

Drautzburg, Thorsten; Fernández-Villaverde, Jesús; Guerrón-Quintana, Pablo;
Oosthuizen, Dick

Working Paper

Filtering with Limited Information

CESifo Working Paper, No. 11243

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Drautzburg, Thorsten; Fernández-Villaverde, Jesús; Guerrón-Quintana, Pablo;
Oosthuizen, Dick (2024) : Filtering with Limited Information, CESifo Working Paper, No. 11243,
CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/302728>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Filtering with Limited Information

*Thorsten Drautzburg, Jesús Fernández-Villaverde, Pablo Guerrón-Quintana,
Dick Oosthuizen*

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Filtering with Limited Information

Abstract

We propose a new tool to filter non-linear dynamic models that does not require the researcher to specify the model fully and can be implemented without solving the model. If two conditions are satisfied, we can use a flexible statistical model and a known measurement equation to back out the hidden states of the dynamic model. The first condition is that the state is sufficiently volatile or persistent to be recoverable. The second condition requires the possibly non-linear measurement to be sufficiently smooth and to map uniquely to the state absent measurement error. We illustrate the method through various simulation studies and an empirical application to a sudden stops model applied to Mexican data.

JEL-Codes: C320, C530, E370, E440, O110.

Keywords: filtering, limited information, non-linear model, dynamic equilibrium model, sudden stops.

Thorsten Drautzburg
Federal Reserve Bank of Philadelphia
Philadelphia / PA / USA
tdrautzburg@gmail.com

Jesús Fernández-Villaverde
University of Pennsylvania
Philadelphia / PA / USA
jesusfv@econ.upenn.edu

Pablo Guerrón-Quintana
Boston College & ESPOL / MA / USA
pguerron@gmail.com

Dick Oosthuizen
University of Pennsylvania
Philadelphia / PA / USA
dickoosthuizen@gmail.com

July 19, 2024

We thank Ed Herbst, Eva Janssens, and Felipe Meza for their useful comments. The views expressed here are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or its Board of Governors. Any errors or omissions are the responsibility of the authors.

1 Introduction

“[You can] do something without having to do everything.” Lars Peter Hansen, 2014.

Economists are often interested in filtering dynamic models, that is, in backing out the states and the shocks that drive the equilibrium path of the economy. There are at least four reasons for this interest. First, the path of shocks (or their distribution) is a reality check of the model. Suppose the shocks required to fit the data are unlikely (i.e., their probability is very low) or implausible (i.e., they contradict other sources of evidence or narrative analytics). In that case, a researcher may consider moving in an alternative direction. Second, economists are interested in historical decompositions: evaluating how much a shock contributed to the observed dynamics. In that way, we can build counterfactuals. Third, knowing the current state of the economy, including the shocks, is necessary when forecasts or optimal policies are state-dependent. Fourth, once you have recovered the shocks conditional on some parameter values, evaluating the likelihood of the model associated with these parameter values is straightforward. This opens the door to the structural estimation of the model either through maximum likelihood or the Bayesian approach.

Filtering a generic, fully specified dynamic equilibrium model can be accomplished with a non-linear filter, such as the sequential Monte Carlo filter described in [Herbst and Schorfheide \(2016, ch. 10\)](#) and [Fernández-Villaverde et al. \(2016\)](#). Unfortunately, this approach suffers from three drawbacks. First, a sequential Monte Carlo filter requires fully specifying the model, including aspects that are not central to the paper, such as auxiliary functional forms, an issue about which various researchers may disagree. Worse, some models’ findings might be fragile with respect to these auxiliary assumptions ([Canova and Ferroni, 2022](#)). Second, a sequential Monte Carlo filter requires solving the full model to specify the transition equations required in the simulation. This can be computationally costly. And, third, the filter suffers from the curse of dimensionality, making it difficult to apply to large models.¹ While new filters like the one in [Farmer \(2021\)](#) are most promising, they still face the same challenges of fully specifying the model, solving it, and a curse of dimensionality.

In this paper, we tackle the problem of non-linear filtering from a different perspective. More concretely, we follow the dictum in [Hansen \(2014, p. 950\)](#): “[Hansen \(1982\)](#) builds on a long tradition in econometrics of ‘doing something without having to do everything.’” This entails the study of partially specified models, that is, models in which only a subset of

¹For the simpler, linear normal case, we can apply the Kalman filter. Some of the concerns (i.e., the curse of dimensionality) are less relevant to this situation. Others still hold (i.e., the need to specify auxiliary parametric forms at least to first and second order). As we will show later, a misspecified linear Kalman filter can yield misleading estimates of the hidden state in our non-linear examples.

economic relations are formally delineated.” In the spirit of partial specification, we refer to our approach as the “partial (information) filter.”

We are not the first to push this idea. [Andreasen et al. \(2018\)](#) present a method for estimating models non-linearly using the population moments generated by a perturbation solution of the model. This approach, however, still requires solving the model with perturbation (which can sometimes be hard when the model has non-differentiabilities such as occasionally binding constraints) and specifying the moments of the shocks that enter into the computation of the model population moments. An early application of the partial filter we present here, published in [Drautzburg et al. \(2021\)](#), filters the shock process in an economy solved using the method of [Andreasen et al. \(2018\)](#). [Gallant et al. \(2017\)](#) also pursue a similar idea of Bayesian estimation of state space models via moment conditions with a partially specified measurement equation. While these authors focus on parameter estimates, we focus on filtering the hidden state variables.

The key advantage of our filtering method is the ease of implementation. Specifically, our approach approximates the state dynamics with a flexible statistical model, such as a vector autoregression (VAR), possibly with non-linear terms. This VAR serves as a plug-in estimate of the model expectations.² Given the statistical model, we can filter out an estimate of the hidden state from the measurement equation. While the VAR is a purely auxiliary model, we obtain the full estimates as a fixed point of the VAR estimation and the filtering stage.

The partial filtering approach requires two main conditions to be satisfied. First, the measurement equation needs to identify the hidden state. This assumption could be violated, for example, with an observation equation that is quadratic in the hidden state when the state space is unrestricted. However, in most applications, natural restrictions on the state space emerge, for example, that prices and quantities are non-negative. Second, the researcher requires a consistent estimator of the model expectations. Using simulations, we show that this condition can be relaxed: our procedure also works in the presence of misspecification in the form of unknown measurement error as long as the measurement error is relatively small or the hidden state is sufficiently persistent.

Using VARs to approximate model expectations is not a new idea. For example, [Campbell \(1991\)](#), [Bernanke and Kuttner \(2005\)](#), and [Chahrour et al. \(2021\)](#) use VARs to approximate expected returns. Closer to us, [Sbordone \(2002\)](#) uses VAR-based expectations to stand in for a firm’s expectations in a New Keynesian price setting equation. Our innovation is that, under regularity conditions, we can use a VAR approximation also for filtering hidden state variables.

²While we have focused on time-invariant VARs, a tractable way to allow for time variation has been proposed by [Petrova \(2019\)](#). Allowing for time variation is a straightforward extension of our approach.

After presenting some theoretical results, we investigate the partial filter’s performance with a battery of Monte Carlo experiments. These exercises confirm that, in many situations, the partial filter provides excellent estimates of the objects of interest for many researchers (e.g., shocks and unobserved states).

Next, we apply the partial filter to a sudden stops model of a small open economy. The model, postulated by [Mendoza \(2010\)](#), introduces a loan-to-value constraint in the small open economy real business cycle model. Because of its inherent non-linearities, this model is challenging to solve. For example, with the *fixed-point iteration (FiPIt)* algorithm developed in [Mendoza and Villalvazo \(2020\)](#), it takes 115 seconds to compute the solution for one set of parameter values.³ Adding a full-information filter, such as a sequential Monte Carlo filter, to this solution pushes the computation time further (depending on the details of the data, around another minute or two). In comparison, our partial information filter takes 47 seconds. While the difference between three/four minutes and 47 seconds might not look large, we may need to repeat the computation for tens of thousands of different parameter values (for example, if we are exploring the fit of the model to the data). More importantly, our sudden stops model is simple; larger non-linear models may take hours to compute, while our partial filter model is much more robust to the curse of dimensionality.

We investigate the performance of our partial filter in this sudden stops model with a Monte Carlo experiment, where we generate pseudo-true samples from the model. After demonstrating that the partial filter works well in this Monte Carlo experiment, we apply the filter to Mexican data from 1980Q1 to 2018Q4. The partial filter identifies well the different financial crises that hit Mexico.

While our application is a model with rational expectations, our framework easily allows for parametric departures from rational expectations, for example, by including the appropriate change of measure in the filtering equation. Think, for instance, about the case of cognitive discounting in [Gabaix \(2020\)](#), where the deviation of the forecast from the balanced growth path is shrunk by a factor that is constant for a given forecast horizon. We could address other departures from rational expectations, such as limited information, by limiting the information set we use to compute the approximating VAR.

Since at least [Roberts \(1995\)](#), a growing literature has used survey data for assessing macroeconomic models or for estimating them. See, for example, [Coibion et al. \(2021\)](#). However, surveys usually only include information about objects readily understood by individuals, such as inflation, but not about structural objects, such as shocks or co-state variables. Our partial filter helps us to learn about these structural objects and their evolution.

³All run times throughout the paper were measured on a computer with an Intel Core i7-8700 @ 3.200GHz with six cores and 12 logical processors, 32.0 GB memory, Windows 10, and Matlab 2021a.

The rest of the paper is structured as follows. Section 2 sets up the general environment and discusses the necessary conditions that ensure that the filter is consistent. To fix ideas, Section 3 analyzes a simple univariate example that features different non-linearities in the state or measurement equations. Section 4 uses a simulation study of a sudden stops model. Section 5 applies the partial filter to the same model using Mexican data. Section 6 concludes.

2 The partial information filter

This section presents the partial information filter. First, we introduce the general environment. Second, we describe our filter. Third, we postulate some regularity conditions required by our procedure. Finally, we outline some extensions.

2.1 Environment

Let us consider a dynamic equilibrium model whose equilibrium conditions, or a subset thereof, are described by an $n_x \times 1$ vector-valued function \mathbf{f} :

$$\mathbf{f}(\mathbf{x}_t, \mathbf{y}_t, \mathbb{E}_t[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)]) = \mathbf{0}, \quad (2.1)$$

where $\mathbf{y}_t = [y_{i,t}]_{i=1}^{n_y}$ is an $n_y \times 1$ vector of observables at period t , \mathbf{x}_t is an $n_x \times 1$ vector of hidden states that we are interested in backing out, \mathbb{E}_t is the conditional expectations operator, and $\mathbf{0}$ is a functional zero. The function $\mathbf{f}(\cdot)$ will usually stack optimality conditions for the agents in the model, policy rules by the government, budget and resource constraints, and stochastic processes for exogenous states. The function $\mathbf{g}(\cdot)$ deals with the part of the model involving future states and observables. Typically, \mathbf{x}_t is only a subset of the entire state vector \mathbf{s}_t that determines the model expectations. The functions $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ are conditional on a vector of states other than \mathbf{x}_t , denoted $\mathbf{s}_{-,t}$, but we omit this conditioning to save notation when no confusion arises.

We highlight three points about equation (2.1). First, $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ do not need to include all equilibrium conditions required for a full solution of the model. Our goal is to accomplish all we need in a particular application by dealing with only a few of them (and, consequently, only a subset of states). Second, at the cost of heavier notation, we could handle cases where some of the states are observed, but nothing of importance is lost by omitting that situation. Similarly, by including lagged values of the states and observables in vectors \mathbf{y}_t and \mathbf{x}_t , we can deal with general timing and stochastic structures. Third, we let \mathbf{x}_t and \mathbf{y}_t enter as arguments of the function $\mathbf{g}(\cdot)$ to encompass situations where the conditional expectations

depend explicitly on them. For instance, in a model with habit persistence, the expected marginal utility tomorrow depends on consumption today.

In what follows, we proceed under the assumption that the filtering problem (2.1) has a unique solution. This high-level assumption could be satisfied because either \mathbf{f} and \mathbf{g} have appropriate properties –such as linearity– or the researcher knows that the hidden states live in a domain $\mathbb{X}_t \subseteq \mathbb{R}^{n_x}$ where the solution is unique. Formally:

Assumption 1. *For some known domain \mathbb{X}_t and given \mathbf{y}_t and $\mathbf{s}_{-,t}$, there is a unique solution $\tilde{\mathbf{x}} \in \mathbb{X}_t$ such that $\mathbf{f}(\tilde{\mathbf{x}}, \mathbf{y}_t, \mathbb{E}[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t) | \tilde{\mathbf{x}}, \mathbf{s}_{-,t}]) = \mathbf{0}$.*

The following scalar example is useful to fix ideas:

Example 1. *Let the model be:*

$$y_t = \mu_y + x_t + \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}^2] \quad (2.2a)$$

$$x_t = \mu_x + \left(\rho_1 + \frac{\eta}{1 + x_{t-1}^2} \right) x_{t-1} + u_t, \quad (2.2b)$$

$$u_t \sim iid(0, \sigma_u^2), \quad (2.2c)$$

We consider two special cases.

(a) *Linear model: $\kappa_2 = \eta = 0$.*

Here $f(\cdot) = y - \mu_y - x - \kappa_1 \mathbb{E}[x'|x]$. Under the population expectation, $\mathbb{E}[x'|x] = \mu_x + \rho_1 x$ so that $f(\cdot) = y - \mu_y - (1 + \kappa_1 \rho_1)x - \kappa_1 \mu_x = 0$. Thus, there is a unique solution in \mathbb{R} given by $x = \frac{y - \mu_y - \kappa_1 \mu_x}{1 + \kappa_1 \rho_1}$ as long as Assumption 1 holds, which implies that $\kappa_1 \rho_1 \neq -1$.

(b) *Linear law of motion and quadratic observations equation: $\eta = 0$.*

Here,

$$y_t = \mu_y + x_t + \kappa_1 \mu_x + \kappa_1 \rho_1 x_t + 2\kappa_2 \rho_1 \mu_x x_t + \kappa_2 (\mu_x^2 + \sigma_u^2) + \kappa_2 \rho_1^2 x_t^2.$$

We require the observation to be monotonically increasing or decreasing in the hidden state ($\frac{\partial y_t}{\partial x_t} \gtrless 0$), which is equivalent to $x_t \gtrless -\frac{1}{2} \frac{1 + \kappa_1 \rho_1 + 2\kappa_2 \rho_1 \mu_x}{\rho_1^2 \kappa_2} \equiv \underline{x}$. Thus, for a given κ_2 , if the observation equation is monotonically increasing, it has a unique solution for x in $\mathbb{X} = (\underline{x}, \infty)$.

If the equations in $\mathbf{f}(\cdot)$ do not involve expectations (i.e., $\mathbf{g}(\cdot)$ is a dummy argument of $\mathbf{f}(\cdot)$), we can back out \mathbf{x}_t directly from (2.1) under Assumption 1. For example, given some observations on input factors, one may extract total factor productivity growth from the production function as in Fernald (2012). Also, in some situations, $\mathbb{E}_t[\mathbf{g}(\cdot)]$ might come from survey data. For example, we may use firms' surveyed inflation expectations in the Phillips

curve when filtering for cost-push shocks. However, survey expectations about structural objects (e.g., hidden state variables) are generally unavailable.

Hence, we will deal more often with situations where $\mathbf{g}(\cdot)$ is nontrivial, $\mathbb{E}_t[\mathbf{g}(\cdot)]$ is not readily available, and we need to back out the hidden state vector by finding the sequence of $\hat{\mathbf{x}}_t$ that solves the equation:

$$\mathbf{f}(\hat{\mathbf{x}}_t, \mathbf{y}_t, \hat{\mathbb{E}}_t[\mathbf{g}(\hat{\mathbf{x}}_{t+1}, \mathbf{y}_{t+1}, \hat{\mathbf{x}}_t, \mathbf{y}_t)]) = \mathbf{0}, \quad (2.3)$$

where $\hat{\mathbb{E}}_t(\cdot)$ is an estimate of the conditional expectation.

To estimate $\hat{\mathbb{E}}_t(\cdot)$, we will proceed in two steps. In the first step, we factorize each of the elements in the m -dimensional function $\mathbf{g}(\cdot)$ as:

$$\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i \equiv g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t) \times g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t). \quad (2.4)$$

Because $g_{i,1}(\cdot)$ could be a constant function, this decomposition is without loss of generality. However, in many applications, we are interested in equilibrium conditions featuring expected discounted values (e.g., in Bellman equations), and we may want to deal with different parts of these expectations separately. For instance, if $\mathbf{g}(\cdot) = \Lambda_{t+1}\Pi_{t+1}$, where Λ_{t+1} is the discount factor, many models allow to substitute the expectation of Λ_{t+1} for the risk-free rate r_t : $\mathbb{E}_t[\Lambda_{t+1}] = r_t^{-1}$. Also, the same model may admit different factorizations of $\mathbf{g}(\cdot)$. Choosing one depends on their suitability for the economic question being investigated or the available data (see [Drautzburg et al., 2021](#), for an example in which the decomposition 2.4 is applied).

Writing the terms in equation (2.4) in expected terms and using that, for any two scalar random variables (x_t, y_t) , we have that $\mathbb{E}_t[x_{t+1}y_{t+1}] = \mathbb{E}_t[x_{t+1}]\mathbb{E}_t[y_{t+1}] + \text{Cov}_t[x_{t+1}, y_{t+1}]$, we find:

$$\begin{aligned} \mathbb{E}_t[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] &\equiv \mathbb{E}_t[g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] \times \mathbb{E}_t[g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] \\ &\quad + \text{Cov}_t[g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i, g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i], \end{aligned} \quad (2.5)$$

where $i = 1, \dots, m$ denotes the rows of \mathbf{g} , \mathbf{g}_1 , and \mathbf{g}_2 that we require.

In the second step, we use an auxiliary statistical model and, depending on the application, additional equilibrium conditions to come up with an estimator for the covariances in equation (2.5). If some of the conditional first moments are also missing after using the equilibrium conditions, the auxiliary statistical model may help us estimate them. With these estimates in hand, we use equation (2.3) to back out the hidden state vector.

Next, we assume that the lack of explicit conditioning on \mathbf{s}_- in our estimates is asymptotically irrelevant:

Assumption 2. As $T \rightarrow \infty$, $\sup_{\mathbf{x}, \mathbf{y}, s_-} |\mathbb{E}[\mathbf{g}(\mathbf{x}', \mathbf{y}', \tilde{\mathbf{x}}, \mathbf{y}) | \tilde{\mathbf{x}}, s_-] - \widehat{\mathbb{E}}[\mathbf{g}(\mathbf{x}', \mathbf{y}', \tilde{\mathbf{x}}, \mathbf{y}) | \tilde{\mathbf{x}}, \mathbf{y}]| \xrightarrow{a.s.} 0$.

Least-squares estimators can satisfy this assumption in our univariate example.

Example 2. Consider again the model presented in (2.2).

(a) *Linear law of motion ($\eta = 0$) with bounded domain:*

Given $\eta = 0$, we have that $\mathbb{E}[g(x') | x, s_-] = \mathbb{E}[g(x') | x] = \text{cons} + (\kappa_1 \rho_1 + 2\kappa_2 \rho_1 \mu_x)x + \kappa_2 \rho_1^2 x^2$. To bound the domain, let $\mathbb{X} = [\underline{x}, \bar{x}]$ and $x \in \mathbb{X}$. Let $\hat{\rho}_1$ be a consistent estimator, such as the OLS estimator $\hat{\rho}_1 \xrightarrow{p} \rho_1$ as $T \rightarrow \infty$. Then:

$$\begin{aligned} & \sup_{x \in \mathbb{X}} |\mathbb{E}[g(x') | x, s_-] - \widehat{\mathbb{E}}[g(x') | x, y]| \\ &= \sup_{x \in \mathbb{X}} |(\kappa_1 \rho_1 + 2\kappa_2 \rho_1 \mu_x)x + \kappa_2 \rho_1^2 x^2 - (\kappa_1 \hat{\rho}_1 + 2\kappa_2 \hat{\rho}_1 \mu_x)x - \kappa_2 \hat{\rho}_1^2 x^2| \\ &\leq \sup_{x \in \mathbb{X}} |(\kappa_1 + 2\kappa_2 \mu_x)(\rho_1 - \hat{\rho}_1)| |x| + \sup_{x \in \mathbb{X}} |\kappa_2(\rho_1^2 - \hat{\rho}_1^2)| x^2 \\ &\leq |(\rho_1 - \hat{\rho}_1)| (\kappa_1 + 2\kappa_2 \mu_x) \max\{-\underline{x}, \bar{x}\} + |(\rho_1^2 - \hat{\rho}_1^2)| \kappa_2 \max\{\underline{x}^2, \bar{x}^2\}. \end{aligned}$$

Because of the continuous mapping theorem, both $\rho_1 - \hat{\rho}_1$ and $\rho_1^2 - \hat{\rho}_1^2$ converge to zero in probability. For any $\epsilon > 0$, $\lim_{T \rightarrow \infty} \Pr\{|\rho_1 - \hat{\rho}_1| > \epsilon\} = 0$ and $\lim_{T \rightarrow \infty} \Pr\{|\rho_1^2 - \hat{\rho}_1^2| > \epsilon\} = 0$. Thus, for any $\tilde{\epsilon} > 0$, we can redefine ϵ such that $\lim_{T \rightarrow \infty} \Pr\{|(\rho_1 - \hat{\rho}_1)| (\kappa_1 + 2\kappa_2 \mu_x) \max\{-\underline{x}, \bar{x}\} > \tilde{\epsilon}/2\} = 0$ and $\lim_{T \rightarrow \infty} \Pr\{|\kappa_2(\rho_1^2 - \hat{\rho}_1^2)| \max\{\underline{x}^2, \bar{x}^2\} > \tilde{\epsilon}/2\} = 0$. Hence, $\lim_{T \rightarrow \infty} \Pr\{|(\rho_1 - \hat{\rho}_1)| (\kappa_1 + 2\kappa_2 \mu_x) \max\{-\underline{x}, \bar{x}\} + |\kappa_2(\rho_1^2 - \hat{\rho}_1^2)| \max\{\underline{x}^2, \bar{x}^2\} > \tilde{\epsilon}\} = 0$ and, thus, $\sup_{x \in \mathbb{X}} \Pr\{|\mathbb{E}[g(x') | x, s_-] - \widehat{\mathbb{E}}[g(x') | x, y]| > \tilde{\epsilon}\} = 0$. Therefore, the expectation converges uniformly on \mathbb{X} .

(b) *Non-linear law of motion with linear observation equation ($\kappa_2 = 0$):*

We still have that $\mathbb{E}[g(x') | x, s_-] = h(x)$, where $h(x) = \kappa_1 \mu_x + \kappa_1 \left(\rho_1 + \frac{\eta}{1+x^2}\right) x$. [Chen and Christensen \(2015, Theorem 2.1\)](#) show that, under regularity conditions, a least-squares series estimator such as a B-spline converges uniformly to the true regression function $h(x)$ and, thus, to $\mathbb{E}[g(x') | x, s_-]$ in our setting.

Besides a uniformly consistent estimator for the expectation process, we also require that the functional equation is continuous in the relevant arguments.

Assumption 3. $\mathbf{f}(\cdot)$ is uniformly continuous in \mathbf{x} and $\mathbb{E}[\mathbf{g}(\cdot) | \cdot]$ and $\mathbb{E}[\mathbf{g}(\cdot) | \mathbf{x}, \cdot]$ is uniformly continuous in \mathbf{x} .

The uniform continuity of f and its arguments along with the uniqueness guarantee consistency of the partial filter.

Lemma 1. *Under Assumptions 1 through 3, if $\hat{\mathbf{x}}^T$ and $\hat{\mathbb{E}}[g(x', y', x, y)|x, y]$ converge with $\hat{\mathbf{x}}_t \in \mathbb{X}_t \forall t$, they converge to \mathbf{x}^T and $\mathbb{E}[g(x', y', x, y)|x, y]$ as $T \rightarrow \infty$.*

Proof. Suppose not. It must be that $\hat{\mathbb{E}}[\cdot]$ converges because otherwise Assumption 2 is violated. Thus, $\hat{\mathbf{x}}^T$ must diverge from \mathbf{x}^T . However, because the expectation converges, a diverging sequence of $\hat{\mathbf{x}}^T$ contradicts Assumption 1: For any $\epsilon > 0$, there exists a T_ϵ such that for all $T > T_\epsilon$ the approximation error is uniformly smaller than ϵ . Now, by uniform continuity (Assumption 3), that bounds the error in $\{\hat{x}_t\}$ by some δ_ϵ . We can pick ϵ to make δ_ϵ vanish. Thus, if there is convergence, it is to the population values. \square

For illustrative purposes, we model the covariance and the first moments in the next subsection using a VAR(1), possibly including non-linear terms. However, other statistical models are conceptually straightforward to use.

2.2 Algorithms

We specify now a VAR(1) in $g_{1,t}, g_{2,t}$ (or their non-degenerate components), (a subset of $n_{\tilde{y}}$ elements of) \mathbf{y}_t , and $\hat{\mathbf{x}}_t$ for $t = 1, \dots, T$ as our auxiliary statistical model. A superscript denotes the whole sequence of a variable. For instance, $\mathbf{y}^T = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T\}$. We collect all the VAR variables in $\boldsymbol{\xi}_t$, a $(2m + n_{\tilde{y}} + n_x) \times 1$ vector. By including the possibly non-linear functions $g_{\cdot,t}$ directly in the VAR, rather than its arguments, even a linear VAR yields an approximation to the expectation of the non-linear function itself. Some applications may, nevertheless, call for the inclusion of non-linear terms, which we collect in the vector $\boldsymbol{\phi}_t = [\phi_{j,t}]_j$, for example $\boldsymbol{\phi}_t = [x_{j,t}^2, y_{\ell,t}^2]$ for some indices j, ℓ to include selected lagged hidden states and observables as predictors. Then, the VAR is given by:

$$\boldsymbol{\xi}_t = \boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_{t-1} + \mathbf{B}\boldsymbol{\phi}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \text{Var}[\boldsymbol{\varepsilon}_t] = \boldsymbol{\Sigma}.$$

Using the companion form of the VAR, we can accommodate the general VAR case with more lags. While, given the structure of many time series, a low-dimensionality VAR will capture their dynamics well, one can easily adapt our steps to other auxiliary statistical models.

To express the expectation in terms of the VAR elements, order $g_{1,t}$ and $g_{2,t}$ as the first two variables of the VAR. Also, let \mathbf{e}_i be a selection vector with a 1 entry in position i and 0 elsewhere. Finally, use that the first $2m$ elements of the VAR characterize g_1 and g_2 . One can then write the components of (2.5) as:

$$\mathbb{E}_t[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)]_i \equiv \mathbf{e}'_i(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\phi}_t)\mathbf{e}'_{i+n_x}(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\phi}_t) + \mathbf{e}'_i\boldsymbol{\Sigma}\mathbf{e}_{i+n_x}. \quad (2.6)$$

With this in hand, the i -th element in equation 2.1 transforms into

$$\mathbf{f}(\mathbf{x}_t, \mathbf{y}_t, \mathbf{e}'_i(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\phi}_t)\mathbf{e}'_{i+n_x}(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\phi}_t) + \mathbf{e}'_i\boldsymbol{\Sigma}\mathbf{e}_{i+n_x})_i = 0. \quad (2.7)$$

Note that \mathbf{x}_t appears inside $\boldsymbol{\xi}_t$ and possibly inside $\boldsymbol{\phi}_t$.

We propose two approaches to filtering with the previous equation. First, we find the sequence $\hat{\mathbf{x}}^T$ that solves equation (2.7) for all i . Because $\hat{\mathbf{x}}_t$ appears in $\boldsymbol{\xi}$ (and possibly in $\boldsymbol{\phi}$) as well as indirectly in the VAR parameter matrices \mathbf{A} , \mathbf{B} , and $\boldsymbol{\Sigma}$, we search for a fixed point in equation (2.7). Second, we find $\hat{\mathbf{x}}^T$ through a Gibbs sampler that quantifies the estimation uncertainty.

Fixed point. We solve for the fixed point of $\{\hat{x}_t\}_{t=1}^T$ and $\{\hat{\mathbb{E}}_t[g(\hat{x}_{t+1}, y_{t+1}, \hat{x}_t, y_t)]\}_{t=1}^T$ numerically. In practice, we have found that initializing $\{\hat{x}_t\}_{t=1}^T$ at a simple starting guess, such as the steady state of our model of interest plus some added noise, and, then, iterating on VAR-parameter estimation and filtering until convergence works well.⁴

Gibbs sampler. A Gibbs sampler allows us to quantify estimation uncertainty. We start the Gibbs sampler at some appropriate guess, such as, for example, the steady state of the model or the fixed point found in the previous algorithm. Then, $d = 1, \dots, D$, we iterate on:

1. Given the previous draw $\{\hat{\mathbf{x}}_t^{(d)}\}_{t=1}^T$, draw the parameters $\boldsymbol{\mu}^{(d)}, \mathbf{A}^{(d)}, \mathbf{B}^{(d)}, \boldsymbol{\Sigma}^{(d)}$ from the posterior for a VAR in $\boldsymbol{\xi}_t^{(d)}$ with controls $\boldsymbol{\phi}_t^{(d)}$.
2. Given \mathbf{y}^T and the VAR parameters, solve for $\{\hat{\mathbf{x}}_t^{(d+1)}\}_{t=1}^T$ using (2.6) in equation (2.3).

In some applications, researchers may have a strong prior over the VAR parameters or even treat them as known. If the researcher wishes to evaluate a specific model that can be simulated, the researcher can estimate the VAR on simulated data. Furthermore, we can add an initial calibration step to either algorithm where we calibrate the first moments to be model-consistent by treating the (log) deviations of the observed series as (log) deviations from the model moments.

2.3 Extensions

In our empirical application, we use a model with an occasionally binding constraint. To ensure that the model multiplier correctly identifies periods with binding constraints, we

⁴The fixed point algorithm is easier to implement. However, in Monte Carlo experiments with misspecified and non-linear DGPs, we found that the fixed point algorithm does not lead to convergence for every sample path because small changes in state estimates can lead to discontinuous changes in parameter estimates and vice versa. The Gibbs sampler below had a more robust performance but required more effort.

introduce and subsequently solve for the time-varying intercepts for expected marginal utility, say, for the period with binding constraints.

Other extensions are conceptually straightforward. For example, one could weaken the identifying assumptions above or allow for *iid* measurement error in the filtering equation (2.1). In the latter case, equation (2.1) would hold only in expectation, and one may need to adjust the estimator for it. For instance, in the linear model with *iid* measurement error, the use of lagged filtered values as an instrument may correct for an attenuation bias. Below, we introduce measurement error in the simulation but treat it as a form of unknown model misspecification.

With measurement error and multiple equations identifying the hidden state, it might be possible to test whether the hidden state is identified consistently by the different equations. Lastly, when we use a misspecified, inconsistent estimator for $\hat{\mathbb{E}}[\mathbf{g}(\cdot)|\cdot]$ with a small but non-zero error, we can use additional assumptions on \mathbf{f} to bound the error on the filtered state.

2.4 Testing

The partial filter makes two key assumptions: First, to back out the hidden states, the model equations must be invertible in population. Second, the population expectations are well approximated by the statistical model. The first assumption might be difficult to assess in applications because the estimated state is constructed from current and past observables only, given parameters. Thus, there is no meaningful test to confirm that we do not need future information to recover the shock – which would make the hidden state only recoverable but not identifiable in the language of [Chahrour and Jurado \(2018\)](#). In applications, external information, such as competing estimates or narrative information, can be used as a joint test of the partial filter assumptions, as in the application in [Drautzburg et al. \(2021\)](#).

However, we can test whether the statistical model of $\mathbb{E}_t[g(x_{t+1}, y_{t+1}, x_t, y_t)]$ is adequate. This requires both the information set to be rich enough and the statistical model to be flexible enough. Our proposed test is to assess whether the VAR residuals for the observables, which include the residuals of $g(\hat{x}_{t+1}, y_{t+1}, \hat{x}_t, y_t)$, are uncorrelated with their lags and the lagged residuals for the other observables. If they exhibit significant correlation, we could improve the model of $\mathbb{E}_t[g(x_{t+1}, y_{t+1}, x_t, y_t)]$ by enlarging the information set or by making the model more flexible.

Specifically, we propose the multivariate [Ljung and Box \(1978\)](#) test statistic of residual autocorrelation and use a wild bootstrap to compute critical values, allowing for heteroskedasticity or stochastic volatility but imposing the Null of zero autocorrelation. As illustrated below, this test has the power to diagnose misspecification.

The univariate [Ljung and Box \(1978\)](#) test statistic for lag length k is:

$$LB_k = T(T+2) \sum_{\ell=1}^k \frac{\hat{r}_\ell^2}{T-\ell}, \quad (2.8)$$

where \hat{r}_ℓ is the ℓ th order autocorrelation of y_t . In multivariate settings, we use the following generalization discussed in [Tsay \(2010, chapter 8.1.4\)](#):

$$LB_k = T^2 \sum_{\ell=1}^k \frac{\text{vec}(\hat{\mathbf{R}}_\ell)' (\hat{\mathbf{R}}_0^{-1} \otimes \hat{\mathbf{R}}_0^{-1}) \text{vec} \hat{\mathbf{R}}_\ell}{T-\ell}, \quad (2.9)$$

where \mathbf{R}_ℓ is the ℓ th order autocorrelation matrix of the observables.

3 A scalar example

In this section, we explore the properties of the partial filter using a Monte Carlo exercise. We also consider possible misspecification in the form of measurement error and a higher-dimensional state vector to assess the robustness of the proposed filter.

As before, there is a single observable y_t that depends on a single hidden state variable x_t that the researcher tries to extract:

$$y_t = \mu_y + x_t + \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}^2] + \sigma_e e_t \quad (3.1a)$$

$$x_t = \mu_x + \left(\rho_1 + \frac{\eta}{1 + x_{t-1}^2} \right) x_{t-1} + \sigma_u u_t \quad (3.1b)$$

where $e_t \sim iid(0, 1)$ and $u_t \sim iid(0, 1)$. Here, e_t has the interpretation of measurement error, while u_t is a shock to the hidden state. In the appendix, we consider ARMA(p,1) processes for x_t as well as a non-linear forcing term for x_t .

We assume that the researcher considers the misspecified observation equation:

$$y_t = \mu_y + x_t + \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}^2]. \quad (3.2)$$

That is, when $\sigma_e > 0$, the model is misspecified. This misspecification vanishes as $\frac{\sigma_e}{\sigma_u} \searrow 0$, akin to the misspecification in [Schorfheide \(2005\)](#). Examples of an observation equation with such an expectation term include a forward-looking Phillips curve or a consumption Euler equation. The observation equation depends on the current state variable x_t and expectations of a function of the future state x_{t+1} , which is non-linear when $\kappa_2 \neq 0$. Since the system

is jointly homogeneous in means, coefficients, and standard deviations, we normalize the coefficient on the current state to unity: $\kappa_1 = 1$.

Following [Gordon et al. \(1993\)](#), the state equation (3.1b) is first-order Markov: either a linear AR(1) (if $\eta = 0$) or a non-linear process (if $\eta \neq 0$). In the appendix, we consider two variations of this equation: first, a variant that includes a cosine function $\gamma_0 \cos(\gamma_1 t)$ of calendar time t , where γ_0 controls the amplitude and γ_1 the frequency, and, second, a variant of the linear case with an ARMA(p,1) structure for x_t .

As stated above, our filtering approach involves three steps. First, we approximate the state dynamics with a (possibly non-linear) VAR. Second, we evaluate the expectations terms in the observation equation using the VAR approximation. Third, we back out the hidden state given the VAR dynamics and the observables. The VAR dynamics can be obtained by simulating a structural model or by estimation (in the simplest case, iterating between state inference and VAR estimation). We describe each step in more detail next.

Step 1. We approximate the dynamics of the hidden state x_t with a VAR(L) in the hidden state and the observable y_t , as well as an optional non-linear function $\phi(x_t, y_t)$. The function $\phi(\cdot)$ can be vector-valued and could include $g(\cdot)$ and other non-linear functions of x_t or y_t :

$$\begin{bmatrix} x_t \\ y_t \\ \phi(x_t, y_t) \end{bmatrix} \approx \mathbf{a}_0 + \sum_{\ell=1}^L \mathbf{A}_\ell \begin{bmatrix} x_{t-\ell} \\ y_{t-\ell} \\ \phi(x_{t-\ell}, y_{t-\ell}) \end{bmatrix} + \boldsymbol{\varepsilon}_t. \quad (3.3)$$

Step 2. We substitute for the expectation using the VAR. Defining $\phi(x_t, y_t) \equiv [x_t^2, x_t/(1+x_t^2)]'$ we have that:

$$\hat{\mathbb{E}}[g(x_{t+1}, y_{t+1}, x_t, y_t) | y^t] = \hat{\mathbb{E}}[\kappa_1 x_{t+1} + \kappa_2 x_{t+1} | y^t] = [\kappa_1, 0, \kappa_2, 0] \sum_{\ell=1}^L \mathbf{A}_\ell \begin{bmatrix} x_{t+1-\ell} \\ y_{t+1-\ell} \\ x_{t+1-\ell}^2 \\ x_{t+1-\ell}/(1+x_{t+1-\ell}^2) \end{bmatrix}. \quad (3.4)$$

Step 3. Plugging the above expression into equation (3.2) yields an equation that can be solved for x_t . We iterate over t , taking $x_{t-\ell}, \ell \geq 1$ as given and only solve for x_t . When $\phi(\cdot)$ includes a rational function or polynomial terms in x_t , the observation equation is a polynomial equation. The coefficients are functions of current and past y_t and x_t . The example above, with $g(x_t, y_t) = x_t^2$ and $\kappa_2 \neq 0$ and zero coefficients on $x_t/(1+x_t^2)$, generally

yields two solutions.⁵ In the case of $\kappa_2 = 0$ and non-zero coefficients on $x_t/(1 + x_t^2)$, the resulting equation can be re-written as a third-order polynomial, which generically has three solutions. While one could use statistical criteria to choose among solutions, in practice, we use economic intuition to choose among solutions, for example, choosing a solution that lies within a certain domain.

The point estimate of the partial filter is a fixed point of the VAR parameters and the state estimates. Below, we solve for this fixed point iteratively. The VAR parameters are estimated in the first step, given state estimates. The state estimates are then updated given the VAR parameter estimates. To quantify estimation uncertainty, this iterative procedure can incorporate a Gibbs sampler by drawing the VAR parameters from the posterior rather than using the point estimates.

To test the model, we add the following step:

Test. Compute the Ljung-Box statistic (2.8) with a lag order at least as high as that of the approximating VAR. We then use the *iid* wild bootstrap to generate D artificial samples from equation (3.3). For each artificial sample, we re-estimate the VAR and compute the Ljung-Box statistic. We reject the partial filter at significance level α if the realized Ljung-Box statistic lies above the $(1 - \alpha) \times 100$ percentile of the bootstrapped distribution.

A number of special cases of the system of equations (3.1) illustrate the challenges and advantages of the partial filter. We focus on the following three cases:

1. Fully linear case: If $\kappa_2 = \eta = 0$, the model is linear. In this case, the Kalman filter applies and is optimal when the disturbances are normal.
2. Quadratic observation equation, linear state equation: If $\eta = 0$, the state equation is linear. In this case, we restrict μ_x and bound the support of v_t to ensure that the model is identified. Otherwise, there are generally two sequences $\{x_t\}$ that are consistent with $\{y_t\}$ and the model is unidentified.
3. Linear observation equation, non-linear state equation: If $\kappa_2 = 0$, we have an observation equation that is linear in the forecast $\mathbb{E}_t[x_{t+1}]$, but we retain the full non-linearity of the state equation. In this case, even the law of motion is generally a correspondence that can have disconnected ergodic sets.

These scenarios demonstrate that our method works well if the misspecification is not severe, i.e., provided that the signal-to-noise ratio $\frac{\sigma_u}{\sigma_e}$ is high enough. In case 1, when the Kalman filter is applicable, we also provide a comparison with the linear Kalman filter.

⁵In the presence of misspecification or approximation error, these solutions need not be real. We then proceed by working with the real component of the solution.

Our Monte Carlo exercise uses $N = 100$ replications. We consider first-order Markov processes with low to high values of persistence $\rho_1 \in \{0.25, 0.5, 0.75, 0.9, 0.95\}$. We normalize the variance of the measurement error $\sigma_e^2 = 1$ and vary the conditional variance of the stochastic component of the hidden state on the following grid: $\{0.252, 1, 2^2, 5^2, 10^2\}$. We use a burn-in in the simulations of 100 periods and a sample of $T = 250$, a typical length for quarterly post-WWII applications. Except in the application with the non-linear observation equations, we assume zero means $\mu_y = \mu_x = 0$ and Gaussian distributions. We test for the adequate lag order of the approximating VAR. In Appendix A.3, we report results using an AR(2), an ARMA(2,1), and an AR(4) process. We show that the specification test correctly rejects lag lengths that are too short and that the partial filter produces reliably high correlations between the filtered state and the actual state.

3.1 Fully linear model

To illustrate the mechanics of the partial filter, we first consider a special case in which the researcher knows the model is correctly specified and imposes this knowledge at the estimation stage. We then consider the misspecified case. Last, we provide a comparison with the Kalman filter.

3.1.1 No misspecification ($\sigma_e = 0$)

We can rewrite the observation equation as:

$$y_t = (1 + \rho_1)x_t, \quad (3.5)$$

where x_t follows a zero mean AR(1) process:


$$x_t = \rho_1 x_{t-1} + v_t. \quad (3.6)$$

Due to the perfect collinearity of x_t and y_t , the partial filter is very simple. We use an AR(1), yielding $\widehat{\mathbb{E}}_t[x_{t+1}^{(d)}] = \widehat{A}_{1,1}[x_t^{(d)}] + \widehat{a}_1$, where \widehat{a}_1 is a constant term included in the AR(1) and that is zero in population. While we could update $\widehat{x}_t^{(1)} = (1 + \widehat{\rho}_1)^{-1}y_t$ given an estimate $\widehat{\rho}_1$ based on an initial guess of $\widehat{x}_t^{(0)} \propto y_t$, here we iterate until convergence on the following two steps for comparison with the more complex cases below:

For $d = 1, 2, \dots$ do until convergence of $\{x_t^{(d)}\}_{t=1}^T$

Step 1. $x_t^{(d)} = y_t - \widehat{A}_{11}x_t^{(d-1)} - \widehat{a}_1^{(d-1)}$

Step 2. Estimate AR(1) coefficients $\hat{A}^{(d)}, \hat{a}_0^{(d)}$ regressing $\begin{bmatrix} x_1^{(d)} \\ \vdots \\ x_T^{(d)} \end{bmatrix}$ on $\begin{bmatrix} x_0^{(d)} & 1 \\ \vdots & \\ x_{T-1}^{(d)} & 1 \end{bmatrix}$.



Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45-degree line.

Figure 1: Partial filter in fully linear case without measurement error: Filtered state \hat{x}_t vs. true x_t for 100 different $T = 250$ samples for varying degrees of persistence ρ_0 .

Figure 1 shows how the fit of the partial filter changes as a function of the persistence of the true state. Specifically, the left panel is a scatter plot when $\rho_1 = 0.25$, the center panel when $\rho_1 = 0.75$, and the right panel when $\rho_1 = 0.95$. Within each panel, the true state x_t is measured on the horizontal axis. Different colors correspond to the 100 different Monte Carlo samples. The dashed lines are the 45-degree lines.

All realizations are close to the 45-degree line. Comparing the left panel to the right panel shows that the filtered state becomes more concentrated and is relatively closer to the true state when the persistence is higher. Given these scatter plots, it is not surprising that the partial filter yields close to perfect correlations and that the relative standard deviations are just slightly above unity. See Appendix A.1 for details on the correlations and relative standard deviations.

3.2 Linear case with misspecification

Focusing on the zero-mean AR(1) as DGP, we can rewrite the observation equation as:

$$y_t = (1 + \rho_1)x_t + \sigma_e e_t. \quad (3.7)$$


The partial filter for this case is simple. In this version, the researcher is unaware of the stochastic term, that is, $\sigma_e > 0$. The estimated VAR is linear. When it has a single lag, the assumed observation equation is $\hat{\mathbb{E}}_t[x_{t+1}^{(d)}] = \hat{\mathbf{A}}_{1,o}[x_t^{(d)}, y_t] + \hat{a}_{0x}$. We iterate until convergence on the following two steps (ignoring the misspecification by treating e_t as zero):

For $d = 1, 2, \dots$ do until convergence of $\{x_t^{(d)}\}_t$

Step 1. $x_t^{(d)} = y_t - \hat{A}_{1x} x_t^{(d-1)} - \hat{A}_{1y}^{(d-1)} y_t - \hat{a}_{0x}^{(d-1)}$

Step 2. Estimate VAR coefficients $\hat{\mathbf{A}}^{(d)}, \hat{\mathbf{a}}_0^{(d)}$ regressing $\begin{bmatrix} y_1 & x_1^{(d)} \\ \vdots & \vdots \\ y_T & x_T^{(d)} \end{bmatrix}$ on $\begin{bmatrix} y_0 & x_0^{(d)} & 1 \\ \vdots & \vdots & \vdots \\ y_{T-1} & x_{T-1}^{(d)} & 1 \end{bmatrix}$.

Figure 2 shows scatter plots for low to high persistence of the state. The results are similar to the case without misspecification. All realizations are close to the 45-degree line. Comparing the left panel to the right panel shows that the filtered state is relatively closer to the true state when the persistence is relatively higher. Compared to the case without misspecification, the scatter plot exhibits more dispersion around the 45-degree line.



Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45-degree line.

Figure 2: Partial filter in the fully linear case with misspecification: Filtered state \hat{x}_t vs. true x_t for 100 different samples for varying degrees of persistence ρ_1 .

Testing the statistical model underlying the partial filter using the Ljung-Box test reveals that, at low levels of the signal-to-noise ratio $\frac{\sigma_u}{\sigma_e}$, a VAR(1) approximation is inadequate. Figure 3 illustrates this for the case of $\rho_1 = 0.75$ and $\sigma_u = \frac{\sigma_e}{4}$ in Panel (a) and $\sigma_u = 10\sigma_e$ in Panel (b). The figure reports the rejection rates by VAR lag length on the horizontal




Figure 3: Partial filter in the fully linear case: Rejection rate of the Ljung-Box test by VAR lag length and test lag order

axis and Ljung-Box lag order on the vertical axis. The shading corresponds to a heat map or rejection rates using p -values of 10% as the cutoff. If the approximating model were correctly specified, we would expect to see dark blue shades. Instead, for $\sigma_u = \frac{\sigma_e}{4}$, Figure 3(a) shows lighter colors corresponding to rejection rates between 38% and 45% for a VAR with one lag. With two or more lags, the rejection rates are close to or lower than 10%. In Panel (b), these rejection rates range from 7% to 12% and do not change much with the lag order of the approximating VAR. We conclude that a VAR(2) is necessary to handle situations with significant misspecification and proceed using this as our approximating model.

Table 1: Partial filter performance and comparison with Kalman filter: Different AR(1) DGPs: $x_t = \rho_1 x_{t-1} + \sigma_u u_t, u_t \sim iid(0, 1)$.

		Correlation			Relative standard deviation		
		Partial Filter	KF $\sigma_e = 0$	Correct KF	Partial Filter	KF $\sigma_e = 0$	Correct KF
$\rho_1 = 0.25$	$\sigma_u = \frac{1}{4}$	0.31 [0.23, 0.36]	0.31	0.32	4.03 [3.65, 4.80]	3.27	0.32
	$\sigma_u = 1$	0.79 [0.74, 0.80]	0.79	0.79	1.40 [1.23, 1.62]	1.26	0.79
	$\sigma_u = 10$	0.99 [0.97, 1.00]	1.00	1.00	1.02 [0.94, 1.12]	1.00	1.00
		Partial Filter	KF $\sigma_e = 0$	Correct KF	Partial Filter	KF $\sigma_e = 0$	Correct KF
$\rho_1 = 0.75$	$\sigma_u = \frac{1}{4}$	0.50 [0.44, 0.55]	0.55	0.65	2.87 [2.53, 3.38]	1.87	0.66
	$\sigma_u = 1$	0.92 [0.88, 0.93]	0.93	0.94	1.19 [1.10, 1.33]	1.07	0.94
	$\sigma_u = 10$	1.00 [0.99, 1.00]	1.00	1.00	1.01 [0.98, 1.07]	1.00	1.00
		Partial Filter	KF $\sigma_e = 0$	Correct KF	Partial Filter	KF $\sigma_e = 0$	Correct KF
$\rho_1 = 0.95$	$\sigma_u = \frac{1}{4}$	0.70 [0.62, 0.76]	0.79	0.89	1.73 [1.44, 2.12]	1.25	0.92
	$\sigma_u = 1$	0.98 [0.95, 0.99]	0.98	0.99	1.07 [1.02, 1.13]	1.02	0.99
	$\sigma_u = 10$	1.00 [1.00, 1.00]	1.00	1.00	1.01 [1.00, 1.04]	1.00	1.00

Note that σ_u in the table corresponds to σ_u/σ_e with $\sigma_e = 1$ in the simulations with misspecification. The partial filter uses a VAR(2).

Table 1 provides summary statistics for the partial filter that uses a VAR(2) approximation. It considers different levels of persistence and signal-to-noise ratios. For each parameter combination, the table reports the median correlation and median relative standard deviation of the filtered state with the true state, along with 68% confidence intervals across the Monte Carlo simulations. The table shows that, as the misspecification vanishes, the correlation approaches one from below, and the relative standard deviation approaches one from above. For example, for $\rho_1 = 0.25$, the median correlation (relative s.d.) across simulations is 0.31 when $\sigma_u = \frac{1}{4}\sigma_e$. This grows to 0.79 when $\sigma_u = \sigma_e$ and is 0.99 when $\sigma_u = 10\sigma_e$. The relative standard deviation falls quickly from 4.03 to 1.40 and 1.02 for these signal-to-noise ratios. The results improve with higher values of persistence. Even when $\sigma_u = \frac{\sigma_e}{4}$, the median correlation rises from 0.31 when $\rho_1 = 0.25$ to 0.70 when $\rho_1 = 0.95$ and the relative standard deviation falls from 4.03 with $\rho_1 = 0.25$ to 1.73 with $\rho_1 = 0.95$.

3.2.1 Kalman filter comparison

Intuitively, the partial filter under misspecification attributes the measurement error to the true state and, thus, overstates its volatility. We now show that a similar behavior arises when the Kalman filter is misspecified.

Using standard results, initialized with the steady state variance, the Kalman filter implies the following state variance $\text{Var}[x_{t+1}|y^t]$ and gain $G(\text{Var}[x_{t+1}|y^t])$:

$$\begin{aligned}\text{Var}[x_{t+1}|y^t] &= \frac{\sigma_u^2}{2} - \frac{1 - \rho_1^2}{(1 + \kappa_1 \rho_1)^2} \frac{\sigma_e^2}{2} + \sqrt{\sigma_u^2 \sigma_e^2 + \frac{1}{4} \left(\frac{1 - \rho_1^2}{(1 + \kappa_1 \rho_1)^2} \sigma_e^2 - \sigma_u^2 \right)^2} \\ &= \sigma_u^2 \quad \text{if } \sigma_e = 0. \\ G(\text{Var}[x_{t+1}|y^t]) &= \frac{(1 + \kappa_1 \rho_1) \text{Var}[x_{t+1}|y^t]}{\sigma_e^2 + (1 + \kappa_1 \rho_1)^2 \text{Var}[x_{t+1}|y^t]} \\ &= \frac{1}{1 + \kappa_1 \rho_1} \quad \text{if } \sigma_e = 0.\end{aligned}$$

This gain, in turn, implies the following recursive state estimate, which becomes static absent measurement error:

$$\begin{aligned}\mathbb{E}[x_t|y^t] &= \rho \mathbb{E}[x_{t-1}|y^{t-1}] + G(\text{Var}[x_{t+1}|y^t])(y_t - (1 + \kappa_1 \rho_1)G(\text{Var}[x_{t+1}|y^t])\rho \mathbb{E}[x_{t-1}|y^{t-1}]) \\ &= \frac{y_t}{1 + \kappa_1 \rho_1} \quad \text{if } \sigma_e = 0.\end{aligned}$$

Interestingly, the fixed point of the partial filter with a lag length of one period yields the same estimate absent measurement error. To see this, consider the true parameter values for the approximating VAR and write:⁶

$$\begin{aligned}\hat{x}_t &= y_t \frac{(1 - \kappa_1 A_{1y})}{1 + \kappa_1 A_{1x}} \\ &= \frac{y_t}{1 + \kappa_1 \rho_1} \quad \text{if } \sigma_e = 0.\end{aligned}$$

Table 1 above also lists the correlation and relative standard deviation from the misspecified Kalman filter (assuming $\sigma_e = 0$, “KF $\sigma_e = 0$ ”) and the correctly specified Kalman filter (“Correct KF”) alongside the partial filter results. Unlike the partial filter, the Kalman filter is given the true parameter values (except for σ_e in the misspecified case). Both the partial filter and the misspecified Kalman filter have similar correlations and relative standard deviations that approach unity from above as σ_u increases. In contrast, the correctly specified Kalman filter has a relative standard deviation that approaches unity from below because it

⁶In general, $A_{1x} \neq \rho_1$ and $A_{1y} \neq 0$ as the measurement error $\sigma_e e_t$ makes discovering x_t impossible.

attributes some variability in the data to the measurement error. When $\sigma_u = 10$, all three models yield correlations and relative standard deviations near unity.

3.3 Quadratic observation equation with misspecification

Next, we allow for a quadratic state equation. The measurement and state equations imply:

$$y_t = \mu_y + x_t + \kappa_1 \mathbb{E}_t[x_{t+1}] + \kappa_2 \mathbb{E}_t[x_{t+1}^2] + e_t \quad (3.8)$$

$$= \mu_y + x_t + \kappa_1 \mu_x + \kappa_1 \rho_1 x_t + 2\kappa_2 \rho_1 \mu_x x_t + \kappa_2 (\mu_x^2 + \sigma_u^2) + \kappa_2 \rho_1^2 x_t^2 + e_t. \quad (3.9)$$

The case of the quadratic observation equation illustrates an important requirement for the partial filter: the model has to be identified absent measurement error, i.e., if $e_t = 0$. Otherwise, the model is only set-identified with solutions generically spread across disjoint regions of the state space. Here, two different values of x_t are consistent with a given realization of y_t absent measurement error. In typical applications, economic insight can be used to rule out regions of the state space that give rise to multiplicity (i.e., values of monetary policy shocks that are wildly different from narrative evidence).

To ensure that our data generating process has a point-identifiable state, we restrict y_t to be increasing in x_t , as in Example 1(b). Since this is only possible with bounded shocks, we now assume that v_t follows a truncated normal distribution.⁷ Retracing the steps of Example 1(b), we require that:

$$\frac{\partial y_t}{\partial x_t} = 2\rho_1^2 \kappa_2 x_t + (1 + \kappa_1 \rho_1 + 2\kappa_2 \rho_1 \mu_x) \geq 0 \quad \Leftrightarrow \quad x_t \geq -\frac{1}{2} \frac{1 + \kappa_1 \rho_1 + 2\kappa_2 \rho_1 \mu_x}{\rho_1^2 \kappa_2}. \quad (3.10)$$

Given bounded shocks, x_t is bounded from below by $\frac{\mu_x + \underline{u}}{1 - \rho_1}$, where \underline{u} is the lower bound on the disturbances. Given $\rho_1 > 0$, the sufficient condition becomes a restriction on the unconditional mean of the state x_t :

$$\mu_x \geq \frac{\rho_0^2 (\kappa_1 - 2\underline{u} \kappa_2) - \kappa_1 \rho_0 + \rho_0 - 1}{2\kappa_2 \rho_0}. \quad (3.11)$$

We focus on a very non-linear DGP with $\kappa_1 = \frac{1}{10}$, $\kappa_2 = 10$, $\mu_y = 0$. We set μ_x so that equation (3.11) holds with equality. For comparison, we also consider an almost linear case with $\kappa_1 = 10$, $\kappa_2 = \frac{1}{10}$, and $\mu_y = 0$ to assess whether the extra flexibility impairs the filter. However, for most sample paths, the performance of the filter continues to be good.

In this setup, given the possible multiplicity of solutions, we need to choose an initial


⁷We truncated the normal distribution to $[-1, 1]$ and scale it to have the desired variance σ_u^2 .

guess with a reasonable scale. Here, we use the following starting guess that assumes that the state variable has a scale compatible with the observable and the known parameters:

$$\hat{x}_t^{(0)} = \frac{1}{\kappa_1 + \kappa_2} \frac{1}{\kappa_2^2 + \kappa_1^2} (\kappa_2^2 \sqrt{y_t} + \kappa_1^2 y_t) + 0.1 \times \mathcal{N}(0, 1). \quad (3.12)$$

We use a VAR with a single lag and a square term in the hidden state x_{t-1} . We initialize the VAR coefficients to have a persistence in the state of 0.9 and zero loadings otherwise.


We illustrate this DGP for various degrees of persistence, $\rho_1 \in \{0.25, 0.75, 0.95\}$. Since the non-linearity only enters via the expected future state, when this state is not persistent, the model is almost linear, as the left panel in Figure 4 shows. For the higher degrees of persistence, however, the non-linearity is noticeable: One gauge of this non-linearity is the relative size of the coefficients in a regression of y_t on the true x_t and x_t^2 – and the departure from the 45-degree line in the figures. For $\rho_1 = 0.25$, the non-linearity matters little: the OLS coefficient on x_t is 0.75 and the LS coefficient on x_t^2 is 0.45. The relative magnitudes are reversed when $\rho_1 = 0.75$ or $\rho_1 = 0.95$.



Note: The dashed line is the 45-degree line; the solid line is the fit of an OLS regression on x_t and x_t^2 . Very non-linear case with $\kappa_1 = \frac{1}{10}$, $\kappa_2 = 10$.

Figure 4: Quadratic observation equation for varying degrees of persistence ρ_1 . True state vs. observation with $T=25,000$ observations.

Figure 5 shows scatter plots of the true vs. filtered state for the near-linear DGP (right column) and the very non-linear DGP (left and center columns). For the very non-linear DGP, we show results for two levels of persistence ($\rho_1 = 0.25$ and $\rho_1 = 0.75$), while $\sigma_u = \sigma_e = 1$ remains constant. For comparison, we also show the near-linear DGP with $\rho_1 = 0.75$. Two features stand out. First, for the higher level of persistence, the partial filter delivers estimates close to the 90-degree line and, thus, close to the truth. Second, the truth is much more volatile than the filtered estimates when $\rho_1 = 0.25$. For almost all sample paths, the filter performs well for both the very non-linear and the near-linear specification.



Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45-degree line.

Figure 5: Partial filter with a quadratic observation equation and misspecification: Filtered state \hat{x}_t vs. true x_t for 100 different samples for varying degrees of persistence ρ_1 .

In this model, the partial filter provides adequate forecasts already with a VAR(1) approximation. Table A.1 in the appendix summarizes the corresponding Ljung-Box tests, again for the near-linear DGP and very non-linear DGP as well as for different levels of persistence ρ_1 in the rows and degrees of misspecification $\frac{\sigma_u}{\sigma_e}$. The absence of rejections of the VAR(1) approximation in this DGP with a quadratic observation equation contrasts with the linear DGP, which generated a noticeable amount of rejections with a VAR(1) approximation. One reason may be that the quadratic observation equation decreases the effective degree of misspecification as more of the observable y_t is driven by x_t .

Table 2: Performance: Quadratic observation equation with different AR(1) DGPs: $x_t = \rho_1 x_{t-1} + \sigma_u u_t$, $y(t) = x(t) + \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_t^2] + e_t$, $u_t, e_t \sim iid$.

		Correlation		Relative standard deviation	
		Near-linear	Very non-linear	Near-linear	Very non-linear
$\rho_1 = 0.25$	$\sigma_u = \frac{1}{4}$	0.67 [0.62, 0.70]	0.10 [0.04, 0.17]	0.61 [0.11, 0.65]	0.77 [0.74, 0.80]
	$\sigma_u = 1$	0.96 [0.96, 0.97]	0.82 [0.80, 0.84]	0.42 [0.41, 0.44]	0.12 [0.11, 0.12]
	$\sigma_u = 10$	1.00 [1.00, 1.00]	0.94 [0.93, 0.95]	0.40 [0.40, 0.41]	0.10 [0.09, 0.10]
		Near-linear	Very non-linear	Near-linear	Very non-linear
$\rho_1 = 0.75$	$\sigma_u = \frac{1}{4}$	0.95 [0.95, 0.96]	0.96 [0.95, 0.96]	0.99 [0.65, 1.02]	0.87 [0.85, 0.90]
	$\sigma_u = 1$	1.00 [1.00, 1.00]	1.00 [1.00, 1.00]	0.94 [0.94, 0.95]	0.79 [0.78, 0.81]
	$\sigma_u = 10$	1.00 [1.00, 1.00]	1.00 [1.00, 1.00]	0.94 [0.93, 0.96]	0.79 [0.77, 0.80]
		Near-linear	Very non-linear	Near-linear	Very non-linear
$\rho_1 = 0.95$	$\sigma_u = \frac{1}{4}$	0.99 [0.99, 0.99]	1.00 [1.00, 1.00]	1.15 [1.14, 1.16]	1.05 [1.05, 1.05]
	$\sigma_u = 1$	1.00 [1.00, 1.00]	1.00 [1.00, 1.00]	1.14 [1.13, 1.15]	1.05 [1.05, 1.05]
	$\sigma_u = 10$	0.99 [0.96, 1.00]	1.00 [1.00, 1.00]	1.12 [1.02, 1.15]	1.05 [1.05, 1.05]

The partial filter uses a VAR(1). Near-linear: $\kappa_1 = 10, \kappa_2 = 0.1$ Very non-linear: $\kappa_1 = 0.1, \kappa_2 = 10$.

Finally, Table 2 summarizes the performance of the partial filter in terms of correlations with the truth and relative standard deviations. The correlations are reported in the two left columns and the relative standard deviations in the two right columns. Confirming the visual evidence from the scatter plots, the partial filter does well except in the case of very low persistence. With $\rho_1 = 0.75$ and $\rho_1 = 0.95$ the correlations are essentially one and the relative standard deviations near one. For a low persistence of $\rho_1 = 0.25$, the partial filter struggles when misspecification is also pervasive ($\sigma_u = \frac{1}{4}\sigma_e$), and more so in the very non-linear case. However, with $\sigma_u = \sigma_e$ the median correlation is already 0.82 even in the very non-linear case. The 68% confidence intervals are tight, indicating that the sample that is a visual outlier in the scatter plots is, indeed, an outlier.


3.4 Linear observation equation, non-linear state equation with misspecification

We now consider a non-linear state equation with non-linear persistence. Specifically:

$$y_t = x_t + \kappa_1 \mathbb{E}_t[x_{t+1}] + e_t \quad (3.13a)$$

$$x_t = \left(\rho_1 + \frac{\eta}{1 + x_{t-1}^2} \right) x_{t-1} + v_t, \quad (3.13b)$$

where $e_t \sim iid(0, \sigma_e^2)$ and $v_t \sim iid(0, \sigma_u^2)$. Here, we set $\eta = 25$. As we illustrate below, these parameter values interact with ρ_1 to give rise to rich non-linearities.




Note: The dashed line is the 45-degree line; the solid linear is the fit of an OLS regression on x_t and x_t^2 .

Figure 6: Linear observation equation, non-linear state equation for varying degrees of persistence ρ_1 . True state vs. observation with $T=25,000$ observations.

Figure 6 illustrates the non-linear dynamics for three different degrees of persistence ρ_1 by plotting the current state x_t against the past state x_{t-1} . Even though $\rho_1 > 0$ in all three cases, when $\rho_1 = 0.25$ the non-linear mean-reversion via $\eta > 0$ dominates the autocorrelation

via ρ_1 . When $\rho_1 = 0.5$, the linear and non-linear force approximately offset each other, while the positive, linear autocorrelation dominates for $\rho_1 = 0.75$.



Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45-degree line.

Figure 7: Partial filter with a non-linear state equation: Filtered state \hat{x}_t vs. true x_t for 100 different samples for varying degrees of persistence ρ_1 .

Turning to the partial filter, we consider two approximating models. First, a linear VAR in terms of x_t, y_t . Second, a non-linear VAR that also includes $x_t/(1 + x_t^2)$ and its lags as observables. Anticipating the results of the specification tests below, each VAR uses two lags. Figure 7 summarizes the results using once more scatter plots of the filtered vs. true state. The left and center panels use a non-linear VAR approximation for the case of $\rho_1 = 0.25$ and $\rho_1 = 0.75$, respectively. The right panel uses a linear VAR approximation for $\rho_1 = 0.75$. The most salient feature of the scatter plots is that the sample paths of the true model lie in two disjoint regions. Importantly, for each of the 100 sample paths, the partial filter correctly identifies the right region of the state and produces estimates near the 90-degree line. The linear and non-linear VAR exhibit broadly similar performance, although the linear VAR has occasional sample paths away from the 90-degree line.

Table 3: Partial filter with non-linear state equation and linear VAR(p): Ljung-Box rejection rates with different lag lengths.

	Non-linear VAR(1)			Non-linear VAR(2)			Linear VAR(2)		
	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 10$	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 10$	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 10$
$\rho = 0.25$	0.20	0.27	0.20	0.17	0.23	0.03	0.55	0.25	0.04
$\rho = 0.5$	0.27	0.30	0.07	0.30	0.30	0.03	0.77	0.32	0.05
$\rho = 0.75$	0.17	0.20	0.20	0.20	0.17	0.20	0.83	0.23	0.04
$\rho = 0.9$	0.60	0.60	0.10	0.67	0.50	0.10	0.77	0.13	0.07
$\rho = 0.95$	1.00	0.67	0.17	0.83	0.63	0.10	0.64	0.15	0.04

Table 3 shows the results of the specification tests for the VAR with one or two lags. The results for the non-linear VAR are in the left panel (VAR(1) case) and the middle panel

(VAR(2)). The right panel shows the results for a linear VAR(2) approximation. Within each panel, the rows represent different degrees of persistence and the columns decreasing degrees of misspecification. Overall, the specification tests reject the models with high degrees of misspecification, and the VAR models with two lags tend to perform better than those with a single lag. For the linear VAR with $\sigma_u = 10\sigma_e$, all rejection rates are below 10%. With the non-linear VAR(2), they are at or below 10% for four out of five levels of persistence ρ_1 , and 20% for $\rho_1 = 0.75$.

Table 4: Partial filter performance: Non-linear law of motion with different DGPs.

		Linear VAR(2)		Non-linear VAR(2)	
		Correlation	Relative s.d.	Correlation	Relative s.d.
$\rho = 0.25$	$\sigma_u = \frac{1}{4}$	0.17 [0.10, 0.24]	3.66 [3.16, 4.23]	0.15 [0.07, 0.19]	3.46 [3.33, 3.82]
	$\sigma_u = 1$	0.52 [0.47, 0.56]	1.20 [1.06, 1.34]	0.54 [0.46, 0.58]	1.10 [1.01, 1.24]
	$\sigma_u = 10$	0.96 [0.95, 0.96]	0.96 [0.92, 1.01]	0.97 [0.97, 0.98]	0.98 [0.94, 1.01]
		Correlation	Relative s.d.	Correlation	Relative s.d.
$\rho = 0.75$	$\sigma_u = \frac{1}{4}$	0.35 [0.31, 0.39]	3.45 [3.24, 4.05]	0.38 [0.33, 0.44]	0.38 [0.33, 0.44]
	$\sigma_u = 1$	0.85 [0.83, 0.87]	1.29 [1.20, 1.37]	0.86 [0.84, 0.87]	0.86 [0.84, 0.87]
	$\sigma_u = 10$	0.99 [0.99, 0.99]	1.00 [0.98, 1.02]	1.00 [1.00, 1.00]	1.00 [1.00, 1.00]
		Correlation	Relative s.d.	Correlation	Relative s.d.
$\rho = 0.95$	$\sigma_u = \frac{1}{4}$	0.58 [0.53, 0.64]	1.92 [1.73, 2.19]	0.67 [0.67, 0.70]	2.06 [1.58, 2.11]
	$\sigma_u = 1$	0.96 [0.94, 0.97]	1.09 [1.04, 1.15]	0.97 [0.96, 0.98]	1.04 [1.03, 1.08]
	$\sigma_u = 10$	1.00 [1.00, 1.00]	1.00 [1.00, 1.02]	1.00 [1.00, 1.00]	1.01 [1.00, 1.02]

Table 4 summarizes the performance of the partial filter in terms of the correlation with the true state and the relative standard deviation. The left panels are for the linear VAR(2), while the right panels are for the non-linear VAR(2). In terms of correlations, both VARs perform similarly well. The correlations decrease with σ_u and are near one when the misspecification is small, independent of ρ_1 . However, when $\rho_1 = 0.25$ and misspecification is significant with $\sigma_u = \frac{1}{4}\sigma_e$, the median correlation is a mere 0.15 with the non-linear VAR, which rises to 0.67 when $\rho_1 = 0.95$. The relative standard deviation tends to be excessively high, with medians as high as 3.66, but approaches unity when the misspecification vanishes. Overall, the results improve with σ_u and the persistence ρ_1 .

4 Application I: A sudden stops model

To evaluate the quantitative performance of the partial information filter, we apply it to a popular non-linear economy: the sudden stops model of [Mendoza \(2010\)](#). The model

introduces a loan-to-value constraint in the workhorse small open economy real business cycle model.

4.1 Model structure

Time t is discrete and goes forever. The economy is inhabited by a representative firm and a representative household. The firm produces a tradable good y_t with a technology $y_t = e^{\epsilon_t^A} A k_t^\gamma L_t^\alpha v_t^\eta$, where k_t is capital, L_t is labor, v_t are imported intermediate goods, ϵ_t^A is a productivity shock, and $\gamma + \alpha + \eta = 1$. The tradable good is sold at a world-determined price (normalized as the numeraire). Imported inputs are purchased at an exogenous price $\log p_t = \log p + \epsilon_t^P$, where ϵ_t^P is a price shock. Working capital loans pay for a fraction ϕ of the cost of imported inputs and labor in advance of sales. These loans are obtained from foreign lenders at the beginning of each period and repaid at the end of the period. Lenders charge the world gross real interest rate $\log R_t = \log R + \epsilon_t^R$ on these loans, where ϵ_t^R is an interest rate shock. The shocks in the economy, $s_t = \{\epsilon_t^A, \epsilon_t^P, \epsilon_t^R\}$, follow a joint first-order Markov process with unconditional zero means.

The household has a utility function over sequences of consumption c_t and labor L_t :

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{\left(c_t - \frac{L_t^\omega}{\omega}\right)^{1-\sigma}}{1-\sigma},$$

where β is the discount factor, ω determines the elasticity of labor supply, and σ controls risk aversion. We follow the specification in [Mendoza and Villalvazo \(2020\)](#) with exogenous discounting instead of the case with endogenous discounting as in [Mendoza \(2010\)](#).

The household can accumulate capital subject to quadratic adjustment costs:

$$k_{t+1} = (1 - \delta)k_t + \tilde{i}_t - \frac{a}{2} \frac{(k_{t+1} - k_t)^2}{k_t}, \quad (4.1)$$

where δ is the depreciation rate, \tilde{i}_t is gross investment, and a is a constant that scales the adjustment costs. Because of the adjustment costs, capital has a market price q_t possibly different from one. The household can also trade a one-period, zero-coupon foreign bond b_{t+1} at a price $q_t^b = R_t^{-1}$ (where $b_t < 0$ means the household is borrowing).

Since the household owns the firm, the household maximizes its preferences subject to the budget constraint:

$$(1 + \tau)c_t + \tilde{i} = y_t - p_t v_t - \phi(R_t - 1)(w_t L_t + p_t v_t) - q_t^b b_{t+1} + b_t, \quad (4.2)$$

where $y_t = \exp(\epsilon_t^A) k_t^{1-\alpha-\eta} L_t^\alpha v_t^\eta$. This budget constraint tells us that the resources of this economy are used for consumption, taxed at a constant rate τ , and gross investment.⁸ The resources come from the output net of imported inputs and the servicing of working capital loans and from the change in the foreign bond position.

In addition, the household is subject to a collateral constraint:

$$\phi R_t(w_t L_t + p_t v_t) - q_t^b b_{t+1} \leq \kappa q_t k_{t+1}. \quad (4.3)$$

Equation (4.3) limits the total debt (working capital loans plus negative positions in bonds) to be less than or equal to a fraction κ of the market value of the end-of-period capital stock. We define a crisis in our model as a period in which the collateral constraint is binding. The occasionally binding nature of the constraint prevents the use of the Kalman filter.

The prices q_t and w_t that appear in equations (4.2) and (4.3) are endogenous market prices taken as given by the household when solving its optimization problem. Since the wage rate must be on the labor supply curve (i.e., it must equal the tax-adjusted marginal disutility of labor), $w_t = L^{\omega-1}(1 + \tau)$. Similarly, the price of capital satisfies $q_t = \frac{\tilde{\partial}_t}{\partial k_{t+1}}$.

The equilibrium conditions of the model boil down to:

$$u_c(t) = R_t \beta \mathbb{E}_t[u_c(t+1)] + \mu_t(1 + \tau) \quad (4.4)$$

$$q_t = \frac{\mathbb{E}_t[M_{t+1}] \mathbb{E}_t[d_{t+1} + q_{t+1}] + \text{Cov}_t(M_{t+1}, d_{t+1} + q_{t+1})}{R_t \mathbb{E}_t[M_{t+1}] + \tilde{\mu}_t[(1 + \tau) - \kappa]} \quad (4.5)$$

$$L_t^{\omega-1}(1 + \tau) = \frac{e^{\epsilon_t^A} A k_t^\gamma L_t^\alpha v_t^\eta}{(1 + \phi(R_t - 1) + \tilde{\mu}_t(1 + \tau)\phi R_t)}, \quad (4.6)$$

where $u_c(t) \equiv \left(c_t - \frac{L_t^\omega}{\omega}\right)^{-\sigma}$ is the marginal utility of consumption, d_t are the dividends from capital (since we substituted $R_{t+1}^q \equiv (d_{t+1} + q_{t+1})/q_t$), $M_{t+1} \equiv \frac{\beta u_c(t+1)}{u_c(t)}$ is the stochastic discount factor, and $\tilde{\mu}_t \equiv \frac{\mu_t}{u_c(t)}$ is the multiplier of the budget constraint rescaled by $u_c(t)$. See [Mendoza and Villalvazo \(2020, p. 86\)](#) for details.

4.2 Calibration and solution

To enhance the comparability of our exercise with previous results in the literature, we borrow all our parameter values from [Mendoza and Villalvazo \(2020\)](#), except the discount factor β ,

⁸A constant consumption tax τ helps us to match the average share of government expenditures in GDP in the data. Given the structure of the utility function, a constant tax does not distort the savings-consumption margin, and it does not create a time-varying distortion on labor supply. The tax revenues are used for government consumption, which does not enter into the utility function. Therefore, we do not need to discuss the government further.

which we lower marginally from 0.92 to 0.918 to yield an average crisis probability of 30.8% (see Table 5 below). The calibration is annual (although in the empirical application, we scale flows to quarterly data).

Parameter	value
Average productivity: A	6.982
Capital share: γ	0.31
Labor share: α	0.59
Imported inputs share: η	0.10
Working capital parameter: ϕ	0.2579
Discount factor: β	0.918
Labor elasticity coefficient: ω	1.8461
Risk aversion coefficient: σ	2.0
Depreciation rate: δ	0.088
Capital adjustment cost: a	2.75
Tax on consumption: τ	0.17
Collateral coefficient: κ	0.20
Productivity shock: persistence	0.537
Productivity shock: s.d.	1.340
Imported input price shock: persistence	0.737
Imported input price shock: s.d.	3.345
Interest rate shock: persistence	0.572
Interest rate shock: s.d.	1.958

Table 5: Model calibration.

While the goal of the partial filter is to allow a researcher to study the data without solving the model, we need to compute the model to generate pseudo-true data and perform a Monte Carlo experiment. We do so following the *fixed-point iteration (FiPIt)* algorithm developed in [Mendoza and Villalvazo \(2020\)](#). The algorithm works by conjecturing the decision rule for bonds and the capital pricing function, deriving a set of implied decision rules that follow from these conjectures, and solving the Euler equations for bonds and capital via fixed-point iteration to find new values of the bonds decision rule and capital pricing function.

Since this algorithm works over a discrete state space, we define discrete grids for the three states of the economy, (b, k, s) . For the endogenous states, we define grids with M nodes for bonds and N nodes for capital, respectively. The grid for the shock triples $s \in \mathbf{S}$ comes from [Mendoza \(2010\)](#), who assumes that \mathbf{S} has eight triples (i.e., each shock has two realizations) and an associated 8×8 Markov transition probability matrix. [Mendoza \(2010\)](#) computes \mathbf{S} and its transition matrix to match the persistence and volatility of total factor productivity, the intermediate input prices, and the world real interest rate. The solution algorithm runs on average in 114 seconds.

4.2.1 The partial filter

Because static observables reveal the shocks, we focus on filtering the unobserved co-state variables. Let us assume that the goal of a researcher is to filter the multiplier on the collateral constraint, expressed in units of consumption $\tilde{\mu}_t = \frac{\mu_t}{u_{c,t}}$, and the price of capital q_t without having to solve for the whole model. Our assumption represents the goals of many researchers: to learn about two variables of key importance for the behavior of the model that is either unobservable, $\tilde{\mu}_t$, or hard to measure, q_t . In fact, the indices of capital prices are subject to much controversy, and the researcher might want to avoid taking a position in the debate. Also, we can implement our filter in two scenarios: when the researcher is willing to assume some parameter values and when the parameter values are unknown. We will focus on the former case but briefly discuss the latter.

We assume that we can measure C_t and L_t (which gives us $u_{c,t}$), R_t , and d_t . These four series are readily available for many countries. For example, d_t can be construed from national income and product accounts. Hence, we define the vector of observables $\mathbf{y}_t = [u_c, L_t, R_t, d_t]$.

The first step is to define the hidden state vector $\mathbf{x}_t = [\tilde{\mu}_t, q_t]$ and select two equilibrium conditions of the model that are likely to be informative about \mathbf{x}_t . We can rewrite the Euler equation (4.4) as:

$$\tilde{\mu}_t = \frac{u_c(t) - R_t \beta \mathbb{E}_t[u_c(t+1)]}{(1+\tau)} \frac{1}{u_c(t)}, \quad (4.7)$$

and the relative price of capital (4.5) becomes:

$$q_t = \frac{\frac{\beta}{u_{c,t}} \mathbb{E}_t[u_{c,t+1}] \mathbb{E}_t[d_{t+1} + q_{t+1}] + \frac{\beta}{u_{c,t}} \text{Cov}_t(u_{c,t+1}, d_{t+1} + q_{t+1})}{R_t \frac{\beta}{u_{c,t}} \mathbb{E}_t[u_{c,t+1}] + \tilde{\mu}_t [(1+\tau) - \kappa]}, \quad (4.8)$$

when we substitute $M_{t+1} \equiv \frac{\beta u_c(t+1)}{u_c(t)}$. This last expression is particularly convenient because we will specify our VAR in terms of the marginal utility of consumption. Notice that we can filter out $\tilde{\mu}_t$ with u_c, L_t, R_t . In contrast, to filter out q_t , we also need d_t .

In addition to equations (4.7) and (4.8), we will also use the slackness condition:

$$\tilde{\mu}_t [q_t^b b_{t+1} - \phi R_t [w_t L_t + p_t v_t] + \kappa q_t k_{t+1}] = 0, \quad (4.9)$$

which we can simplify further by taking advantage of the fact that, in equilibrium the (measurable) fraction $\frac{\phi R_t}{1+\phi(R_t-1)-\eta}$ of output is spent on labor and imported inputs (see Appendix B.2). We use this condition to adjust our state estimates. Specifically, when the borrowing constraint is binding, but $\tilde{\mu}_t$ is not estimated to be positive, we lower our estimate of the expected future utility, and vice versa when the borrowing is slack and yet the estimated $\tilde{\mu}_t$ is inferred

to be positive.⁹

Mapping the equilibrium conditions into the notation of the partial information filter, we have a 2-dimensional g -function:

$$g(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t) = [\mathbf{y}_{1,t+1}; \mathbf{y}_{4,t+1} + \mathbf{x}_{2,t+1}; \beta \mathbf{y}_{1,t+1} / \mathbf{y}_{1,t}] \quad (4.10)$$

and an \mathbf{f} -function:

$$\mathbf{f}(\mathbf{x}_t, \mathbf{y}_t, \mathbb{E}_t[\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t]) = \begin{bmatrix} -\tilde{\mu}_t + \frac{u_c(t) - R_t \beta \mathbb{E}_t[u_{c,t+1}]}{(1+\tau)} \frac{1}{u_c(t)} \\ -q_t + \frac{\frac{\beta}{u_{c,t}} \mathbb{E}_t[u_{c,t+1}] \mathbb{E}_t[d_{t+1} + q_{t+1}] + \frac{\beta}{u_{c,t}} \text{Cov}_t(u_{c,t+1}, d_{t+1} + q_{t+1})}{R_t \frac{\beta}{u_{c,t}} \mathbb{E}_t[u_{c,t+1}] + \tilde{\mu}_t [(1+\tau) - \kappa]} \end{bmatrix} = \mathbf{0}. \quad (4.11)$$

When the borrowing constraint is slack, we solve \mathbf{f} for $[\tilde{\mu}_t, q_t]$. When the borrowing constraint is binding (or violated), but $\tilde{\mu}_t = 0$, we add a time-varying intercept in the VAR that lowers $\mathbb{E}_t[u_{c,t+1}]$ to be consistent with the binding constraint at time t . Similarly, when the constraint is not binding but $\tilde{\mu}_t > 0$, we introduce a time-varying intercept that raises $\mathbb{E}_t[u_{c,t+1}]$ to lower the inferred $\tilde{\mu}_t$.

The second step in implementing the partial filter is to run a VAR(x) on $\mathbf{y}_t, \mathbf{x}_t$, and use it to approximate the expectations and covariances in $\mathbf{f}(\cdot)$. In the VAR(1) case, we collect the VAR variables in the vector $\boldsymbol{\xi}_t^{(d)} = [u_{c,t}, L_t, R_t, d_t, \tilde{\mu}_t^{(d)}, q_t^{(d)}]$. The case of more lags can be handled via the companion form of the VAR.

Then, we use the fact that $\boldsymbol{\xi}_t^{(d)} = \mathbf{a}^{(d)} + \mathbf{A}^{(d)} \boldsymbol{\xi}_{t-1}^{(d)} + \boldsymbol{\varepsilon}_t^{(d)}$ and $\boldsymbol{\Sigma}^{(d)} = \widehat{\text{Var}}[\boldsymbol{\varepsilon}_t^{(d)}]$ to write:

$$f_1(\mathbf{x}_t^{(d)}, \mathbf{y}_t, \widehat{E}_t^{(d)}[g(\mathbf{x}_{t+1}^{(d-1)}, \mathbf{y}_{t+1}, \mathbf{x}_t^{(d-1)}, \mathbf{y}_t)]) \approx -\mathbf{x}_{1,t}^{(d)} + \frac{u_c(t) - R_t \beta \mathbf{e}'_1(\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \mathbf{x}_t^{(d-1)})}{(1+\tau)} \frac{1}{u_c(t)}$$

and

$$\begin{aligned} f_2(\mathbf{x}_t^{(d)}, \mathbf{y}_t, \widehat{E}_t^{(d)}[g(\mathbf{x}_{t+1}^{(d-1)}, \mathbf{y}_{t+1}, \mathbf{x}_t^{(d-1)}, \mathbf{y}_t)]) \approx & -\mathbf{x}_{2,t}^{(d)} \\ & + \frac{\frac{\beta}{u_c(t)} \mathbf{e}'_1(\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \mathbf{x}_t^{(d-1)}) (\mathbf{e}'_4 + \mathbf{e}'_6) (\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \mathbf{x}_t^{(d-1)})}{R_t \frac{\beta}{u_c(t)} \mathbf{e}'_1(\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \mathbf{x}_t^{(d-1)}) + \tilde{\mu}_t^{(d)} [(1+\tau) - \kappa]} \\ & + \frac{\frac{\beta}{u_c(t)} [(\boldsymbol{\Sigma}_{1,4}^{(d)} + \boldsymbol{\Sigma}_{1,6}^{(d)})]}{R_t \frac{\beta}{u_c(t)} \mathbf{e}'_1(\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \mathbf{x}_t^{(d-1)}) + \tilde{\mu}_t^{(d)} [(1+\tau) - \kappa]}. \end{aligned}$$

For numerical stability, we do not immediately update the state estimates. Instead, we

⁹We do not use this equation to filter Tobin's q because of measurement error. The observed debt is averaged across maturities and is not risk-free, causing a discrepancy between the model and the data. In the data for Mexico, we would otherwise infer an implausible increase in Tobin's q when the borrowing constraint binds. Given observables, only an increased collateral value could otherwise justify the observed borrowing.

use a third step to update the guesses slowly. Specifically, we update the state estimates for \hat{q}_t and $\hat{\mu}_t$, using iterative rules of the form $\hat{q}_t^{(d)} = (1 - s)\hat{q}_t^{(d-1)} + s\mathbf{x}_{2,t}^{(d)}$. Updating the state estimates slowly improves convergence. Because in the model $\tilde{\mu}_t \geq 0$, we also censor $\hat{\mu}_t$ at $-5e^{-4}$ and, for numerical stability, we restrict \hat{q}_t to lie within $[0.9, 1.1]$ and to average 1. These bounds have minor effects on the simulated data but improve the numerical stability in the empirical application. See Appendix B.3 for a detailed algorithm.

All the steps we have taken so far follow directly from the partial information filter approach described earlier: (i) we selected two equilibrium conditions plus the slackness condition of the model based on their informativeness about the states we wanted to learn about and performed basic algebraic manipulations with them; (ii) we ran a VAR on four observables and used the results to substitute variances and covariances on the three equations from step (i); and (iii) we slowly updated our estimates for \hat{q}_t and $\hat{\mu}_t$. None of these steps requires much computational effort. The total run time of the algorithm is 47 seconds.

4.3 Simulation study

We simulate the model 100 times for 250 periods using the solution method presented in Subsection 4.2. Since our experiments with univariate processes in Section 3 suggested that the partial filter recovers the hidden state more reliably when the persistence of the state is higher, we look at sample paths from our economy with different properties.

More concretely, we retain only simulated samples when we are in a crisis (e.g., $\mu_t > 0$) between 10 and 15% of the periods, when we are in a crisis between 30 and 35% of the periods, and when we are in a crisis between 50 and 55% of the periods until we have accumulated 100 samples each. These are the most useful samples to test our partial filter. Since in all simulated samples we are using the same calibration, this choice means that we are implicitly selecting the shock realizations that give us the patterns that we are looking for.¹⁰ Then, we apply our partial filter as outlined above to the selected simulated samples.

4.3.1 Results

We begin by assessing the adequacy of the auxiliary statistical model. Figure 8 shows the results of the autocorrelation test for the residuals of the observed variables. Panel (a) shows the case with a sudden stop frequency of 10%–15%. Panel (b) does the same for the case with

¹⁰Average technology declines, and the average import price increases as the frequency of crises increases. Technology shocks average 1.002 when the crisis frequency is in $[10\%, 15\%]$, but they only average 0.9995 when the crisis frequency is in $[50\%, 55\%]$. Import price shocks average 1.0250 with a low crisis probability and 1.032 with a high crisis frequency. While the second moments remain stable, some higher moments also differ. For example, in the high crisis frequency sample, the interest rate shock has a skewness of 0.36, as opposed to -0.09 in the low crisis simulation.

a sudden stop frequency of 50%–55%. The horizontal axis of each panel shows the number of lags of the approximating VAR; the vertical axis shows the number of lags of the test statistic. The colors represent the fraction of simulations for which the bootstrapped Ljung-Box test failed to reject the Null of no autocorrelation. Whereas the dark blue color represents values near zero, shades of green values near 0.5, and bright yellow signify that we failed to reject in (nearly) all simulations. Both panels show that VAR models with few lags are rejected, while models with two or more lags capture even higher-order autocorrelation.




Figure 8: Rejection rate for the Null of zero residual autocorrelation as a function of the lags of the approximating VAR and the lags in the test

Since the data-generating process features both shocks and decision rules that are first-order Markov, our simulation indicates that the non-linear nature of the model or the discretization of the shock processes push us toward using a higher-order VAR. In what follows, we focus on the VAR approximation with two lags.

Sudden stop freq. [%]	$\text{corr}(\hat{q}_t, q_t)$	$\text{corr}(\hat{\mu}_t, \tilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\tilde{\mu}_t$
[10.0, 15.0]	0.77 [0.64, 0.87]	0.85 [0.66, 0.91]	1.29 [1.01, 1.57]	1.66 [0.79, 2.88]
[30.0, 35.0]	0.74 [0.55, 0.87]	0.84 [0.71, 0.92]	1.32 [1.12, 1.66]	1.40 [0.96, 1.84]
[50.0, 55.0]	0.70 [0.45, 0.81]	0.84 [0.68, 0.91]	1.18 [1.00, 1.52]	1.28 [1.02, 1.74]

Table 6: Medians [68% CI] of correlations, relative standard deviations, and classification errors

Table 6 reports the correlation and relative standard deviation for the two filtered variables for each of the three crisis frequency scenarios. The partial filter does well in recovering the price of capital q_t , independently of the crisis frequency in the sample: the median correlation across simulations ranges from 0.70 to 0.77 in the three scenarios. The median standard deviation of the estimated \hat{q}_t relative to the true q_t ranges from 1.18 to 1.32. The inner 68% confidence intervals for the correlations are skewed to the left, while the confidence intervals for relative standard deviations are skewed to the right. The performance is comparable for $\tilde{\mu}_t$,

with correlations around 0.85. The relative standard deviation of the filtered $\tilde{\mu}_t$ is higher, with the median ranging from 1.66 to 1.28. Again, the confidence intervals are skewed.¹¹ With a Minnesota prior rather than a flat prior, the correlations improve for Tobin’s q without impacting those for the crisis indicator much. The relative standard deviations for Tobin’s q rise, while those for the crisis indicator fall. See Table B.4 in the Appendix.

Sudden stop	Crisis classification		
Frequency	False positives	False negatives	False(any)
[10.0, 15.0]	3.3 [1.4, 5.4]	0.0 [0.0, 8.1]	3.7 [2.0, 5.3]
[30.0, 35.0]	5.5 [3.1, 8.1]	0.0 [0.0, 7.8]	4.9 [2.8, 6.8]
[50.0, 55.0]	8.7 [5.3, 11.6]	0.8 [0.0, 5.5]	4.9 [3.3, 6.9]

Table 7: Crisis classification errors

It might be more relevant for policymakers to gauge how well our partial filter does at learning when the borrowing constraint binds. Table 7 lists the false positive and false negative rates when classifying crises, i.e., how often we estimate that the constraint is binding when it is not and, vice versa, how often we miss that the constraint is actually binding. Here, our partial filter displays a great performance, with median false positives between 3.3%, when the true crisis frequency is 10%-15%, and 8.7%, when the true crisis frequency is between 50% and 55%. The false negative rates are near zero. The probability of any false classification is 4 to 5%. With a Minnesota prior, the results are broadly similar, but with a lower false positive rate at the cost of a higher false negative rate (see Table B.5 in the appendix). Table 7 shows that the partial information filter performs very well in classifying crises correctly.

Figure 9 shows the data underlying the previous tables.¹² The top panels show the time path of the true (blue) and estimated (orange) time series for the sample with the median correlation. In line with Table 6, the two lines track each other well in all three cases. There appears to be less excess volatility for $\tilde{\mu}_t$ when the crisis frequency is higher, in line with a relative standard deviation closer to one. Comparing the simulated time series with the filtered counterparts shows that the partial filter reliably recovers the large movements in q_t and $\tilde{\mu}_t$. This includes even short-lived crises, such as the one around period 230 in the middle panel of the top row.

The bottom panel shows scatter plots of the relationship between the true and the filtered series, with different colors indicating different sample paths. While there is some dispersion around the 45-degree line and the filtered series tend to have a wider range, the scatter plots

¹¹Out of a total of 300 simulations for the three crisis frequencies, all but six converge. We drop the simulations that have not converged. The results remain virtually unchanged, however, if they are included.

¹²Figures B.9 and B.10 in the Appendix show the sample paths for true and filtered q and $\tilde{\mu}$ for all three crisis frequencies.

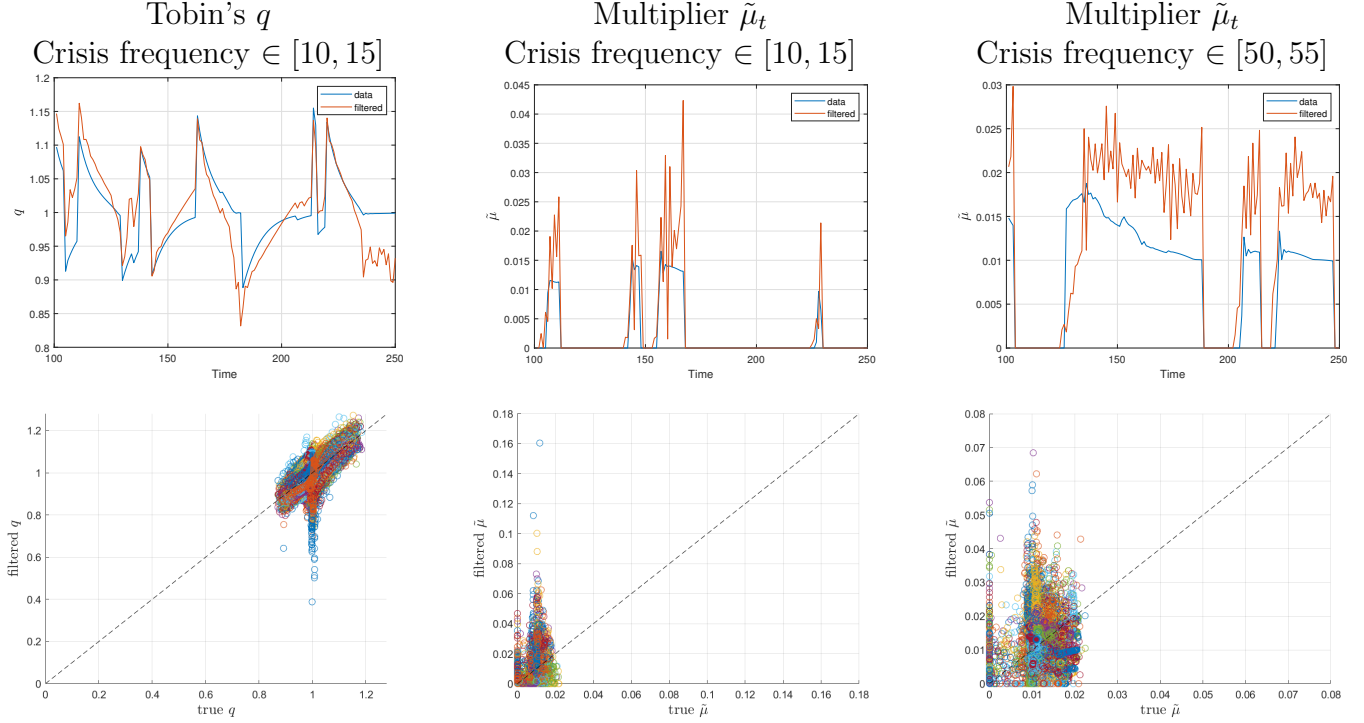


Figure 9: Sample paths for q_t and $\tilde{\mu}_t$ with the median correlation (top panels) and scatter plot pooling all samples (bottom panels) for different in-sample crisis frequencies.

show that the partial filter recovers the right pattern of q . For $\tilde{\mu}_t$, the relationship is noisier, but the excess volatility decreases with the crisis frequency, resulting in a more compact distribution that is also reflected in narrower axes. While there is some mass on the axes, which represent classification errors, a substantial mass is inside, reflecting the correct crisis classification.

5 Application II: Using Mexican data

Having established that the filter performs well in a Monte Carlo experiment, we apply it to Mexican quarterly data ([OECD](#); [Banco de México](#); [CEPALSTAT](#)) from 1980Q1 to 2018Q4.

Data and partial filter specification. To account for trends outside the model, we detrend the variables (in logs) using the HP filter and re-center them at the model-implied means (in levels). We use data on real consumption and hours worked to construct marginal utility. Our measure of the world real interest rate is the LIBOR ([ICE Benchmark Administration](#)) spliced with the 3-month U.S. T-bill rate ([Board of Governors of the Federal Reserve System](#)) and converted to real terms using the U.S. CPI ([U.S. BLS](#)). We take dividends from the S&P/IFCG M MEXICO index, spliced after 2008Q3 with the MSCI MEXICO index

([Datastream International](#)). We construct the annual capital stock using the perpetual inventory method and linearly interpolate it to obtain a quarterly series. See Appendix B.2 for details.

We use the same specification of the partial filter as in our earlier simulation study. This algorithm is a fixed point algorithm that ignores uncertainty about the parameters of the approximating VAR. The step that updates parameters uses point estimates. Here, we also provide results using the Gibbs sampler, which takes a draw from the posterior distribution of the approximating VAR. For each parameter draw, we then solve for filtered sequences.¹³ We initialize $q_t = 1$ and $\tilde{\mu}_t^{(0)} = 0.01$ and add white noise to each series to avoid singularity in the initial VAR estimates. We then set $\mathbf{a}^{(0)}, \mathbf{A}^{(0)}, \Sigma^{(0)}$ to the OLS estimates associated with $\xi_t^{(0)} = [u_{c,t}, L_t, R_t, d_t, \tilde{\mu}_t^{(0)}, q_t^{(0)}]$.¹⁴

Specification tests. We report p-values of the Ljung-Box result for residual autocorrelation in Panel (a) of Figure 10, while Panel (b) reports the results of the Kolmogorov-Smirnov test. Both sets of tests suggest that a VAR with three or more lags adequately models the behavior of the VAR observables.




Figure 10: Bootstrapped specification tests for residual autocorrelation and equality of distribution

Consider the tests for multivariate residual autocorrelation first. The figure reports p -values for one through six lags of the VAR for Ljung-Box test lags also ranging from one to six periods. For a VAR with a single lag, the p -values range from less than 0.01 (one lag in the test, dark blue) to 0.14 (lighter blue, two lags in the test). Starting with a VAR of lag length three, we cannot reject the Null of zero residual autocorrelation at any horizon considered,

¹³The filtering step is again parameterized the same as in the simulation study. For successive parameter draws, we initialize the algorithm with the filtered sequence from the previous draws, unless in the rare event that this leads to a divergent filtering sequence. In the case where the current initial condition leads to divergence, we re-initialize the filtering sequence with the initial draw.

¹⁴As before, we use the slackness condition to shift the future expected marginal utility up or down to ensure the slackness condition binds during a reasonable fraction of the sample. Specifically, we lower (raise) future expected marginal utility if the borrowing constraint is violated and $\tilde{\mu}_t = 0$ (if the slackness condition is above its 15th percentile and $\tilde{\mu}_t > 0$).

as evidenced by the light blue to yellow colors in the figure, which indicate p -values of 0.29 and higher.

The Kolmogorov-Smirnov test provides an auxiliary test of the adequacy of the approximating statistical model. Under the Null of equality of the unconditional distribution from the approximating VAR, the bootstrapped p -values should be distributed uniformly on $[0, 1]$, and about 10% of them should be below 0.1. Figure 10 (b) reports this fraction. With a single lag, between 19% and 49% of p -values are below 0.1, indicating meaningful misspecification. This performance improves, and with four lags, between 5% and 23% of p -values are below 0.1 (or 13% when averaged across variables). We conclude that a VAR with sufficiently many lags captures both the dynamics and the distribution adequately.

We thus focus on a VAR with four lags below. For comparison, we report the results with various numbers of lags. Overall, the results are insensitive to the lag length, but at lower lag lengths, the volatility of $\tilde{\mu}_t$ is understated.

Results. Does our partial filter recover the Mexican debt crises during our sample? Yes. Panel (a) of Figure 11 shows the filtered multiplier on the borrowing constraint, $\tilde{\mu}_t$, and plots a number of economic events that we suspect could lead to a binding borrowing constraint. These events are: (1) the debt crisis of 1982Q3; (2) the collapse of oil prices in 1986 that predated Black Friday in 1987Q3; (3) the “Tequila Crisis” of 1994Q4; (4) the Asian crisis of 1997Q3; (5) the Russian crisis and LTCM crisis of 1998Q3; and (6) the collapse of Lehman Brothers in 2008Q3. The partial filter is successful in identifying all these crisis events except the Asian crisis of 1997Q3. Interestingly, it is unclear whether the Asian crisis should have much of an impact on Mexico. The fixed point of the algorithm identifies a crisis event around 2003 that might be a false positive. Indeed, when taking parameter uncertainty into account, this event disappears (see Panel (a) of Figure 12).

The initial guess, a 0-1 dummy for whether the collateral constraint would be violated given a counterfactual $q_t = 1$, does well in capturing the crises. Furthermore, the filtered results do not inherit the false positive from the initial guess. Indeed, starting from a completely random initial guess yields a similar crisis classification. See Figure B.12.

Panel (b) of Figure 11 shows the corresponding movement in q_t . The blue line represents the filtered value for q_t . The same crisis events that are associated with a binding borrowing constraint are also associated with a lower price of capital (the binding borrowing constraint drives up the rate of return the representative household requires). Moreover, the filter shows that our initial guess does not drive the inference about Tobin’s q .

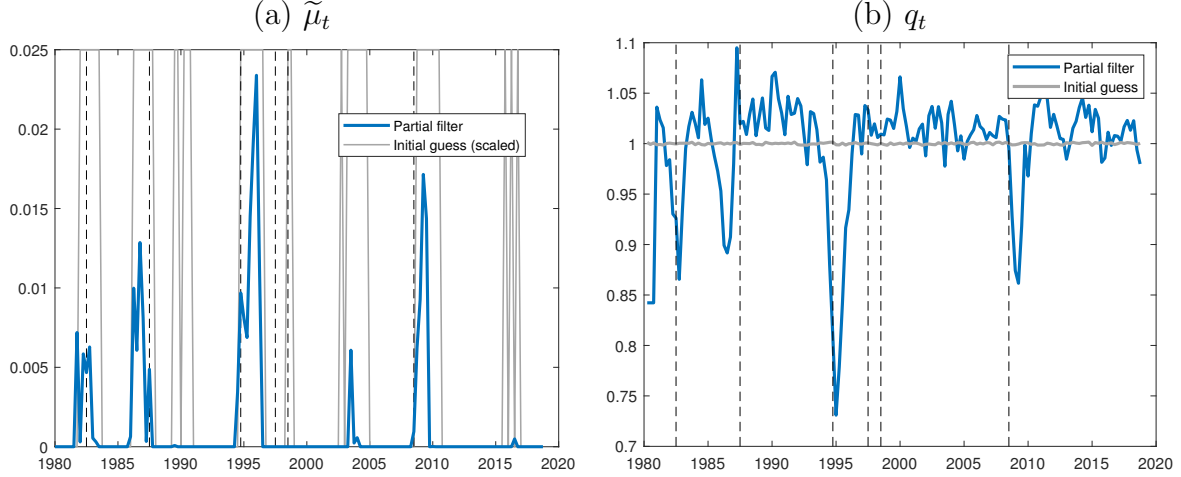
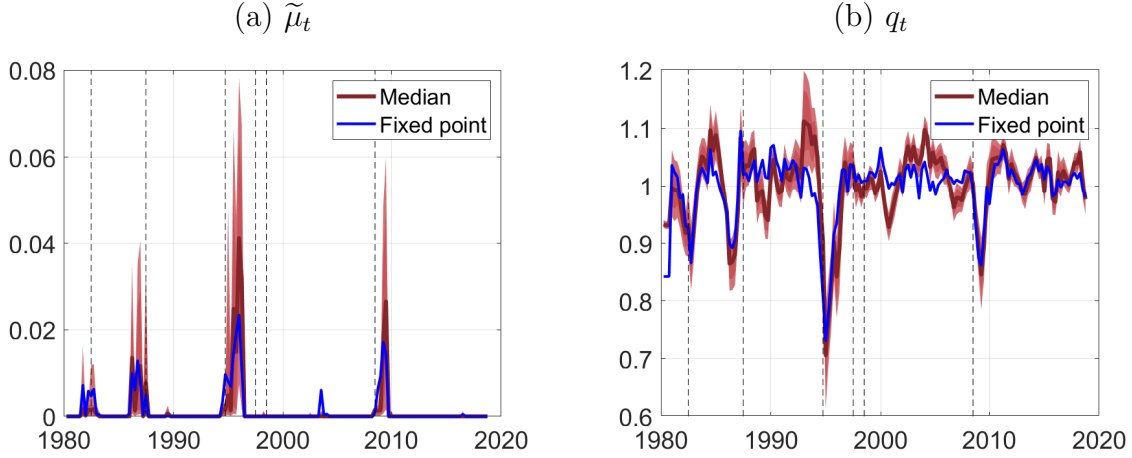


Figure 11: Filtered values of $\tilde{\mu}_t$ and q_t .

Parameter uncertainty. While we treated the estimated expectations as certain in the analysis so far, we can quantify the effect of parameter uncertainty on expectations and thus the filtered estimates using a Gibbs sampler. We iterate between estimating the VAR with a Minnesota prior and filtering $\tilde{\mu}_t$ and q_t , which, in turn, enter the VAR as data.¹⁵



Posterior median and 68% and 90% credible sets

Figure 12: Filtered values of $\tilde{\mu}_t$ and q_t .

Allowing for parameter uncertainty broadly confirms the results from the fixed point algorithm obtained so far. First, the partial filter with parameter uncertainty identifies almost the same crises as the fixed point algorithm. Panel (a) of Figure 12 shows 68% and 90% credible sets, the posterior median, and the fixed point estimates of the multiplier on the borrowing constraint. While the credible sets are wide in some periods, the filter

¹⁵Our prior for the VAR parameters centers the VAR parameters at independent, stationary AR(1)s with the persistence of 0.9 and standard deviations given by the training sample.

rarely identifies positive multipliers outside crisis times. During these times, the multiplier is positive with at least 95% posterior probability. Unlike the fixed point algorithm, now we do not see significant false positives. The other crisis events appear to be both statistically significant and economically meaningful.

Another feature of the posterior uncertainty is its asymmetry. The asymmetry is more noticeable in the multiplier estimates. However, large right tails in the estimates of $\tilde{\mu}_t$ also manifest themselves as asymmetric credible sets for Tobin's q (Panel (b) of Figure 12), for example during the 2007-2008 financial crisis. Overall, the effects of parameter uncertainty are small.

6 Conclusion

The partial information filter allows researchers to evaluate the historical implications of their models without solving them or even fully specifying them. The latter point is important because non-linear dynamic models are often matched to certain moments but without specifying enough shock processes and shock dynamics required for full information methods. Even when full information non-linear filters are available, the computational cost of these filters may discourage researchers from evaluating the historical implications of their model. Our partial information filter thus has the potential to allow for better scrutiny of dynamic macroeconomic models. It is, therefore, an important complement to likelihood-free methods in model estimation.

Our analysis points to two requirements for the partial information filter. First, the measurement equations need to identify the state absent measurement error uniquely. Second, the results are robust to misspecification if the measurement is not too noisy, and the state needs to exhibit enough persistence. Based on our experience, we conjecture that a large class of models of interest satisfy these two requirements.

References

- Andreasen, M., Fernández-Villaverde, J., and Rubio-Ramírez, J. (2018). The pruned state-space system for non-linear DSGE models: Theory and empirical applications. Review of Economic Studies, 85:1–49.
- Banco de México (2022). Deuda neta total del sector público, saldo promedio y final - (cg7). <https://www.banxico.org.mx/SieInternet/consultarDirectorioInternetAction.do?sector=9&accion=consultarCuadro&idCuadro=CG7&locale=es> (Accessed January 2022).
- Bernanke, B. S. and Kuttner, K. N. (2005). What explains the stock market’s reaction to Federal Reserve policy? Journal of Finance, 60(3):1221–1257.
- Board of Governors of the Federal Reserve System (2020). 3-month Treasury bill secondary market rate, discount basis [dtb3]. Retrieved from FRED, Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/DTB3> (Accessed July 2020).
- Campbell, J. Y. (1991). A variance decomposition for stock returns. Economic Journal, 101(405):157–179.
- Canova, F. and Ferroni, F. (2022). Mind the gap! Stylized dynamic facts and structural models. American Economic Journal: Macroeconomics, 14(4):104–35.
- CEPALSTAT (2022). Producto interno bruto anual (pib) por objeto del gasto a precios corrientes en moneda nacional. Individual URLs provided in the appendix.
- Chahrour, R., Cormun, V., Leo, P. D., Guerron-Quintana, P., and Valchev, R. (2021). Exchange rate disconnect redux. Boston College Working Papers in Economics 1041, Boston College Department of Economics.
- Chahrour, R. and Jurado, K. (2018). News or noise? The missing link. American Economic Review, 108(7):1702–1736.
- Chen, X. and Christensen, T. M. (2015). Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions. Journal of Econometrics, 188(2):447–465.
- Coibion, O., Gorodnichenko, Y., Kumar, S., and Ryngaert, J. (2021). Do you know that I know that you know...? Higher-order beliefs in survey data. Quarterly Journal of Economics, 136(3):1387–1446.

- Datastream International (2021). S&P/IFCG M MEXICO (IFGMMX) index and MSCI MEXICO index (MSMEXFL). Available: Refinitiv Workspace. (September 30, 2021).
- Drautzburg, T., Fernández-Villaverde, J., and Guerrón-Quintana, P. (2021). Bargaining shocks and aggregate fluctuations. Journal of Economic Dynamics and Control, 127:104121.
- Farmer, L. E. (2021). The discretization filter: A simple way to estimate nonlinear state space models. Quantitative Economics, 12(1):41–76.
- Fernald, J. G. (2012). A quarterly, utilization-adjusted series on total factor productivity. Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- Fernández-Villaverde, J., Rubio-Ramírez, J. F., and Schorfheide, F. (2016). Solution and estimation methods for DSGE models. In Taylor, J. B. and Uhlig, H., editors, Handbook of Macroeconomics, volume 2, pages 527–724. Elsevier.
- Gabaix, X. (2020). A behavioral New Keynesian model. American Economic Review, 110(8):2271–2327.
- Gallant, A. R., Giacomini, R., and Ragusa, G. (2017). Bayesian estimation of state space models using moment conditions. Journal of Econometrics, 201(2):198 – 211.
- Gordon, N., Salmond, D., and Smith, A. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings F (Radar and Signal Processing), 140:107–113(6).
- Hansen, L. P. (1982). Large sample properties of Generalized Method of Moments estimators. Econometrica, 50(4):1029–1054.
- Hansen, L. P. (2014). Nobel Lecture: Uncertainty outside and inside economic models. Journal of Political Economy, 122(5):945–987.
- Herbst, E. P. and Schorfheide, F. (2016). Bayesian Estimation of DSGE Models. Number 10612 in Economics Books. Princeton University Press.
- ICE Benchmark Administration (2020). 3-Month London Interbank Offered Rate (LIBOR), based on U.S. Dollar [USD3MTD156N]. Retrieved from FRED, Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/USD3MTD156N> (Accessed July 2020).
- Ljung, G. M. and Box, G. E. P. (1978). On a measure of lack of fit in time series models. Biometrika, 65(2):297–303.

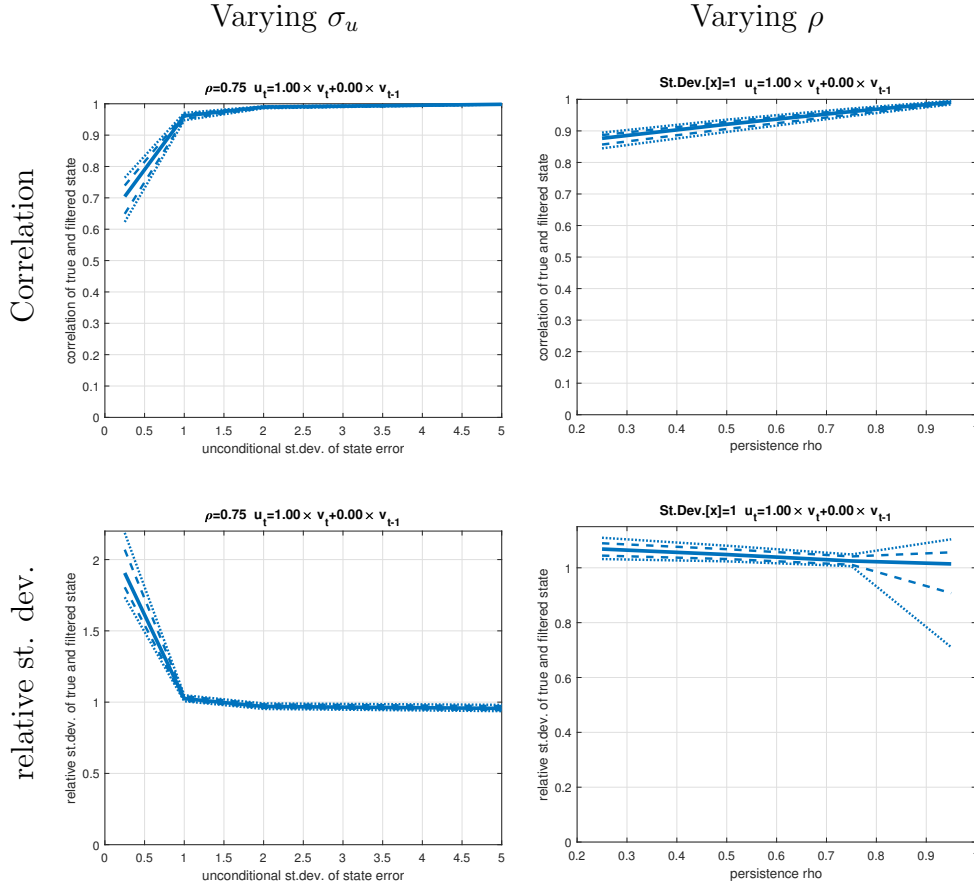
- Mendoza, E. and Villalvazo, S. (2020). FiPIt: A simple, fast global method for solving models with two endogenous states and occasionally binding constraints. Review of Economic Dynamics, 37:81–102.
- Mendoza, E. G. (2010). Sudden stops, financial crises, and leverage. American Economic Review, 100(5):1941–66.
- OECD (2021). Macroeconomic time series for Mexico. Retrieved from FRED, Federal Reserve Bank of St. Louis. Individual URLs provided in the appendix.
- Penn World Table 10.01 (2021). Average annual hours worked by persons engaged for Mexico [avhwpemxa065nrug]. Retrieved from FRED, Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/AVHWPEMXA065NRUG> (Accessed November 2021).
- Petrova, K. (2019). A quasi-Bayesian local likelihood approach to time varying parameter VAR models. Journal of Econometrics, 212(1):286 – 306.
- Roberts, J. M. (1995). New Keynesian economics and the Phillips curve. Journal of Money, Credit and Banking, 27(4):975–984.
- Sbordone, A. M. (2002). Prices and unit labor costs: a new test of price stickiness. Journal of Monetary Economics, 49(2):265–292.
- Schorfheide, F. (2005). VAR forecasting under misspecification. Journal of Econometrics, 128(1):99–136.
- Tsay, R. (2010). Analysis of Financial Time Series. Wiley.
- U.S. BLS (2021). Macroeconomic time series for the U.S.A. Retrieved from FRED, Federal Reserve Bank of St. Louis. Individual URLs provided in the appendix.

Appendix

A Monte Carlo study in the scalar case

A.1 No measurement error

Here, we provide additional correlation and standard deviation plots for the linear case without measurement error.



Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median.

Figure A.1: Partial filter in the linear case without measurement error. Correlation and relative standard deviation of the filtered state \hat{x}_t vs. true x_t for varying ρ_0 .

Figure A.1 plots the median correlation between the true and filtered states (along with 68% and 90% confidence intervals) in the top row, and the relative standard deviations of the filtered state relative to the true state in the bottom row. The two panels on the left vary σ_u , for fixed $\rho_1 = 0.75$. The two panels on the right fix $\sigma_u = 1$ and vary ρ_1 . The top row shows that the median correlation is close to one in all cases shown and always above

0.95 with tight confidence intervals. The bottom row shows that the standard deviation of the estimated state is slightly above that of the true state for all levels of σ_u and ρ_0 .

A.2 Additional results: Quadratic observation equation

	Near-linear observation equation					Very non-linear observation equation				
	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 2$	$\sigma_u = 5$	$\sigma_u = 10$	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 2$	$\sigma_u = 5$	$\sigma_u = 10$
$\rho = 0.25$	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00
$\rho = 0.5$	0.00	0.00	0.02	0.00	0.01	0.00	0.00	0.00	0.00	0.00
$\rho = 0.75$	0.00	0.00	0.00	0.01	0.03	0.00	0.00	0.00	0.00	0.00
$\rho = 0.9$	0.00	0.01	0.02	0.00	0.01	0.00	0.00	0.00	0.00	0.00
$\rho = 0.95$	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table A.1: VAR(1) rejection rates in different DGPs. Near-linear: $\kappa_1 = 10, \kappa_2 = 0.1$ Very non-linear: $\kappa_1 = 0.1, \kappa_2 = 10$.

A.3 Linear model with ARMA(p,q) laws of motion

We now generalize the linear law of motion (3.1) with $\eta = 0$ to an ARMA(p,1) model:

$$y_t = \mu_y + x_t + \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}^2] + \sigma_e e_t \quad (\text{A.1a})$$

$$x_t = \mu_x + \sum_{\ell=1}^p \rho_\ell x_{t-\ell} + \sigma_u u_t \quad (\text{A.1b})$$

$$u_t = \frac{\phi_0 v_t + \phi_1 v_{t-1}}{\sqrt{\phi_0^2 + \phi_1^2}}, \quad (\text{A.1c})$$

Here we summarize the performance of four different linear DGPs for the case of the linear observation equation. In all cases, $y_t = x_t + \mathbb{E}_t[x_{t+1}] + e_t$ and $[e_t, u_t]' \stackrel{iid}{\sim} \mathcal{N}(0, I_2)$.

1. Figure A.2 considers the case in the main text, where $x_t = 0.95x_{t-1} + \sigma_u u_t$.
2. Figure A.3 considers the case in the main text, where $x_t = 0.5x_{t-1} + 0.4x_{t-2} + \sigma_u u_t$.
3. Figure A.4 considers the case in the main text, where $x_t = 0.5x_{t-1} + 0.4x_{t-2} + \sigma_u v_t$ and $v_t = \frac{1}{\| [1, 4] \|} u_t + \frac{4}{\| [1, 4] \|} u_{t-1}$.
4. Figure A.5 considers the case in the main text, where $x_t = 0.9x_{t-4} + \sigma_u u_t$.

Within each figure, the top panel shows the fraction of rejected draws according to a Ljung-Box test with a significance level of 10% (with a lag order equal to the maximum lag of the AR coefficients in the DGP) as a function of the lags of the approximating VAR. Within the top panel, the different lines reflect decreasing degrees of misspecification of the approximating model by increasing the standard deviation σ_u of the state x_t relative to the measurement error. In each figure, the rejection rate falls as the lag order of the VAR increases. The Ljung-Box test successfully rejects models with lag lengths that are too short. However, in the case of the back-loaded AR(4), the test is conservative and rejected about 60% of the time.

The center and bottom panels of each figure show the correlation of the filtered and the true state across simulations with the median and 68% confidence intervals as a function of the lag length of the approximating VAR, separately for accepted and rejected simulations. The center panel shows the result for the highest degree of misspecification and the bottom for the lowest degree of misspecification. While there are a few differences in the performance of the partial filter for accepted and rejected draws, the correlations are stable across lag lengths of the approximating model. For the high degrees of misspecification, the median correlations range from 0.4 with the AR(4) as DGP to 0.85 in the ARMA(2,1). With the low degree of misspecification, the correlations are between 0.8 in the case of the AR(4) and almost one in the other cases.

Overall, the specification test has bite, and the partial filter performance with the linear observation equation is robust to different DGPs.




Figure A.2: AR(1) with $\rho = 0.95$. y_t has measurement error.

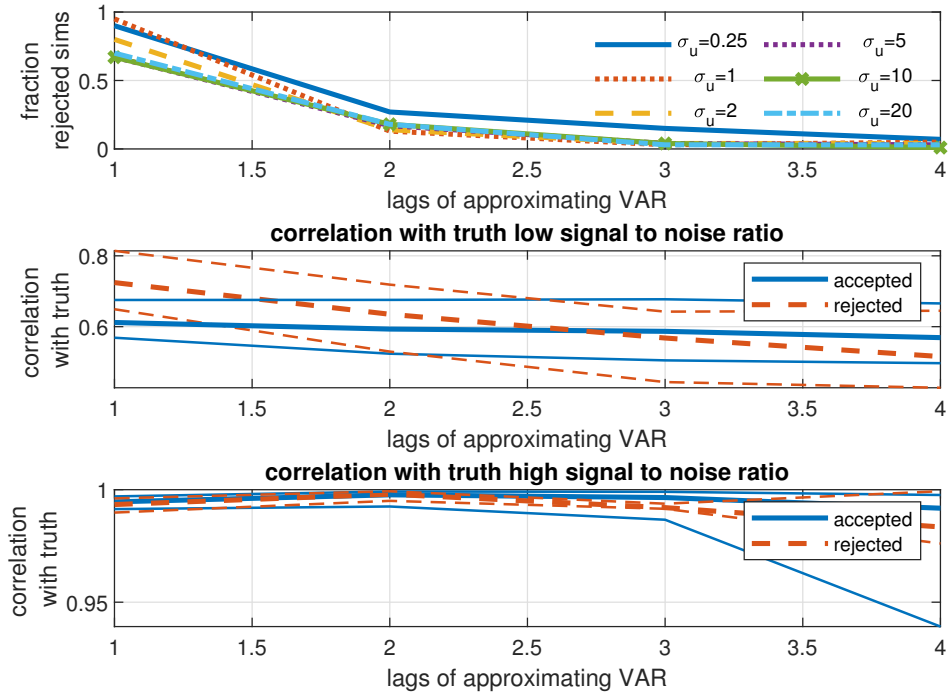


Figure A.3: x_t AR(2) process with $\rho_1 = 0.5$, $\rho_2 = 0.4$. y_t has measurement error.




Figure A.4: x_t ARMA(2,1) process with $\rho_1 = 0.5, \rho_2 = 0.4, \phi_0 = \frac{1}{|[1,4]|}, \phi_1 = \frac{4}{|[1,4]|}$. y_t has measurement error.




Figure A.5: x_t AR(4) process with $\rho_1 = \rho_2 = \rho_3 = 0, \rho_4 = 0.9, \phi_0 = \frac{1}{|[1,4]|}, \phi_1 = \frac{4}{|[1,4]|}$. y_t has measurement error.


A.4 Non-linear law of motion with cosine forcing term

To generate more persistence and non-linearity, we introduce a cosine term:

$$y_t = x_t + \kappa_1 \mathbb{E}_t[x_{t+1}] + e_t \quad (\text{A.2a})$$

$$x_t = \gamma_0 \cos(\gamma_1(t-1)) + \left(\rho_0 + \frac{\rho_1}{1+x_{t-1}^2} \right) x_{t-1} + u_t, \quad (\text{A.2b})$$


Below, we use $\gamma_0 = 4, \gamma_1 = 1.2$.



Note: The dashed line is the 45-degree line; the solid linear is the fit of an OLS regression on x_t and x_t^2 .

Figure A.6: Linear observation equation, non-linear state equation with cosine forcing term for varying degrees of persistence ρ_1 . True state vs. observation with $T=25,000$ observations.

Figure A.6 shows a long simulation from the DGP for x_{t+1} vs. x_t . It shows that the hidden state x_t can live in nearly disjoint regions. With $\rho_1 = 0.25$, within each region, the persistence is negative, but it is positive across regions. For $\rho_1 = 0.75$, the DGP generates a donut-shaped relationship between x_t and x_{t+1} .



Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45-degree line.

Figure A.7: Partial filter with a non-linear state equation with cosine forcing term: Filtered state \hat{x}_t vs. true x_t for 100 different samples for varying degrees of persistence ρ_1 .

Despite this challenging environment, the performance of the filter is still acceptable; see

Figure A.7. However, the partial filter does not consistently pass the specification tests; see Table A.2.

Table A.2: Partial filter with non-linear state equation with cosine forcing term and non-linear VAR(p): Ljung-Box rejection rates with different lag lengths.

	Non-linear VAR(1) approximation					Non-linear VAR(2) approximation				
	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 2$	$\sigma_u = 5$	$\sigma_u = 10$	$\sigma_u = \frac{1}{4}$	$\sigma_u = 1$	$\sigma_u = 2$	$\sigma_u = 5$	$\sigma_u = 10$
$\rho = 0.25$	0.10	0.12	1.00	0.94	1.00	0.12	0.20	0.98	1.00	1.00
$\rho = 0.5$	0.92	0.92	1.00	0.34	0.84	0.90	0.90	1.00	0.46	0.84
$\rho = 0.75$	1.00	0.96	0.18	0.64	0.06	0.98	1.00	0.20	0.32	0.14
$\rho = 0.9$	0.98	0.16	0.08	0.78	0.38	1.00	0.56	0.42	0.30	0.24
$\rho = 0.95$	0.96	0.14	0.22	0.42	0.48	0.98	0.44	0.20	0.14	0.12




Figure A.8: Rejection rates and correlation of filtered state with truth in non-linear law of motion with cosine forcing term. ($\rho_1 = 0.9$)

B Sudden stops model

This appendix presents additional results for the simulation as well as for the empirical filtering exercise on the sudden stops model.

Table A.3: Partial filter performance: Non-linear law of motion with cosine forcing term.

Non-linear VAR(2)				
		Correlation	Relative s.d.	
$\rho = 0.25$	$\sigma_u = \frac{1}{4}$	0.43 [0.37, 0.47]	1.58 [1.50, 1.80]	
	$\sigma_u = 1$	0.74 [0.71, 0.95]	0.80 [0.75, 0.96]	
	$\sigma_u = 10$	NaN [NaN, NaN]	NaN [NaN, NaN]	
		Correlation	Relative s.d.	
$\rho = 0.75$	$\sigma_u = \frac{1}{4}$	0.76 [0.76, 0.76]	0.76 [0.76, 0.76]	
	$\sigma_u = 1$	NaN [NaN, NaN]	NaN [NaN, NaN]	
	$\sigma_u = 10$	0.99 [0.99, 1.00]	0.99 [0.99, 1.00]	
		Correlation	Relative s.d.	
$\rho = 0.95$	$\sigma_u = \frac{1}{4}$	0.80 [0.80, 0.80]	1.27 [1.27, 1.27]	
	$\sigma_u = 1$	0.96 [0.95, 0.98]	0.93 [0.92, 0.95]	
	$\sigma_u = 10$	0.99 [0.99, 1.00]	0.98 [0.96, 0.99]	

B.1 Additional simulation results

Sudden stop freq. [%]	$\text{corr}(\hat{q}_t, q_t)$	$\text{corr}(\hat{\mu}_t, \tilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\tilde{\mu}_t$
[10.0, 15.0]	0.90 [0.80, 0.96]	0.81 [0.53, 0.93]	1.35 [1.19, 1.61]	0.63 [0.39, 1.12]
[30.0, 35.0]	0.90 [0.82, 0.95]	0.88 [0.79, 0.94]	1.39 [1.20, 1.57]	0.86 [0.57, 1.20]
[50.0, 55.0]	0.81 [0.76, 0.89]	0.90 [0.81, 0.96]	1.39 [1.13, 1.64]	1.06 [0.83, 1.33]

Table B.4: Medians [68% CI] of correlations, relative standard deviations, and classification errors. VAR with Minnesota prior.

Sudden stop frequency	Crisis classification		
	False positives	False negatives	False(any)
[10.0, 15.0]	1.0 [0.0, 3.1]	10.9 [0.0, 40.3]	3.0 [1.2, 6.1]
[30.0, 35.0]	4.3 [1.2, 7.5]	1.3 [0.0, 10.1]	3.7 [2.4, 5.9]
[50.0, 55.0]	7.7 [2.5, 11.7]	0.0 [0.0, 2.7]	4.1 [2.4, 5.7]

Table B.5: Crisis classification errors. VAR with Minnesota prior.

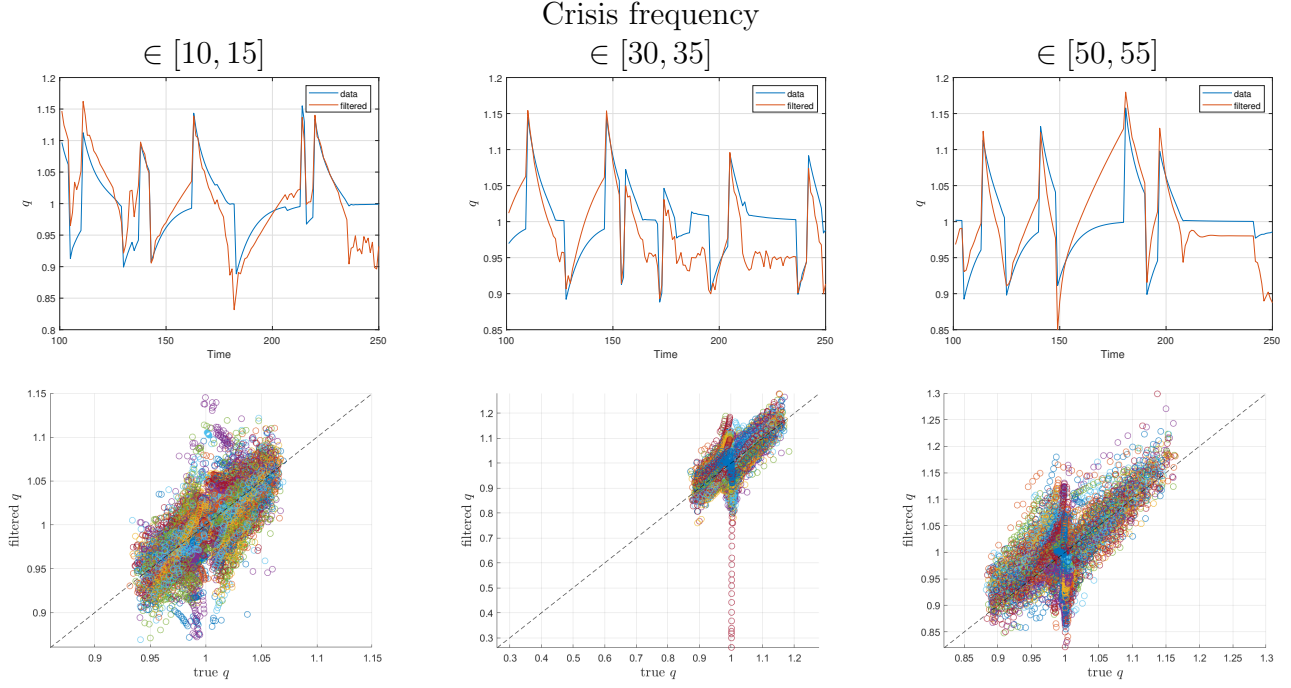


Figure B.9: Sample paths for q_t with the median correlation (top panels) and scatter plot pooling all samples (bottom panels) as a function of the in-sample crisis frequency.

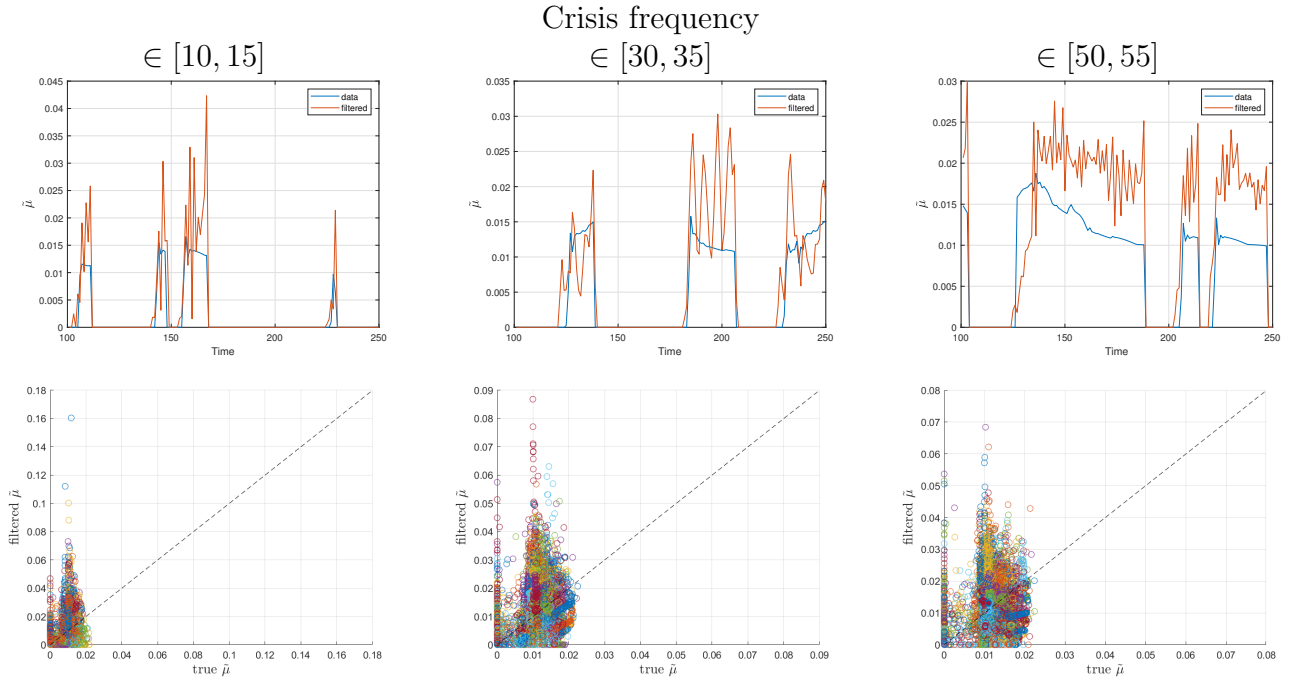


Figure B.10: Sample paths for the multiplier on the borrowing constraint $\tilde{\mu}_t$ for the sample paths with the median correlation (top panels) and scatter plot pooling all samples (bottom panels) as a function of the in-sample crisis frequency.

B.2 Data

Mapping data to the model. [Mendoza \(2010\)](#) defines GDP as output minus intermediate inputs.

$$GDP \equiv \exp \epsilon^A F(k, L, v) - pv.$$

Thus, F is gross output, not GDP:

$$y \equiv \exp \epsilon^A F(k, L, v) = GDP + pv.$$

The static optimality conditions yield that:

$$\eta \frac{y}{v} = p(1 + \phi(R - 1)) \quad \Leftrightarrow \quad pv = \eta \frac{y}{1 + \phi(R - 1)}.$$

Plugging in:

$$GDP = y - pv = y \frac{1 + \phi(R - 1) - \eta}{1 + \phi(R - 1)} \quad \Leftrightarrow \quad y = \frac{1 + \phi(R - 1)}{1 + \phi(R - 1) - \eta} GDP$$

The static optimality conditions yield that:

$$\begin{aligned} \alpha \frac{y}{L} &= w(1 + \phi(R - 1)) \\ \eta \frac{y}{v} &= p(1 + \phi(R - 1)) \end{aligned}$$

so that:

$$(wL + pv)(1 + \phi(R - 1)) = (\alpha + \eta)y \quad \Leftrightarrow \quad (1 + \phi(R - 1)) \frac{wL + pv}{y} = (\alpha + \eta).$$

Plugging in for y in terms of GDP:

$$(1 + \phi(R - 1) - \eta) \frac{wL + pv}{GDP} = (\alpha + \eta).$$

Dividends are net of depreciation:

$$d_t = (1 - \alpha - \eta) \frac{y_t}{k_t} - \delta + \left(\frac{k_{t+1} - k_t}{k_t} \right)^2 \Psi' \left(\frac{k_{t+1} - k_t}{k_t} \right).$$

In terms of GDP:

$$d_t = \frac{1 + \phi(R - 1)}{1 + \phi(R - 1) - \eta} (1 - \alpha - \eta) \frac{GDP_t}{k_t} - \delta + \left(\frac{k_{t+1} - k_t}{k_t} \right)^2 \Psi' \left(\frac{k_{t+1} - k_t}{k_t} \right).$$

Then, the total dividends in the economy are:

$$k_t d_t = \frac{1 + \phi(R - 1)}{1 + \phi(R - 1) - \eta} (1 - \alpha - \eta) GDP_t - \delta k_t + \left(\frac{k_{t+1} - k_t}{k_t} \right)^2 \Psi' \left(\frac{k_{t+1} - k_t}{k_t} \right) k_t$$

The slackness condition depends on $\phi R_t(w_t L_t + p_t v_t)$. Relative to GDP:

$$\frac{q_t^b b_{t+1}}{GDP_t} - \frac{\phi R_t}{1 + \phi(R_t - 1) - \eta} (\alpha + \eta) + \kappa \frac{q_t k_{t+1}}{GDP_t} \geq 0.$$

Data sources. When we filter the model with Mexican data, our measure for real consumption and real GDP are also taken from FRED (mnemonics NAEXKP02MXQ661S and NAEXKP01MXQ661S; [OECD](#)). See Table B.6 for access information.

Our measure for hours worked is an index of monthly hours worked in the Mexican manufacturing sector (FRED mnemonic: HOHWMN03MXQ661N; [OECD](#)). We convert the measure to logs and regress it on quarterly dummies to remove seasonalities. Because manufacturing hours are about twice as volatile as overall hours, we divide the fluctuations by a factor of two.¹⁶

With respect to the interest rate, when available, we use the LIBOR as our measure of the world interest rate ([ICE Benchmark Administration](#); since 2022 this series is no longer available on FRED). Prior to the availability, we use the 3-month U.S. Treasury rate ([Board of Governors of the Federal Reserve System](#)), shifted so that the two measures agree during the first period of overlap. To convert these nominal measures into real rates, we subtract the one-year realized change in the log of the U.S. CPI ([U.S. BLS](#)).

We construct the dividend series using the dividends from the S&P/IFCG M MEXICO (IFGMMX) index, appended after 2008Q3 with the MSCI MEXICO index (MSMEXFL; [Datastream International](#)). We rescale the latter series to have the same mean as the IFGMMX series, calculated over the 1988Q1 to 2008Q3 sample. We then denominate dividends in 1980Q1 dollars, take the log, use the x13 filter to deseasonalize them, and then detrend

¹⁶In the U.S., the ratio of standard deviations of average weekly hours in manufacturing relative to the total private sector is 2.51 at a quarterly frequency and 1.95 at an annual frequency from 1965 to 2021 (based on the series with mnemonics AWHMAN and AWHNONAG; [U.S. BLS](#)). For Mexico, where we can compare it to annual data only from 1995 onward [Penn World Table 10.01](#), the ratio is 1.52. While the annual data from Mexico indicate lower volatility, applying the same ratio of annual to quarterly volatility from the U.S. to Mexican data suggests an adjustment by a factor of 1.96: $\frac{sd_{mfq,US}^q}{sd_{mfq,US}^a} sd_{mfq,MX} = \frac{2.51}{1.95} 1.52 = 1.96$.

Source	Title	Link
OECD	Gross Domestic Product by Expenditure in Constant Prices: Private Final Consumption Expenditure for Mexico [NAEXKP02MXQ661S]	https://fred.stlouisfed.org/series/NAEXKP02MXQ661S (Accessed November 2021).
OECD	Gross Domestic Product by Expenditure in Constant Prices: Total Gross Domestic Product for Mexico [NAEXKP01MXQ661S]	https://fred.stlouisfed.org/series/NAEXKP01MXQ661S (Accessed November 2021).
OECD	Monthly Hours Worked: Manufacturing for Mexico [HOHWMN03MXQ661N]	https://fred.stlouisfed.org/series/HOHWMN03MXQ661N (Accessed November 2021).
U.S. BLS	Average Weekly Hours of Production and Nonsupervisory Employees, Manufacturing [AWHMAN]	https://fred.stlouisfed.org/series/AWHMAN (Accessed November 2021).
U.S. BLS	Average Weekly Hours of Production and Nonsupervisory Employees, Total Private [AWHNONAG]	https://fred.stlouisfed.org/series/AWHNONAG (Accessed November 2021).
U.S. BLS	Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCNS]	https://fred.stlouisfed.org/series/CPIAUCNS (Accessed November 2021).
OECD	Average Annual Hours Worked by Persons Engaged for Mexico [AVHWPEMXA065NRUG],	https://fred.stlouisfed.org/series/AVHWPEMXA065NRUG (Accessed November 2021).
ICE Benchmark Administration	3-Month London Interbank Offered Rate (LIBOR), based on U.S. Dollar [USD3MTD156N]	https://fred.stlouisfed.org/series/USD3MTD156N (Accessed July 2020 – no longer available).
Board of Governors of the Federal Reserve System	3-Month Treasury Bill Secondary Market Rate, Discount Basis [DTB3]	https://fred.stlouisfed.org/series/DTB3 (Accessed July 2020).

Note: All data retrieved from FRED, Federal Reserve Bank of St. Louis.

Table B.6: Mexican and international data retrieved from FRED: Sources with links.

log real dividends using the (quarterly) HP filter. Last, we rescale the mean to match the model mean marginal product of capital.

To evaluate the slackness condition, we need data on GDP and capital formation. We use annual data from [CEPALSTAT](#); see Table B.7. Specifically, we splice Annual Gross Domestic Product (GDP) by Expenditure at constant prices in National Currency [various base years] and similarly for Gross Fixed Capital Formation. To construct capital, we assume the capital stock starts at the same average value of GDP in 1960 as in the model, and then we add the investment series to the capital stock estimate, using depreciation of 8.8%. This is correct

Source	Title	Link
Banco de México	Deuda Neta Total del Sector Público, saldo promedio y final - (CG7)	https://www.banxico.org.mx/SieInternet/consultarDirectorioInternetAction.do?sector=9&accion=consultarCuadro&idCuadro=CG7&locale=es (Accessed January 2022).
CEPALSTAT	Producto interno bruto anual (PIB) por objeto del gasto a precios corrientes en moneda nacional: Producto interno bruto (PIB), México [año base 1960, ..., 2013]	https://statistics.cepal.org/portal/cepalstat/dashboard.html?theme=2&lang=es (Accessed January 2022).
CEPALSTAT	Producto interno bruto anual (PIB) por objeto del gasto a precios corrientes en moneda nacional: Formación bruta de capital fijo, México [año base 1960, ..., 2013]	https://statistics.cepal.org/portal/cepalstat/dashboard.html?theme=2&lang=es (Accessed January 2022).

Table B.7: Mexican public debt and annual national account data: Sources with links.

according to the model up to a first-order approximation.

We obtain data for debt from [Banco de México](#).¹⁷ We compute the quarterly debt in current prices of pesos using the end-of-quarter value. We convert the resultant values to a 2013 base year like the other variables in the slackness condition.

We center all variables at the model-implied means before detrending with the HP filter. Results with the Baxter-King filter are similar.

¹⁷See Table B.7 for access information.

B.3 Details of the implementation of the filter

Algorithm 1 Fixed-point algorithm for the sudden stops model

- Step 0 Initialize the starting guesses as $q_t^{(0)} = 1 + \mathcal{U}[-600^{-1}, 600^{-1}]$ and $\tilde{\mu}_t$ as a dummy for the event of a violated borrowing constraint at $q_t = 1$, i.e., as a dummy for $\phi R_t(w_t L_t + p_t v_t) - q_t^b b_{t+1} > \kappa k_{t+1}$. Set $d = 0$ and an upper bound on iterations D .
- Step 1 Set $d = d + 1$.
- Step 2 Estimate the VAR parameters $\hat{\mathbf{a}}^{(d)}, \hat{\mathbf{A}}^{(d)}, \hat{\mathbf{\Sigma}}^{(d)}$ using OLS on the data and the scaled guesses for $[q_t^{(d-1)}/\bar{q}^{(d-1)}, \tilde{\mu}_t^{(d-1)}]$.
- Step 3 Solve for the latent variables $[q_t^{(d)}, \tilde{\mu}_t^{(d)}]^T$ given $\hat{\mathbf{a}}^{(d)}, \hat{\mathbf{A}}^{(d)}, \hat{\mathbf{\Sigma}}^{(d)}$.
- (a) Solve equations (4.10) and (4.11) for $[q_{*t}^{(d)}, \tilde{\mu}_{*t}^{(d)}]^T$ with $\mathbb{E}_t[u_{c,t+1}] = \hat{\mathbb{E}}_t^{(d)}[u_{c,t+1}] + \Delta_t^{(d)} = \mathbf{e}'_{u,c}(\mathbf{a}^{(d)} + \mathbf{A}^{(d)} \hat{x}_t^{(d-1)}) + \Delta_t^{(d-1)}$.
 - (b) Update $\Delta_t^{(d)}$: If the implied solution violates the borrowing constraint at time s but $\tilde{\mu}_{*t}^{(d)} = 0 (\leq 0)$, set $\Delta_s^{(d)} = \Delta_s^{(d-1)} - 4e - 6$ for all s with binding constraints. If $\tilde{\mu}_{*t}^{(d)} > 0$, but the borrowing constraint is higher than in p th percentile in sample, set $\Delta_s^{(d)} = \Delta_s^{(d-1)} + 4e - 6$. In the simulation, we set p equal to the target in-sample frequency plus 10p.p. (i.e., $p \in \{0.25, 0.45, 0.55\}$). In the empirical application, we set $p = 0.15$.
 - (c) Set $q_t^{(d)} = (1 - s_q^{(d-1)})q_t^{(d-1)} + s_q^{(d-1)}q_{*t}^{(d)}$ and $\tilde{\mu}_t^{(d)} = \max\{(1 - s_\mu^{(d-1)})\tilde{\mu}_t^{(d-1)} + s_\mu^{(d-1)}\tilde{\mu}_{*t}^{(d)}, -5e - 4\}$. Here, $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.01$ for $d = 1, \dots, D/10$, $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.02$ for $d = D/10 + 1, \dots, D/5$, and $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.0025$ for $d = D/5, \dots, D$.
 - (d) For $d \leq D/20$ truncate $q_t^{(d)}$ at 1 ± 0.1 and $q_t^{(d)} = \frac{q_t^{(d)}}{T^{-1} \sum_s q_s^{(d)}}$.
- Step 4 If $T^{-1} \sum_{t=1}^T |q_t^{(d)} - q_t^*| > 1e^{-6}$ or the number of periods with positive multipliers $\tilde{\mu}_t$ changes within the last 100 iterations, go to Step 1. Otherwise, save and report the estimates.
-

Panels (a) and (b) of Figure B.11 show the inferred $\tilde{\mu}_t$ and q_t over time for a VAR with a flat prior and lag lengths ranging from one to six. The VAR with four lags is highlighted as the thicker, solid, and purple line. In both cases, VARs with more than one lag give a consistent answer. The VAR with a single lag yields smaller values for $\tilde{\mu}_t$ and different dynamics for q_t . Figure B.13 suggests that this is due to the dividend dynamics, which the VAR(1) may not adequately capture. Panels (c) and (d) of Figure B.11 show the analogous results for a VAR estimated with a weak Minnesota prior. The results are very similar to the

results with a flat prior. Overall, the results suggest that once the VAR is sufficiently rich, the results are insensitive to the lag length.

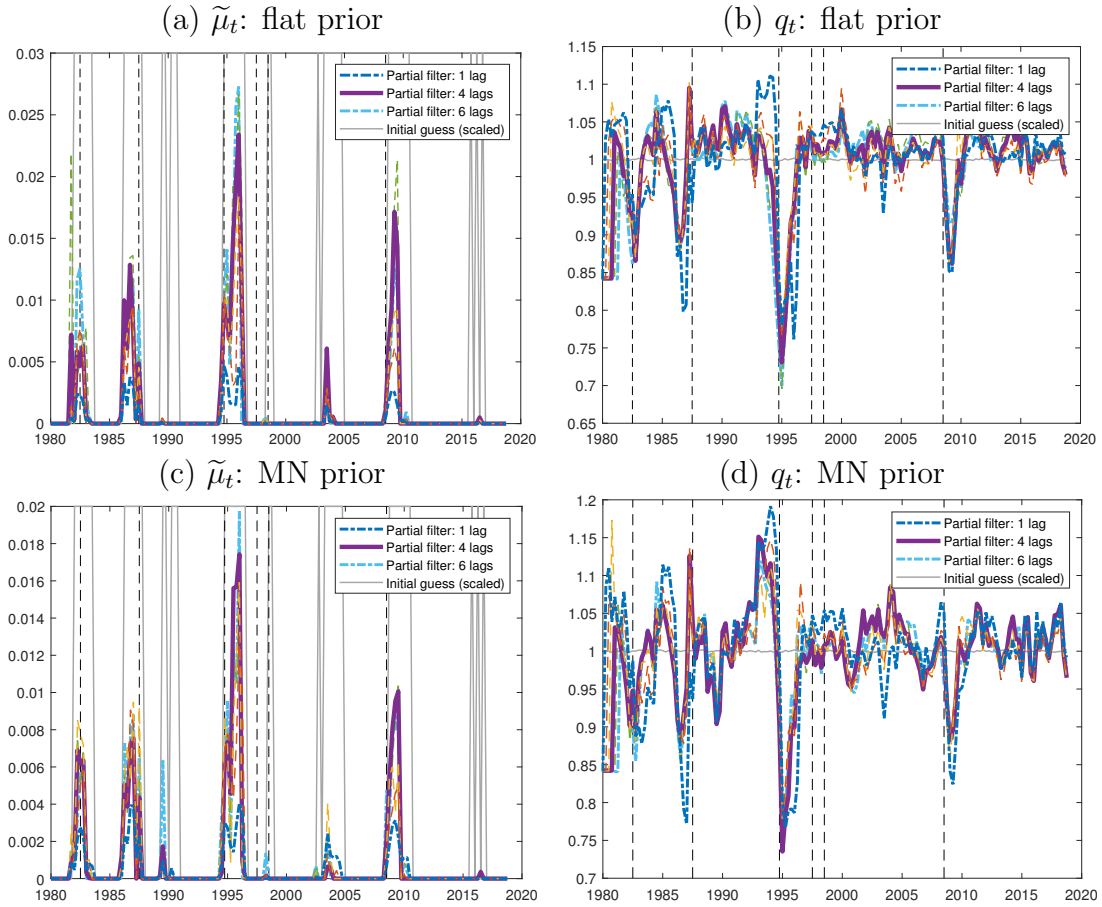


Figure B.11: Filtered values of $\tilde{\mu}_t$ and q_t for different VAR lag lengths.

Panels (a) and (b) of Figure B.12 show that the initial period aside, the inference on the crisis periods does not hinge on the initial guess. Still, there are some differences in the inferred severity of the crisis. These, in turn, have some effects on the inferred values of Tobin's q during crisis times, such as in the period following 1995. Overall, however, the inference depends little on the initial conditions.

Figure B.13 shows that the filtered Tobin's q closely tracks the fluctuations in measured dividends, crisis periods aside.

Figure B.14 shows that the Gibbs posterior has converged. It compares the 90% and 68% credible sets as well as the medians for the hidden co-state variables q_t and $\tilde{\mu}_t$ for the initial 500 draws (after discarding a burn-in; shown as area plots) and the subsequent 500 draws (shown as dashed lines). The credible sets are hard to tell apart, and only small inaccuracies are visible for the 90% credible set. For example, the lower panel shows that, in the 2007-2008 financial crisis, the 95th percentile is slightly higher in the second half of the posterior than

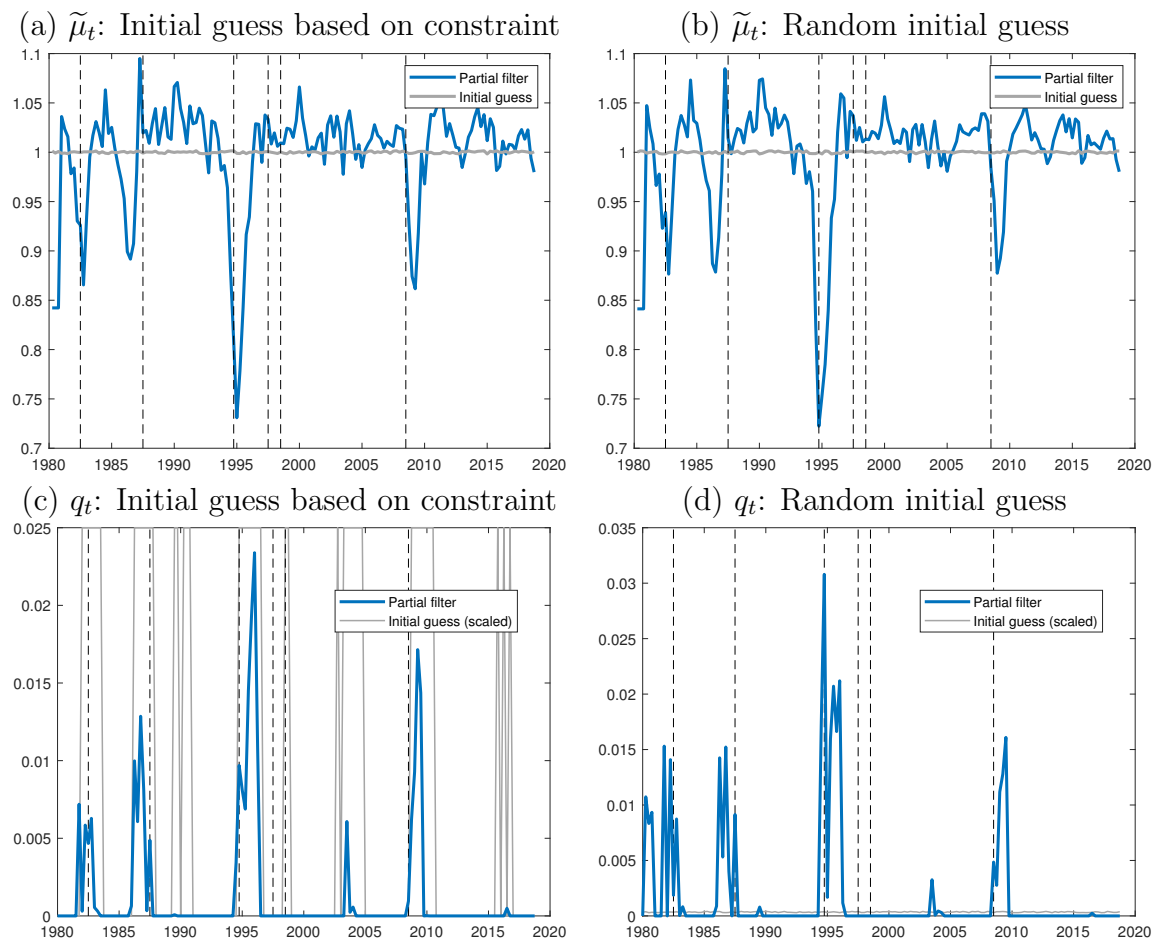


Figure B.12: Filtered values of $\tilde{\mu}_t$ (top) and q_t (bottom) given different initial guesses.

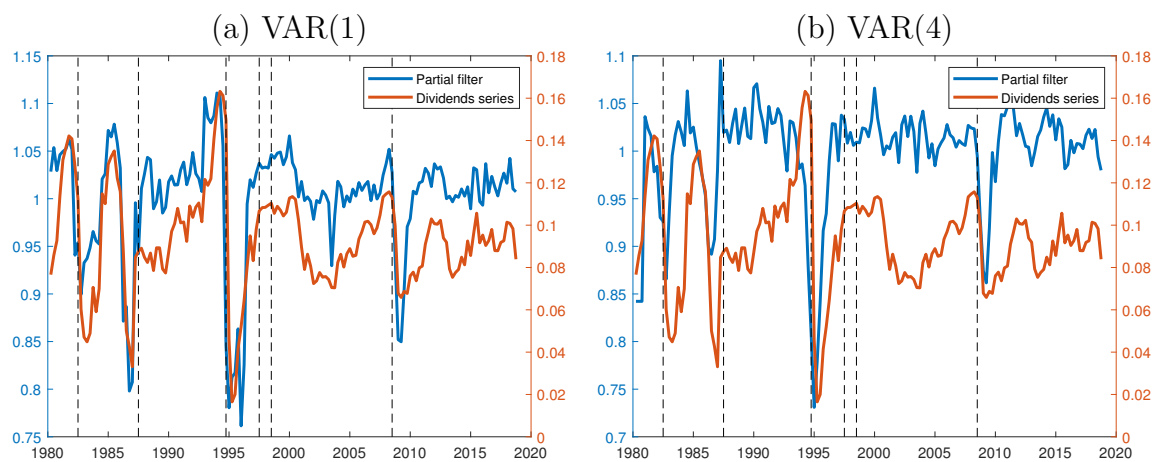
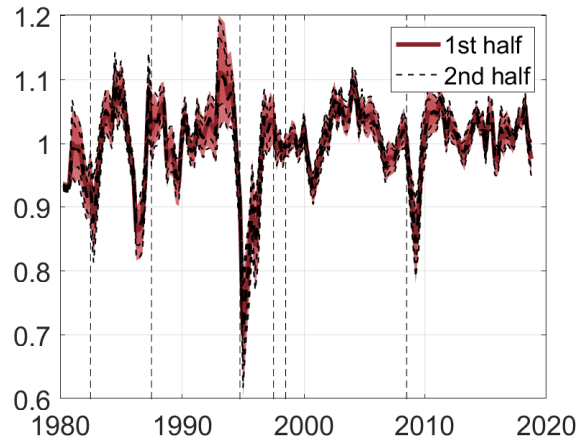


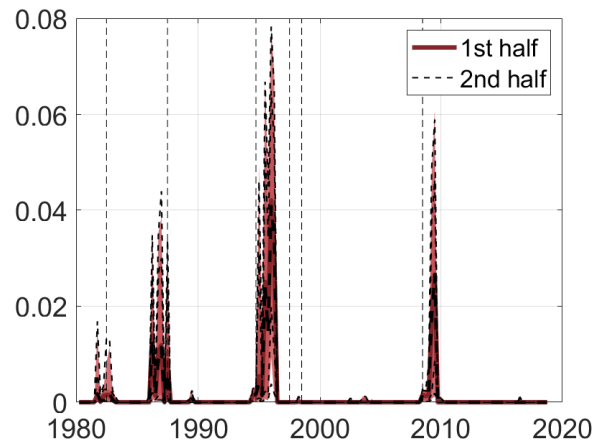
Figure B.13: Comparison of q_t in the partial filter and the dividends series.

in the first half. However, even for other local extremes in the Lagrange multiplier, such as in the mid-1990s, the credible sets are hard to distinguish. The estimates for Tobin's q line up even more closely than those for $\tilde{\mu}$.

(a) Tobin's q



(b) Lagrange multiplier $\tilde{\mu}$



Shown are the median, 68% and 90% credible sets for the first half and the second half of the posterior draws.

Figure B.14: Comparison of the posterior based on the first half and the second half of the Gibbs sampler draws.