



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/234808/>

Version: Published Version

Article:

Polito, V. (2025) Optimal macroeconomic policy in nonlinear models: a VSTAR perspective. *Studies in Nonlinear Dynamics and Econometrics*. ISSN: 1081-1826

<https://doi.org/10.1515/snde-2024-0100>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Vito Polito*

Optimal Macroeconomic Policy in Nonlinear Models: A VSTAR Perspective

<https://doi.org/10.1515/sn-de-2024-0100>

Received September 16, 2024; accepted September 17, 2025; published online October 9, 2025

Abstract: The paper describes a method to transform vector smooth transition autoregressions in a form that is particularly suitable for policy analysis, because it is of low-dimension and retains certainty-equivalence. Optimal rules are calculated with the state dependent coefficients approach, which allows linear methods to solve nonlinear problems. The methodology is applied to revisit interest rate and quantitative easing (QE) monetary policy in the United States during 1979–2018. The actual size of QE and duration of the zero lower bound are found close to those prescribed by the optimal policy, but not the timing and composition of QE. The methodology compares favourably against alternative optimization approaches based on linearization or numerical search.

Keywords: smooth transition models; nonlinear quadratic regulator; zero lower bound; quantitative easing; optimal monetary policy

JEL Classification: C32; C61; E52

1 Introduction

This paper describes a methodology for the analysis of optimal macroeconomic policy when the interaction between the economy and the policy sector is described by a widely-used class of nonlinear models: vector smooth transition autoregressions (VSTARs). VSTAR models have become increasingly popular in macroeconomics given their ability to better describe data compared to linear models. In applied works VSTARs are often used for forecasting and to account for history dependence in the response of macroeconomic variables to shocks. While the econometric literature has long established methods for specification, estimation and structural analysis with VSTAR, no attention has been devoted so far to the analysis of policy. The present paper attempts to fill this methodological gap.

The proposed methodology is applied to evaluate the actual conduct and effects of conventional (interest rate) monetary policy and large scale asset purchase programs, frequently referred to as quantitative easing (QE), undertaken by the Fed since 2008 against the framework obtained from the optimal policy. This seems a pertinent application of the VSTAR, given the extent of nonlinearity displayed by macroeconomic data of the United States over the last 40 years and, particularly, by the federal funds rate and the Fed's balance sheet in the decade after the onset of the Great Recession. According to this analysis, the size of QE undertaken by the

I thank Juan Paez-Farrell, Alessia Paccagnini, Michele Piffer, Peter Spencer, Mike Wickens, Timo Teräsvirta, Yunyi Zhang for numerous conversations on the topic, comments and suggestions. I also thank for comments participants at the 2019 International Workshop on Financial Markets and Nonlinear Dynamics in Paris, and at seminars held at the Universities of Brunel, Cardiff, Sheffield, Newcastle and York.

*Corresponding author: Vito Polito, School of Economics, University of Sheffield, S10 2TN Sheffield, UK; CESifo, Munich, Germany; and NETSPAR, Tilburg, The Netherlands, E-mail: v.polito@sheffield.ac.uk

Fed and the duration of the zero lower bound (ZLB) are not too dissimilar from those prescribed by the optimal policy, but there are differences in terms of the timing and composition of QE.

Given a welfare function and a structural model of the economy, monetary policy could be evaluated considering commitment to a rule derived by maximising welfare subject to the model. Quantitative policy evaluation based on structural models is however not problem free. Above all it is often difficult to get agreement on how to specify a structural model of the economy. This is true even for well-established structural models such as the Smets and Wouters (2003) model, as argued by Chari et al. (2009). Challenges compound when dealing with QE because structural models tend to abstract from large changes in the economy, assuming parameters and sources of shocks to be constant over time and known by policy makers. These issues are particularly noteworthy for quantitative analyses dealing with the significant instabilities and jumps in the data of the last 25 years, as well as the asymmetric nature of the ZLB. An alternative solution is to quantify policy based on a nonlinear vector autoregression model of the economy. Its main attraction is that, whatever the structural model, a close approximation is given by a vector autoregression representation. Although the resulting policy would not be first best compared with using a correct structural model, it would be a simple and transparent way to proceed. It would also provide a useful benchmark against which to compare any other policy decisions, and it has the advantage of being easily automated.

Against this backdrop, the paper makes two methodological contributions. To better highlight these, it is useful to recap how the effects of systematic changes in policy are evaluated when using a linear VAR. In a VAR all (policy and non-policy) variables are endogenous. To quantify the impact of a systematic change in policy, the VAR needs to be transformed into a VARX in which the policy variables are exogenous. This is achieved through a block decomposition separating the vector of policy instruments from the remaining variables in the VAR. Depending on whether the policy variables are ordered first or last, at least two possible VARX representations can be derived, see for example Sack (2000) and Stock and Watson (2001). These VARX representations differ in terms of the timing assumption about the economy response to a policy change, which can be either with a lag (A1) or contemporaneous (A2). If the objective function of the decision maker is quadratic, optimal policy can then be calculated as a linear quadratic regulator (LQR) problem using either of these two VARX models. Crucially, both A1 and A2 satisfy certainty equivalence, solving the LQR problem regardless of the particular sequence of disturbances. For this reason solutions can be computed using linear dynamic programming based on Riccati equation iteration, as shown in Ljungqvist and Sargent (2018).

The first methodological contribution refers to the derivation of a VSTARX model suitable for optimal policy analysis from a reduced-form VSTAR. It is shown that when using either A1 or A2 in a VSTAR, differences between the two VSTARX representations go well beyond the timing of interaction among variables. A1 gives rise to a nonlinear quadratic regulator (NLQR) problem that is high dimensional and no longer consistent with certainty equivalence. In contrast, A2 has the advantage of delivering a small dimensional NLQR problem compatible with certainty equivalence.

The second methodological contribution refers to the calculation of optimal decision rules that solve the NLQR problem. This consists of employing state dependent coefficient (SDC) factorization to transform the nonlinear VSTARX into an affine model with SDC matrices. As a result, the NLQR problem becomes isomorphic to the LQR. Under A2, the problem is also low dimensional and certainty equivalent, and it can be solved with linear dynamic programming methods. The solution yields an optimal feedback rule with time-varying coefficients that can be combined with the VSTARX to construct the reduced-form VSTAR under the optimal rule.

Of course, any nonlinear model could be linearized and then analyzed with LQR. However, as well as providing solutions that are not globally valid, linearization assumes away crucial features of macroeconomic data, such as asymmetries, threshold effects and large transitional changes. A VSTAR could well capture these features, but the LQR solution would disregard them, potentially resulting in incorrect and/or inefficient decision making. Alternatively, policy analysis could be carried out with numerical methods. While these can deliver globally valid solutions without disregarding nonlinearity, their main drawback is that they are computationally intensive and their effectiveness in finding the optimal solution is often impeded when nonlinearity is significant. The proposed SDC method has the advantage of accounting for nonlinearity while being computationally simple and

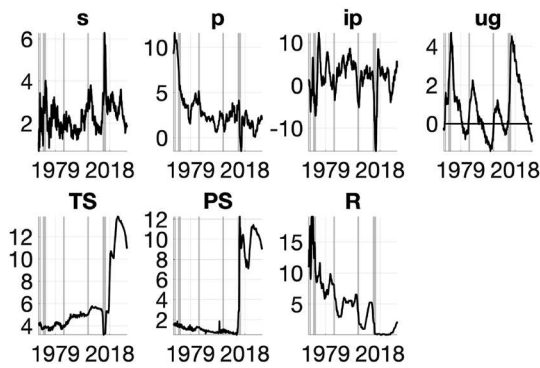


Figure 1: Macroeconomic data, United States 1979:8-2018:10. NBER recessions in grey. All variables are in percentage.

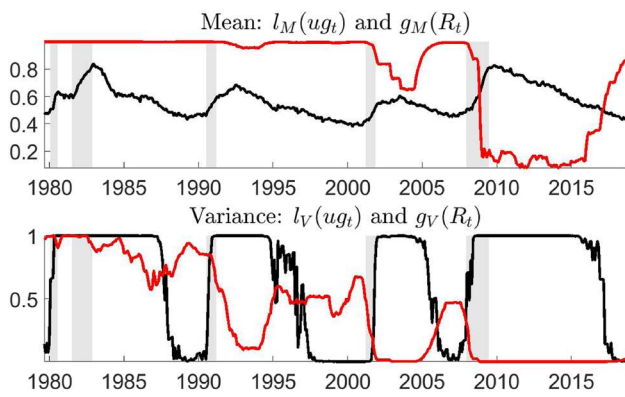


Figure 2: Transition probabilities $l_i(ug_t)$ (black) and $g_i(R_t)$ (red), in the mean $i = M$ and covariance matrix $i = V$ of the estimated VSTAR. NBER recessions in grey.

easy to understand given his similarity to the LQR. It also allows the use of dynamic programming, one of the most powerful tools for solving decision problems in economics.

The proposed methodology is applied using a VSTAR estimated with maximum likelihood on monthly data for the United States during 1979:8-2018:10 (Figure 1). In the benchmark specification, the model includes seven variables: credit spread, inflation, industrial production, unemployment rate, the federal funds rate, and total assets held by the Fed split between Treasury and privately-held securities.¹ The policy analysis therefore accounts for the macroeconomic effects of changes in both the size and composition of the Fed's balance sheet as well as the ZLB. Nonlinearity in the mean and variance of the VSTAR is captured by allowing the coefficients to change during periods of economic slack and when the federal funds rate is near to the ZLB. It is shown that the estimated VSTAR can well capture nonlinearities in the data (Figure 2) and provides a plausible description of the response of the three policy instruments to shocks at different times of the United States's monetary history (Figure 3).

The estimated VSTAR is used to evaluate how far from optimal were the actual interest rate policy and QE undertaken by the Fed in the decade after the onset of the Great Recession. By focusing exclusively on the post-2008 period, the policy analysis avoids the substantial macroeconomic shifts that occurred between the pre- and post-Great Recession periods. The optimal VSTAR is obtained from the joint optimization of all monetary policy instruments under alternative specifications of the Fed's preferences. The evaluation considers three different dimensions: the gains from optimization; the response of monetary policy to shocks; and counterfactual

¹ This split is chosen to conform with the portfolio-balancing and the credit channels highlighted by the literature on the transmission mechanisms of QE, see for example Joyce et al. (2012). Following Kuttner (2018), throughout the paper asset purchases are referred to as QE regardless of whether these entail public or private securities.

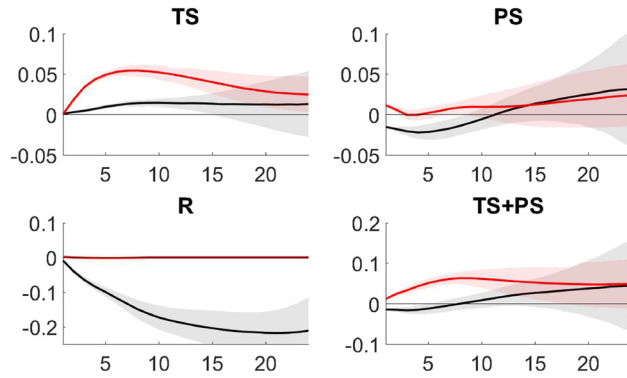


Figure 3: Median and two standard deviation confidence bands of IRFs of policy instruments to a credit shock during normal times (black) and the ZLB period (red). All variables are in percentage.

Table 1: Gains from monetary policy optimization computed from the VSTAR.

	Items in the loss function					Loss	Stabilization	
	$(p_t - \bar{p})^2$	$u g_t^2$	ΔTS_t^2	ΔPS_t^2	ΔR_t^2	V	G	\hat{u}
Policy	Baseline: $v_p = v_{ug} = v_{\Delta TS} = v_{\Delta PS} = v_{\Delta R} = 1$							
Actual	1.24	5.46	0.02	0.05	0.38	7.15		
Optimal	1.05	5.33	0.05	0.25	0.39	7.08	0.98	0.26 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	8.05	11.15	0.95 ^b
	Weights II: $v_p = 0.5, v_{ug} = v_{\Delta TS} = v_{\Delta PS} = v_{\Delta R} = 1$							
Actual	1.24	5.46	0.02	0.05	0.38	6.53		
Optimal	1.14	5.29	0.06	0.19	0.37	6.47	0.84	0.23 ^a
NoQE	2.47	5.55	0.02	0.00	0.00	6.81	4.11	0.53 ^b
	Weights III: $v_{ug} = 0.5, v_p = v_{\Delta TS} = v_{\Delta PS} = v_{\Delta R} = 1$							
Actual	1.24	5.46	0.02	0.05	0.38	4.42		
Optimal	1.03	5.40	0.05	0.22	0.38	4.39	0.67	0.24 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	5.27	16.18	1.31 ^b
	Weights IV: $v_{ug} = v_p = 1, v_{\Delta TS} = v_{\Delta PS} = v_{\Delta R} = 0.5$							
Actual	1.24	5.46	0.02	0.05	0.38	6.92		
Optimal	0.97	5.28	0.06	0.38	0.46	6.70	3.16	0.47 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	8.03	13.82	1.05 ^b
	Weights V: $v_{ug} = v_p = v_{\Delta R} = 1, v_{\Delta TS} = v_{\Delta PS} = 0.5$							
Actual	1.24	5.46	0.02	0.05	0.38	6.92		
Optimal	0.97	5.28	0.06	0.38	0.46	6.70	3.16	0.47 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	8.03	13.82	1.05 ^b
	Weights VI: $v_{ug} = v_p = v_{\Delta TS} = v_{\Delta PS} = 1, v_{\Delta R} = 0.5$							
Actual	1.24	5.46	0.02	0.05	0.38	7.14		
Optimal	1.04	5.32	0.07	0.25	0.39	7.03	1.44	0.32 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	8.03	11.16	0.95 ^b
	Average (across weights)							
Actual	1.24	5.46	0.02	0.05	0.38	6.51		
Optimal	1.03	5.32	0.06	0.28	0.41	6.40	1.71	0.33 ^a
No-QE	2.47	5.55	0.02	0.00	0.00	7.37	11.71	0.97 ^b

G is the stabilization gain; \hat{u} the unemployment compensation; ^a indicates comparison between actual and optimal policy; ^b between actual and no-QE policy. *Source:* Author's calculations. See main text for more details.

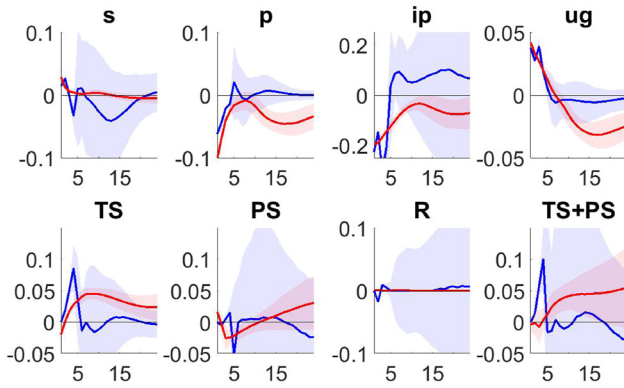


Figure 4: Median and two standard deviation confidence bands of IRFs to a credit shock during the ZLB period under the actual (red) and the optimal (blue) policy. All variables are in percentage.

simulations. According to these results, the actual QE and ZLB policy were not too far from those otherwise prescribed by the optimal analysis. In particular, the joint optimization of the monetary policy instruments since 2008 would have resulted in: (i) Reduction of the average unemployment rate by 0.33%. This is about one-third of that measured for the actual QE policy relative to a scenario of no-QE (0.97%) (Table 1); (ii) Responses of QE to credit shocks very similar to those obtained under the actual policy, particularly beyond the very short-run horizon (three months) (Figures 4 and 5); (iii) Increase in size of the Fed’s balance sheet as a proportion to GDP (between 17.98 and 19.33%) in line with that observed from the data (17.96). There are however differences in the timing and composition of QE, as according to the optimal policy larger purchases of Treasury securities should have occurred during the first phase of QE (Figure 6 and Table 3); (iv) Average ZLB duration (between 68 and 73 months) close to that observed from the data (96 months), in contrast with the view that the ZLB should have terminated much earlier (after 26 months) supported by the standard Taylor rule, see Williamson (2015) (Table 2).

To gauge some indication on the quality of the quantitative performance of the proposed methodology, the results are compared with those from three alternative optimization approaches. The first is based on linearization. The other two, model predictive control and direct search, are nonlinear optimizers that rely upon numerical search using the VSTARX from A1. The proposed methodology is found to deliver faster and larger optimization gains compared to any of these alternatives (Table 5).

Two possible caveats on the quantitative results should be noted. The first concerns the specification of the Fed preferences. Economic theory offers little guidance on this since structural models with QE include heterogeneous households, a feature which prevents the identification of a unique micro-founded objective function for the central bank. Optimal policy is computed in this paper assuming a standard quadratic loss function in inflation and unemployment, as in Sims et al. (2023) and consistently with the Fed’s dual mandate of

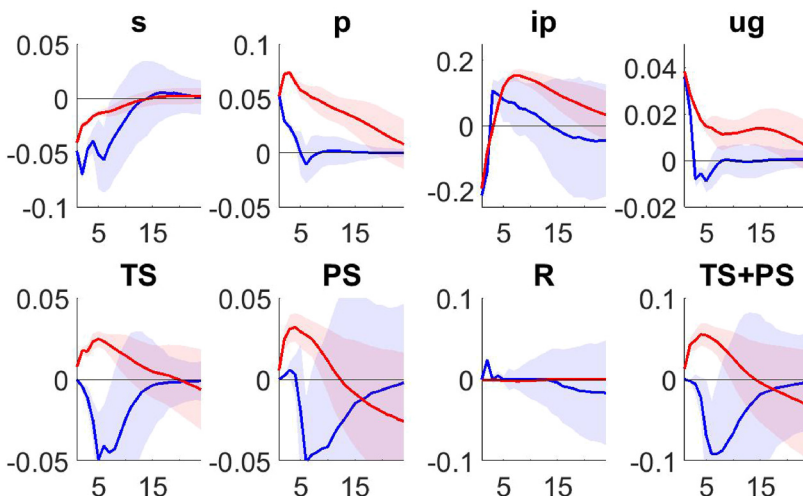


Figure 5: Median and two standard deviation confidence bands of IRFs to a supply-type shock during the ZLB period under the actual (red) and the optimal (blue) policy. All variables are in percentage.

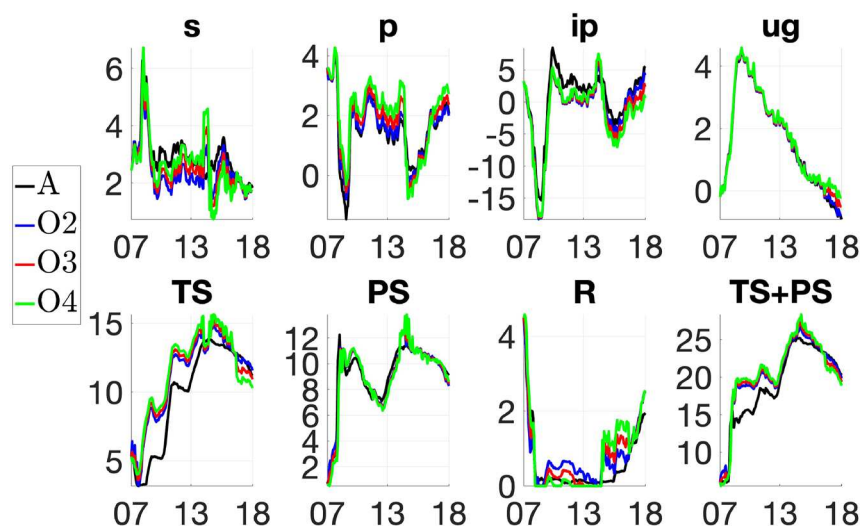


Figure 6: Counterfactual simulations of interest rate and QE policy in the United States. ‘A’ denotes the actual paths; ‘O2’, ‘O3’ and ‘O4’ the counterfactual paths under the inflation targets of 2, 3 and 4 %, respectively. All variables are in percentage.

Table 2: Duration of the ZLB period under the actual and optimal policy.

	Number of months when			
	R less than:			
	0.25	0.5	0.75	100
	85	97	100	104
$\bar{\pi} = 2$	26	55	87	104
$\bar{\pi} = 3$	51	72	74	84
$\bar{\pi} = 4$	71	73	73	75

Source: Author’s calculations. See main text for more details.

price stability and full employment. For robustness, alternative specifications are appraised by varying the loss function parameters. The second caveat is that, in principle, counterfactual analysis with a VSTAR enters the territory of the Lucas (1976) critique. While its quantitative significance is often debated, see Rudebusch (2005) and Sims and Zha (2006), to guard against it the loss function is expanded including penalties for changes in the policy instruments. This ensures that the counterfactual policy simulations are close to the actual data, making them less susceptible to the critique. It is also shown that the simulated paths under the optimal policy are within the range of macroeconomic forecasts from professional survey data (Table 4). In addition, the analysis mainly focuses on the post-2008 period, thus avoiding the change in policy regime relative to the pre-2008 period.²

1.1 Related Literature

The paper makes several contributions to the literature. The methodological part is related to the research on nonlinear time series modelling, which has long established methods for the specification, estimation and structural analysis of VSTAR, see Granger et al. (1993), Dijk et al. (2002) and Martin et al. (2013). The paper complements this literature, tackling a different task, the analysis of systematic policy. The methodological part also

² Prominent examples of counterfactual monetary policy analysis based on reduced-form models are Sack (2000), Stock and Watson (2001), Sims and Zha (2006), Polito and Spencer (2016) for conventional monetary policy; Lenza et al. (2010), Kapetanios et al. (2012), Gambacorta et al. (2014), Dahlhaus et al. (2018), for QE. Many of these works analyze counterfactual policy deviations, such as the no-QE policy scenario, much larger than those considered in this paper.

contributes to the macroeconomic literature on optimal decision analysis. As mentioned above, the predominant paradigm here is still the LQR, which applies to linear economic environments, see Ljungqvist and Sargent (2018). The approach adopted in the paper illustrates how to bring the same LQR techniques to nonlinear economic models.

The paper also contributes to three strands of applied econometrics. The first is the analysis of the wider effects of QE and the ZLB based on VAR or semi-structural models. Examples include Lenza et al. (2010), Kapetanios et al. (2012), Gambacorta et al. (2014), Dahlhaus et al. (2018). These works evaluate what the unconventional monetary policy undertaken by the Fed since the Great Recession *managed to* achieve relative to a scenario of no-intervention. The paper complements this literature evaluating what the actual QE and ZLB policy *could have* achieved, using as reference the framework from the optimal policy. The second is the applied literature on macroeconomic asymmetries and the transmission channels of QE and the ZLB, typified by the works of Guerrieri and Iacoviello (2017) and Sims and Wu (2020); Sims et al. (2023). The present paper shares with these works a similar emphasis on the dual source of nonlinearity, from the economy and the policy sectors, but for the purpose of empirical analysis. The third is the literature using VSTAR-type models to account for nonlinearity in macroeconomic time series, whose applications spans from business cycle, Terasvirta and Anderson (1992), international spillovers, Gefang and Strachan (2009), fiscal policy, Auerbach and Gorodnichenko (2012) (conventional) monetary policy, Dahlhaus (2017), and credit markets, Atanasova (2003). The paper contributes to this literature using a VSTAR for the analysis of QE and the ZLB.

There is a growing literature that develops dynamic stochastic general equilibrium (DSGE) models to study monetary policy when the short term interest rate is constrained by the zero lower bound and the Fed embarks on QE. This literature considers actual and/or optimal QE intervention targeting either private assets, through outright purchases or collateralized short-term loans, as for example in Del Negro et al. (2017), or long term government bonds, as for example in Sims and Wu (2021). The present paper contributes to this literature considering in a reduced form framework the design of QE programmes that simultaneously involve private securities and government bonds purchases. Being reduced form, the analysis is unable to explain the economic mechanisms that may explain why each type of QE is required and desired. It is also silent about the distortions that these purchases may induce on private sector decisions. The main advantage of taking a reduced form approach is that it facilitates the empirical analysis of the effects of conventional and unconventional monetary policy despite the large data instabilities occurred during both the build-up of QE and its first phase of unwinding since 2015, allowing to compare and contrast the effects of different types of asset purchases, how these policies interact with each other and quantify their outcomes in terms of stabilization gains.

This paper is also related to a recent literature that uses DSGE models to conduct counterfactual policy analysis in reduced-form or semi-structural frameworks that aim to be robust to the Lucas critique, see, Barnichon and Mesters (2023), Beraja (2023) and McKay and Wolf (2023). While the methodology in this paper does not explicitly aim to insulate policy analysis from the Lucas critique, it is implemented in a way that attempts to minimize the quantitative relevance of the critique. The methodologies in this recent literature have also limited applicability. Beraja (2023)'s approach only applies to DSGE models that can be linearized, thereby excluding nonlinearities induced by ZLB or leverage constraints that typically features DSGE models with QE. Barnichon and Mesters (2023)'s sufficient statistics cannot be used to compute the optimal policy if there is not enough empirical evidence to estimate the responses to all policy news shocks (The authors also discuss circumstances in which their optimal policy perturbation approach is not applicable to nonlinear models). The identification result of McKay and Wolf (2023) also relies on linearity and is not suited to study change in the steady state. In contrast, the approach proposed in this paper has the advantage of being applicable to a much broader set of nonlinearities and policy experiments, as demonstrated in the empirical application.

The paper is organized as follows. Section 2 describes the reduced-form VSTAR and the two VSTARX representations under A1 and A2. Section 3 illustrates the SDC method and derives the VSTAR under control. Sections 4 and 5 present the results from the estimated VSTAR and the analysis of optimal policy, respectively. Section 6 concludes with a summary. Supplementary material and robustness results are in Appendix.

2 The VSTAR Model

2.1 Reduced Form

Let $\mathbf{y}_t = \begin{bmatrix} \mathbf{x}'_t & \mathbf{u}'_t \end{bmatrix}'$ be a $n \times 1$ vector, with \mathbf{x}_t denoting a $p \times 1$ vector of variables describing the economy and \mathbf{u}_t a $m \times 1$ vector of policy instruments. The dynamic of \mathbf{y}_t for $t \geq 1$ is described by the VSTAR:

$$\begin{aligned} \mathbf{y}_{t+1} = & [1 - g_M(\mathbf{s}'_u \mathbf{y}_t)] \left\{ \begin{array}{l} [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] [\lambda_1 + \Lambda_1(L) \mathbf{y}_t] \\ + l_M(\mathbf{s}'_x \mathbf{y}_t) [\lambda_2 + \Lambda_2(L) \mathbf{y}_t] \end{array} \right\} \\ & + g_M(\mathbf{s}'_u \mathbf{y}_t) \left\{ \begin{array}{l} [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] [\lambda_3 + \Lambda_3(L) \mathbf{y}_t] \\ + l_M(\mathbf{s}'_x \mathbf{y}_t) [\lambda_4 + \Lambda_4(L) \mathbf{y}_t] \end{array} \right\} + \mathbf{v}_{t+1}, \end{aligned} \quad (1)$$

where \mathbf{y}_0 is given; λ_j are vectors of constant coefficients of dimensions $n \times 1$; $\Lambda_j(L)$ denote q lags of the vector \mathbf{y}_t , i.e. $\Lambda_j(L) = \sum_{k=1}^q \Lambda_{jk} L^k$ with each Λ_{jk} being a $n \times n$ matrix of coefficients and L the lag operator; \mathbf{v}_t is a vector of independently and identically distributed reduced-form disturbances with covariance matrix Ω_t given by:

$$\begin{aligned} \Omega_t = & [1 - g_V(\mathbf{s}'_u \mathbf{y}_t)] \{ [1 - l_V(\mathbf{s}'_x \mathbf{y}_t)] \Omega_1 + l_V(\mathbf{s}'_x \mathbf{y}_t) \Omega_2 \} \\ & + g_V(\mathbf{s}'_u \mathbf{y}_t) \{ [1 - l_V(\mathbf{s}'_x \mathbf{y}_t)] \Omega_3 + l_V(\mathbf{s}'_x \mathbf{y}_t) \Omega_4 \}, \end{aligned} \quad (2)$$

where Ω_j are symmetric matrices of coefficients, $j = 1, \dots, 4$. The terms $g_i(\mathbf{s}'_i \mathbf{y}_t)$ and $l_i(\mathbf{s}'_i \mathbf{y}_t)$ denote continuous (transition) functions, bounded between zero and one, capturing the effect of nonlinearity on the transmission mechanism of shocks in the mean, $i = M$, and covariance matrix, $i = V$, of the VSTAR in (1) and (2), \mathbf{s}_x and \mathbf{s}_u being selection vectors identifying transition variables from \mathbf{x}_t and \mathbf{u}_t , respectively. In particular, $l_i(\mathbf{s}'_i \mathbf{y}_t)$ links variations in the VSTAR coefficients to changes in the state of the economy, while $g(\mathbf{s}'_u \mathbf{y}_t)$ links variations in the VSTAR coefficients to changes in the policy stance.

The VSTAR model in (1)-(2) provides a natural framework to study monetary policy in non-linear macroeconomic environments. The idea of considering dual nonlinearity, from the economy and the policy sectors, is not new. It features in recent works of Guerrieri and Iacoviello (2017) and Sims and Wu (2020); Sims et al. (2023) based on DSGE models where nonlinearity stems from the private sector, due to collateral constraints on borrowing, and from the policy sector, due to the ZLB. The specified VSTAR shares with these works a similar emphasis on the dual nonlinearity, with the purpose however of evaluating its significance for empirical analysis.

2.2 VSTARX Form

The VSTAR in (1)-(2) is a reduced form, suitable for econometric estimation and dynamic analysis like forecasting. The coefficients in the equations for the policy instruments describe the so-called actual policy rules. The coefficient in the equations for the economy vector describe how the economy evolves over time subject to the actual policy rule.

To quantify the impact of a change in policy, the coefficients in the equations for the instruments need to be replaced with those from another policy rule, either ad-hoc or optimal in some sense. This requires transforming the VSTAR into a VSTARX where the policy variables are treated as exogenous. As in a VAR, this transformation requires only a zero block factorization (sufficient condition) of the reduced-form covariance matrix, to separate the policy vector from the economy vector. Thus, there are (at least) two possible representations depending on whether the policy variables are ordered first or last.³

³ Other representations could be obtained by further restricting the covariance matrix of the equations for the variables in the economy vector. These further restrictions are however not necessary when using a VAR (or a VSTAR) to study the effects of changes in the policy rules.

To elucidate the structure of, and the differences between, these two VSTARX forms, consider partitioning the mean equation (1) to separate the economy from the policy instruments.⁴ The covariance matrix (2) can be conformably partitioned as:

$$\mathbf{\Omega}_t = \begin{bmatrix} \mathbf{\Omega}_{xxt} & \mathbf{\Omega}_{xut} \\ \mathbf{\Omega}_{uxt} & \mathbf{\Omega}_{uut} \end{bmatrix}. \quad (3)$$

Under A1, the economy reacts with a lag to a change in policy, requiring $\mathbf{\Omega}_{xut} = \mathbf{0}$ in (3).

Under A2, the economy reacts within the same period, thus requiring $\mathbf{\Omega}_{uxt} = \mathbf{0}$.

To start, consider applying A1 and A2 to the VAR obtained after all transition functions in the VSTAR are set equal to zero. This yields:

$$A1: \mathbf{y}_{t+1} = \tilde{\lambda}_1 + \tilde{\Lambda}_1(L)\mathbf{y}_t + \tilde{\Gamma}_1\mathbf{u}_t + \tilde{\mathbf{v}}_{t+1}, \tilde{\mathbf{v}}_t \sim i.i.d.(\mathbf{0}, \tilde{\mathbf{\Omega}})$$

$$A2: \mathbf{y}_{t+1} = \hat{\lambda}_1 + \hat{\Lambda}_1(L)\mathbf{y}_t + \hat{\Gamma}_1\mathbf{u}_{t+1} + \hat{\mathbf{v}}_{t+1}, \hat{\mathbf{v}}_t \sim i.i.d.(\mathbf{0}, \hat{\mathbf{\Omega}}).$$

The above are two VARX models in which the vector of policy instruments \mathbf{u} enters on the right side as an exogenous variable.⁵ These two VARX differ in terms of the interaction between the economy and the policy instruments, being with a lag under A1 and contemporaneous under A2. Using either of them as constraint for optimization makes no difference in terms dimensionality, since a change in the policy vector has a direct effect on the endogenous variables only through the time-invariant coefficients $\tilde{\Gamma}_1$ and $\hat{\Gamma}_1$ under A1 and A2, respectively. Further, disturbances are homoscedastic in both models. Thus a change in policy has not effect on the covariance matrices $\tilde{\mathbf{\Omega}}$ and $\hat{\mathbf{\Omega}}$, implying that certainty-equivalent policy rules can be calculated using both A1 and A2.

Applying A1 to the VSTAR, yields:

$$\mathbf{y}_{t+1} = [1 - g_M(\mathbf{s}'_u\mathbf{y}_t)] \left\{ \begin{array}{l} [1-l_M(\mathbf{s}'_x\mathbf{y}_t)][\phi_1 + \mathbf{\Phi}_1(L)\mathbf{y}_t + \mathbf{\Gamma}_1\mathbf{u}_t] \\ + l_M(\mathbf{s}'_x\mathbf{y}_t)(\phi_2 + \mathbf{\Phi}_2(L)\mathbf{y}_t + \mathbf{\Gamma}_2\mathbf{u}_t) \end{array} \right\} \\ + g_M(\mathbf{s}'_u\mathbf{y}_t) \left\{ \begin{array}{l} [1-l_M(\mathbf{s}'_x\mathbf{y}_t)][\phi_3 + \mathbf{\Phi}_3\mathbf{y}_t + \mathbf{\Gamma}_3(L)\mathbf{u}_t] \\ + l_M(\mathbf{s}'_x\mathbf{y}_t)[\phi_4 + \mathbf{\Phi}_4(L)\mathbf{y}_t + \mathbf{\Gamma}_4\mathbf{u}_t] \end{array} \right\} + \mathbf{e}_{t+1}, \quad (4)$$

$$\mathbf{\Sigma}_t = \begin{bmatrix} \mathbf{\Omega}_{xxt} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (5)$$

where $\phi_j = [\lambda'_{xj} \quad \mathbf{0}']'$; $\mathbf{\Phi}_j(L) = \begin{bmatrix} \mathbf{\Lambda}_{xxj}(L) & \mathbf{\Lambda}_{xuj}(L) - \mathbf{\Lambda}_{xulj} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$, with (i) $\mathbf{\Lambda}_{xxj}(L)$ and $\mathbf{\Lambda}_{xuj}(L)$ denoting partitions of the responses of the economy variables, $\mathbf{\Lambda}_{xj}(L)$, to their own lagged values and to the policy instruments, respectively, and (ii) $\mathbf{\Lambda}_{xulj}(L)$ the responses of the economy variables to the first lag of the policy instruments; $\mathbf{\Gamma}_j = [\mathbf{\Lambda}'_{xuj1} \quad \mathbf{0}']'$ is the response of the economy variables to the first lag of the policy instruments; $j = 1, \dots, 4$; and $\mathbf{e}_t = [\mathbf{v}'_{xt} \quad \mathbf{0}']' \sim i.i.d.(\mathbf{0}, \mathbf{\Sigma}_t)$.

Conversely, applying A2 to the VSTAR yields:

$$\mathbf{y}_{t+1} = [1 - g_M(\mathbf{s}'_u\mathbf{y}_t)] \left\{ \begin{array}{l} [1-l_M(\mathbf{s}'_x\mathbf{y}_t)][\phi_{1t} + \mathbf{\Phi}_{1t}(L)\mathbf{y}_t] \\ + l_M(\mathbf{s}'_x\mathbf{y}_t)[\phi_{2t} + \mathbf{\Phi}_{2t}(L)\mathbf{y}_t] \end{array} \right\} \\ + g_M(\mathbf{s}'_u\mathbf{y}_t) \left\{ \begin{array}{l} [1-l_M(\mathbf{s}'_x\mathbf{y}_t)][\phi_{3t} + \mathbf{\Phi}_{3t}(L)\mathbf{y}_t] \\ + l_M(\mathbf{s}'_x\mathbf{y}_t)[\phi_{4t} + \mathbf{\Phi}_{4t}(L)\mathbf{y}_t] \end{array} \right\} + \mathbf{\Gamma}_t\mathbf{u}_{t+1} + \mathbf{e}_{t+1}, \quad (6)$$

⁴ Step-by-step derivation of all results presented in this section is provided in Appendix A in Supplementary Material.

⁵ While by definition \mathbf{y}_t includes \mathbf{u}_t , the coefficients corresponding to the equations of the variables in \mathbf{u}_t are all zero in the two VARX. Thus \mathbf{u}_t is effectively exogenous in both models.

$$\Sigma = \begin{bmatrix} \sigma_x^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \Gamma_t = \begin{bmatrix} \mathbf{G}'_{12t} & \mathbf{I} \end{bmatrix}', \mathbf{G}_{12t} = \mathbf{\Omega}_{xut} \mathbf{\Omega}_{uut}^{-1}, \quad (7)$$

where $\phi_{jt} = [(\lambda_{xj} - \mathbf{G}_{12t} \lambda_{uj})' \quad \mathbf{0}']'$, $\Phi_{jt}(L) = \begin{bmatrix} \Lambda_{xj}(L) - \mathbf{G}_{12t} \Lambda_{uj}(L) \\ \mathbf{0} \end{bmatrix}$, $j = 1, \dots, 4$, $\mathbf{e}_t = [(\mathbf{v}_{xt} - \mathbf{G}_{12t} \mathbf{v}_{ut})' \quad \mathbf{0}']' \sim i.i.d.(\mathbf{0}, \Sigma)$, and σ_x^2 denotes the variance of $\mathbf{v}_{xt} - \mathbf{G}_{12t} \mathbf{v}_{ut}$.

The nonlinear dynamic of \mathbf{y}_t under both (4) and (6) describes a VSTAR with exogenous variables, i.e. a VSTARX, since the policy vector \mathbf{u}_t enters on the right side of both models as an exogenous variable.⁶ Therefore, both (4) and (6) can be used to represent the nonlinear economy constraint faced by a regulator in charge of setting \mathbf{u}_t . When the objective of the regulator is approximated by a quadratic function, the resulting problem is referred to as the NLQR problem.

It is useful to point out the differences between (4)-(5) and (6)-(7) due A1 and A2. The VSTARX in (4)-(5) corresponds to the equations for the economy vector \mathbf{x}_t in the reduced-form VSTAR (1)-(2). In contrast, the VSTARX in (6)-(7) is obtained by extrapolating from the reduced-form covariance matrix (3) the contemporaneous response of the economy to policy and then incorporating this into the systematic part of \mathbf{x}_t .

There are three further differences, specific to the nonlinear nature of the VSTAR. The first refers to the response coefficient of the economy variables \mathbf{x}_t to changes in the policy vector \mathbf{u}_t . In (4) this reflects the nonlinearity in the mean of the VSTAR as it is given by $\Gamma_j = [\Lambda'_{xuj1} \quad \mathbf{0}']'$, $j = 1, \dots, 4$. In (6), this is determined by the nonlinearity in the covariance matrix of the VSTAR as it is given by Γ_t in (7). The second refers to the coefficients ϕ , Φ and Γ . In (4) these are fixed, and it is their nonlinear combination that varies across states, reflecting the state-dependent nature of the VSTAR coefficients. In (6), these coefficients are time varying and written as ϕ_{jt} , Φ_{jt} and Γ_{jt} , $j = 1, \dots, 4$, since they are re-calculated to incorporate time variation from the reduced-form covariance matrix. The third difference refers to the structure of the covariance matrix of the disturbances in the two representations: this is time varying in (5), but constant in (7).

Both (4)-(5) and (6)-(7) can be employed as constraints in a NLQR problem. The VSTARX in (4)-(5) has however two main drawbacks. First, the NLQR solution becomes highly dimensional. This is because the coefficients in (4) depend on the lagged value of the policy vector, the variable set by the regulator. Thus changes in \mathbf{u}_t affect on impact the entire system through Γ_j and $g_M(\mathbf{s}'_t \mathbf{y}_t)$. Second, the solution no longer satisfies certainty equivalence, since changes in the policy vector also affect the covariance matrix in (5). The VSTARX in (6)-(7) yields a solution that has lower dimensionality compared to (4)-(5). This is because changes in the policy vector \mathbf{u}_{t+1} affect the economy through Γ_j but not $g_M(\mathbf{s}'_t \mathbf{y}_t)$, which is pre-determined when policy choices are made. In addition, the VSTARX in (6)-(7) satisfies the certainty-equivalence principle, because the disturbances in (6) have constant covariance given by (7). In other words, the representation in (6)-(7) results in a low-dimension and certainty-equivalent NLQR.

The empirical evidence on optimal monetary policy analysis based VAR, finds that the two timing protocols of A1 and A2 in general do not result in large quantitative differences, see Polito and Wickens (2012). Given the methodological advantages highlighted above, the NLQR problem is thus specified and solved in the next section using the VSTARX based on A2 in (6)-(7). A further reason for preferring A2 is that a reduced-form policy experiment could miss the timing of the policy impact on the economy, as forward-looking agents in the private sector can be expected to incorporate information about the implementation of policy at the time of announcement, see Dahlhaus et al. (2018). As also argued by Stock and Watson (2001), A2 minimizes this shortcoming by at least allowing the economy to react within the same period of a change in policy.⁷

⁶ The coefficients in the equations corresponding to the policy instruments are all zero in (4) and (6).

⁷ For comparison, the VSTARX derived using A1 is employed in two different types of nonlinear optimization exercises in the numerical analysis. These results are presented in Section 5.4.

3 Optimal Policy

3.1 Specification

The problem of setting optimal policy is specified as a closed-loop regulator determining the sequence of policy instruments $\{\mathbf{u}_t\}_{t=1}^{\infty}$ that minimizes the quadratic loss function

$$V_0 = E_0 \sum_{t=0}^{\infty} \beta^t \left[(\mathbf{y}_t - \bar{\mathbf{y}})' \mathbf{Q} (\mathbf{y}_t - \bar{\mathbf{y}}) \right], \quad (8)$$

where E_0 denotes mathematical expectation conditional on information at $t = 0$; $\beta \in (0, 1)$ is a discount factor; $\bar{\mathbf{y}}$ is a vector of targets; \mathbf{Q} is a symmetric positive semidefinite matrix of coefficients. Given \mathbf{y}_0 , the regulator sets $\{\mathbf{u}_t\}_{t=1}^{\infty}$ subject to the constraints in (6)-(7).⁸

The above NLQR problem has not closed-form solution. One option is to linearize (6) so that the NLQR problem effectively becomes a LQR and a closed-form solution can be calculated with linear dynamic programming, for example using Riccati equation iteration as in Ljungqvist and Sargent (2018). The alternative is to solve the problem numerically using the nonlinear model of the economy as it is or a higher-order approximation. Linearization would somehow be inconsistent with the aim of this paper, i.e. studying the significance of macroeconomic nonlinearity for optimal policy analysis. The main drawback of numerical methods is that they can be computationally intensive and there is not guarantee about their optimization performance in highly nonlinear models. This is because direct search methods can easily default on local minima while gradient methods can quickly loose efficiency once differentiation becomes either intractable or very high dimensional due to nonlinearity. Luukkonen et al. (1988) show how to employ Taylor-series expansion of the transition function of the VSTAR to derive a polynomial approximation that can be used to test for nonlinearity in the data. With this approach, however, the economy constraint of the NLQR would remain still highly nonlinear, even when using a first-order approximation of the transition function.

The SDC method extends Riccati equation iteration used in LQR to the nonlinear case. This is done by approximating the NLQR problem into a sequence of one-period LQR problems that can be solved independently. The SDC method has the advantage of being numerically simpler than many other nonlinear optimization techniques, as well as being closely related to the well-understood Riccati equation method used for LQR. The SDC method is widely applied in engineering and data science.⁹ The next subsections describe how to adapt it for optimal macroeconomic policy analysis with a VSTAR.

3.2 The SDC Method

Applying the SDC method to solve the NLQR problem of minimizing (8) subject to (6)-(7) requires two steps. The first consists of employing SDC factorization to transform the nonlinear model (6) into an affine structure with SDC matrices. As a result, the NLQR problem can be written in a form that is isomorphic to the LQR problem and it can be solved with linear dynamic programming techniques. The second step consists of solving the NLQR problem iterating on the resulting SDC solutions. This yields a feedback rule for the policy instruments with time-varying response coefficients that can be combined with the original constraint to derive the VSTAR under control.

⁸ This is equivalent to calculate the solution under commitment. The analysis can be easily extended to compute solutions under discretion, where the regulator minimizes the loss function only in the current period, $(\mathbf{y}_t - \bar{\mathbf{y}})' \mathbf{Q} (\mathbf{y}_t - \bar{\mathbf{y}})$.

⁹ The method was pioneered by the early works of Pearson (1962) and Wernli and Cook (1975). Kaiser et al. (2017) illustrate a recent application in data science. Guerrieri and Iacoviello (2017) use SDC factorization to solve a nonlinear DSGE model as a VAR with SDCs. An interesting avenue of research, outside the scope of this paper, could be to investigate whether a mapping then exists into either a VSTAR or other nonlinear VAR representations.

3.2.1 Factorization

SDC factorization is typically used, outside economics, for solving (i) continuous time, (ii) deterministic NLQR problems, (iii) using the method of variation. The approach is adapted here for solving (i) discrete time, (ii) stochastic NLQR problems, (iii) using dynamic programming. To illustrate the logic of SDC factorization, consider the differential equation for an autonomous system, with fully observable state, nonlinear in the state and affine in the control: $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}(t)$, $\mathbf{x}(0) = \mathbf{x}_0$, where $\mathbf{x} \in \mathbb{R}^n$ is the state vector, $\mathbf{u} \in \mathbb{R}^m$ is the control vector, $\mathbf{f}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ are continuous functions, $\mathbf{f}(\mathbf{0}) = \mathbf{0}$, $\mathbf{g}(\mathbf{x}) \neq \mathbf{0} \forall \mathbf{x}$, and $t \in [0, \infty]$. The two nonlinear functions can be factorized as $\mathbf{f}(\mathbf{x}) = \mathbf{A}(\mathbf{x})\mathbf{x}$ and $\mathbf{g}(\mathbf{x}) = \mathbf{B}(\mathbf{x})$ where $\mathbf{A}(\mathbf{x})$ and $\mathbf{B}(\mathbf{x})$ are matrices of coefficients that depend on the state vector. Thus the nonlinear differential equation can be written into an affine form with SDC matrices: $\dot{\mathbf{x}} = \mathbf{A}(\mathbf{x})\mathbf{x} + \mathbf{B}(\mathbf{x})\mathbf{u}$.¹⁰

Proceeding in a similar fashion, the nonlinear system in equation (6) can be factorized in terms of SDC matrices and rewritten as

$$\mathbf{y}_{t+1} = \mathbf{c}(\mathbf{y}_t) + \mathbf{A}(\mathbf{y}_t)\mathbf{y}_t + \mathbf{B}(\mathbf{y}_t)\mathbf{u}_{t+1} + \mathbf{e}_{t+1} \quad (9)$$

for $t \geq 0$, using:

$$\begin{aligned} \mathbf{c}(\mathbf{y}_t) &= [1 - g_M(\mathbf{s}'_u \mathbf{y}_t)] \{ [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] \phi_{1t} + l_M(\mathbf{s}'_x \mathbf{y}_t) \phi_{2t} \} \\ &\quad + g_M(\mathbf{s}'_u \mathbf{y}_t) \{ [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] \phi_{3t} + l_M(\mathbf{s}'_x \mathbf{y}_t) \phi_{4t} \}, \\ \mathbf{A}(\mathbf{y}_t) &= [1 - g_M(\mathbf{s}'_u \mathbf{y}_t)] \{ [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] \Phi_{1t}(L) + l_M(\mathbf{s}'_x \mathbf{y}_t) \Phi_{2t}(L) \} \\ &\quad + g_M(\mathbf{s}'_u \mathbf{y}_t) \{ [1 - l_M(\mathbf{s}'_x \mathbf{y}_t)] \Phi_{3t}(L) + l_M(\mathbf{s}'_x \mathbf{y}_t) \Phi_{4t}(L) \}, \\ \mathbf{B}(\mathbf{y}_t) &= \Gamma_t, \end{aligned}$$

with the upper part of Γ_t , i.e. \mathbf{G}_{12t} , being given by (7).¹¹ Three results about the above SDC factorisation are worth highlighting. First, by construction the SDC factorization gives a representation of the model of the economy mathematically equivalent to VSTARX model in (6). Second, the SDC vectors $\mathbf{c}(\mathbf{y}_t)$ and matrices $\mathbf{A}(\mathbf{y}_t)$ are constructed adding through weighted coefficients across the four states, with weights given by $[1 - g_M(\mathbf{s}'_u \mathbf{y}_t)][1 - l_M(\mathbf{s}'_x \mathbf{y}_t)]$, $[1 - g_M(\mathbf{s}'_u \mathbf{y}_t)]l_M(\mathbf{s}'_x \mathbf{y}_t)$, $g_M(\mathbf{s}'_u \mathbf{y}_t)[1 - l_M(\mathbf{s}'_x \mathbf{y}_t)]$ and $g_M(\mathbf{s}'_u \mathbf{y}_t)l_M(\mathbf{s}'_x \mathbf{y}_t)$ for state $j = 1, 2, 3$ and 4 , respectively. However, the SDC matrix $\mathbf{B}(\mathbf{y}_t)$ encapsulates the highly nonlinear structure of (7). Third, all SDC matrices in (9) depend on the previous period state vector. For this reason, these are unaffected in the current period by a change in policy.¹²

3.2.2 Solution

Following the SDC factorization, the closed-loop regulator problem consists of finding the sequence $\{\mathbf{u}_{t+1}\}_{t=0}^{\infty}$ that minimizes (8) subject to (9) given \mathbf{y}_0 . For convenience, this is re-written below:

$$\begin{aligned} \min_{\{\mathbf{u}_{t+1}\}_{t=0}^{\infty}} V_0 &= E_0 \sum_{t=0}^{\infty} \beta^t \left[(\mathbf{y}_t - \bar{\mathbf{y}})' \mathbf{Q} (\mathbf{y}_t - \bar{\mathbf{y}}) \right] \\ \text{s.t. : } \mathbf{y}_{t+1} &= \mathbf{c}(\mathbf{y}_t) + \mathbf{A}(\mathbf{y}_t)\mathbf{y}_t + \mathbf{B}(\mathbf{y}_t)\mathbf{u}_{t+1} + \mathbf{e}_{t+1}, \\ \mathbf{e}_{t+1} &\sim \text{i.i.d.}(\mathbf{0}, \Sigma) \text{ and } \mathbf{y}_0 \text{ given.} \end{aligned}$$

¹⁰ Wernli and Cook (1975) extend SDC factorization to the general model $\dot{\mathbf{x}} = f(x, u, t)$. This yields $f(x, u, t) = A(x, u, t)x + B(x, u, t)u$, where the matrices $A(x, u, t)$ and $B(x, u, t)$ depend on x , u and t .

¹¹ SDC factorization is very general, and applicable to many other nonlinear systems, including artificial neural networks, threshold and bilinear models.

¹² An issue often highlighted is that SDC factorizations are not unique, since there are many nonlinear models that can be encapsulated into the SDC matrices in (9), see, e.g., Beeler et al. (2000). This occurs when SDC factorization is used to approximate an unknown nonlinear model. However, this issue does not apply here, as the starting point of the factorization is the known system in (6).

This formulation of the NLQR problem is isomorphic to the LQR except for the objects $\mathbf{c}(\mathbf{y}_t)$, $\mathbf{A}(\mathbf{y}_t)$ and $\mathbf{B}(\mathbf{y}_t)$, whose elements are functions of the state vector \mathbf{y}_t . The SDC method treats the matrices $\mathbf{c}(\mathbf{y}_t)$, $\mathbf{A}(\mathbf{y}_t)$ and $\mathbf{B}(\mathbf{y}_t)$ as fixed in any given period $t \geq 0$ and solves the NLQR problem in that period as a LQR. As a result, the closed-loop solution can be derived as:

$$\mathbf{u}_{t+1} = \mathbf{k}_t + \mathbf{K}_t \mathbf{y}_t, \quad (10)$$

which is a linear feedback rule with time-varying coefficients, \mathbf{k}_t and \mathbf{K}_t .¹³ This can be combined with the equations for the economy in SDC form in (9) as:

$$\begin{bmatrix} \mathbf{I} & -\mathbf{B}(\mathbf{y}_t) \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{t+1} \\ \mathbf{u}_{t+1} \end{bmatrix} = \begin{bmatrix} \mathbf{c}(\mathbf{y}_t) \\ \mathbf{k}_t \end{bmatrix} + \begin{bmatrix} \mathbf{A}(\mathbf{y}_t) \\ \mathbf{K}_t \end{bmatrix} \mathbf{y}_t + \begin{bmatrix} \mathbf{e}_{t+1} \\ \mathbf{0} \end{bmatrix} \quad (11)$$

to recover a reduced-form model with time-varying coefficients:

$$\mathbf{y}_{t+1} = \lambda_t^* + \Lambda_t^*(L) \mathbf{y}_t + \mathbf{e}_{t+1}^*, \quad (12)$$

where

$$\lambda_t^* = \begin{bmatrix} \mathbf{c}(\mathbf{y}_t) + \mathbf{B}(\mathbf{y}_t) \mathbf{k}_t \\ \mathbf{k}_t \end{bmatrix}, \quad \Lambda_t^*(L) = \begin{bmatrix} \mathbf{A}(\mathbf{y}_t) + \mathbf{B}(\mathbf{y}_t) \mathbf{K}_t \\ \mathbf{K}_t \end{bmatrix}, \quad \mathbf{e}_t^* = \begin{bmatrix} \mathbf{e}_t \\ \mathbf{0} \end{bmatrix}$$

and $\mathbf{e}_t \sim (\mathbf{0}, \Sigma)$, with Σ being defined in (7). Equation (12) is the VSTAR under the optimal feedback rule in SDC form. This differs from the reduced-form VSTAR in (1)-(2) in three respects. First, the coefficients of the equations for the economy in the VSTAR in (1) are replaced with a linear combination of the corresponding coefficients from the VSTARX model in (6)-(7) and the policy rule coefficients (see the upper parts of λ_t^* and $\Lambda_t^*(L)$). Second, the coefficients of the equations for the policy vector in the VSTAR in (1) are replaced with those in (10). Third, the covariance matrix of the VSTAR in (12) is constant, while it is time-varying and state-dependent in the original VSTAR. It is worth noting that the system (12) can be used to study the effects of any policy rule, including one with ad-hoc coefficients. To this end it is sufficient to replace \mathbf{k}_t and \mathbf{K}_t in (11) with any given values and then derive the corresponding reduced-form system as in (12).

4 Estimated VSTAR

4.1 Data

The VSTAR is estimated using aggregate monthly data for the United States from 1979:8 to 2018:10.¹⁴ In the benchmark specification, the endogenous vector comprises seven variables: an indicator of credit risk, the spread between the Baa corporate rate and the 10-year treasury rate (s); the inflation rate (p), measured as the annual change in the personal consumption expenditures deflator; the annual growth rate of industrial production (ip); the difference between the actual and the natural rate of unemployment, the unemployment gap (ug); Treasury securities held by the Fed as a proportion to GDP (TS); private securities held by the Fed as a proportion to GDP (PS); and the federal fund rate (R).¹⁵ The selected measure of credit risk is used as a proxy for financial conditions in the macroeconomy.¹⁶ Inflation and the unemployment gap are used as indicators of the Fed targets of price

¹³ Appendix B in Supplementary Material shows how to adapt the method of Riccati equation iteration described in Ljungqvist and Sargent (2018) to solve the above NLQR problem and the algorithm required for the computation of the solution. Appendix C in Supplementary Material describes how stability is ensured upon the optimal solution.

¹⁴ The use of monthly data is quite common in VSTAR analysis, as it ensures that the number of available observations is large enough for the estimation of the model parameters.

¹⁵ Data sources and transformations are described in Appendix D in Supplementary Material.

¹⁶ Galvão and Owyang (2018) use a factor VSTAR to derive an index of financial conditions that is highly correlated with macroeconomic variables of the United States. The indicator s has an estimated correlation of 98 % with their index. An alternative proxy for credit risk is the VIX, but using this would significantly curtail the sample size as it is only available from the 1990s onward.

stability and maximum employment, respectively, as in Williamson (2015). Industrial production is used as a proxy for aggregate output, following Sack (2000). The sum of *TS* and *PS* equals total assets held in the Fed's balance sheet. Thus the VSTAR includes conventional and QE monetary policy tools, and can account for variation in both the size of the Fed's balance sheet and the composition of its assets portfolio.

As in Gambacorta et al. (2014), the total size of the Fed's balance sheet is used to proxy QE. The distinction between *TS* and *PS* captures the effects of QE on aggregate demand through the portfolio-balance and the credit channels, respectively, highlighted by the economic theory. Kuttner (2018) supports this distinction on the ground that it characterizes the conduct of QE in the United States since 2008, which involved large purchases of *PS* in the immediate aftermath of the Great Recession, and of *TS* afterwards.¹⁷

Figure 1 plots the data, with NBER recession periods highlighted in grey for reference. with NBER recessions highlighted in grey. Several clear irregularities are visible over the sample period, such as the large spikes in the credit risk, inflation, industrial production and unemployment gap; the asymmetric dynamics of the unemployment gap above and below zero; the large transitional changes in the two QE variables around the Great Recession; and the flat path of the federal fund rate while at the ZLB. It is the presence of these irregularities that motivates the use of a VSTAR.¹⁸

4.2 Specification

Variables are ordered as $\mathbf{y}'_t = [s_t \ p_t \ ip_t \ ug_t \ TS_t \ PS_t \ R_t]$. The transition functions are specified as follows. $l_i(\mathbf{s}'_x \mathbf{y}_t)$ is set to capture nonlinearity due to the amount of slack in the economy, measured by the unemployment gap. Thus $\mathbf{s}'_x = [0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0]$ and $l_i(\mathbf{s}'_x \mathbf{y}_t) = l_i(ug_t)$. This is similar to Ramey and Zubairy (2018), except that in their work the natural unemployment rate is constant, while it is time varying here following the definition of the unemployment gap of Williamson (2015).¹⁹ $g_i(\mathbf{s}'_u \mathbf{y}_t)$ is set to capture changes in the VSTAR parameters depending on whether or not the economy is near to the ZLB. Thus $\mathbf{s}'_u = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1]$ and $g_i(\mathbf{s}'_u \mathbf{y}_t) = g_i(R_t)$.

Transition across states due to nonlinearity in the economy is modelled using a logistic function $l_i(ug_t): (-\infty, +\infty) \rightarrow [0, 1]$ specified as: $l_i(ug_t) = \{1 + \exp[-\gamma_{li}(ug_t - c_i)]\}^{-1}$, where $\gamma_{li} > 0$ denotes the speed of adjustment across states for the parameters of \mathbf{y}_t , $i = M$, and $\mathbf{\Omega}_t$, $i = V$; c_i is a threshold parameter. Since the federal funds rate cannot be negative, transition across states due to nonlinearity in the policy vector is instead modelled with the incomplete gamma function $g_i(R_t): (0, +\infty) \rightarrow [0, 1]$ described as: $g_i(R_t) = \frac{1}{\Gamma(\gamma_{gi})} \int_0^{R_t} e^{-t} t^{\gamma_{gi}-1} dt$, where $\Gamma(\gamma_{gi})$ is the gamma function with shape parameter γ_{gi} for \mathbf{y}_t , $i = M$, and $\mathbf{\Omega}_t$, $i = V$; and the implied scale parameter has been set equal to one. Logistic functions are widely employed in the applied macroeconomic literature to capture nonlinearity stemming from the economy, see, e.g., Auerbach and Gorodnichenko (2012), Galvão and Owyang (2018), Ramey and Zubairy (2018). The incomplete gamma function is used to capture nonlinearity driven by a variable that is nonnegative, being for example volatility, as in Lanne and Saikkonen (2005), or the nominal rate of interest, as in Hurn et al. (2022). To discipline the economic interpretation, the threshold parameters are set equal to zero, i.e. $c_i = 0$, $i = M, V$. This means that nonlinearity stemming from the economy vector depends entirely on whether the actual rate of unemployment is above or below the natural rate, i.e. the amount of slack in the economy.

To ensure positive definiteness on the variance covariance matrix in (2) and on its determinant, each matrix of coefficients $\mathbf{\Omega}_j$ is decomposed as

$$\mathbf{\Omega}_j = \mathbf{D}'_j \mathbf{D}_j, \quad (13)$$

where \mathbf{D}_j is a lower triangular matrix, $j = 1, \dots, 4$.

¹⁷ Dahlhaus et al. (2018) use instead the sum of *TS* and *PS* with maturity longer than five years. This alternative proxy for QE displays very similar dynamics to *TS* + *PS*, since the share of assets with maturity of five or more years in the Fed's balance sheet ranges between 60 and 80 % between 2010 and 2015.

¹⁸ Luukkonen et al. (1988)'s LM test rejects linearity in each of these time series at 95 % confidence.

¹⁹ The unemployment gap is not the only possible measure of the state of the economy. Auerbach and Gorodnichenko (2012), for example, employ the moving average of GDP growth. However, GDP growth is not explicitly targeted by the Fed.

4.3 Estimation

Given a sample of $t = 1, \dots, T$ observations and the i.i.d. assumption on the disturbances \mathbf{v}_t , the log-likelihood function of the VSTAR in (1)-(2) conditional on the first q observations is formed as:

$$\ln L_T(\mathbf{\Pi}) = -[2(T - q)]^{-1} \sum_{t=q+1}^T \left(n \ln 2\pi + \ln |\mathbf{\Omega}_t| + \mathbf{v}_t' \mathbf{\Omega}_t^{-1} \mathbf{v}_t \right), \quad (14)$$

where $\mathbf{\Pi}$ is a vector collecting all VSTAR parameters in equations (1), (2) and (13). It is convenient to partition the vector of parameters as $\mathbf{\Pi} = [\mathbf{\Pi}_1, \mathbf{\Pi}_2]$, where $\mathbf{\Pi}_1 = \{\lambda_1, \mathbf{\Lambda}_1(L), \lambda_2, \mathbf{\Lambda}_2(L), \lambda_3, \mathbf{\Lambda}_3(L), \lambda_4, \mathbf{\Lambda}_4(L)\}$ includes the linear coefficients in the conditional mean (1), while $\mathbf{\Pi}_2 = \{\mathbf{D}_1, \mathbf{D}_2, \mathbf{D}_3, \mathbf{D}_4, \gamma_{IM}, \gamma_{IV}, \gamma_{gM}, \gamma_{gV}\}$ includes the coefficients pertinent to the covariance matrix and to the transition functions.

To establish the optimal lag length and estimate the model parameters, an algorithm similar to that used by Auerbach and Gorodnichenko (2012) is adapted to the VSTAR in equations (1), (2) and (13).²⁰ The algorithm exploits the possibility of writing the model in a linear form conditional on a subset of the parameters, thereby estimating the remaining parameters analytically. Thus, estimation involves repeating until convergence two sequential steps. The first consists of determining a draw of $\mathbf{\Pi}_2$. Conditional on $\mathbf{\Pi}_2$, the VSTAR is linear in the remaining parameters over the four states. In the second step the remaining parameters in $\mathbf{\Pi}_1$ can be determined analytically via generalized least squares and the log-likelihood can be evaluated from (14). The estimation is undertaken including up to 12 lags and the model fit evaluated using information criteria such as, Akaike (AIC), Hannan (HIC) and Schwarz (SIC), in both standard and normalised form. The model with one lag is the preferred specification according to almost any of the information criteria considered.²¹ In Appendix E in Supplementary Material it is also shown that the proposed VSTAR specification yields a better fit of the data compared to several other VSTAR specifications often used in the applied macroeconomic literature.

Figure 2 plots the transition probabilities estimated from the VSTAR for the mean (top panel) and the covariance matrix (bottom panel) parameters. The black lines show the probability of the economy being, $l_i(ug_t) \simeq 1$, or not being, $l_i(ug_t) \simeq 0$, in a slack state. The red lines the probability of policy being away from, $g_i(R_t) \simeq 1$, or near to, $g_i(R_t) \simeq 0$, the ZLB. The top panel shows four sharp shifts in the mean coefficients linked to the economy vector occurring around recessions.²² After 2001, nonlinearity from the policy sector becomes significant too, in particular during 2009–2015, as the federal funds rate reaches the ZLB. The bottom panel shows shifts in the coefficients of the covariance matrix around recessions, whereas changes driven by R_t reflect predominantly the downward trend in the federal funds rate observable in Figure 1. The results in Figure 2 highlight the flexibility of the estimated VSTAR, given that it can simultaneously capture gradual changes in the mean coefficients linked with changes in the amount of slack in the economy and the more sudden variation in the volatility coefficients.²³

4.4 Impulse Response Functions

IRF analysis with VSTAR is typically based on the algorithm of Koop et al. (1996), with shocks identified through Cholesky factorization of the covariance matrix $\mathbf{\Omega}_t$. In a small-scale VAR the Cholesky decomposition yields structural shocks that can be easily reconciled with an economic interpretation. This is not necessarily the case

²⁰ Further details on the estimation methodology and lag selection results are in Appendix E in Supplementary Material.

²¹ VSTAR models with one lag are not uncommon in applied macroeconomics, because the nonlinear nature of the model on its own can capture a significant part of the data dynamics, see for example Dahlhaus (2017) and Galvão and Owyang (2018).

²² These four regimes are also documented by Liu et al. (2019) using a change-point VAR model estimated on macroeconomic data of the United States since the 1960s.

²³ The estimated transition probabilities are somewhat between those presented by Auerbach and Gorodnichenko (2012), that assume gradual changes, and those in Ramey and Zubairy (2018), that assume sudden changes. Overall, the gradual changes in the probabilities from the estimated VSTAR is in part driven by the use of monthly data, while these two studies use quarterly and annual data, respectively.

in larger (and nonlinear) models like the estimated VSTAR. For this reason, IRF analysis is carried out in this paper by adapting the algorithm of Koop et al. (1996) to include sign restrictions.²⁴ To formulate a benchmark, the restrictions are set to identify a credit-type shock that on impact yields macroeconomic dynamics similar to those observed in the United States after 2008, leaving entirely unrestricted the policy variables. This shock is such that both the spread and the unemployment gap increase while inflation and industrial production decrease at the one month horizon.²⁵ To evaluate the different responses of the policy instruments during periods when either conventional monetary policy or QE is active, the IRFs are computed using as starting values for y_t the averages over two sub-samples of equal size. The first is from 1994:6 to 2001:5. This is referred to as the normal times, covering a period when the federal funds rate is well above the ZLB. The second includes observations for the ZLB period, from 2009:1 to 2015:12.

Figure 3 displays the responses of the three policy instruments plus the total size of the Fed's balance sheet ($TS + PS$) to the identified credit shock over 24 months. In each panel, the black lines denote the responses during normal times, the red lines the responses in the ZLB period and the shaded areas denote two standard deviations confidence bands. The response of the policy instruments mimics that of expansionary monetary policy intervention at different periods of the United States' economic history. In normal times the shock leads to a reduction of the federal funds rate, while QE shows little dynamics. During the ZLB the monetary expansion occurs through QE. It is remarkable how the path of the IRFs for the two QE variables reproduces changes in the Fed assets' portfolio similar to those observed after QE1-QE2, as visible from the initial expansion of PS holding, subsequently supported by a large increase of TS . These dynamics of the Fed's balance sheet, consisting of a large initial purchase of PS followed by TS , are also visible from the evolution of TS and PS around the Great Recession in Figure 1.²⁶

The IRF analysis indicates that the estimated VSTAR returns a plausible description of the monetary policy instrument responses to macroeconomic shocks at different times of the United States' monetary history, thereby providing a reliable tool for the analysis of optimal policy in the next section. In Appendix F in Supplementary Material it is shown that these results are robust to different specifications of the VSTAR, obtained by replacing industrial production with either GDP or public debt, and different types of economic disturbances, such as demand and supply shocks.

5 Optimal VSTAR

Quantitative analysis of optimal macroeconomic policy is now carried applying the methodology for optimal policy in Sections 2 and 3 to the estimated VSTAR in Section 4. In particular, the evaluation focuses on: the gains from the optimization of policy, the economy and policy response to shocks under the optimal policy, and counterfactual simulations to compare actual data dynamics against the optimal ones. Once optimal rules are computed, macroeconomic dynamics are simulated using the VSTAR under the optimal policy in (12). This is used to derive a reference framework against which to compare the dynamics observed from the data, that reflect the actual monetary policy undertaken.

The reference framework for the optimal policy is not unique. It depends on the optimization protocol followed by the Fed (whether commitment or discretion); the constraints faced by Fed; the preferences of the Fed in terms of objective function and targets; and the number of available instruments. The analysis here considers solutions under commitment for joint optimization (optimal coordination) of both the federal funds rate and

²⁴ Appendix F in Supplementary Material describes in details the algorithm for computing the IRFs.

²⁵ These sign restrictions are also motivated by the observation that the sample correlation coefficients of the spread with p , ip , and ug are -0.278 , -0.607 and 0.468 , respectively.

²⁶ QE1 lasted 17 months. The Fed initially purchased private securities, consisting of \$500 billion in agency mortgage-backed securities (MBS) and \$100 billion in agency debt from Fannie Mae, Freddie Mac and Federal Home Loan Banks. Later the Fed purchased an additional \$750 billion of MBS, \$100 billion in agency debt, and \$300 billion in Treasuries. QE2 lasted for 8 months. During this phase, the Fed purchased a total of \$767 billion of Treasuries. See Kuttner (2018), for further details.

the QE instruments. Thus the economy vector is specified as $\mathbf{x}'_t = [s_t \ p_t \ ip_t \ ug_t]$ and the policy vector as $\mathbf{u}_t = [TS_t \ PS_t \ R_t]$. As described in Section 2.2, the coefficients from the estimated VSTAR are used to recover the VSTARX under the timing assumption A2. Thus, the constraint faced by the Fed is equivalent to that in (6)-(7).

The preferences of the Fed are specified in terms of minimization of a discounted intertemporal quadratic function that includes two main components. The first penalizes deviations of the inflation rate and the unemployment gap from their respective targets. This is consistent with the Fed dual mandate of price stability and maximum employment. The second penalizes changes in the monetary policy instruments. This allows to control for different degrees of gradualism in the use of each policy instrument, mimicking for example the gradual movements of the federal funds rate observed in several periods over 1979–2018, the sluggish evolution of the QE variables before the Great Recession, and gradual normalization programs often undertaken by central banks in advanced economies. It also ensures that the paths of the economy and policy instruments under the optimal rules do not deviate too much from those observed in the data, thus containing the quantitative importance of Lucas critique on the results.²⁷

Given the above, the Fed's objective function is specified as:

$$V_0 = E_0(1 - \beta) \sum_{t=0}^{\infty} \beta^t \left[v_p (p_t - \bar{p})^2 + v_{ug} (ug_t - \bar{ug})^2 + v_{\Delta TS} \Delta TS_t^2 + v_{\Delta PS} \Delta PS_t^2 + v_{\Delta R} \Delta R_t^2 \right], \quad (15)$$

where \bar{p} and \bar{ug} denote inflation and unemployment gap targets, respectively; v_z , $z = p, ug, \Delta TS, \Delta PS$ and ΔR , denote weights on the stabilization of inflation, unemployment gap and the three policy instruments; Δ is the lag operator.²⁸

The parameters of the objective function are calibrated as follows. The discount factor β is set to 0.996, corresponding to an annual rate of interest of 4.83 %, equal to the sample average for R_t . The results are not significantly affected by other reasonable values of the discount factor. The unemployment target is set as $\bar{ug} = 0$, since the unemployment gap is measured as the deviation of the actual rate of unemployment from the natural rate, and the Fed's mandate includes full employment. The calibration of the remaining parameters (\bar{p} , v_p , v_{ug} , $v_{\Delta TS}$, $v_{\Delta PS}$, and $v_{\Delta R}$) is less clear cut as it is reasonable to expect that the Fed preferences in terms of inflation target and weights on policy instruments may have changed during 1979–2018. To establish a benchmark the inflation target is set equal to the sample average over the simulation horizon and all policy weights are given equal value, $v_p = v_{ug} = v_{\Delta TS} = v_{\Delta PS} = v_{\Delta R} = 1$. When necessary, the robustness of the quantitative results that follow is evaluated against other reasonable choices of these parameters.

5.1 Gain from Optimization

What would have been the gain in terms of macroeconomic stabilization if QE had been optimally coordinated with interest rate policy? The answer to this question is first quantified considering macroeconomic dynamics in the 10 years from November 2007, one year before the Fed announcement of QE1, to November 2017, the date from which the Fed's balance sheet is allowed to shrink gradually.²⁹ Accordingly, the policy analysis is based on data that exclude the pre-2008 policy regime change, thereby limiting the quantitative impact of the Lucas critique on the results. Following Sack (2000), the gain from monetary policy optimization is evaluated comparing two scenarios. The first, referred to as the *actual policy*, is based on the observed dynamics of the economy and policy instruments over the considered sub-period. The second, the *optimal policy*, is constructed assuming that the Fed chooses the optimal combination of conventional and QE policy in each month and this optimal

²⁷ Appendix G in Supplementary Material shows that control for the volatility of QE in the welfare loss function can be micro-founded from a New Keynesian model.

²⁸ The welfare loss function in equation (15) is obtained from the general welfare loss in equation (8) through appropriate calibration of the matrix \mathbf{Q} .

²⁹ The starting date is chosen to account for possible aggregate responses in the economy before the announcement of the first QE program, due to the private sector's anticipation of future policy intervention. The results do not change significantly if the starting date for the optimization is moved closer to that of the Fed's announcement, November 2008.

policy is implemented given the true state of the economy in that month. The economy response to the optimal policy in any given month is computed including the actual shocks occurring in that month, using the VSTAR in (12). In addition, the inflation target in the loss function (15) is set as $\bar{p} = 1.55$, corresponding to the average inflation rate from November 2007 onward. As a further reference, a *no-policy* scenario is also simulated.³⁰ This is based on the implied economy and policy dynamics obtained if no-change in QE had been implemented since November 2007.

Following Dennis and Söderström (2006), two measures of the optimization gain are computed. The first is the (percentage) change in the loss due to implementation of a new policy, V' , relative to the loss measured from another policy, V , i.e. $G = 100 \times [1 - V'/V]$. The second is the unemployment-equivalent compensation. To clarify this, consider the unemployment gap term in (15). This can be written as $v_{ug}ug_t^2 = v_{ug}(u_t - \bar{u}_t)^2$, where u_t and \bar{u}_t denote the actual and the natural rate of unemployment, respectively. The unemployment-equivalent compensation is defined as the permanent deviation $\hat{u} \neq 0$ of the actual unemployment rate from the natural rate that results in a change of the loss by $(1 - \beta) \sum_{t=0}^{\infty} \beta^t v_{ug} \hat{u}^2$ such that V equals V' . It follows that $\hat{u} = \sqrt{(V - V')/v_{ug}}$. Thus G is a direct measure of the macroeconomic stabilization differential between two policies, whereas \hat{u} quantifies the implied permanent change in the rate of unemployment.

Table 1 presents the results. Columns two to six report the volatility of the five items in the loss function (15); column seven the implied losses; whereas the last two columns report the calculated measures of the optimization gain. For robustness, six different specifications of the weights in (15) are considered. The first (*Baseline*) is the benchmark which gives equal weight to each item in (15). The remaining five evaluate the effects of lower weight to the stabilization of either inflation (*II*), unemployment (*III*), all policy instruments (*IV*), the two QE instruments (*V*), or the federal funds rate (*VI*). For each set of policy weights the loss is computed as the undiscounted weighted average of the five terms in (15).³¹ The stabilization gains and unemployment-equivalent compensations are measured first to compare monetary policy optimization relative to the actual policy undertaken by the Fed, and then the actual policy relative to the no-QE scenario.

Two main results emerge from Table 1. First, policy optimization would have increased the stabilization gain during 2007–2017. On average across different weight specifications, optimal policy would have reduced macroeconomic volatility by about 2 % relative to the actual policy. In terms of unemployment-equivalent compensation, policy optimization would have resulted on average in a further reduction of the unemployment rate of about 0.33 % compared to the actual policy. Second, the gains from switching from the actual to the optimal policy are not as large as those achieved by the actual implementation of QE relative to the no-QE policy scenario. On average across the different weight specifications, macroeconomic volatility and the unemployment rate would have been about 12 and 1 % higher, respectively, had QE not been implemented during 2007–2017.

Looking across the different specifications of the weights, it can be observed that the gains from either the actual or the optimal policy increase when the Fed targets unemployment stabilization more aggressively than inflation stabilization (compare case II with either I or III). This is because unemployment is relative more volatile than inflation during 2007–2017, in turn implying that assigning a higher weight to this variable brings larger gains. Reducing the weight attached to any policy instrument in the objective function produces higher stabilization gains (compare case I with cases IV, V and VI). This is because the further reduction in the volatility of inflation and unemployment more than compensates the increase in the volatility of the instruments.

5.2 Impulse Response Functions

How the economy and the policy instruments respond to shocks at the ZLB once changes to QE and the federal funds rate are optimally coordinated? To answer this question, the IRFs are calculated from the VSTAR under

³⁰ Counterfactual analyses based on a no-policy scenario are frequent in the recent literature on the macroeconomic effects of QE, see, e.g., Lenza et al. (2010), Kapetanios et al. (2012), Dahlhaus et al. (2018).

³¹ For β close to 1, the discounted quadratic loss is approximately equal to the weighted average of the squared deviations from target of the variables included in it.

the optimal policy in (12) assuming that the economy and the policy instruments (\mathbf{y}_t) start from the ZLB period.³² The preferences of the Fed are set as under the baseline calibration and the inflation target is set equal to the ZLB sample average of 1.29 %.

Figure 4 shows the responses from the VSTAR under the optimal policy rules (blue) to a credit shock during the ZLB period. For comparison, the corresponding responses from the estimated VSTAR, which reflect the dynamics under the actual Fed monetary policy are also reported (red).³³ According to these results, when the federal funds rate is constrained by the ZLB, the Fed responds under the optimal policy to a credit shock by increasing the size of the balance sheet more sharply than under the actual policy, at least in the short run (first 3 months). At the same time there is significant portfolio rebalancing towards Treasury securities, since *TS* holdings increase while *PS* holdings decrease in the short run horizon. The response of the Fed's asset portfolio under the optimal policy show less persistence than those from the actual policy. This is because the optimal policy dampens very rapidly the impact of the shock on inflation and unemployment.

To illustrate how the optimal policy changes in response to different macroeconomic conditions, Figure 5 shows the responses under the actual (red) and optimal (blue) policy to a supply-type shock during the ZLB.³⁴ The federal funds rate remains constrained by the ZLB under both policies. The responses of QE are however different, being expansionary under the actual policy and contractionary under the optimal. The optimal QE policy response to the supply-type shock shows little rebalancing of the Fed portfolio contrary to what is observed from the responses to the credit shock displayed above.

In summary, QE under the optimal policy results in increase of the size of the balance sheet and significant short-run portfolio rebalancing toward *TS* in response to a credit-type shock. The responses to the supply-type shock are expansionary under the actual policy and contractionary under the optimal, with little evidence of portfolio rebalancing.

5.3 Counterfactual Monetary Policy

What would have been the dynamics of the economy and policy instruments had the federal funds rate and QE been coordinated optimally since 2008? To answer this question, the VSTAR under the optimal policy is used for three counterfactual experiments. The simulations are carried out starting from November 2007, using the specification of the Fed preferences with equal weights.

Each counterfactual simulation employs a different value for the inflation target, 2, 3 and 4 %. These do not reflect any specific view on the actual preferences of the Fed, but are selected for gauging indication on the sensitivity of the results to the choice of the inflation target.³⁵ It is important to clarify the role of the inflation target in these simulations. All other things equal, the optimal policy would require expansionary (contractionary) QE whenever actual inflation is below (above) target. As the unemployment gap is positive during most of the post-2007 period, the response of monetary policy to unemployment stabilization is stronger the lower is actual inflation relative to the target.

Figure 6 illustrates the results from these counterfactuals. Black lines (A) show the actual policy instruments dynamics as observed from data; whereas blue (O2), red (O3) and green (O4) lines denote the dynamics of the same instruments simulated from the optimal VSTAR under the three alternative inflation targets of 2, 3 and 4 %, respectively. The simulations are carried out assuming that in each period the Fed implements the optimal policy given the observed state of the economy.

³² As described in Appendix F in Supplementary Material, the IRFs under the optimal policy are computed using simulation from the VSTAR in (12). The history for the ZLB period is again the average of \mathbf{y}_t during 2009:1–2015:12.

³³ The responses for the policy instruments under the actual policy are the same as those in Figure 3.

³⁴ As described in Appendix F2 in Supplementary Material, the supply-type shock is constructed such that both the unemployment gap and inflation increase at the one month horizon.

³⁵ The 2 % inflation target was officially announced by the Fed as a long-run objective in January 2012. Ball (2014) argues in favour of a 4 % inflation target. The 3 % target is used as an intermediate case.

Three main results can be observed from Figure 6. First, the dynamics of the economy are fairly similar across the three counterfactual simulations. Compared to the data, the main differences are that credit risk would have been lower and inflation higher during 2008–2013 under the optimal policy regardless of the choice of the inflation target. The dynamics of industrial production and the unemployment gap are remarkably similar to those observed from the data, under any of the three counterfactual simulations. Second, the path of the federal funds rate appears close to that observed from the data under any of the counterfactual simulations, particularly when using the 4 % inflation target. Third, the dynamics of the Fed's balance sheet are qualitatively similar to the observed ones under any of the three counterfactual scenarios. The most noticeable exception is the period 2009–2015 period, when the size of the simulated Fed's balance sheet is larger than the actual one. This is due to the more rapid increase in *TS* purchases under any of the counterfactual relative to the observed purchases.

Table 2 gives a breakdown of the duration of the ZLB period under the actual and optimal policy, based on the counterfactual analysis. In the data, the federal funds rate never reaches zero and the identification of the periods in which the ZLB is binding is somewhat arbitrary, depending on the chosen threshold below which the policy rate is deemed to be at the ZLB. To guard against this, the number of ZLB periods in Table 2 is counted considering four different thresholds, corresponding to months in which the federal funds rate is below to either 25, 50, 75 or 100 basis points. According to the results in the table, the actual duration of the ZLB ranges between 85 and 104 months, depending on the selected interest rate threshold. The duration of the ZLB from the counterfactual experiments is closer to the observed duration as the inflation target and the cut-off rate for the ZLB increase. It is worth highlighting that the estimated duration of the ZLB period is considerably longer than that otherwise predicted by conventional evaluations of monetary policy. For example, Williamson (2015) estimates that the ZLB should have terminated around the beginning of 2011 (26 months) had monetary policy during the Great Recession being conducted according to a standard Taylor rule.

The counterfactual simulations can be used to evaluate different segments of the QE intervention undertaken by the Fed since November 2007. The Fed embarked in three large-scale asset purchase programs, commonly known as QE1, QE2, QE3 and the Maturity Extension Program (MEP).³⁶ Table 3 reports how *TS* holding, *PS* holding and the total size of the Fed's balance sheet change during each of these programs under the actual policy and the three counterfactual simulations. Under the actual QE policy, the total increase in the balance sheet was about 17.96 % of GDP. The size of the change of the Fed's balance sheet as a percentage of GDP under the optimal policy is close to the actual one, ranging between 17.98 and 19.33 % depending on the inflation target. Comparing across programs, QE1 should have been almost a third larger, regardless of the inflation target. Consequently, this would have resulted in QE2 and QE3 programs of smaller size under the optimal policy compared to the actual. It would have also resulted in a reduction of the size of the balance sheet (relative to GDP) during the MEP under the optimal policy about four times larger than the actual reduction observed in the data.

As noted above, the counterfactual simulations under the optimal policy entail only small changes from the actual data. To provide some indication of the extent of these small changes, Table 4 reports in the top row the mean of the absolute deviation from the actual data of the macroeconomic forecasts from the Survey of Professional Forecasters of the Federal Reserve Bank of Philadelphia during 2007:11-2018:10, whereas the remaining rows report the corresponding deviations for the simulated data used for the results in Figure 6 and Table 1.³⁷ The table does not include *TS* and *PS* since the Survey covers all macroeconomic data used in the empirical analysis except for the balance sheet. The results in Table 4 provide some reassurance with regard to the closeness of the simulated data to the actual data. The counterfactual simulations involve smaller deviations from the actual data compared to the professional forecasts, for inflation, industrial production and unemployment, whereas the deviations for the spread the federal funds rate are marginally larger.

³⁶ The MEP consisted of purchases of long-term Treasuries, accompanied by the sale of the same quantity of short term securities, leaving the overall size of the balance sheet unchanged.

³⁷ The Professional Forecasts database reports for each indicator the 25th and 75th percentile of the professional forecasts in the current and next three quarters. The data reported in the first row of numbers in Table 4 refer to the mean of the two percentiles computed from the next quarter forecasts. These are interpolated to compute the implied next month forecasts.

Table 3: Change in the Fed's asset portfolio in percentage of GDP under actual and optimal policy over different QE phases.

	QE1	QE2	MEP	QE3	Total	QE1	QE2	MEP	QE3	Total
	Actual policy					Optimal policy, $\bar{\pi}=2$				
<i>TS</i>	2.02	5.10	-0.48	3.69		4.04	3.52	-0.66	1.37	
<i>PS</i>	5.35	-1.75	-0.08	4.12		7.37	-1.83	-1.51	5.68	
Total	7.37	3.35	-0.56	7.81	17.96	11.42	1.68	-2.16	7.04	17.98
	Optimal policy, $\bar{\pi}=3$					Optimal policy, $\bar{\pi}=4$				
<i>TS</i>	3.97	3.54	-0.71	0.96		3.89	3.56	-0.76	0.55	
<i>PS</i>	8.06	-1.82	-1.63	6.30		8.74	-1.81	-1.76	6.92	
Total	12.02	1.72	-2.34	7.26	18.66	12.63	1.76	-2.52	7.47	19.33

Source: Author's calculations. See main text for more details.

Table 4: Mean of the absolute deviation from actual data of professional forecasts and simulated data from optimal policy, November 2007 – October 2018.

	<i>s</i>	<i>p</i>	<i>ip</i>	<i>ug</i>	<i>R</i>
	Professional forecasts				
	0.178	0.680	2.752	0.123	0.117
	Optimal policy				
$\bar{\pi} = 2$	0.590	0.180	1.785	0.048	0.323
$\bar{\pi} = 3$	0.562	0.356	2.108	0.083	0.313
$\bar{\pi} = 4$	0.592	0.604	2.447	0.139	0.377
1.55					
(Baseline)	0.630	0.163	1.785	0.046	0.359
(Weights II)	0.673	0.164	1.813	0.048	0.380
(Weights III)	0.442	0.162	1.292	0.038	0.288
(Weights VI)	0.844	0.191	2.401	0.058	0.449
(Weights V)	0.777	0.164	2.276	0.052	0.364
(Weights VI)	0.686	0.200	1.893	0.052	0.446

Source: Author's calculations. See main text for more details.

In summary, according to this analysis the scale of total QE intervention observed in the data appears to be close to that prescribed by the optimal QE policy under conventional calibration of the Fed preferences. The main differences between actual and optimal QE are in terms of timing and composition, since the optimal QE would have entailed earlier increases in the Fed's holding of *TS* compared to the actual policy. This would have resulted in lower spread and higher inflation in the aftermath of the Great Recession. Overall the counterfactual simulations indicate that the observed duration of the ZLB is in line with that otherwise prescribed by the optimal policy under most calibrations.

5.4 Alternative Optimization

The results presented above are the outcomes of two important requirements of the proposed optimization methodology: transformation of the VSTAR into a VSTARX using A2 and optimization with the SDC method. To evaluate the performance of the proposed optimization methodology, the results are compared with those from three possible alternative approaches.³⁸ The first, labelled as LIN, is based on the linearized-equivalent of the VSTAR, i.e. the VAR model estimated over the same sample period, whose performance against the VSTAR is

³⁸ Appendix H in Supplementary Material gives a detailed description of each of these three approaches.

highlighted in Appendix E in Supplementary Material. As it is a reduced-form model, for the purpose of this exercise the VAR is transformed into a VARX using A2. The optimization problem then boils down to a LQR which is solved with standard Riccati equation iteration. The second and third approach both employ the VSTARX obtained from the estimated VSTAR under A1. For this reason optimization can only be undertaken with nonlinear methods. In the second, labelled as MPC, this is pursued using a method that mimics the logic of the model predictive control algorithm, whose application in a number of nonlinear decision problems in economics has been recently illustrated by Grüne et al. (2015). This consists of positing a time-varying coefficients feedback rule $\mathbf{u}_t = \mathbf{k}_t^{(MPC)} + \mathbf{K}_t^{(MPC)} \mathbf{y}_{t-1}$ in conjunction with the system of equation (4) and then using the quasi-Newton method to find the coefficients $\mathbf{k}_t^{(MPC)}$ and $\mathbf{K}_t^{(MPC)}$ such that the one period ahead forecast of \mathbf{y}_t minimizes the loss (15). In the third approach, labelled as DNS, optimization is pursued using direct numerical search over the policy rule coefficients. This consists of using the time-varying coefficients feedback rule $\mathbf{u}_t = \mathbf{k}_t^{(DNS)} + \mathbf{K}_t^{(DNS)} \mathbf{x}_{t-1}$ to derive estimates of $\hat{\mathbf{u}}_t$, and then forming $\mathbf{y}_t = [\mathbf{x}_t' \hat{\mathbf{u}}_t']'$. Again the quasi-Newton method is employed to find the coefficients $\mathbf{k}_t^{(DNS)}$ and $\mathbf{K}_t^{(DNS)}$ that minimize the loss (15). The results from LIN highlight the effects of ironing out (through linearization) all the nonlinear features in the data prior to optimization. The results from MPC and DNS highlight instead the effects of doing optimization under the alternative transformation protocol, A1, therefore not being able to use dynamic programming but having to rely upon numerical search. For the sake of comparison, the stabilization gains are re-computed under each method and the quality of the performance is evaluated considering three dimensions: feasibility (ability to find an optimal solution), efficiency (ability to find the solution with the lowest loss) and computational speed (time required to find a solution).

Table 5 gives a summary of the outcomes of these numerical experiments, considering optimization for the post-2007 period and the full sample. For each approach and sample, the optimization is carried using the protocol and each of the six weight specifications described in Section 5.1. The first eighth columns of Table 5 report the average across all specifications of each argument in the loss function (15), the value of the loss V , the stabilization gain G and the unemployment-equivalent compensation \hat{u} , respectively. The final column reports an estimate of the speed of optimization (in seconds) obtained under the specification of the loss function with equal weights from 10 optimization trials under each approach. To facilitate comparison, the results from the actual and optimal policy under the SDC method are also reported.

According to the results in Table 5, each optimization approach can deliver a feasible solution, i.e. a solution that leads to a loss lower than under the actual policy. Regarding efficiency, the results from these numerical

Table 5: Average gains under alternative optimization methods.

	$(p_t - \bar{p})^2$	ug_t^2	ΔTS_t^2	ΔPS_t^2	ΔR_t^2	V	G	\hat{u}	Speed (sec.)
Sample: 2008:8–2018:9									
Actual policy	1.24	5.46	0.02	0.05	0.38	6.51			
Optimal policy									
SDC	1.03	5.32	0.06	0.28	0.41	6.40	1.71	0.33	0.494
LIN	1.02	5.10	0.05	0.05	0.32	5.93	8.94	0.80	0.503
MPC	1.14	5.44	0.08	0.05	0.38	6.45	0.93	0.25	6.580
DNS	1.14	5.44	0.08	0.05	0.38	6.45	1.01	0.27	4.399
Sample: 1979:8–2018:9									
Actual policy	4.75	2.66	0.32	0.02	0.11	7.16			
Optimal policy									
SDC	3.40	2.38	0.36	0.14	0.20	5.82	18.42	1.22	1.539
LIN	3.91	2.48	0.31	0.02	0.19	6.24	12.58	1.01	1.768
MPC	4.62	2.64	0.30	0.02	0.11	7.01	2.06	0.41	21.592
DNS	4.62	2.64	0.46	0.02	0.11	7.14	0.365	0.12	15.351

Source: Author's calculations. See main text for more details.

experiments suggest that SDC and LIN provide superior performance compared to the two numerical methods, as indicated from the lower losses (and therefore stabilization gains and unemployment compensations) that these two approaches deliver regardless of the sample choice. The comparison between SDC and LIN needs careful interpretation. This is because the analysis of optimal policy with LIN is based on the linear VAR model, which gives a very poor fit of the data dynamics compared to the VSTAR which is used by SDC, as shown in Appendix in Supplementary Material, Table E.2. Further, the performance of LIN relative to SDC deteriorates as the simulation horizon is extended to cover the whole sample, thereby exacerbating the importance of nonlinearity in the data. This points towards the importance of accounting for nonlinearity in the quantitative analysis of monetary policy, highlighting the difference it makes using a VAR rather than a nonlinear model like the VSTAR. Finally, optimization with SDC and LIN is way much faster than with the two numerical methods. This is not surprising since iteration until convergence of the Riccati equation in SDC and LIN is known to be faster than numerical search.

6 Conclusions

Macroeconomic data display many nonlinear features (asymmetries, thresholds and large swings) that often make linear models inadequate for reliable quantitative analysis. VSTAR is a popular tool employed in econometrics and applied macroeconomics for inference, structural analysis and forecasting when data are nonlinear. This paper uses VSTAR for a different task, the analysis of optimal policy.

At least two possible VSTARX can be employed to calculate optimal policy from a reduced-form VSTAR. One of these makes the analysis particularly tractable, since it preserves certainty equivalence and ensures low dimensionality of the decision making problem. The optimization employs SDC factorization to transform the VSTARX into a VAR with SDC matrices. This makes the NLQR problem isomorphic the LQR. Consequently, a nonlinear optimization problem can be solved with the same dynamic programming techniques used for LQR.

The methodology is applied to revisit the coordination between interest rate monetary policy and QE in the United States since 2008 and evaluate its effects on wider economy. To this end, the empirical analysis first establishes an estimated VSTAR that can be regarded as a reliable starting point for the subsequent policy analysis. According to the results found in the empirical analysis, the quantity of assets purchased the Fed and the duration of the ZLB in the decade after the onset of the Great Recession are not too dissimilar from those otherwise prescribed by the optimal policy. While the ZLB duration prescribed under the optimal policy is in line with the actual duration, there are however differences in terms of the timing and composition of QE. It is shown that the proposed methodology compares favourably, in terms of feasibility, efficiency and computational speed, with three alternative optimization approaches based on linear and nonlinear methods.

References

- Atanasova, C. 2003. "Credit Market Imperfections and Business Cycle Dynamics: A Nonlinear Approach." *Studies in Nonlinear Dynamics & Econometrics* 7 (4). <https://doi.org/10.2202/1558-3708.1112>.
- Auerbach, A. J., and Y. Gorodnichenko. 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy* 4 (2): 1–27.
- Ball, L. M. 2014. *The Case for a Long-run Inflation Target of Four Percent*. IMF Working Papers 2014, 092. Washington, D.C.: International Monetary Fund.
- Barnichon, R., and G. Mesters. 2023. "A Sufficient Statistics Approach for Macro Policy." *American Economic Review* 113 (11): 2809–45.
- Beeler, S. C., H. T. Tran, and H. T. Banks. 2000. "Feedback Control Methodologies for Nonlinear Systems." *Journal of Optimization Theory and Applications* 107 (1): 1–33. <https://doi.org/10.1023/a:1004607114958>.
- Beraja, M. 2023. "A Semistructural Methodology for Policy Counterfactuals." *Journal of Political Economy* 131 (1): 190–201.
- Chari, V. V., P. J. Kehoe, and E. R. McGrattan. 2009. "New Keynesian Models: Not yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics* 1 (1): 242–66.
- Dahlhaus, T. 2017. "Conventional Monetary Policy Transmission During Financial Crises: An Empirical Analysis." *Journal of Applied Econometrics* 32 (2): 401–21.

- Dahlhaus, T., K. Hess, and A. Reza. 2018. "International Transmission Channels of US Quantitative Easing: Evidence from Canada." *Journal of Money, Credit and Banking* 50 (2-3): 545–63.
- Del Negro, M., G. Eggertsson, A. Ferrero, and N. Kiyotaki. 2017. "The Great Escape? A Quantitative Evaluation of the Fed's Liquidity Facilities." *American Economic Review* 107 (3): 824–57.
- Dennis, R., and U. Söderström. 2006. "How Important is Precommitment for Monetary Policy?" *Journal of Money, Credit and Banking*: 847–72. <https://doi.org/10.1353/mcb.2006.0054>.
- Dijk, D. V., T. Teräsvirta, and P. H. Franses. 2002. "Smooth Transition Autoregressive Models - A Survey of Recent Developments." *Econometric Reviews* 21 (1): 1–47.
- Galvão, A. B., and M. T. Owyang. 2018. "Financial Stress Regimes and the Macroeconomy." *Journal of Money, Credit and Banking* 50 (7): 1479–505.
- Gambacorta, L., B. Hofmann, and G. Peersman. 2014. "The Effectiveness of Unconventional Monetary Policy at the Zero Lower Bound: A Cross-Country Analysis." *Journal of Money, Credit and Banking* 46 (4): 615–42.
- Gefang, D., and R. Strachan. 2009. "Nonlinear Impacts of International Business Cycles on the UK - a Bayesian Smooth Transition VAR Approach." *Studies in Nonlinear Dynamics & Econometrics* 14 (1). <https://doi.org/10.2202/1558-3708.1677>.
- Granger, C. W., and T. Teräsvirta. 1993. *Modelling Non-Linear Economic Relationships*. Oxford, UK: Oxford University Press.
- Grüne, L., W. Semmler, and M. Stieler. 2015. "Using Nonlinear Model Predictive Control for Dynamic Decision Problems in Economics." *Journal of Economic Dynamics and Control* 60: 112–33.
- Guerrieri, L., and M. Iacoviello. 2017. "Collateral Constraints and Macroeconomic Asymmetries." *Journal of Monetary Economics* 90: 28–49.
- Hurn, S., N. Johnson, A. Silvennoinen, and T. Teräsvirta. 2022. "Transition from the Taylor Rule to the Zero Lower Bound." *Studies in Nonlinear Dynamics & Econometrics* 26 (5): 635–47.
- Joyce, M., D. Miles, A. Scott, and D. Vayanos. 2012. "Quantitative Easing and Unconventional Monetary Policy - an Introduction." *The Economic Journal* 122 (564): F271–88.
- Kaiser, E., J. N. Kutz, and S. L. Brunton. 2017. "Data-Driven Discovery of Koopman Eigenfunctions for Control." arXiv preprint arXiv:1707.01146.
- Kapetanios, G., H. Mumtaz, I. Stevens, and K. Theodoridis. 2012. "Assessing the Economy-Wide Effects of Quantitative Easing." *The Economic Journal* 122 (564): F316–47.
- Koop, G., M. H. Pesaran, and S. M. Potter. 1996. "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (1): 119–47.
- Kuttner, K. N. 2018. "Outside the Box: Unconventional Monetary Policy in the Great Recession and Beyond." *Journal of Economic Perspectives* 32 (4): 121–46.
- Lanne, M., and P. Saikkonen. 2005. "Non-Linear GARCH Models for Highly Persistent Volatility." *The Econometrics Journal* 8 (2): 251–76.
- Lenza, M., H. Pill, and L. Reichlin. 2010. "Monetary Policy in Exceptional Times." *Economic Policy* 25 (62): 295–339.
- Liu, P., K. Theodoridis, H. Mumtaz, and F. Zanetti. 2019. "Changing Macroeconomic Dynamics at the Zero Lower Bound." *Journal of Business & Economic Statistics* 37 (3): 391–404.
- Ljungqvist, L., and T. J. Sargent. 2018. *Recursive Macroeconomic Theory*. Cambridge, Massachusetts: MIT press.
- Lucas, R. E. 1976. "Econometric Policy Evaluation: A Critique." In *Carnegie-Rochester Conference Series on Public Policy*, Vol. 1, 19–46. North-Holland. <https://www.sciencedirect.com/science/article/pii/S0167223176800036>.
- Luukkonen, R., P. Saikkonen, and T. Teräsvirta. 1988. "Testing Linearity Against Smooth Transition Autoregressive Models." *Biometrika* 75 (3): 491–9.
- Martin, V., S. Hurn, and D. Harris. 2013. *Econometric Modelling with Time Series: Specification, Estimation and Testing*. Cambridge, UK: Cambridge University Press.
- McKay, A., and C. K. Wolf. 2023. "What Can time-series Regressions Tell us About Policy Counterfactuals?" *Econometrica* 91 (5): 1695–725.
- Pearson, J. 1962. "Approximation Methods in Optimal Control I. Sub-Optimal Control." *International Journal of Electronics* 13 (5): 453–69.
- Polito, V., and P. Spencer. 2016. "Optimal Control of Heteroscedastic Macroeconomic Models." *Journal of Applied Econometrics* 31 (7): 1430–44.
- Polito, V., and M. Wickens. 2012. "Optimal Monetary Policy Using an Unrestricted VAR." *Journal of Applied Econometrics* 27 (4): 525–53.
- Ramey, V. A., and S. Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy* 126 (2): 850–901.
- Rudebusch, G. D. 2005. "Assessing the Lucas Critique in Monetary Policy Models." *Journal of Money, Credit and Banking* 37 (2): 245–72.
- Sack, B. 2000. "Does the Fed Act Gradually? A VAR Analysis." *Journal of Monetary Economics* 46 (1): 229–56.
- Sims, E. R., and J. C. Wu. 2020. "Are QE and Conventional Monetary Policy Substitutable?" *International Journal of Central Banking* 16 (1): 195–230.
- Sims, E., and J. C. Wu. 2021. "Evaluating Central Banks' Tool Kit: Past, Present, and Future." *Journal of Monetary Economics* 118: 135–60.
- Sims, C. A., and T. Zha. 2006. "Were There Regime Switches in US Monetary Policy?" *American Economic Review* 96 (1): 54–81.
- Sims, E., J. C. Wu, and J. Zhang. 2023. "The Four-Equation New Keynesian Model." *Review of Economics and Statistics* 105 (4): 931–47.
- Smets, F., and R. Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.
- Stock, J. H., and M. W. Watson. 2001. "Vector Autoregressions." *Journal of Economic Perspectives* 15 (4): 101–15.

- Terasvirta, T., and H. M. Anderson. 1992. "Characterizing Nonlinearities in Business Cycles Using Smooth Transition Autoregressive Models." *Journal of Applied Econometrics* 7 (S1): S119–S136.
- Wernli, A., and G. Cook. 1975. "Suboptimal Control for the Nonlinear Quadratic Regulator Problem." *Automatica* 11 (1): 75–84.
- Williamson, S. 2015. "The Road to Normal: New Directions in Monetary Policy." *Annual Report*: 6–23.

Supplementary Material: This article contains supplementary material (<https://doi.org/10.1515/snde-2024-0100>).