RESUMO EXPANDIDO (PÚBLICO GERAL/NÃO ASSOCIADOS AKB) / EXTENDED ABSTRACTS (NON AKB MEMBERS) - ÁREA 3. MÉTODO, TEORIAS ECONÔMICAS HETERODOXAS E HISTÓRIA ECONÔMICA

A NOVEL APPROACH TO ESTIMATING THE DISTRIBUTIONAL EFFECTS OF MONETARY POLICY IN PANEL DATA: A DISTRIBUTED-LAG TWO-WAY FIXED EFFECTS (TWFE) MODEL WITH MULTIPLE ROBUSTNESS CHECK

Rafael Pahim (mi.oj@hotmail.com)

Sylvio Antonio Kappes (sylvio.kappes@feac.ufal.br)

The empirical and theoretical literature on the distributional effects of monetary policy has expanded substantially over the past two decades. Despite this growth, no consensus has emerged regarding the direction or magnitude of these effects. Existing studies diverge along several dimensions, including the choice of estimator, the identification strategy for monetary policy shocks, and the temporal and geographic coverage of the data. Against this backdrop, the present paper proposes a novel empirical approach to estimating the distributional effects of monetary policy, contributing an additional set of estimates that may help reconcile the mixed findings reported in prior work or, alternatively, lend further support to the non-neutrality view emphasized by Kappes (2023).

The research problem motivating this study stems from the observation that, although the difference-in-differences (DiD) design has become one of the most widely adopted tools for causal inference in economics, its application to the

analysis of monetary policy — and, in particular, to its distributional effects — has been largely overlooked. The primary reason for this omission lies in a fundamental limitation of the standard DiD framework: treatment variables are typically modeled as binary indicators, whereas monetary policy shocks are inherently continuous, varying in both magnitude and direction. This mismatch has favored the adoption of alternative specifications, such as vector autoregressions (VARs) and local projections (LPs), which more naturally accommodate continuous treatments. The present paper seeks to bridge this methodological gap by developing an identification and estimation strategy compatible with the DiD design while preserving the defining features of monetary policy shocks.

The paper pursues three main objectives, developed in a pedagogically oriented manner throughout the text. First, it introduces a novel monetary policy shock identification methodology and integrates it into a distributed-lag two-way fixed effects (TWFE) model. Second, it discusses a set of robustness checks applicable to the proposed specification, drawing on alternative approaches well established in the literature. Third, it conducts a Monte Carlo simulation to evaluate the internal consistency of the proposed estimator, providing evidence on its reliability under controlled conditions.

The methodology builds upon a conceptual discussion of macroeconomic shocks and their applicability to the monetary policy context, drawing on the identification criteria set out by Ramey (2016). According to these criteria, a well-defined shock must be exogenous (orthogonal) with respect to all current and lagged endogenous variables in the system, uncorrelated with other exogenous shocks, and must represent either an unanticipated movement in an exogenous variable or, within an intertemporal framework, new information regarding future movements in exogenous variables. Building on Nunes et al. (2025), the proposed identification strategy extends their framework along two key dimensions: (i) aligning the shock measure more closely with the criteria established by Ramey (2016), and (ii) incorporating features central to the monetary policy literature, most notably the persistence of policy actions.

The construction of the shock proceeds in successive stages. First, deviations of the central bank policy rate from its sample mean are computed via Z-score normalization. Recognizing that distinct monetary policy regimes prevailed in different subperiods — including the contractionary regimes of the 1970s and early 1980s associated with the oil shocks and the Volcker disinflation, the Great Moderation, and the post-Global Financial Crisis period — the paper implements a regime-adjusted normalization, identifying regime breaks via the multiple structural breaks methodology of Bai and Perron (1998). The adjusted Z-score is then purged of its predictable components by regressing it on unemployment, core inflation, and its own lag, within a random-effects panel model with time effects. The residuals from this regression are interpreted as monetary policy shocks — that is, as intra-regime deviations orthogonal to the systematic component of the central banks' reaction function. These shocks are subsequently classified into discrete categories according to their direction (contractionary or expansionary) and magnitude (low, medium, large, and very large), rendering them naturally compatible with the binary structure required by the DiD framework.

Regarding the estimator, the paper develops, in a progressive and pedagogical manner, the derivation of the distributed-lag TWFE model starting from the canonical 2 × 2 DiD design. The exposition clarifies the role of the parallel trends assumption, the two-way within transformation (double-demeaning), and the interpretation of the TWFE estimator as a weighted average of implicit 2 × 2 DiD comparisons across units and time periods in staggered designs. The inclusion of lagged shock indicators accommodates the persistence typical of monetary policy effects, which the literature documents as concentrated within the first two years following implementation, though potentially extending to longer horizons according to recent evidence. The baseline specification includes two lags of each shock indicator; alternative specifications with three, four, and ten lags are employed as robustness checks to assess sensitivity to medium- and long-horizon dynamics. From the estimated coefficients, cumulative impulse-response functions are constructed for each shock category. The empirical illustration relies on a panel of 16 advanced economies covering the period 1950–2019.

As additional robustness checks, the paper implements two widely established econometric approaches: a VAR model identified through Cholesky decomposition and a local projection (LP) à la Jordà (2005). For the LP specification, a novel adaptation of the forecast-error identification strategy proposed by Auerbach and Gorodnichenko (2013) and Furceri et al. (2018) is introduced, incorporating an additional step that removes the autoregressive component of the shock series — bringing the measure into closer compliance with the orthogonality requirement of Ramey (2016). The identification is operationalized using the OECD Economic Outlook database (1986–2019), which enables long-horizon panel estimation while avoiding the data access constraints associated with proprietary sources.

Two sets of results merit particular emphasis. The first concerns the properties of the proposed shock measure. The descriptive statistics indicate that the identified shocks are relatively infrequent and broadly symmetric across positive and negative realizations, with most observations concentrated in the low- and medium-intensity categories. Large and very large shocks account for less than 4 percent and less than 2 percent of the observations, respectively. This distributional pattern is consistent with an interpretation of the shocks as transitory and unexpected movements relative to the systematic conduct of monetary policy, rather than as persistent shifts in policy stance or regime-level changes. Furthermore, a comparison between the proposed measure and the benchmark series of Romer and Romer (2004) reveals a moderate correlation and broadly similar temporal patterns, suggesting external coherence of the identification strategy.

The second set of results pertains to the Monte Carlo simulation, conducted with 1,000 independent data-generating processes (DGPs) calibrated using empirical sample moments and allowing for alternative error structures (cross-sectional correlation, serial correlation, and the combination of both). Simulations were performed for specifications with two, four, and ten lags. The results, summarized graphically, show that the average estimated coefficients closely track the true parameter values embedded in the DGP, with the latter consistently falling within the 95 percent confidence intervals constructed from the distribution of the estimates. This pattern holds across all lag and error structures, as well as across all shock categories, encompassing both high-

frequency (small) and low-frequency (large) shocks. The absence of systematic deviations between estimated and true coefficients indicates that, under correct specification, the distributed-lag TWFE estimator recovers the true parameters of interest on average, exhibiting unbiasedness in finite samples and appropriate uncertainty quantification.

In sum, the paper makes two novel methodological contributions to shock identification, develops a difference-in-differences framework tailored to the requirements of monetary policy estimation in panel data, presents two robustness checks well established in the literature, and demonstrates, through Monte Carlo simulation, the internal consistency of the proposed estimator. The natural next step is to apply the proposed framework to empirical inequality data, thereby providing the literature with an additional set of estimates that may help characterize the response of distributional outcomes to monetary policy shocks

Palavras-chave: monetary policy; difference-in-difference; shock indentification; empirical.