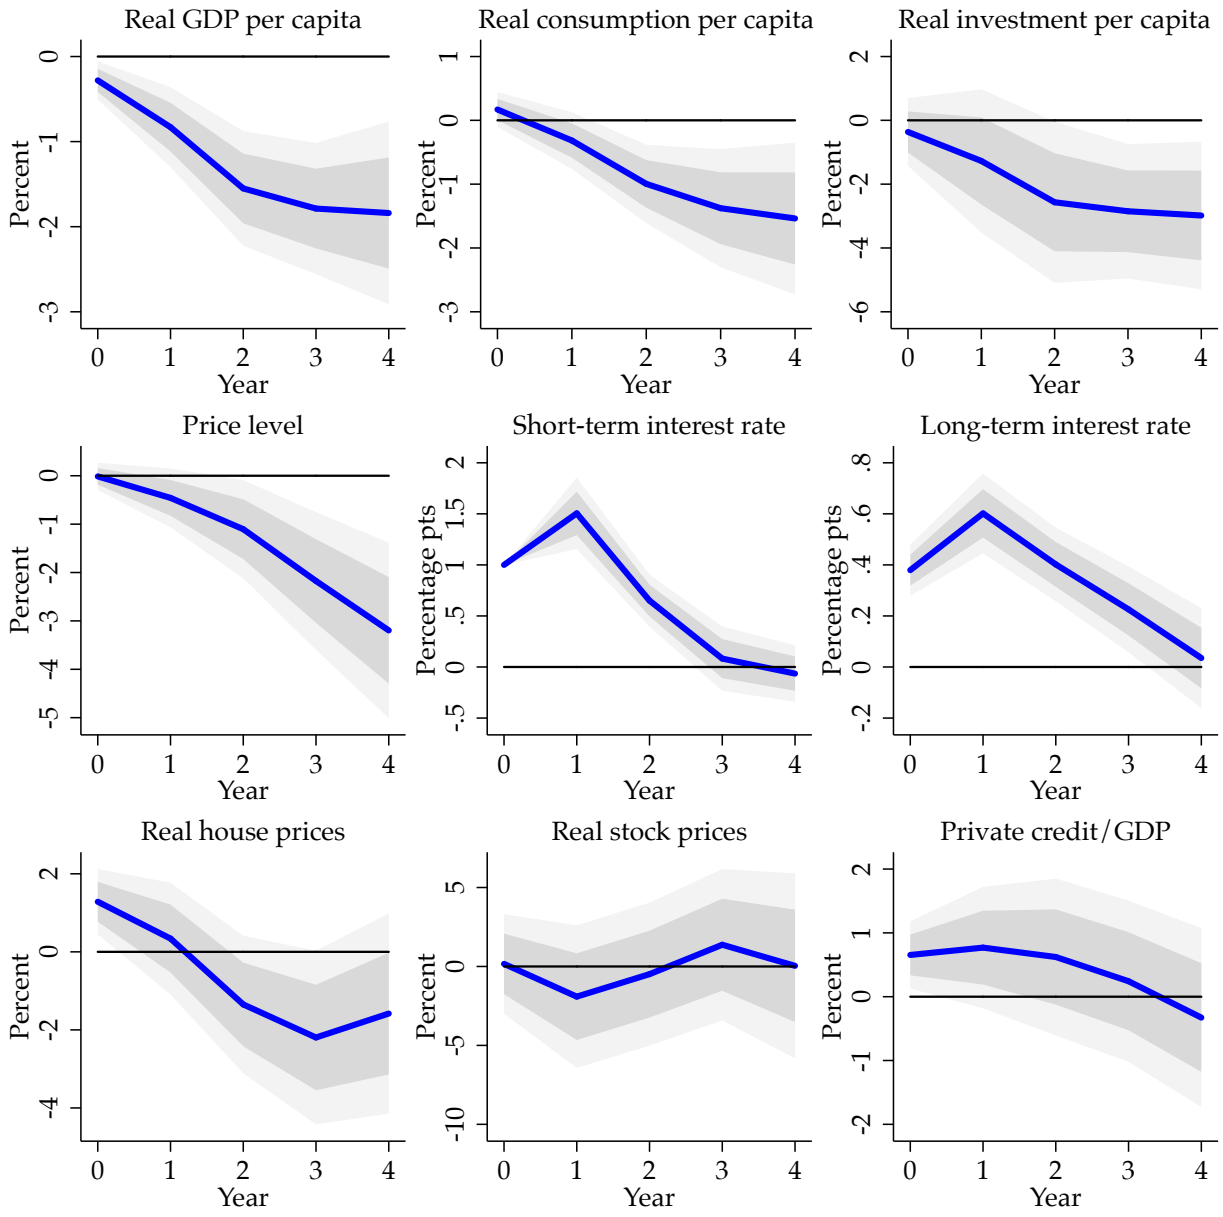
In this section we briefly present a comprehensive study of the causal response of macroeconomic outcomes to a short-term interest rate increase of 100 bps. Figures 3 and 4, summarize responses for our full set of variables for the Full and Post-WW2 samples, respectively.

Since the two sets of results are qualitatively and quantitatively very similar, we discuss only the full sample results. In these figures, a +100 bps increase in the short rate *causally* leads to a 2% decline in real GDP per capita (or about 0.5% per anuum), a 1%–1.5% decline in real consumption per capita, a 3%–4% decline in real investment per capita, and a 3% decline in the price level (second row, first column), where all effects are relative to the no-change policy counterfactual and the measurements are cumulative over the horizon of 4 years.

Turning to the second row of charts in Figure 3 we look first at the own response of short-term interest rates to a +100 bps rate rise in year 0 (row 2, column 2). This path reflects the intrinsic persistence of changes in interest rates. In this case, short-term interest rates increase by +150 bps in year 1, drop back to +75 bps in year 2, and then decline to effectively zero in both years 3 and 4. The next chart (row 2, column 3) shows the response of long-term interest rates, which are, as is well-known, more subdued in amplitude than short rates; the long-term interest rate moves about half as much. A +100 bps rise in the short rate *causally* leads to the long rate rising +40 bps in year 1, rising to +60 bps in year 2, and then falling back towards zero by year 4.

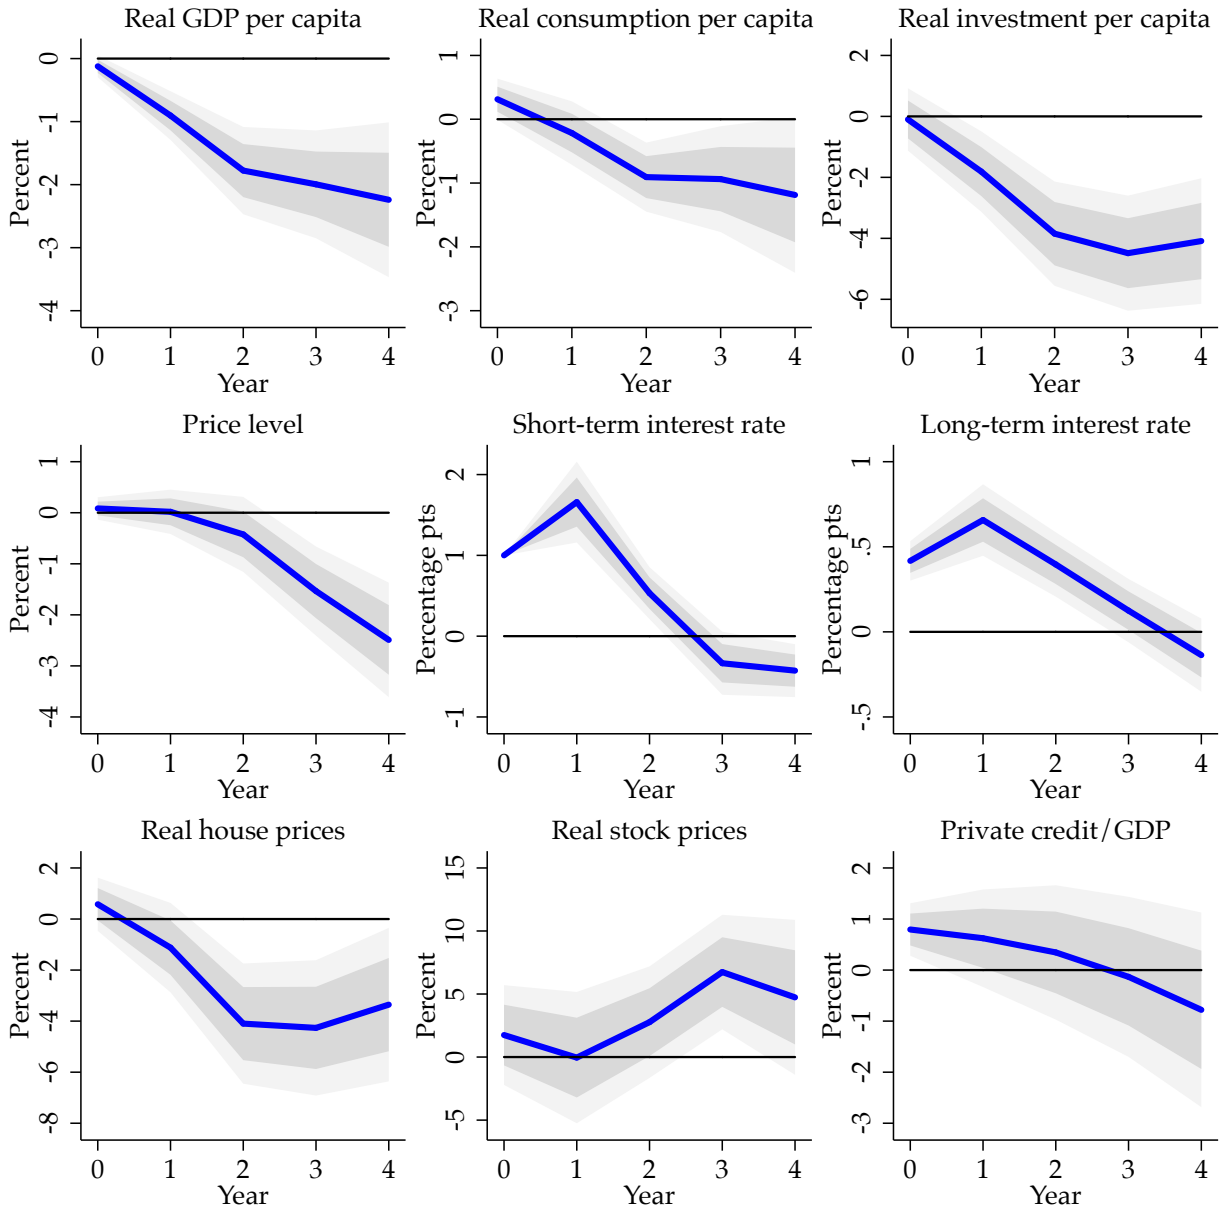
The last row (columns 1 and 2) show the responses of two key asset prices: a +100 bps rise in the short rate *causally* leads to a cumulative 4% decline in real house prices and a cumulative 5% decline in real stock prices over 4 years, again as compared to the no-change policy counterfactual. These

Figure 3: *Full baseline results. Full sample*



Notes: Full sample: 1870–2013 excluding WW1: 1914–1919 and WW2: 1939-1947. LP-IV estimates displayed with a solid blue line and 1 S.D. and 90% confidence bands. See text.

Figure 4: *Full baseline results. Post-WW2 sample*



Notes: Post-WW2 sample: 1948–2013. LP-IV estimates displayed with a solid blue line and 1 S.D. and 90% confidence bands. See text.

responses are consistent with a significant wealth-effect channel for monetary policy, alongside the more often noted income-effect channel visible in the path real GDP.

Finally, in the last row and column of this chart, we display the causal response of aggregate credit (bank lending to the nonfinancial sector relative to GDP). This chart shows that a +100 bps rise in the short rate has a relatively muted effect on the ratio of bank loans to GDP cumulated over four years, although this effect is in the end consistent with models where contractionary monetary policy leads to less demand and/or supply of credit. If anything, the effect is slightly more pronounced when using the Post-WW2 sample although the responses are very similar.

The takeaway from these impulse responses is clear. An exogenous shock to interest rates has sizable effects on real variables (larger than those measured using conventional VARs), but along the lines predicted by most monetary models with rigidities. Term structure responses conform very well with standard results in the literature. Nominal variables decline strongly. Perhaps the only variable that appears to be somewhat unresponsive to interest rates is the credit to GDP variable. Loans decline in response to a shock to interest rates, but their rate of decline matches closely the rate of decline in real economic activity. The next step is to investigate the stability of these results when we allow them to depend on the state of the economy.

8. State dependence

In this section, we extend the linear baseline model to address major issues that have emerged in the research literature as well as in current policy debates. We will stratify the analysis by whether the economy is in a boom or a bust, and by periods of low versus high inflation.

One concern is the possibility that policy responses are asymmetric. Angrist, Jordà, and Kuersteiner (forthcoming) argue that monetary policy loosening may have very different impacts on macroeconomic outcomes than monetary policy tightening. Barnichon and Matthes (2016), using identification via Gaussian basis functions, also find that expansionary monetary policy has virtually no effect, but contractionary policy is powerful. Earlier work by Hamilton and Jordà (2002) investigated a similar asymmetry, based on the nature of the surprise component of changes in the federal funds rate target. In their setting, the asymmetry was motivated less by whether the central bank was in overall tightening or loosening mood, but whether the public perceived the central bank to be more or less hawkish than current economic conditions would indicate.

Adopting an alternative identification method, based on the Romer and Romer (2004) monetary policy instrumental variable used in our analysis, Tenreyro and Thwaites (2016) reach somewhat similar conclusions, but this time by estimating state-dependent responses according to the state of the business cycle with a stratification based on recessions and expansions; they find that monetary policy has a big effect in expansions, but a smaller effect in recessions. They also find contractionary shocks may be more powerful than expansionary shocks, but argue that this is not the entire story since (perhaps surprisingly) in their data contractionary shocks are no more likely in expansions than in recessions. Along similar lines, in a very recent paper, Paul (2015) finds evidence for state

dependence in monetary policy when dividing the cycle up into booms and slump phases, in this case employing the "high frequency" identification method of Gertler and Karadi (2015).[5]

We explore these ideas in the context of a state-dependent analysis that echoes the analysis in Auerbach and Gorodnichecko (2013b) and Jordà and Taylor (2016) for fiscal policy, namely, that stabilization policy can have different effects in the boom (when the output gap is large) versus the slump (when the output gap is small). We feel that some of the questions raised in the earlier literature can be better structured using a similar stratification for interest rates.

An alternative stratification based on the state of the economy links with current debates about the effectiveness of monetary policy in an environment of unusually low inflation. Despite their belief in the "divine coincidence" (Blanchard and Galí 2005), the returns to ultra-loose monetary policy are debatable in this setting. Thus, as a follow up to our stratification based on the output gap, we investigate state dependence based on low inflation environments (defined by a 2% CPI inflation cutoff) relative to all others.

8.1. State dependence in boom and bust episodes

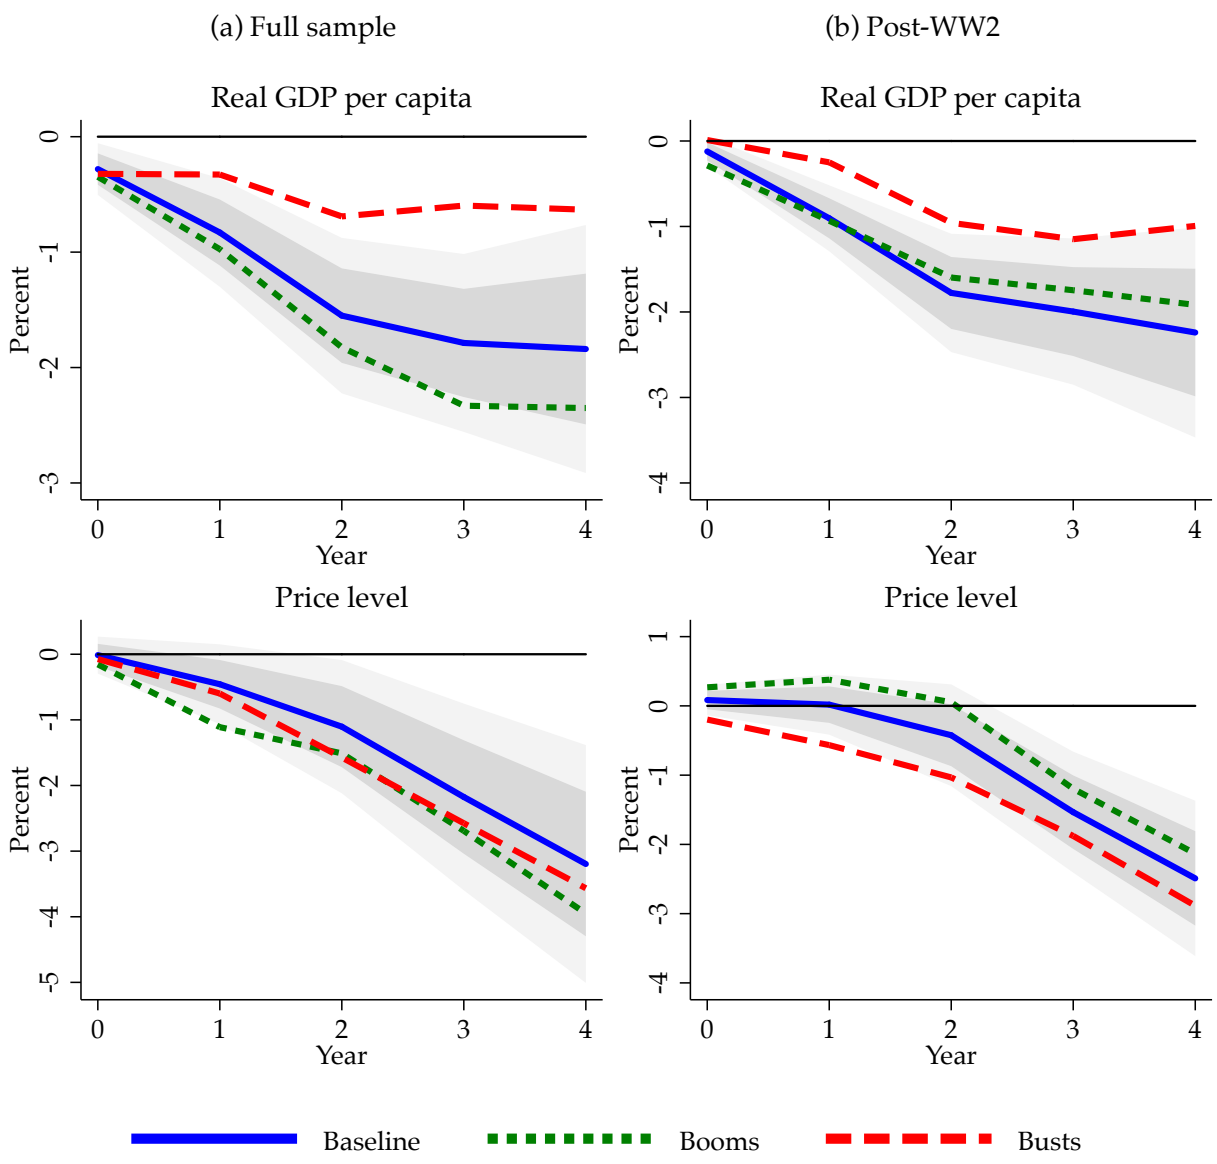
The first set of policy experiments shown in Figure 5 investigates the hypothesis that interest rates may have different effects on outcomes in booms versus busts. To examine this nonlinearity we introduce an indicator state variable equal to 1 in booms and 0 in busts. Booms (busts) are periods when the actual level of log real GDP per capita is above (respectively, below) the level of its long-run country-specific trend component as measured by an HP filtered series with a very low-pass setting (λ=100, with annual data).

These results are strongly indicative of an asymmetric macroeconomic response to interest rates. The experiment is always normalized to be a +100 bps increase in the short-term rate. This is done to facilitate impulse-response comparisons across states even though, of course, tight policy in a bust is unlikely. The response of real GDP to monetary policy appears to be quite strong in booms (about -2.5% by year 4 in the full sample), but considerably weaker in busts (about -0.5% by year 4 in the full sample).

This difference of nearly 2 percentage points (but closer to only 1% Post-WW2) is broadly consistent with Angrist, Jordà, and Kuersteiner (forthcoming), Tenreyro and Thwaites (2016), and Barnichon and Matthes (2016). Evidence of asymmetry is less clear for inflation. Differences using the full sample are very minor and not very different from the average response. The responses calculated for the Post-WW2 sample suggest that if the monetary authority tried to stimulate the economy (opposite to what is displayed but consistent with the outlook in many economies today), it would take 2 years before the effects would be felt in prices, although after 4 years the effects would be similar in size to those estimated for the boom state.

[5]And even before this new wave of studies, which exploit the latest LP methods and identification techniques, there was an older tradition of studies which explored asymmetric monetary policy responses (see, inter alia, Cover 1992; Weise 1999; Ravn and Sola 2004; Lo and Piger 2005).

Figure 5: *State dependence: monetary policy has a weaker effect on output in busts*



Notes: Full sample: 1870–2013 excluding world wars (1914–1919 and 1939–1947). Post-WW2 sample: 1948–2013. Linear LP-IV estimates displayed with a solid blue line and 1 S.D. and 90% confidence bands. Estimates stratified by the boom displayed with a green dotted line whereas estimates in the bust are displayed with a red dashed line. See text.

To sum up, our output-state asymmetry finding says that a central bank rate tightening of +100 bps in a boom would have causal effects on output that are strongly contractionary on output going forward. In contrast, a central bank rate loosening of −100 bps in a bust would have causal effects on output that are, proportionately, weakly expansionary going forward.

8.2. State dependence in "lowflation" episodes

Our final set of experiments, in Figure 6, concerns the hypothesis that monetary policy may have different effects in times of low inflation, a topic of rising interest since the advanced economies entered an era of "lowflation" as the Great Recession wore on after 2008. To examine this nonlinearity we set an indicator state variable equal to 1 when inflation is low (at or below 2%) and 0 otherwise.

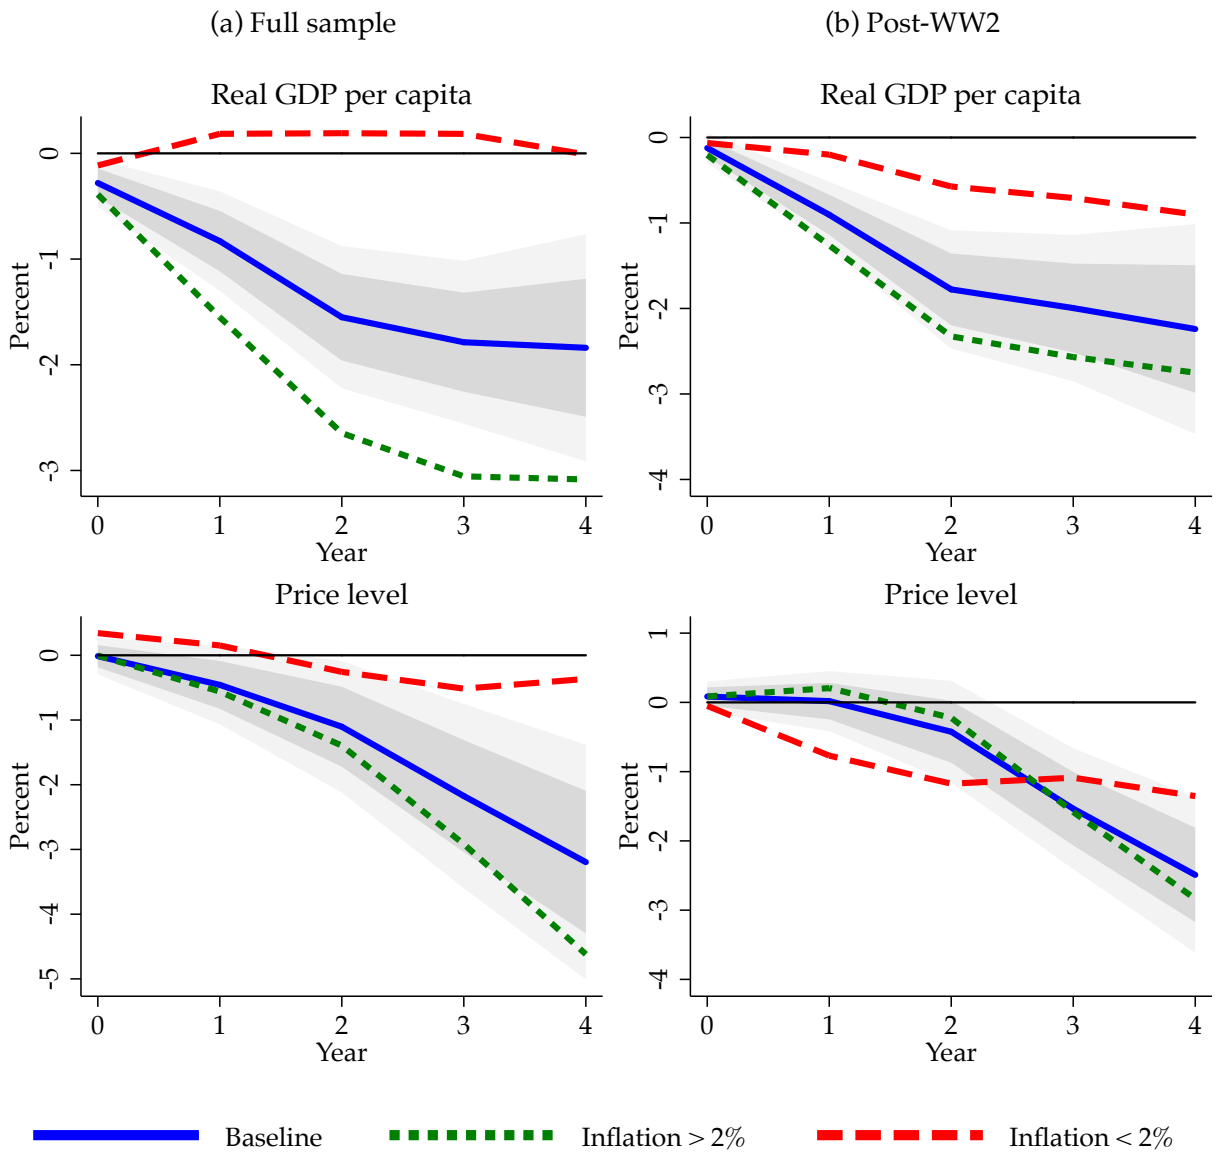
These results reinforce some findings in the previous section. Just as in Figure 5, we normalize the responses to a +100 bps increase in interest rates, while fully recognizing rates are unlikely to go up when there is lowflation. However, the normalization greatly facilitates the comparison across states. The response of real GDP to monetary policy appears to be quite strong when inflation is above 2%. In that scenario, the cumulated response at year 4 is slightly above −3%, somewhat stronger than the response in booms reported in figure 5, which was −2.5%. However, monetary policy loses most of its traction when the inflation rate dips below the 2% threshold and economies tip into a lowflation state. In the full sample results, we find that there is no effect on output although in the Post-WW2 sample, the effect on output is somewhat more visible (around −1%).

To sum up, our inflation-state asymmetry finding says that a central bank rate loosening of −100 bps in times when in inflation is above 2% would have causal effects on output that are strongly expansionary on output going forward. In contrast, a central bank rate loosening of −100 bps in times when in inflation is below 2% would have causal effects on output that are, proportionately, weakly expansionary going forward.

8.3. Asymmetry and monetary policy in a depressed-lowflation trap

The two sets of results we have presented in this section reinforce each other. The concerns of many economists and policymakers reflect the idea that policy may be less effective in some states of the world, such as now, and in a manner that is consistent with several recently emerging theories, and this appears to be well supported by our findings. Our analysis suggests that we may face challenges to policy implementation if we find ourselves in states of the world where inflation is very low. These can also often be times of slow growth or stagnation. In these conditions, the ability to stimulate the economy out of its torpor is that much more difficult; neutral real rates are low, and nominal interest rates are of course likely to be dragged ever closer to the zero lower bound (Williams 2016). Taken together, these conditions—less bang for a given policy move, and less room to make a move—could make it considerably more difficult for monetary policy to effect any change in its principal target variables, output and inflation.

Figure 6: *State dependence: monetary policy has weaker effects when there is lowflation*



Notes: Full sample: 1870–2013 excluding world wars (1914–1919 and 1939–1947). Post-WW2 sample: 1948–2013. Linear LP-IV estimates displayed with a solid blue line and 1 S.D. and 90% confidence bands. Estimates stratified by the lowflation regime displayed with a red dashed line whereas estimates when inflation is above 2% are displayed with a green dotted line. See text.

9. Conclusion

The effectiveness of monetary stabilization policy is not only a major policy issue but also an important matter of some controversy among both theoretical and empirical macroeconomists. We calculate that interest rates have a bigger effect on macroeconomic outcomes than previously measured. The source of the attenuation bias that we report suggests that common identification assumptions prevalent in the literature (based on regression control) are insufficient. Using a quasi-natural experiment in international finance and novel empirical methods, we show why this bias occurs and how to resolve it.

Monetary stabilization policy turns out to be state-dependent, a critical observation in the context of the Great Recession. Policymakers have faced a situation of consistent undershoot relative to their stated objectives and forecasts. A slower than expected growth trajectory has been seen in the U.S. economy, and worse yet in the U.K., Europe, and Japan, and persistent sub-2% inflation afflicts all of these economies. Amid worries of deflation risk and secular stagnation, the extensive and unconventional application of monetary policy tools have failed to shunt the macroeconomic locomotive onto a faster track. Debate centers on why the central banks got so derailed, and the subsequent lack of policy traction. This may be nothing new, and we find economic responses to be smaller when the economy is running below potential. Moreover, monetary stabilization seems to be especially weak in lowflation environments, such as those currently experienced in a number of western economies. These results therefore have profound implications for how today's monetary models are formulated and applied.

References

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. Identification of Causal Effects Using Instrumental Variables. *Journal of the American statistical Association* 91(434): 444–55.

Angrist, Joshua D., Òscar Jordà, and Guido Kuersteiner. Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited. *Journal of Business and Economic Statistics.* Forthcoming.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2013a. Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy* 4(2): 1–27.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2013b. Fiscal Multipliers in Recession and Expansion. In *Fiscal Policy After the Financial Crisis* edited by Alberto Alesina and Francesco Giavazzi. Chicago: University of Chicago Press, pp. 63–98.

Barnichon, Regis, and Christian Matthes. 2016. Gaussian Mixture Approximations of Impulse Responses and The Non-Linear Effects of Monetary Shocks. CEPR Discussion Paper 11374.

Barro, Robert J., and José F. Ursúa. 2008. Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity* 39(1): 255–335.

Blanchard, Olivier, and Jordi Galí. Real Wage Rigidities and the New Keynesian Model. 2007. *Journal of Money, Credit and Banking* 39(S1): 35–65.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. Monetary Policy Shocks: What Have We Learned And To What End? *Handbook of Macroeconomics*, vol. 1, edited by John B.

Taylor and Michael Woodford. Amsterdam: Elsevier, pp. 65–148.
Cloyne, James, and Patrick Hürtgen. The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom. Bank of England Working Paper 493.
John H. Cochrane. 2016. Do Higher Interest Rates Raise or Lower Inflation? University of Chicago. Unpublished.
Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. Plausibly Exogenous. *Review of Economics and Statistics* 94(1): 260–272.
Cover, James Peery. 1992. Asymmetric Effects of Positive and Negative Money-Supply Shocks. *Quarterly Journal of Economics* 107(4): 1261–82.
Cox, David R. 1958. *Planning of Experiments*. New York: John Wiley.
di Giovanni, Julian, Justin McCrary, and Till von Wachter. 2009. Following Germany's Lead: Using International Monetary Linkages to Estimate the Effect of Monetary Policy on the Economy. *Review of Economics and Statistics* 91(2): 315–331.
di Giovanni, Julian, and Shambaugh, Jay C. 2008. The Impact of Foreign Interest Rates on the Economy: The Role of the Exchange Rate Regime. *Journal of International Economics* 74(2): 341–61.
Eichenbaum, Martin. 1992. Comments 'Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy' by Christopher Sims. *European Economic Review* 36(5): 1001–11.
Faust, Jon, Eric T. Swanson, and Jonathan H. Wright. 2004. Identifying VARs Based on High Frequency Futures Data. *Journal of Monetary Economics* 51(6): 1107–31.
Gertler, Mark, and Peter Karadi. 2015. Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics* 7(1): 44–76.
Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 2005. The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models. *American Economic Review* 95(1): 425–36.
Hamilton, James, and Òscar Jordà. 2002. A Model for the Federal Funds Rate Target. *Journal of Political Economy* 5(110): 1135–67.
Hanson, Michael S. 2004. The "Price Puzzle" Reconsidered. *Journal of Monetary Economics* 51(7): 1385–1413.
Imbens, Guido W. 2014. Instrumental Variables: An Econometrician's Perspective. NBER Working Paper 19983.
Imbens, Guido W., and Joshua Angrist. 1994. Identification and Estimation of Local Average Treatment Effects. *Econometrica* 61(2): 467–76.
Ilzetzki, Ethan, Enrique G. Mendoza, and Carlos A. Végh. 2013. How Big (Small?) are Fiscal Multipliers? *Journal of Monetary Economics* 60(2): 239–54.
Jordà, Òscar. 2005. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95(1): 161–82.
Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2015. Betting the House. *Journal of International Economics* 96(S1): S2–S18.
Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. *NBER Macroeconomics Annual 2016*, no. 31, edited by Martin Eichenbaum and Jonathan A. Parker. Chicago: University of Chicago Press. Forthcoming.
Jordà, Òscar, and Alan M. Taylor. 2016. The Time for Austerity: Estimating the Average Treatment Effect of Fiscal Policy. *Economic Journal* 126(590): 219–55.
Klein, Michael W., and Jay C. Shambaugh. 2013. Rounding the Corners of the Policy Trilemma: Sources of Monetary Policy Autonomy. NBER Working Paper 19461.
Knoll, Katharina, Moritz Schularick, and Thomas Steger. No Price Like Home: Global House Prices, 1870–2012. *American Economic Review*. Forthcoming.

Kuttner, Kenneth N. 2001. Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market. *Journal of Monetary Economics* 47(3): 523–44.

Lo, Ming Chien, and Jeremy Piger. 2005. Is the Response of Output to Monetary Policy Asymmetric? Evidence from a Regime-Switching Coefficients Model. *Journal of Money, Credit and Banking* 37(5): 865–86.

Nakamura, Emi, and Jón Steinsson. 2013. High Frequency Identification of Monetary Non-Neutrality. NBER Working Paper 19260.

Obstfeld, Maurice, and Kenneth Rogoff. 1995. The Mirage of Fixed Exchange Rates. *Journal of Economic Perspectives* 9(4): 73–96.

Obstfeld, Maurice, Jay C. Shambaugh, and Alan M. Taylor. 2004 Monetary Sovereignty, Exchange Rates, and Capital Controls: The Trilemma in the Interwar Period. *IMF Staff Papers* 51(S): 75–108.

Obstfeld, Maurice, Jay C. Shambaugh, and Alan M. Taylor. 2005. The Trilemma in History: Tradeoffs among Exchange Rates, Monetary Policies, and Capital Mobility. *Review of Economics and Statistics* 87(3): 423–38.

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. *American Economic Review* 103(3): 129–34.

Paul, Pascal. 2015. The Time Varying Transmission of Monetary Policy Surprises. University of Oxford. Unpublished.

Quinn, Dennis P., Martin Schindler, and A. Maria Toyoda. 2011. Assessing Measures of Financial Openness and Integration. *IMF Economic Review* 59(3): 488–522.

Ramey, Valerie A., and Sarah Zubairy. 2014. Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. NBER Working Paper 20719.

Ravn, Morten O., and Martin Sola. 2004. Asymmetric Effects of Monetary Policy in the United States. *Federal Reserve Bank of St. Louis Review* 86(5): 41–60.

Romer, Christina D. and David H. Romer. 2004. A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review* 94(4): 1055–84.

Rubin, Donald B. 1974. Estimating Causal Effects of Treatments in Randomized and Nonradomized Studies. *Journal of Educational Psychology* 66(5): 688–701.

Rubin, Donald B. 1978. Bayesian Inference for Causal Effects: The Role of Randomization. *Annals of Statistics* 6(1): 34–58.

Rubin, Donald B. 2005. Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. *Journal of the American Statistical Association* 100(469): 322–331.

Shambaugh, Jay C. 2004. The Effect of Fixed Exchange Rates on Monetary Policy. *Quarterly Journal of Economics* 119(1): 301–352.

Sims, Christopher A. 1992. Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy. *European Economic Review* 36(5): 975–1000.

Swanson, Eric T., and John C. Williams. 2014. Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates. *American Economic Review* 104(10): 3154–85.

Tenreyro, Silvana, and Gregory Thwaites. 2016. Pushing on a String: US Monetary Policy is Less Powerful in Recessions. *American Economic Journal: Macroeconomics*. 8(4): 43–74.

Weise, Charles L. 1999. The Asymmetric Effects of Monetary Policy: A Nonlinear Vector Autoregression Approach. *Journal of Money, Credit and Banking* 31(1): 85–108.

Williams, John C. 2016. Monetary Policy in a Low R-star World. FRBSF Economic Letter 2016-23.

Williamson, Stephen D. 2016. Neo-Fisherism: A Radical Idea, or the Most Obvious Solution to the Low-Inflation Problem? *The Regional Economist*. Federal Reserve Bank of St. Louis, July.

[Appendix: To be completed.]