Filtering with limited information

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

October 21, 2022

Abstract

We propose a new tool to filter non-linear dynamic models that does not require us to fully specify the model 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 a hidden state. 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 small open economy model with an occasionally binding constraint.

*Drautzburg: Federal Reserve Bank of Philadelphia, tdrautzburg@gmail.com. Fernández-Villaverde: UPenn, jesusfv@econ.upenn.edu. Guerrón-Quintana: Boston College and ESPOL, pguerron@gmail.com. Oosthuizen: UPenn, dickoos@sas.upenn.edu. We thank Felipe Meza for useful comments. The views expressed herein are our own views only. They do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or its Board of Governors.

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 prominent reasons for this interest. First, the path of shocks (or their distribution) are a reality check of the model. If 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 sharp narrative analytics), a researcher may want to consider revisiting the specification of the model or 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 perform variance decompositions and build counterfactuals. Third, knowing the current state of the economy, including the shocks, is important when forecasts or optimal policy 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 can be completed, for a generic, fully-specified dynamic equilibrium model, with a non-linear filter, such as the sequential Monte Carlo filter described in 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 of it such as some auxiliary functional forms (adjustments costs, evolution of some exogenous variables) that are not central to the message of the model and that various researchers may disagree on. Furthermore, some models' findings are fragile with respect to these assumptions (see Canova and Ferroni, 2018) and we may want to have findings that do not suffer from this fragility. 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 very large models.[1]

In this paper, we tackle the problem of non-linear filtering from a different perspective. More concretely, we follow (Hansen, 2014, p. 950)'s dictum: "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 economic relations are formally delineated." (Hansen, 2014, p. 950). 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 solving and 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

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

(which can sometimes be hard or even impossible 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 (how many moments are required depends on the order of the perturbation). An early application, 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 they focus on parameter estimates, we focus on filtering the hidden state variables. An advantage of our filtering method is the ease of implementation.

Specifically, our approach is to approximate the state dynamics with a flexible statistical model, such as a Vector-Autoregression (VAR). This VAR may include non-linear terms and the 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 economic 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. In 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 new in macroeconomics. For example, Campbell (1991), Bernanke and Kuttner (2005), and Chahrour et al. (2021) use VARs to approximate expected returns. Closer to the structural nature of our paper, Sbordone (2002) uses VAR-based expectations to stand-in for firm's expectations in a New Keynesian price setting equation in her test of price stickiness. Our innovation is that, under regularity conditions, we can use a VAR-approximation also for filtering hidden state variables.

Models with departures from rational expectations are increasingly common, see, for example, Gabaix (2020). While our application focuses on rational expectations, it is straightforward to allow for parametric departures from rational expectations in our framework, for example by including the appropriate change of measure in the filtering equation. In the case of the cognitive discounting in Gabaix (2020), for example, the deviation of the forecast from the balanced growth path is shrunk by a factor that is constant for a given forecast horizon. Other departures, such as limited information could be addressed by limiting the information set we use to compute the approximating the VAR.

[2]While 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 in this form is a straightforward extension of our model.

Going back to at least Roberts (1995), a growing literature has used survey data for assessing macroeconomic models, see, for example, Coibion et al. (2021). With representative agents with model-consistent expectations,[3] one can use survey expectations to substitute our for expectations in economic models – for example, as a measure of inflation expectations in the New Keynesian Phillips Curve (Roberts (1995)). While we focus on models with a single set of model-consistent expectations, the use of surveys is still limited to accounting for expectations over readily understood objects such as inflation, but precludes structural objects such as shocks or co-state variables. For these structural objects, eliciting expectations is not feasible, and a statistical approach such as ours is needed.

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 first uses a simulation study of a small-open economy real business cycle model with capital adjustment costs and an occasionally binding borrowing constraint to filter Tobin's Q and the multiplier on the borrowing constraint. Section 4 also applies the partial filter to the same model for Mexican data from 1980 to 2018. Section 5 concludes.

[3]See Mavroeidis et al. (2014) for a discussion.

2 The partial information filter

In this section, we present how the partial information filter works. First, we introduce the general environment. Second, we describe the two algorithms to accomplish the filtering. Third, we postulate some regularity conditions required by our procedure. Fourth, we point out some directions for extensions.

2.1 Environment

Let us consider a dynamic equilibrium model whose equilibrium conditions, or a subset thereof, can be described by:

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

where $\mathbf{y}_t = [y_{i,t}]_{i=1}^{n_y}$ is a $n_y \times 1$ vector of observables at period t, $\mathbf{x}_t$ a $n_x \times 1$ 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. Strictly speaking, 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 it is not important for expositional purposes. The deterministic steady state of the model, $\{\bar{\mathbf{x}}, \bar{\mathbf{y}}\}$, is defined by:

$$\mathbf{f}(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \mathbf{g}(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{x}}, \bar{\mathbf{y}})]) = \mathbf{0}.$$

We highlight three points about equation (2.1). First, the functions $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ do not need to include all equilibrium conditions required for a full solution of the model. Our goal is, precisely, to get away with handling 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 very 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 deal with 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 $\widetilde{\mathbf{x}} \in \mathbb{X}_t$ such that $\mathbf{f}(\widetilde{\mathbf{x}}, \mathbf{y}_t, \mathbb{E}[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)|\widetilde{\mathbf{x}}, \mathbf{s}_{-,t}]) = \mathbf{0}$.*

The following scalar example is useful to fix ideas:

Example 1. *Let the model be:*

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

$$x_t = \mu_x + \left(\rho_0 + \frac{\rho_1}{1 + x_{t-1}^2}\right) x_{t-1} + u_t, \tag{2.2b}$$

$$u_t \sim iid(0, \sigma_u^2), \tag{2.2c}$$

We consider two special cases.

(a) Linear model: $\kappa_2 = \rho_1 = 0$.
Here $f(\cdot) = y - \mu_y - x - \kappa_1 \mathbb{E}[x'|x]$. *Under the population expectation,* $\mathbb{E}[x'|x] = \mu_x + \rho_0 x$ *so that* $f(\cdot) = y - \mu_y - (1 + \kappa_1\rho_0)x - \kappa_1\mu_x = 0$. *Thus, there is a unique solution in* $\mathbb{R}$ *given by* $x = \frac{y - \mu_y - \kappa_1\mu_x}{1 + \kappa_1\rho_0}$.

(b) Linear law of motion and quadratic observations equation: $\rho_1 = 0$.
Here,

$$y_t = \mu_y + x_t + \kappa_1\rho_0 x_t + 2\kappa_2\rho_0\mu_x x_t + \kappa_2(\mu_x^2 + \sigma_u^2) + \kappa_2\rho_0^2 x_t^2.$$

We require the observation to be monotonically increasing or decreasing in the hidden state:

$$\frac{\partial y_t}{\partial x_t} = 2\rho_0^2\kappa_2 x_t + (1 + \kappa_1\rho_0 + 2\kappa_2\rho_0\mu_x) \gtrless 0 \quad \Leftrightarrow \quad x_t \gtrless -\frac{1}{2}\frac{1 + \kappa_1\rho_0 + 2\kappa_2\rho_0\mu_x}{\rho_0^2\kappa_2} \equiv \underline{x}. \tag{2.3}$$

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 simply back out $\mathbf{x}_t$ directly from (2.1). 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 directly from survey data. We can consider the latter case as a particular application of our procedure when the expectation is about a concept general enough to be captured by surveys.[4] Survey expectations are, however, unavailable if the filtering exercise requires expectations about structural objects, such as hidden state variables.

Most commonly, therefore, we will deal with situations where $\mathbf{g}(\cdot)$ matters and $\mathbb{E}_t[\mathbf{g}(\cdot)]$ is not readily available to the researcher. Thus, we need to back out the hidden state vector by finding the sequence of $\widehat{\mathbf{x}}_t$ that solve the equation:

$$\mathbf{f}(\widehat{\mathbf{x}}_t, \mathbf{y}_t, \widehat{\mathbb{E}}_t[\mathbf{g}(\widehat{\mathbf{x}}_{t+1}, \mathbf{y}_{t+1}, \widehat{\mathbf{x}}_t, \mathbf{y}_t)]) = \mathbf{0}, \tag{2.4}$$

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

[4] For example, we may use firms' surveyed inflation expectations in the Phillips Curve when filtering for cost-push shocks.

To tackle the estimation of this conditional expectation, we will proceed in two steps. In the first step, we factorize the m-dimensional function $\mathbf{g}(\cdot)$ as:

$$\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i \equiv g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t) \times g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t). \tag{2.5}$$

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, for example in Bellman equations, and we may want to deal with different parts of these expectations separately. Also, the same model may admit different natural factorizations of $\mathbf{g}(\cdot)$. Choosing one depends on their suitability for the economic question being investigated or the available data.

Writing the terms in equation (2.5) 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:

$$\begin{aligned}\mathbb{E}_t[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] \equiv\ & \mathbb{E}_t[g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] \times \mathbb{E}_t[g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i] \\ & + \text{Cov}_t[g_{i,1}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i, g_{i,2}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)_i],\end{aligned} \tag{2.6}$$

where $i = 1, \ldots, m$ denotes the rows of $\mathbf{g}$, $\mathbf{g}_1$, and $\mathbf{g}_2$ that we require. This factorization is convenient because, in many cases, the equilibrium conditions of the model allow us to write the product of the conditional expectations in terms of time t observables.

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 (2.6). If some of the conditional first moments are also missing after using equilibrium conditions, the auxiliary statistical model may help us estimate them. With this evaluation (or estimation) of the expectations and the covariances estimate, we return to (2.4) and back out the hidden state vector.

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

Assumption 2. *As* $T \to \infty$, $\sup_{\mathbf{x},\mathbf{y},\mathbf{s}_-} |\mathbb{E}[\mathbf{g}(\mathbf{x}', \mathbf{y}', \widetilde{\mathbf{x}}, \mathbf{y})|\widetilde{\mathbf{x}}, \mathbf{s}_-] - \widehat{\mathbb{E}}[\mathbf{g}(\mathbf{x}', \mathbf{y}', \widetilde{\mathbf{x}}, \mathbf{y})|\widetilde{\mathbf{x}}, \mathbf{y}]| = \mathbf{0}$.

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

Example 2. *Consider again the previous example summarized in* (2.2).

(a) Linear law of motion ($\rho_1 = 0$) *with bounded domain:*
Given $\rho_1 = 0$, *we have that* $\mathbb{E}[g(x')|x, s_-] = \mathbb{E}[g(x')|x] = cons + (\kappa_1\rho_0 + 2\kappa_2\rho_0\mu_x)x + \kappa_2\rho_0^2x^2$. *To bound the domain, let* $\mathbb{X} = [\underline{x}, \bar{x}]$ *and* $x \in \mathbb{X}$. *Let* $\widehat{\rho}$ *be a consistent estimator, such as the OLS estimator* $\widehat{\rho} \xrightarrow{p} \rho_0$ *as* $T \to \infty$. *Then:*

$$\begin{aligned}&\sup_{x\in\mathbb{X}} |\mathbb{E}[g(x')|x, s_-] - \widehat{\mathbb{E}}[g(x')|x, y]| \\ =&\sup_{x\in\mathbb{X}} |(\kappa_1\rho_0 + 2\kappa_2\rho_0\mu_x)x + \kappa_2\rho_0^2x^2 - (\kappa_1\widehat{\rho} + 2\kappa_2\widehat{\rho}\mu_x)x - \kappa_2\widehat{\rho}^2x^2|\end{aligned}$$

$$
\begin{aligned}
&= \sup_{x\in\mathbb{X}} |(\kappa_1 + 2\kappa_2\mu_x)x(\rho_0 - \widehat{\rho}) + \kappa_2 x^2(\rho_0^2 - \widehat{\rho}^2)| \\
&\leq \sup_{x\in\mathbb{X}} |(\kappa_1 + 2\kappa_2\mu_x)(\rho_0 - \widehat{\rho})||x| + \sup_{x\in\mathbb{X}} |\kappa_2(\rho_0^2 - \widehat{\rho})|x^2 \\
&= |(\kappa_1 + 2\kappa_2\mu_x)(\rho_0 - \widehat{\rho})|\max\{-\underline{x}, \bar{x}\} + |\kappa_2(\rho_0^2 - \widehat{\rho}^2)|\max\{\underline{x}^2, \bar{x}^2\} \\
&\leq |(\rho_0 - \widehat{\rho})||(\kappa_1 + 2\kappa_2\mu_x)|\max\{-\underline{x}, \bar{x}\} + |(\rho_0^2 - \widehat{\rho}^2)||\kappa_2|\max\{\underline{x}^2, \bar{x}^2\}
\end{aligned}
$$

But because of the continuous mapping theorem, we have that both $\rho_0 - \widehat{\rho}$ *and* $\rho_0^2 - \widehat{\rho}^2$ *converge to zero in probability. For any* $\epsilon > 0$*, we have that* $\lim_{T\to\infty} \Pr\{|\rho_0 - \widehat{\rho}| > \epsilon\} = 0$ *and* $\lim_{T\to\infty} \Pr\{|\rho_0^2 - \widehat{\rho}^2| > \epsilon\} = 0$*. Thus, for any* $\tilde{\epsilon} > 0$*, we can redefine* ϵ *such that* $\lim_{T\to\infty} \Pr\{|(\rho_0 - \widehat{\rho})||(\kappa_1 + 2\kappa_2\mu_x)|\max\{-\underline{x}, \bar{x}\} > \tilde{\epsilon}/2\} = 0$ *and* $\lim_{T\to\infty} \Pr\{|(\rho_0^2 - \widehat{\rho}^2)||\kappa_2|\max\{\underline{x}^2, \bar{x}^2\} > \tilde{\epsilon}/2\} = 0$*. This implies that* $\lim_{T\to\infty} \Pr\{|(\rho_0 - \widehat{\rho})||(\kappa_1 + 2\kappa_2\mu_x)|\max\{-\underline{x}, \bar{x}\} + |(\rho_0^2 - \widehat{\rho}^2)||\kappa_2|\max\{\underline{x}^2, \bar{x}^2\} > \tilde{\epsilon}\} = 0$ *and therefore* $\sup_{x\in\mathbb{X}} \Pr\{|\mathbb{E}[g(x')|x, s_-] - \widehat{\mathbb{E}}[g(x')|x, y]| > \tilde{\epsilon}\} = 0$*. Hence, 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* $h(x) = \kappa_1\mu_x + \kappa_1\left(\rho_0 + \frac{\rho_1}{1+x^2}\right)x$*. Chen and Christensen (2015, Theorem 2.1) show that, under regularity conditions in a time-series setting, 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. $\mathbf{f}$ *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* $\widehat{\mathbf{x}}^T$ *and* $\widehat{\mathbb{E}}[g(x', y', x, y)|x, y]$ *converge with* $\widehat{\mathbf{x}}_t \in \mathbb{X}_t \forall t$*, they converge to* $\mathbf{x}^T$ *and* $\mathbb{E}[g(x', y', x, y)|x, y]$ *as* $T \to \infty$*.*

Proof. Suppose not. It must be that $\widehat{\mathbb{E}}[\cdot]$ converges, because otherwise Assumption 2 is violated. Thus, $\widehat{\mathbf{x}}^T$ must diverge from $\mathbf{x}^T$. However, because the expectation converges, a diverging sequence of $\widehat{\mathbf{x}}^T$ 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 $\{\widehat{x}_t\}$ by some δ_ϵ. We can pick ϵ to make δ_ϵ vanish. Thus, if there is convergence, it is to the population values. □

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 $g_{1,t}, g_{2,t}$, (a subset of $n_{\widetilde{y}}$ elements of) $\mathbf{y}_t$, and $\widehat{\mathbf{x}}_t$ for $t = 1, ..., T$ as our auxiliary statistical model. We will denote the whole sequence of a variable by a superscript. For instance, $\mathbf{y}^T = \{\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_T\}$. We collect all the VAR variables in $\boldsymbol{\xi}_t$, a $(2m + n_{\widetilde{y}} + n_x) \times 1$ vector. Some applications may call for the inclusion of non-linear terms, which we collect in the vector $\boldsymbol{\zeta}_t = [\zeta_{j,t}]_j$, for example $\boldsymbol{\zeta}_t = [x_{j,t}^2, y_{\ell,t}^2]$ for some indices j, ℓ to include selected lagged hidden states and observables as predictors. We use the self-explanatory notation:

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

We can extend the notation to the general VAR case with more lags using the companion form of the VAR. While, given the structure of many time series, a low dimensionality VAR will capture their dynamics well, the reader can adapt our steps to her favorite auxiliary statistical model.

To express the expectation in terms of the VAR objects, 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 zero everywhere else, and using that the first $2m$ elements of the VAR characterizing g_1 and g_2, we can write the components of (2.6) as:

$$\mathbb{E}_t[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)]_i \equiv \mathbf{e}_i'(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\zeta}_t)\mathbf{e}_{i+m}(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\zeta}_t) + \mathbf{e}_i'\boldsymbol{\Sigma}\mathbf{e}_{i+m} \tag{2.7}$$

and get

$$\mathbf{f}(\mathbf{x}_t, \mathbf{y}_t, \mathbf{e}_i'(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\zeta}_t)\mathbf{e}_{i+m}(\boldsymbol{\mu} + \mathbf{A}\boldsymbol{\xi}_t + \mathbf{B}\boldsymbol{\zeta}_t) + \mathbf{e}_i'\boldsymbol{\Sigma}\mathbf{e}_{i+m})_i = 0 \tag{2.8}$$

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

We propose two approaches to filtering with the previous equation. First, we consider finding the sequence $\widehat{\mathbf{x}}^T$ that solves equation (2.8) for all i. Because $\widehat{\mathbf{x}}_t$ appears in $\boldsymbol{\xi}$ (and possibly in $\boldsymbol{\zeta}$) as well as indirectly in the VAR parameter matrices $\mathbf{A}$, $\mathbf{B}$, and $\boldsymbol{\Sigma}$, we are searching for a fixed point in (2.8). Second, we consider finding $\widehat{\mathbf{x}}^T$ through a Gibbs sampler that quantifies the estimation uncertainty.

Fixed point. When the measurement equation maps to a unique state, the fixed point of equation (2.4) is unique. With finite data, we can, in principle, solve numerically for the finite-dimensional $\widehat{\mathbf{x}}^T$ and the associated estimator satisfying that condition. In practice, we have found that initializing $\widehat{\mathbb{E}}_t^{(0)}[\mathbf{g}(\cdot)]$ based on $\widehat{\mathbf{x}}_t^{(0)} = \bar{\mathbf{x}}\ \forall t$ and $\mathbf{A}^{(0)} = \mathbf{0}$, $\boldsymbol{\mu}^{(0)} = \bar{\boldsymbol{\xi}}_T$ and $\boldsymbol{\Sigma}^{(0)} = \mathbf{0}$ and then iterating on the VAR estimation the filtering until convergence works well.

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

1. Given the previous draw $\{\widehat{\mathbf{x}}_t^{(d)}\}_{t=1}^T$, 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 $\boldsymbol{\zeta}_t^{(d)}$.

2. Given $\mathbf{y}^T$ and the VAR parameters, solve for $\{\widehat{\mathbf{x}}_t^{(d+1)}\}_{t=1}^T$ using (2.7) in equation (2.4).

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 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 fully exploit the model information, we employ the complementary slackness condition to test if the inequality constraint would be violated. We can then introduce and solve for time-varying intercepts for expected marginal utility, say, for the period with binding constraints to ensure that the model multiplier correctly identifies periods with a binding constraint.

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 on expectation and one may need to adjust the estimator for the expectations. For instance, in the linear model with *iid* measurement error, the use of lagged filtered values as an instrument may correct for otherwise present attenuation bias. Below, we introduce measurement error in the simulation but treat it as a form of model misspecification unknown to the researcher.

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

3 A scalar example

We now return to the simple model (2.2) and explore the properties of the partial filter in a Monte Carlo exercise. We also consider possible misspecification in the form of measurement error and an MA(1) structure of the innovations to assess the robustness of the proposed filter.[5]

[5]In principle, we could augment the state vector of the partial filter to capture these elements of the DGP in a correctly specified model. However, we would then need to introduce use extra information to recover these states.

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:

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

$$x_t = \mu_x + \gamma_0 \cos(\gamma_1(t-1)) + \left(\rho_0 + \frac{\rho_1}{1+x_{t-1}^2}\right) x_{t-1} + u_t, \tag{3.1b}$$

$$u_t = \frac{\phi_0 v_t + \phi_1 v_{t-1}}{\sqrt{\phi_0^2 + \phi_1^2}}, \tag{3.1c}$$

where $e_t \sim iid(0, \sigma_e^2)$ and $v_t \sim iid(0, \sigma_u^2)$. Here, e_t has the interpretation of measurement error, while v_t is a shock to the hidden state.

We treat the observation equation (3.1a) as known by the researcher. Examples with such an expectation terms 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: $\kappa_1 = 1$.

The state equation has a deterministic and a stochastic component. The stochastic component is first-order Markov, but non-linear if $\rho_1 \neq 0$. Since we are after applications with incompletely specified models, we allow for a forcing term and ARMA dynamics that we deliberately ignore in the partial filter setup to investigate the behavior under misspecification or partial specification. The forcing term is a cosine function of calendar time t, where γ_0 controls the amplitude and γ_1 the frequency. We allow for the innovation to x_t to follow an MA(1) process. When $\phi_1 \neq 0$, there are two hidden state variables: x_t and v_t.

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 VAR dynamics and observables. The VAR dynamics could be known from simulating a structural model, or estimated - in the simplest case, iterating between state inference and VAR estimation. We discuss each step next.

Step 1. We approximate the dynamics of the hidden state x_t with a VAR in the hidden state and the observables. It includes non-linear terms:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} \approx \mathbf{a}_0 + \sum_{\ell=1}^{L} \left(\mathbf{A}_\ell \begin{bmatrix} x_{t-\ell} \\ y_{t-\ell} \end{bmatrix} + \sum_{p=2}^{P} \mathbf{b}_{\ell,p} x_{t-\ell}^p \right) + \boldsymbol{\varepsilon}_t. \tag{3.2}$$

This VAR uses L lags and allows for a $P-$th order polynomial in the hidden state. While we impose that the state variable captures the non-linear dynamics, we could test or relax this assumption. That is, we could check whether $\mathbf{b}_{\ell,p}$ is statistically different from zero.

Step 2. We apply the statistical identity $\mathbb{E}[Z^2] = \mathbb{E}[Z]^2 + \text{Var}[Z]$ to the measurement equation:

$$\mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}] = \mathbb{E}_t[\kappa_1 x_{t+1} + \kappa_2 x_{t+1}^2] + \kappa_2^2 \,\text{Var}_t[x_{t+1}]. \tag{3.3}$$

Step 3. We plug the VAR dynamics (3.2) into the expression for the expectation terms in the observation equation and solve for the state, treating the measurement error e_t as zero. Given the non-linearity, this yields, generally, P solutions if $\kappa_2 = 0$ and $2P$ solutions when $\kappa_2 \neq 0$. Among these candidates, we choose solutions based on their predictive density. We could also use economic intuition to rule out solutions.

The point estimate of the partial filter is a fixed point of VAR parameters and 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 VAR parameter estimates. If one is interested in quantifying estimation uncertainty, one renders this iterative procedure into a Gibbs sampler by drawing the VAR parameters from the posterior, rather than using the point estimates.

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

1. Fully linear case: If $\kappa_2 = \gamma_0 = \rho_1 = 0$, the model is linear. In this case, the Kalman filter applies and is optimal when the disturbances are Gaussian.

2. Quadratic observation equation, linear state equation: If $\rho_1 = \gamma_0 = 0$, the state equation is linear. In this case, we need to 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.

By analyzing these three cases, we can show that our simple procedure performs well even in the presence of misspecification. The filter does well as long as the misspecification is not severe, i.e., provided that the signal-to-noise ratio is high enough. Also, we provide a comparison with the linear Kalman filter. Additionally, we consider a special case without misspecification.

More concretely, we conduct a Monte Carlo exercise using $N = 100$ replications. Throughout, we consider low to high values for of persistence $\rho_0 \in \{0.25, 0.5, 0.75, 0.9, 0.95\}$. We normalize the variance of the measurement error $\sigma_e^2 = 1$ and vary the unconditional variance of the stochastic component of the hidden state on the following grid: $\{4^{-2}, 1, 2^2, 5^2\}$.[6] We set $\phi_0 = 4 - \phi_1$, and

[6]The variance adjustment is for the linear law of motion for x_t, or for the first-order approximation to the non-linear law of motion.

vary the importance of the backward-looking MA component by letting $\phi_1 \in \{0, 1, 2, 3\}$. We use a burn-in in the simulations of 100 periods and a sample period of $T = 250$, a typical sample 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. Throughout, we assume a single lag in the VAR.

3.1 Linear model without misspecification

Before discussing the case of misspecification, we first examine a special case with neither measurement error nor MA component.

This is equivalent to increasing the signal to noise ratio to infinity in the case with measurement error.

Compared to the fully linear case under misspecification, in the case of no measurement error we can set the unconditional standard deviation of the process directly with σ_u^2. In addition, in the linear case ($\rho_1 = 0$) without measurement error, we have that x_t and y_t are perfectly collinear.

We can rewrite the observation equation as:

$$y_t = (1 + \rho_0)x_t, \tag{3.4}$$

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

$$x_t = \rho_0 x_{t-1} + v_t. \tag{3.5}$$

Due to the perfect collinearity of x_t and y_t, here the partial filter is very simple. We consider 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 $\widehat{x}_t^{(d)}$ directly, we simply iterate until convergence on the following two steps:

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

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 $\widehat{A}^{(d)}, \widehat{a}_0^{(d)}$ regressing $\begin{bmatrix} x_1^{(d)} \\ \vdots \\ x_T^{(d)} \end{bmatrix}$ on $\begin{bmatrix} x_0^{(d)} & 1 \\ \vdots & \\ x_{T-1}^{(d)} & 1 \end{bmatrix}$.

We proceed iteratively for comparison with the more complicated cases. In this simple model, we could also use a simple two-step procedure. Since for any initial guess of $\rho_0 \neq -1$ we have that $y_t \propto x_t$, we could use y_t or the initial guess for x_t to estimate $\widehat{\rho}_0$. We would then update $x_t = (1 + \widehat{\rho}_0)^{-1} y_t$.

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 $\rho_0 = 0.25$, the center panel when $\rho_0 = 0.75$, and the right panel when $\rho_0 = 0.95$. Within each panel, the true state x_t is measured on the

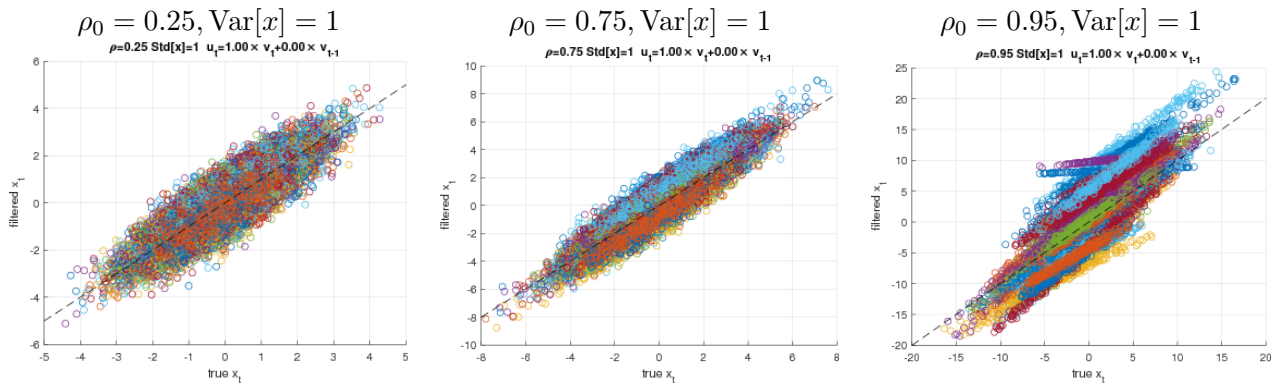

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

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

horizontal axis. Different colors correspond to the 100 different Monte Carlo samples. The dashed line is 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 scatterplots, it is not surprising that the partial filter here yields close to perfect correlations and that the relative standard deviations are just slightly above unity. See Appendix 5 for more details.

3.2 Linear case with misspecification

Given the MA(1) error term and knowing the true DGP, we can rewrite the observation equation as:

$$y_t = (1 + \rho_0)x_t + \frac{\phi_1}{\sqrt{\phi_0^2 + \phi_1^2}} v_t + e_t, \tag{3.6}$$

where x_t follows a zero mean ARMA(1,1) process.

The partial filter here is simple. We consider a linear VAR with a single lag, yielding $\widehat{\mathbb{E}}_t[x_{t+1}^{(d)}] = \widehat{\mathbf{A}}_{1,\cdot}[x_t^{(d)}, y_t^{(d)}] + \widehat{\mathbf{a}}_1$. While we could update $\widehat{x}_t^{(d)}$ directly, we simply iterate until convergence on the following two steps (treating e_t as zero):

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

Step 1. $x_t^{(d)} = y_t - \widehat{A}_{11} x_t^{(d-1)} - \widehat{A}_{12}^{(d-1)} y_t - \widehat{a}_1^{(d-1)}$

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

In this example as well as in the other two examples, we show two types of results. First, fixing the parameters of the DGP, we show scatter plots of the filtered series $\widehat{x}_t$ against the true state x_t.

While ignoring the temporal pattern, it flexibly shows the overall fit and shape of the filtered series. Second, we vary the parameters of the DGP and show how the correlation and relative standard deviation of the filtered vs the true series change with the DGP.

Figure 2 shows scatter plots for low to moderate persistence of the state with $\phi_1 = 0$.[7] The three panels differ in their persistence: The left panel has a low persistence of $\rho_0 = 0.25$, the center panel of $\rho_0 = 0.5$, and the right panel has a persistence of $\rho_0 = 0.75$. 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 line is 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 is relatively closer to the true state when the persistence is relatively higher.

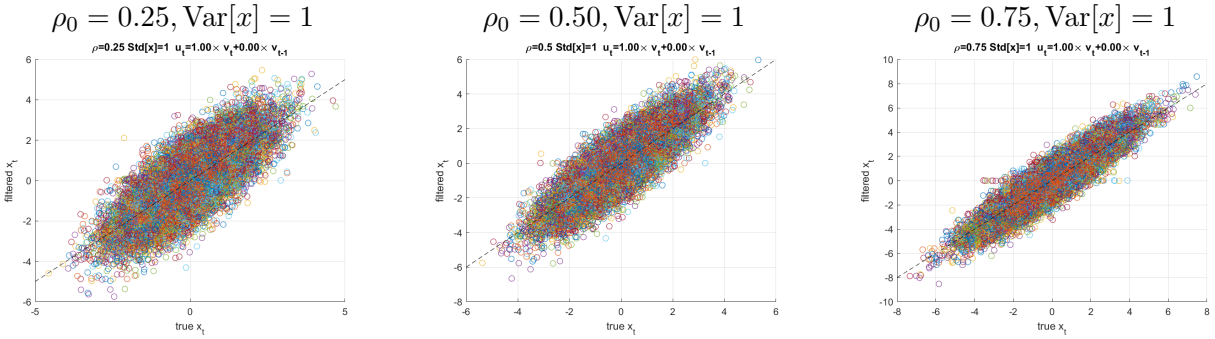

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 2: Partial filter in the fully linear case: Filtered state $\widehat{x}_t$ vs true x_t for 100 different samples for varying degrees of persistence ρ_0.

Table 1, Panels A1 and B1, evaluates the performance of the partial filter in terms of the correlation and standard deviation of the filtered and true state. The upper panel fixes the signal-to-noise ratio (at equal variances) and varies the persistence ρ_0, while the lower panel fixes the persistence at $\rho_0 = 0.75$ and varies the signal-to-noise-ratio. Both panels show median correlations and relative standard deviations, along with the 68% confidence intervals. Overall, there is little variation across Monte Carlo simulations, as indicated by the narrow confidence intervals relative to the median.[8] The left column of Panel B1 in the bottom row shows that the median correlations with a persistence of $\rho_0 = 0.5$ are above 0.9, unless the unconditional standard deviation of x_t is $\frac{1}{4}$, in which case it drops to 0.56. The left column in the top row shows that, holding fixed the unconditional variance of the state to equal the unit variance of the measurement error, the correlation ranges from 0.77 to above 0.98.

The relative standard deviation of the filtered and the true state (right columns of panels A1 and B1 in Table 1) behave similar to the correlations: Except for the signal to noise ratio below unity, the relative standard deviations are below 1.5. For high signal to noise ratios or very persistent states, the relative standard deviations are close to unity.

[7] The results for non-zero ϕ_1 are similar, but the higher persistence marginally increases the correlations between the filtered and true series; see Table A.1 in the Appendix.

[8] The dispersion across Monte Carlo draws falls with the median correlation.

(A) Varying the persistence ρ_0 of the state given $\text{Var}[e_t] = \text{Var}[x_t]$

	A1. Linear case		A2. Quadratic obs. eq.		A3. Non-linear state eq.	
ρ_0	Correlation	Rel st. dev.	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.77 [0.74,0.80]	1.31 [1.25,1.38]	0.80 [0.78,0.82]	0.12 [0.12,0.13]	0.90 [0.89,0.90]	1.21 [1.18,1.24]
0.50	0.85 [0.82,0.87]	1.21 [1.16,1.26]	0.97 [0.96,0.98]	0.43 [0.41,0.45]	0.93 [0.92,0.93]	1.19 [1.17,1.21]
0.75	0.92 [0.91,0.94]	1.11 [1.08,1.14]	0.99 [0.99,1.00]	0.97 [0.93,1.00]	0.66 [0.65,0.67]	1.39 [1.38,1.40]
0.90	0.97 [0.95,0.98]	1.04 [1.03,1.06]	1.00 [1.00,1.00]	1.03 [1.02,1.05]	0.76 [0.75,0.77]	1.73 [1.72,1.76]
0.95	0.98 [0.97,0.99]	1.03 [1.02,1.05]	1.00 [1.00,1.00]	1.06 [1.04,1.07]	0.79 [0.79,0.80]	1.93 [1.91,1.95]

(B) Varying the relative standard deviation of measurement error e_t vs the state x_t given $\rho_0 = 0.75$

$\text{Var}[x_t]^{1/2}$	B1. Linear case		B2. Quadratic obs. eq.		B3. Non-linear state eq.	
$/\text{Var}[x_t]^{1/2}$	Correlation	Rel st. dev.	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.56 [0.51,0.62]	2.70 [2.47,3.06]	0.83 [0.81,0.86]	1.17 [1.13,1.23]	0.65 [0.64,0.65]	1.41 [1.40,1.42]
1.00	0.92 [0.91,0.94]	1.11 [1.08,1.14]	0.99 [0.99,1.00]	0.97 [0.93,1.00]	0.66 [0.65,0.67]	1.39 [1.38,1.40]
2.00	0.98 [0.97,0.98]	0.99 [0.98,1.01]	1.00 [0.99,1.00]	0.95 [0.91,0.99]	0.74 [0.67,0.95]	1.35 [1.12,1.38]
5.00	0.99 [0.99,1.00]	0.96 [0.95,0.97]	1.00 [0.99,1.00]	0.93 [0.90,0.96]	0.97 [0.97,0.98]	1.13 [1.11,1.17]
10.00	1.00 [0.99,1.00]	0.95 [0.94,0.97]	1.00 [0.99,1.00]	0.91 [0.88,0.95]	0.98 [0.97,0.98]	1.11 [1.08,1.15]

Shown are the median correlations and relative standard deviations [68% confidence interval] across simulations.

Table 1: Correlation and relative standard deviation of the filtered vs true state as a function of the persistence of the hidden state (top panel) and the signal to noise ratio (bottom panel). Model with measurement error and ARMA(1,1) shocks.

3.3 Quadratic observation equation with misspecification

Next, we consider the linear state equation, but allow for a quadratic state equation. With *iid* innovations ($\phi_1 = 0$), the measurement and state equations imply:

$$y_t = \mu_y + x_t + \kappa_1 \mathbb{E}_t[x_{t+1}] + \kappa_2 \mathbb{E}_t[x_{t+1}^2] + e_t \tag{3.7}$$

$$= \mu_y + x_t + \kappa_1 \rho_0 x_t + 2\kappa_2 \rho_0 \mu_x x_t + \kappa_2(\mu_x^2 + \sigma_u^2) + \kappa_2 \rho_0^2 x_t^2 + e_t \tag{3.8}$$

The case of the quadratic observation equation serves to illustrate 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, however, economic insight can be used to rule out regions of the state space that give rise to multiplicity.

To ensure that our data generating process has a point-identifiable state, we restrict the measurement 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.[9] Retracing the steps of example 1(b), we require that:

$$\frac{\partial y_t}{\partial x_t} = 2\rho_0^2 \kappa_2 x_t + (1 + \kappa_1 \rho_0 + 2\kappa_2 \rho_0 \mu_x) \geq 0 \quad \Leftrightarrow \quad x_t \geq -\frac{1}{2} \frac{1 + \kappa_1 \rho_0 + 2\kappa_2 \rho_0 \mu_x}{\rho_0^2 \kappa_2} \tag{3.9}$$

Given bounded shocks, x_t is bounded from below by $\frac{\mu_x + \underline{u}}{1-\rho_0}$, where $\underline{u}$ is the lower bound on the

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

disturbances. Given $\rho_0 > 0$, the sufficient condition becomes a restriction on the unconditional mean of the state x_t:

$$\mu_x \geq \frac{\rho_0^2(\kappa_1 - 2\underline{u}\kappa_2) - \kappa_1\rho_0 + \rho_0 - 1}{2\kappa_2\rho_0} \tag{3.10}$$

In the simulations, we consider two different combinations of κ_1 and κ_2:

1. Almost linear: $\kappa_1 = 10, \kappa_2 = \frac{1}{10}, \mu_y = 0$. μ_x set so that (3.10) holds with equality
2. Very non-linear: $\kappa_1 = \frac{1}{10}, \kappa_2 = 10, \mu_y = 0$. μ_x set so that (3.10) holds with equality

Since we already treated a linear case above, considering a nearly linear case may seem superfluous because we have documented that the filter works in the linear case with a linear VAR. However, in non-linear applications we also consider a more flexible approximating model. A priori, the extra flexibility could impair the filter performance. However, we show that this is not the case even in the near-linear case.

In this setup, given the possible multiplicity of solutions, we need to choose an initial guess with a reasonable scale. Here, we use the following guess, based only on observables and the known parameters of the observation equation:

$$\widehat{x}_t^{(0)} = \frac{1}{\kappa_1 + \kappa_2} \frac{1}{\kappa_2^2 + \kappa_1^2} \left(\kappa_2^2\sqrt{y_t} + \kappa_1^2 y_t\right) + 0.1 \times \mathcal{N}(0,1)$$

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.

Before turning to the results, we illustrate the DGP for various degrees of persistence, $\rho_0 \in \{0.25, 0.75, 0.95\}$. 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 3 shows. For the higher degrees of persistence, however, the non-linearity is noticeable: One gauge of that measure 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 $\rho_0 = 0.25$, the non-linearity matters little – the OLS coefficient on x_t is .75, whereas the LS coefficient on x_t^2 is .45 only. The relative magnitudes are reversed when $\rho_0 = .75$ or $\rho_0 = .95$.

Figure 4 uses a scatter plot to compare the filtered state to the true state for 100 different sample paths. Each panel plots the true x_t on the horizontal axis and the filtered $\widehat{x}_t$ on the vertical axis. Different colors indicate different sample paths. The dashed line is the 45-degree line. The three panels correspond again to $\rho_0 \in \{0.25, 0.75, 0.95\}$, with different colors in each panel corresponding to different simulated sample paths. With low persistence (left panel), there are two out of 100 samples paths when the partial filter estimates sequence of states outside the true parameter region. These are shown as orange and purple clouds in the bottom of the graph and correspond to the "wrong root" of the non-linear system. In practice, researchers may be able to

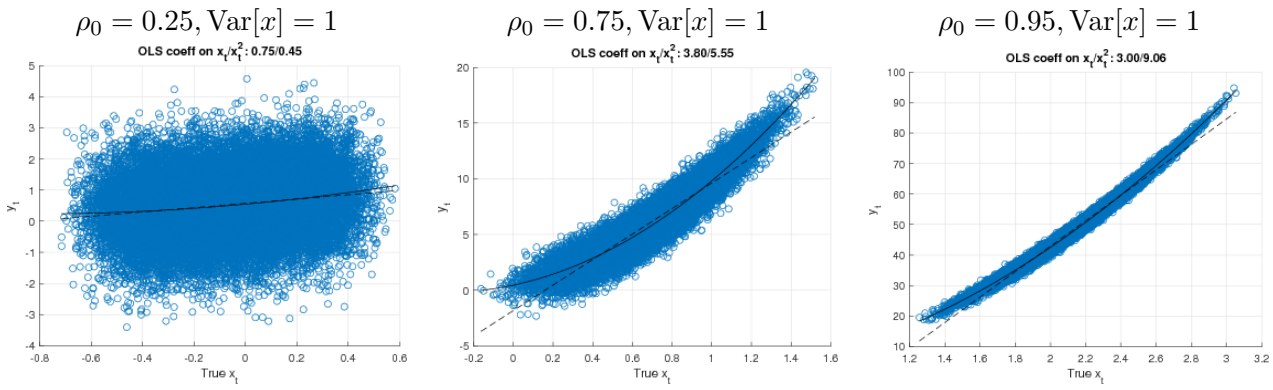

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 3: Very-non linear case ($\kappa_1 = \frac{1}{10}, \kappa_2 = 10$) for varying degrees of persistence ρ_0. True state vs observation with T=25,000 observations.

rule this case out using knowledge of the application. In all other sample paths, the partial filter has a strong correlation with the truth, but exhibits excess smoothness: The filtered state varies only between .5 and 1.5, as opposed to -2.5 to 2.5 in the true DGP. As we discuss below, however, the correlation of the filtered values with the truth is still above .78 in 84% of samples. With higher persistence (center and right panel), the visual fit is remarkably high across all sample paths and the correlations exceeds .96 in 84% of sample paths.

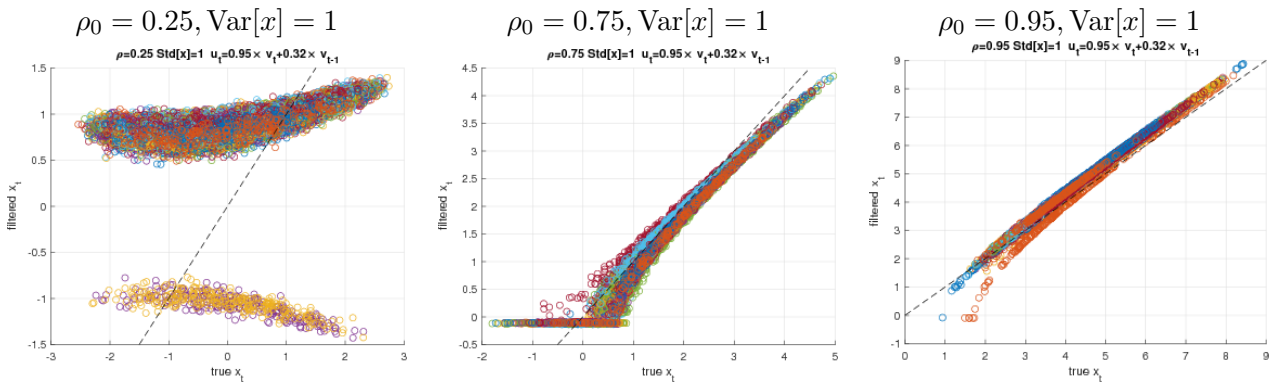

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 4: Very-non linear case ($\kappa_1 = \frac{1}{10}, \kappa_2 = 10$) for varying degrees of persistence ρ_0. True state vs partial filter in the fully linear case (bottom row): Filtered state $\widehat{x}_t$ vs true x_t for 100 different $T = 250$ samples.

Table 1, panels A2 and B2, shows how the fit of the partial filter changes as a function of the signal-to-noise ratio and the persistence of the true state. Specifically, it lists the median correlation between the true and filtered states (along with 68% confidence intervals) in the left column, and the relative standard deviations of the filtered state relative to the true state in the right column. The bottom panel varies the variance of x_t, for fixed $\rho_0 = .75$. The top panel fixes the variance of x_t and varies ρ_0. Turning to the results, the left column shows that the median correlation is above 0.8 in all cases shown, and often well above 0.95 with tight confidence intervals. The right

column shows that, the standard deviation of the estimated state is near that of the true state for a persistence of $\rho_0 \geq 0.75$. However, as indicated in the scatter plots in Figure 4, the filter estimate is excessively smooth when the persistence is low. With $\rho_0 = 0.25$, the relative standard deviation is only 0.12. It rises to 0.43 with $\rho_0 = 0.5$ and approaches unity quickly for $\rho_0 = 0.75$ or higher.

While the previous figures considered the very non-linear DGP, the results are even stronger with a DGP that is close to linear. In principle, the approximating model is overparameterized with a linear DGP. Table A.2 in the Appendix compares the results for correlations and relative standard deviations for the very non-linear and the near-linear cases when we vary the signal-to-noise ratio and the persistence as before. The partial filter produces higher correlations and relative standard deviations even closer to one than in the very non-linear case summarized in Table 1: The correlations are now .9 or higher, as opposed to .80 or higher in the very non-linear case. The worst relative standard deviation is now .39 to .12 in the very non-linear case. The confidence intervals remain tight.

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

We now consider a different form of non-linearity: A non-linear state equation with non-linear persistence and a cosine forcing term. Specifically, the model is now:

$$y_t = x_t + \kappa_1 \mathbb{E}_t[x_{t+1}] + e_t \tag{3.11a}$$

$$x_t = \gamma_0 \cos(\gamma_1 (t-1)) + \left(\rho_0 + \frac{\rho_1}{1+x_{t-1}^2}\right) x_{t-1} + u_t, \tag{3.11b}$$

$$u_t = \frac{\phi_0 v_t + \phi_1 v_{t-1}}{\sqrt{\phi_0^2 + \phi_1^2}}, \tag{3.11c}$$

where $e_t \sim iid(0, \sigma_e^2)$ and $v_t \sim iid(0, \sigma_u^2)$. Here, we set $\rho_1 = 25, \gamma_0 = 8, \gamma_1 = 1.2$. As we illustrate below, these parameter values interact with ρ_0 to give rise to rich non-linearities.

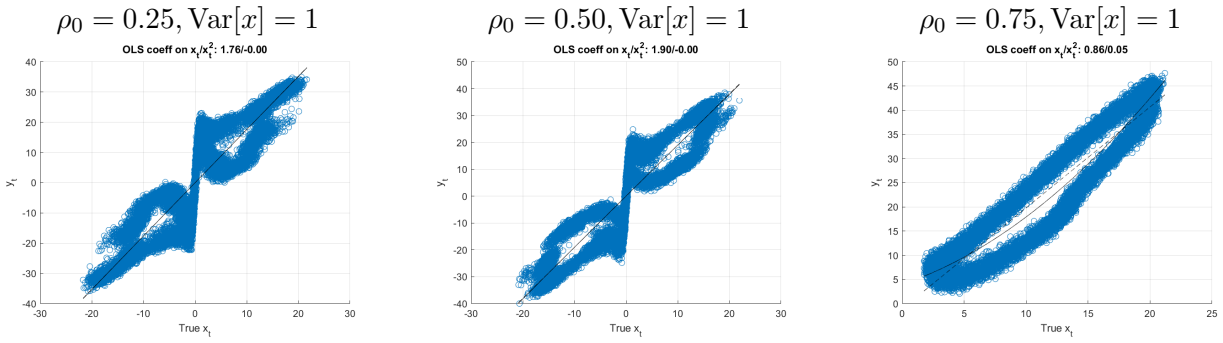

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 5: Linear observation equation, non-linear state equation for varying degrees of persistence ρ_0. True state vs observation with T=25,000 observations.

To illustrate this data-generating process, we once again turn to scatter plots. Figure 5 shows the relationship between the hidden state and the measurement for $\phi_0 = 0$. Even though we have

that $\kappa_2 = 0$, the measurement is related to the state non-linearly because of the expectations term, which varies non-linearly. In all cases, the state space contains regions where one value of x_t can be associated with different measurements. This is due to the non-linear state dynamics: The current state x_t is related to the state yesterday x_{t-1} in a similar non-linear fashion as the measurement y_t is related to the current state (not shown).

Despite the significant non-linearities, the partial filter does exceedingly well. Figure 6 shows that across all 100 sample paths, the filtered estimates are close to the 45 degree line. Even in the one sample path for $\rho_0 = 0.75$ where the state dynamics are in the negative region, a region of the state space outside the restricted domain in the data generating process and not visited during any other simulation, the filter successfully captures this.

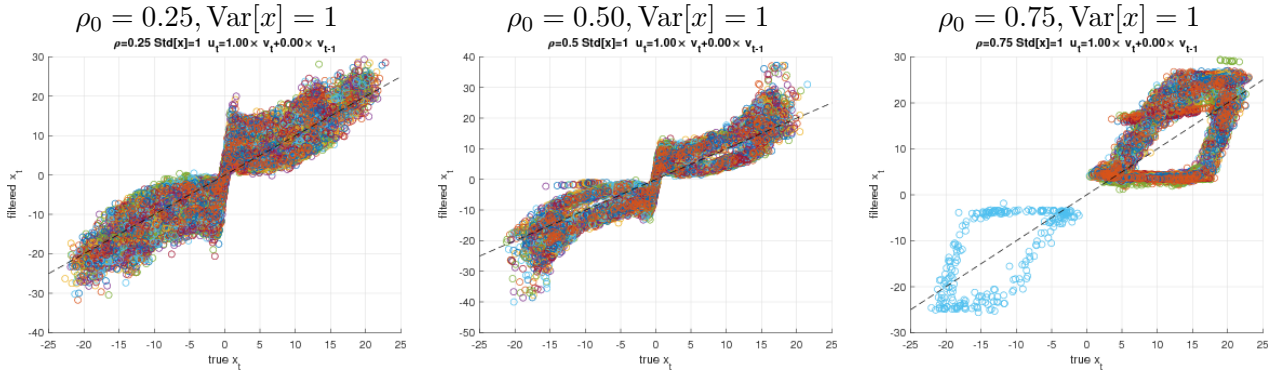

Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 6: Linear observation equation, non-linear state equation for varying degrees of persistence ρ_0. Filtered state $\widehat{x}_t$ vs true x_t for 100 different $T = 250$ samples.

Table 1 panels A3 and B3 quantify the performance of the partial filter with the non-linear law of motion. As before, the left column lists the median correlation of the filtered estimates and the truth along with 68% confidence intervals. The right column shows the relative standard deviation of the filtered estimate. For an intermediate persistence $\rho_0 = 0.75$, the correlation ranges from 0.65 to 0.98, as the variance of the state increases relative to the measurement error variance. Fixing the variance of the hidden state shows that the correlation can dip from above 0.9 to 0.65 when $\rho_0 = 0.75$, but is otherwise always above 0.75. The relative standard deviation in the right column is slightly high, but of the right order of magnitude.

3.5 Comparison with the Kalman Filter

While the Kalman Filter is not applicable in all of the scalar examples that we are considering, it is still a useful benchmark. Because it is widely used for filtering dynamic economies, applied researchers may consider it as a natural starting point even when it is misspecified.

We document that both the partial filter and the Kalman filter typically yield state estimates that are high relative to the truth. However, the partial filter is more robust in the non-linear

applications and yields relative standard deviations that are more consistently close to unity. See A.4.

4 Application: A Sudden Stops Model

To evaluate 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 (LTV) constraint in the workhorse RBC small-open economy 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:

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

where β is the discount factor, ω determines the elasticity of labor supply, and σ controls risk aversion. Note that we follow the specification in Mendoza and Villalvazo (2020) with a standard time-separable expected utility with exogenous discounting, instead of the case with endogenous discounting as in Mendoza (2010).

The household can accumulate capital subject to quadratic adjustment costs:

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

where δ is the depreciation rate, $\widetilde{i}_t$ is gross investment, and a is a constant scaling 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:

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

The left-hand-side of this budget constraint tells us that the resources of this economy are used for consumption, taxed at a constant rate τ and gross investment.[10] The resources on the right-hand side come from output net of imported inputs and the servicing of working capital loans, and the change in the foreign bond position.

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

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

Equation (4.3) limits the total debt (working capital loans plus negative positions in bonds) to be less 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 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), we require that $w_t = L^{\omega-1}(1+\tau)$. Similarly, the price of capital must satisfy $q_t = \frac{\partial \widetilde{i}_t}{\partial k_{t+1}}$.

The equilibrium conditions of the model boil down to:

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

$$q_t = \frac{\mathbb{E}_t\left[M_{t+1}\right]\mathbb{E}_t\left[d_{t+1}+q_{t+1}\right] + \mathrm{Cov}_t(M_{t+1}, d_{t+1}+q_{t+1})}{R_t \mathbb{E}_t\left[M_{t+1}\right] + \widetilde{\mu}_t\left[(1+\tau)-\kappa\right]} \tag{4.5}$$

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

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 $\widetilde{\mu}_t \equiv \frac{\mu_t}{u_c(t)}$ 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, which we lower marginally from 0.92 to 0.918 to yield an average crisis probability of 30.8% (see Table 2

[10] While the model does not include a government, a time-invariant consumption tax τ help 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.

below). The calibration is annual. However, since the ratio of flows to stocks scales in the period length, we can apply the model to quarterly data below without much effort.

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 2: Model calibration.

While the whole point of the partial filter is to allow a researcher to study the data without solving the model, we need to compute the model to perform our Monte Carlo experiment that demonstrates that such a task can be accomplished. We do so following *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.

4.2.1 The partial filter

Since here the shocks are revealed by static observables, we focus instead 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 $\widetilde{\mu}_t = \frac{\mu_t}{u_{c,t}}$, 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 are either unobservable, $\widetilde{\mu}_t$, or hard to measure, q_t (the indices of capital prices are subject to much discussion). 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.

In comparison, we assume that we can measure C_t and L_t (which gives us $u_{c,t}$, recall we have fixed already parameter values), R_t, and d_t. For many countries, these four series are readily available (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 in implementing the partial filter is to define the hidden state vector $\mathbf{x}_t = [\widetilde{\mu}_t, q_t]$ and select two equilibrium conditions of the model that are likely to be informative about $\mathbf{x}_t$, the Euler equation (4.4) rewritten as:

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

and the relative price of capital (4.5) becomes:

$$q_t = \frac{\frac{\beta}{u_{c,t}} \mathbb{E}_t\left[u_{c,t+1}\right] \mathbb{E}_t\left[d_{t+1} + q_{t+1}\right] + \frac{\beta}{u_{c,t}} \mathrm{Cov}_t(u_{c,t+1}, d_{t+1} + q_{t+1})}{R_t \frac{\beta}{u_{c,t}} \mathbb{E}_t\left[u_{c,t+1}\right] + \widetilde{\mu}_t\left[(1+\tau) - \kappa\right]}, \tag{4.8}$$

when we substitute $M_{t+1} \equiv \frac{\beta u_c(t+1)}{u_c(t)}$. This last expression is particularly convenient because, below, we will specify our VAR in terms of the marginal utility of consumption. Notice that we can filter out $\widetilde{\mu}_t$ 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:

$$\widetilde{\mu}_t \left[q_t^b b_{t+1} - \phi R_t \left[w_t L_t + p_t v_t \right] + \kappa q_t k_{t+1} \right] = 0, \tag{4.9}$$

which we can simplify further by taking advantage of the fact that, in equilibrium $w_t L_t + p_t v_t = (\alpha + \eta) y_t$, i.e., a constant fraction of output is spent on labor and imported inputs (this result comes directly from having competitive markets and a Cobb-Douglas production function). We use this condition to adjust our state estimates. When the borrowing constraint is binding, but $\widetilde{\mu}_t$ 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 $\widetilde{\mu}_t$ is inferred to be positive.[11]

Thus, we have a 2-dimensional g-function:

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

[11] We refrain from using this equation to filter Tobin's q directly, because of measurement error: The observed debt is averaged across maturities and 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.

and an f-function:

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

When the borrowing constraint is slack, we solve $\mathbf{f}$ for $[\widetilde{\mu}_t, q_t]$. When the borrowing constraint is binding (or violated) but $\widetilde{\mu}_t = 0$, we add a time-varying intercepts 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 $\widetilde{\mu}_t > 0$, we introduce a time-varying intercept that raises $\mathbb{E}_t[u_{c,t+1}]$ to lower the inferred $\widetilde{\mu}_t$.

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 $\boldsymbol{\xi}_t^{(d)} = [u_{c,t}, L_t, R_t, d_t, \widetilde{\mu}_t^{(d)}, q_t^{(d)}]$ – 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 $\boldsymbol{\Sigma}^{(d)} = \widehat{\mathrm{Var}}[\boldsymbol{\varepsilon}_t^{(d)}]$ 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

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

For numerical stability, we do not immediately update the state estimates. Instead, we use a third step to slowly update the guesses. Specifically, we update the state estimates for $\widehat{q}_t$ and $\widehat{\widetilde{\mu}}_t$, using iterative rules of the form $\widehat{q}_t^{(d)} = (1-s)\widehat{q}_t^{(d-1)} + s\mathbf{x}_{2,t}^{(d)}$. Updating the state estimates slowly improves convergence. Because in the model $\widetilde{\mu}_t \geq 0$, we also censor $\widehat{\widetilde{\mu}}_t$ at a small negative number. See Appendix B.2 for details.

Everything we have done so far is quite straightforward: (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 run a VAR on four observables and used the results to substitute variances and covariances on the three equations from step (i); and (iii) slowly updated our estimates for $\widehat{q}_t$ and $\widehat{\widetilde{\mu}}_t$. None of these steps requires much computational effort.

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. 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.[12] Then, we apply our partial filter as outlined above to the selected simulated samples.

4.3.1 Results

Table 3 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, independent of the crisis frequency in the sample: the median correlation across simulations is above 0.8 in all three scenarios. The median standard deviation of the estimated $\widehat{q}_t$ relative to the true q_t ranges from 1.03 to 1.16. 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 weaker for $\widetilde{\mu}_t$, but the estimates are solid throughout, with correlations rising from 0.32 to 0.51 with the crisis frequency. The relative standard deviation of the filtered $\widetilde{\mu}_t$ is lower, with the median falling slightly from 0.80 to 0.68. One reason may be that higher sudden stop frequencies make the non-linearities in the model more pronounced and thus complicate the filtering exercise.[13]

Sudden stop freq. [%]	$\text{corr}(\widehat{q}_t, q_t)$	$\text{corr}(\widehat{\widetilde{\mu}}_t, \widetilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\widetilde{\mu}_t$
[10.0, 15.0]	0.84 [0.71, 0.94]	0.32 [0.12, 0.50]	1.11 [0.96, 1.37]	0.80 [0.35, 1.46]
[30.0, 35.0]	0.81 [0.65, 0.94]	0.35 [0.11, 0.53]	1.16 [0.97, 1.42]	0.68 [0.38, 1.07]
[50.0, 55.0]	0.85 [0.65, 0.95]	0.51 [0.25, 0.71]	1.03 [0.91, 1.38]	0.68 [0.50, 0.91]

Table 3: Medians [68% CI] of correlations, relative standard deviations, and classifications errors

The fact that the performance for $\widetilde{\mu}_t$ is weaker is not a surprise: $\widetilde{\mu}_t$ has a mass-point at zero, a non-linearity that our extremely simple filter is unlikely to be able to capture well. But, how much

[12] Average technology declines and the average import price increases with the frequency of crisis. 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.

[13] Out 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.

economic information does a researcher lose in practice? We argue that it is much more relevant to gauge how well our partial filter does at learning when the borrowing constraint binds, the most relevant information for policymakers. Table 4 lists the false positive and false negative rates when classifying crises, that is how often we estimate that the constraint is binding when it is not and vice versa, how often we miss that the constraint is binding. Here, our partial filter displays a great performance, with median false positives between 4.0% when the true crisis frequency is 10%-15% and 12.5% when the true crisis frequency is between 50% and 55%. The false negative rates , in contrast, tend to be higher when the crisis frequency is lower. When the crisis frequency is around 10%, the median false negative rate is 13.8%; it only measures 7.1% when the crisis frequency is around 50%. We read the results of Table 4 as a very solid performance of our extremely simple filter.[14]

Sudden stop	Crisis classification	
frequency	False positives	False negatives
[10.0, 15.0]	4.0 [1.6, 6.0]	13.8 [5.6, 20.8]
[30.0, 35.0]	9.4 [6.2, 12.9]	15.7 [8.2, 22.0]
[50.0, 55.0]	12.5 [10.2, 15.9]	7.1 [3.2, 15.1]

Table 4: Crisis classification errors

Figure 7 shows the data underlying the previous tables. Each column corresponds to one of the three targeted in-sample crisis frequencies. The top panels show the time path of the true (blue) and estimated (orange) time series for $\widehat{q}_t$ for the sample with the median correlation. In line with Table 3, the two lines track each other well in all three cases. The amplitude is somewhat more closely aligned when the crisis frequency is higher, reflecting the higher relative standard deviation. Comparing the simulated time series for q_t with the filtered counterparts shows that the partial filter reliably recovers the large movements in q_t, but misses smaller changes. The bottom panel shows scatter plots of the relationship between the true and the filtered q_t, with different colors indicating different sample paths. While there is some dispersion around the 45-degree line and the filtered q tends to have a wider range, the scatter plots show that the partial filter recovers the right pattern of q overall.

Figure 8 shows representative sample paths and overall scatter plots for $\widetilde{\mu}_t$. As before, the top panels show the sample paths for the filtered and actual multipliers for the median correlation. Clearly, as the crisis frequency increases, the top panels are less sparse. When the crisis frequency is low, the partial filter successfully identifies most crises correctly: Four crises, including the most severe and most persistent crisis around period 140, are called correctly. There is a false negative around period 130, and a false positive shortly after period 200. With the less sparsely populated

[14] We checked that when we do not use the slackness condition to infer $\widetilde{\mu}_t$, the correlation increases with the crisis frequency from a median of 0.36 to a median of 0.75. The median correlation for Tobin's q is 0.98 throughout and the relative standard deviations are closer to uity for both q and $\widetilde{\mu}$. However, the false positive rate is much higher, with the median at about one third. See Appendix B.3. Moreover, the filtering algorithm did not converge for 13% of simulations, as opposed to just 2% when using the slackness condition.

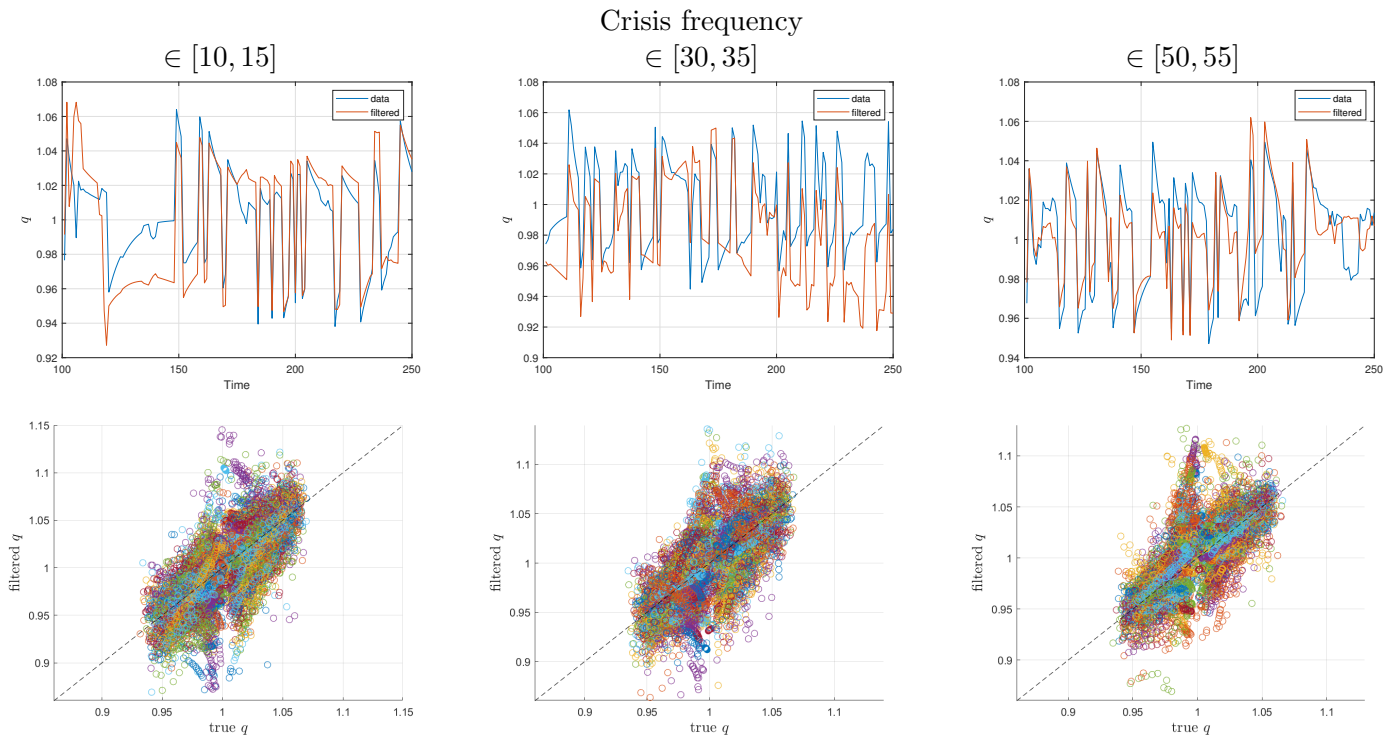

Figure 7: Sample paths for q_t with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis frequency.

panels with a higher crisis frequency, the partial filter also misses few crises, but fails to correctly call the severity of the worst crises, shown as tall blue spikes. The scatter plots in the bottom panels similarly show that there is less mass on the horizontal axes. Even conditional on being positive, the true $\widetilde{\mu}$ are higher (located further to the right) as the crisis frequency rises, indicating that the crises are, on average, not only more frequent but also more severe. The scatter plots with the higher crisis frequency more clearly show the positive correlation between the true and filtered values. However, the mass is concentrated below the 45-degree line, in line with the relative standard deviation that is lower for the filtered $\widetilde{\mu}$.

4.4 Application to Mexico

Having established that the filter performs well in a Monte Carlo experiment, we apply it to Mexican quarterly data 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 spliced with the 3-months U.S. T-bill rate and converted to real terms using the U.S. CPI. We take dividends from the S&P/IFCG M MEXICO index, spliced after 2008Q3 with the MSCI MEXICO index. We construct the annual capital stock needed to

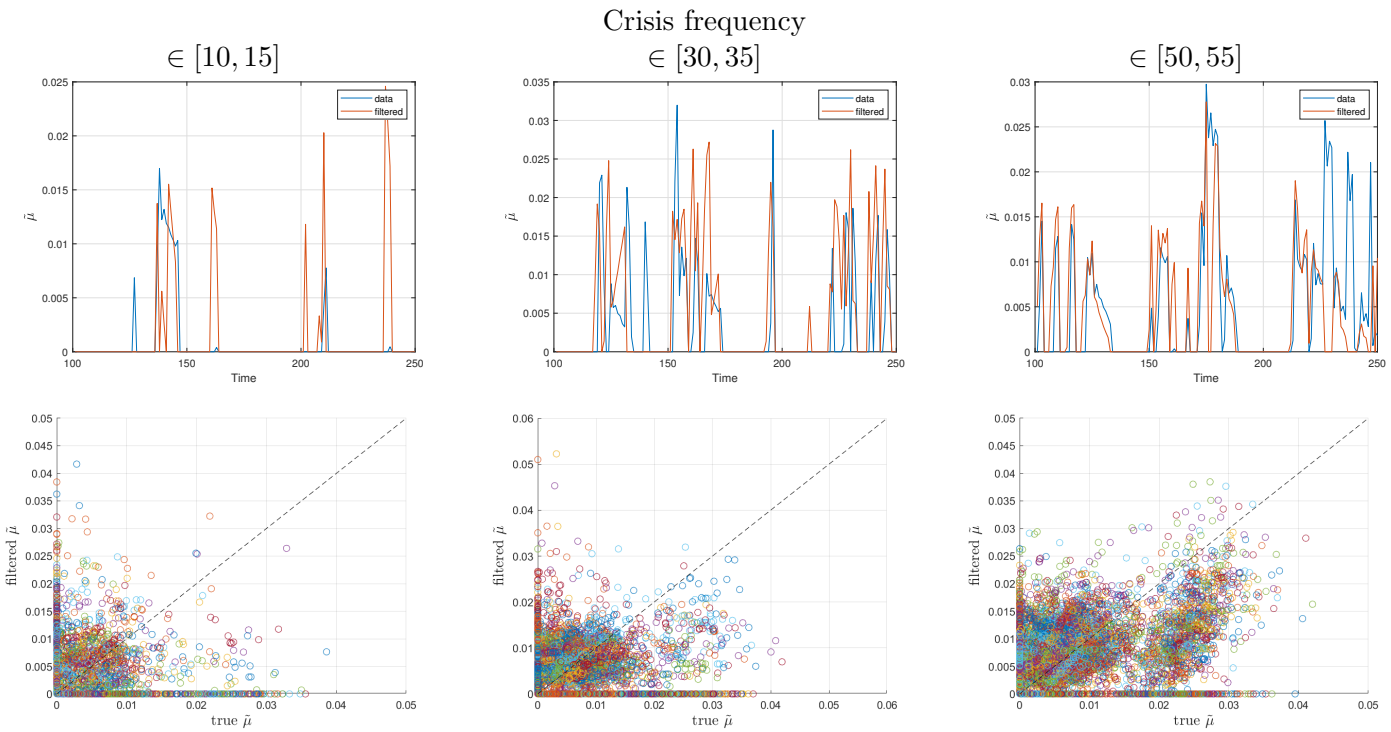

Figure 8: Sample paths for the multiplier on the borrowing constraint $\widetilde{\mu}_t$ for the sample paths with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis frequency.

evaluate the borrowing constraint using the perpetual inventory method and linearly interpolate to obtain a quarterly series. See Appendix B.1 for details.

We use the same specification of the partial filter as in our simulation study before. This algorithm is a fixed point algorithm that ignores uncertainty about the parameters of the approximating VAR: In the step that updates parameters, it uses the 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.[15] 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 $\mathbf{a}^{(0)}, \mathbf{A}^{(0)}, \mathbf{\Sigma}^{(0)}$ to the OLS estimates associated with $\boldsymbol{\xi}_t^{(0)} = [u_{c,t}, L_t, R_t, d_t, \tilde{\mu}_t^{(0)}, q_t^{(0)}]$. [16]

Results. Does our partial filter recover the Mexican debt crises during our sample? Yes. Figure 9(a) shows the filtered multiplier on the borrowing constraint, $\widetilde{\mu}_t$, and plots a number of economic events that we suspect could lead to a binding borrowing constraint. These events are:

[15] The 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 the initial condition leads to divergence, we initialize the filtering sequence instead of in the initial draw.

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

(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 Global Financial Crisis associated with the collapse of Lehman Brothers in 2008q3. The partial filter is successful in identifying these crisis events but one: the Asian crisis of 1997q3. However, it is unclear whether this crisis should have a direct impact on Mexico and the filter identifies the subsequent financial turmoil in 1998. While the duration of these crisis events is unclear, there are no clear false positives.

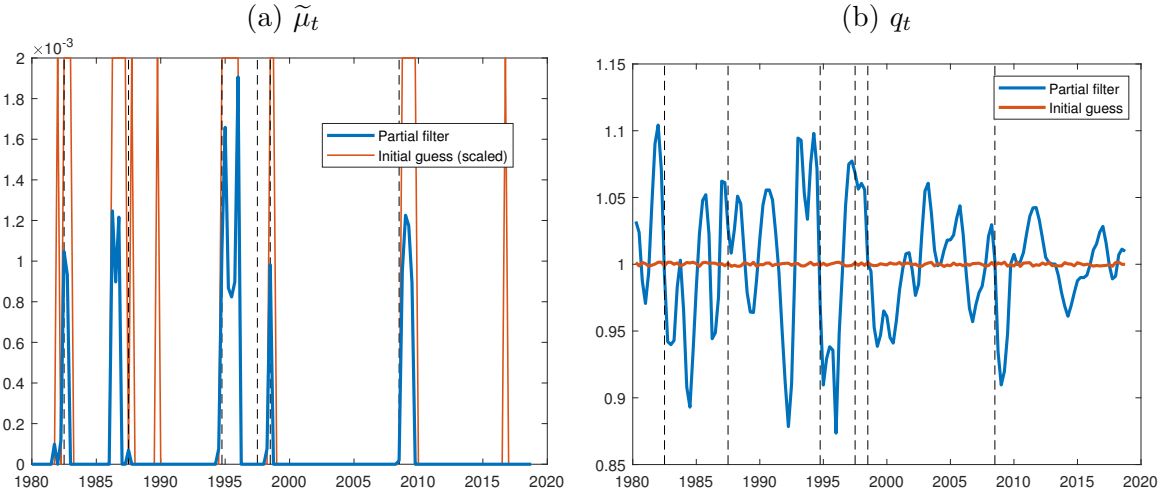

Figure 9: Filtered values of $\widetilde{\mu}_t$ and q_t.

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. But 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.1.

Figure 9(b) 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 – which is natural because the binding borrowing constraint drives up the rate of return the representative household requires. Moreover, the filter shows that the inference about Tobin's q is not driven by our initial guess, displayed in orange.

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-type prior and filtering $\widetilde{\mu}_t$ and q_t, which, in turn, enter the VAR as data.[17]

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 exactly the same crises as the fixed point algorithm. Figure 10(a) shows 68% and 90% credible sets, the posterior

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

median, and the fixed point estimates of the multiplier on the borrowing constraint. While the credible sets are wide in some periods, the filter rarely identifies positive multipliers outside crisis times. During these times, the multiplier is positive with at least 95% posterior probability. The posterior median shows one false positive in 1989.

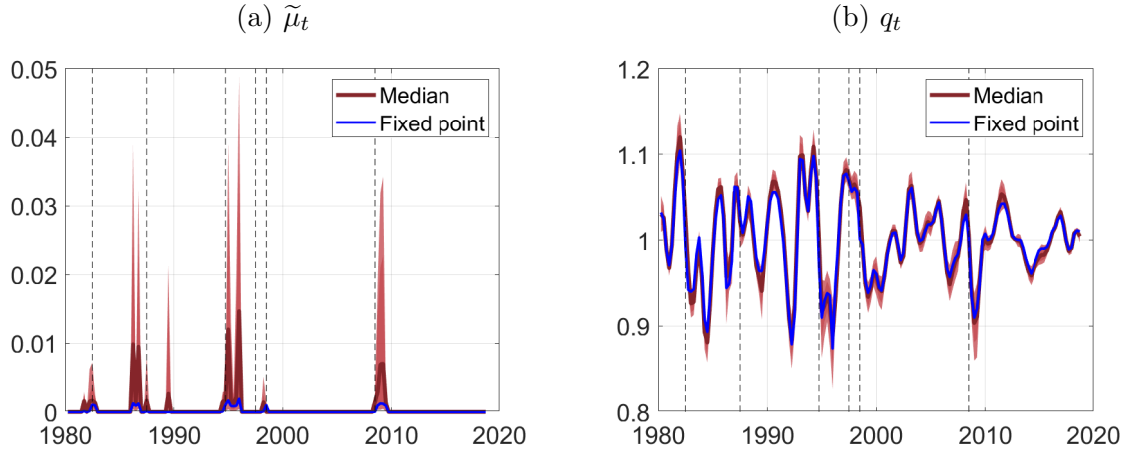

Posterior median and 68% and 90% credible sets

Figure 10: Filtered values of $\widetilde{\mu}_t$ and q_t.

Another feature of the posterior uncertainty is its asymmetry. The assymmetry is more noticeable in the multiplier estimates, although large right tails in the estimates of $\widetilde{\mu}_t$ also manifest themselves as asymmetric credible sets for Tobin's q, for example during the Global Financial Crisis. Overall, the effects of parameter uncertainty are relatively small.

5 Conclusion

The partial information filter allows applied researchers to evaluate the historical implications of their models, without fully specifying the model – and even solving the model. This 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. And 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 macro models. It may thus be an important complement to likelihood-free methods in model estimation.

Our analysis study points to two requirements for the partial information filter. First, the measurement equations need to uniquely identify the state absent measurement error. Second, the results are robust to misspecification if the measurement is not too noisy, and the state needs to exhibit enough persistence.

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.

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.

Campbell, J. Y. (1991). A variance decomposition for stock returns. Economic Journal, 101(405):157–179.

Canova, F. and Ferroni, F. (2018). Mind the gap! Stylized dynamic facts and structural models. Working Papers No 13/2018, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School.

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.

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. The Quarterly Journal of Economics, 136(3):1387–1446.

Drautzburg, T., Fernández-Villaverde, J., and Guerrón-Quintana, P. (2021). Bargaining shocks and aggregate fluctuations. Journal of Economic Dynamics and Control, 127:104121.

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.

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.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. Journal of Economic Literature, 52(1):124–88.

Mendoza, E. and Villalvazo, S. (2020). FiPIt: A Simple, Fast Global Method for Solving Models with Two Endogenous States & 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.

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.

Appendix A: Scalar model

A.1 No measurement error

Here we provide additional correlation and standard deviation plots for the linear case without measurement error. In Figure A.1 we plot 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.

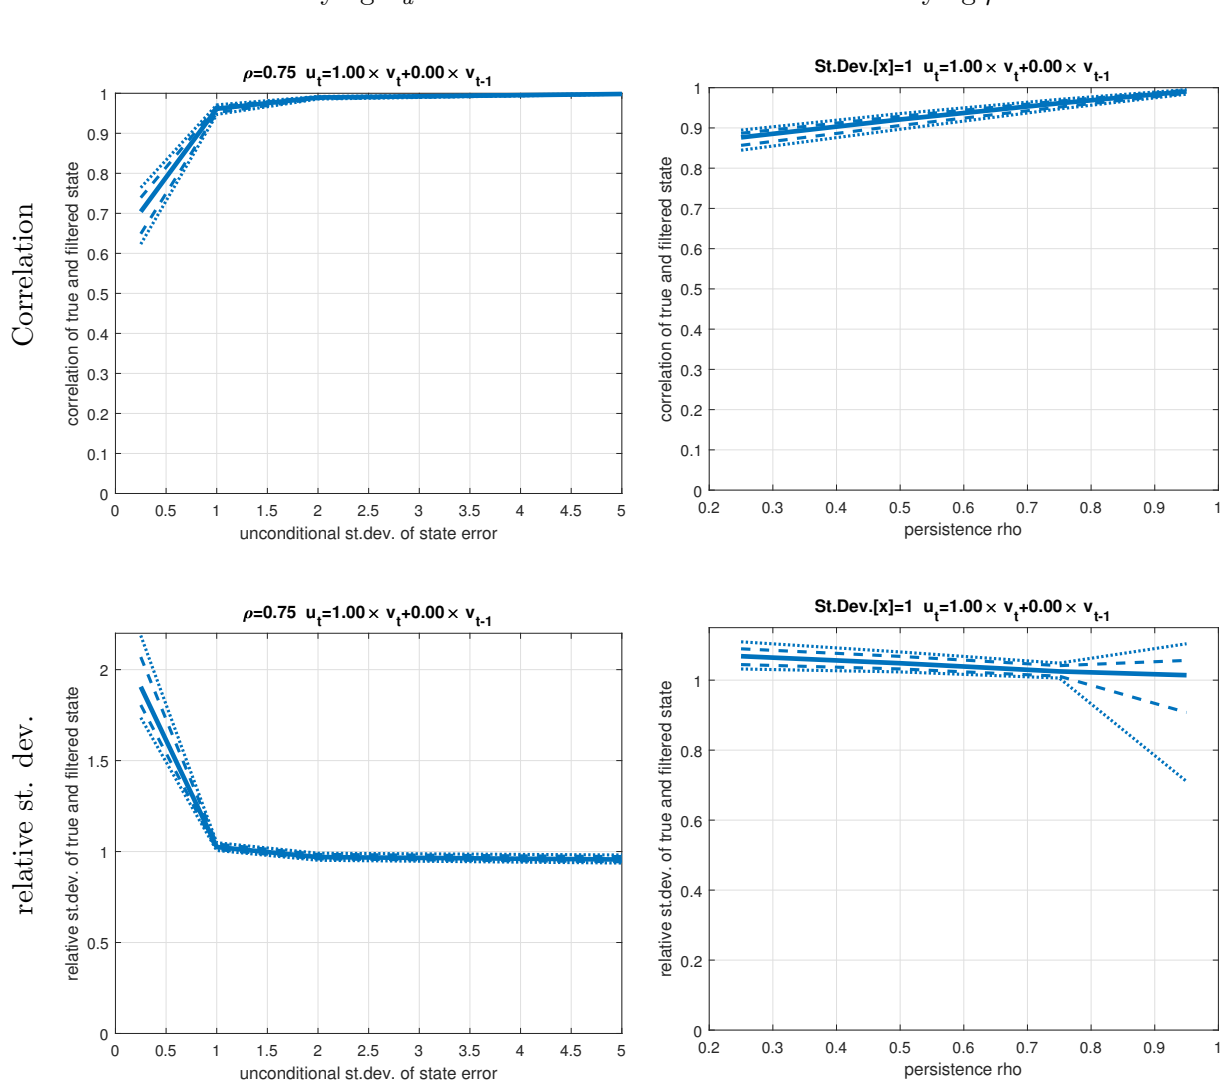

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median.

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

The two panels on the left vary σ_u, for fixed $\rho_0 = .75$. The two panels on the right fix $\sigma_u = 1$ and vary ρ_0. Turning to the results, the top row shows that the median correlation is close to 1 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 Filtering with the non-linear VAR

For now, we proceed in a simple special case: $\rho_1 = \gamma = \phi_1 = 0$. In this case, the observation equation becomes:

$$y_t = x_t(1 + \kappa_1\rho_0) + \kappa_2\rho_0^2 x_t^2 + e_t \tag{3.1a'}$$

Now assume that we estimate the restricted non-linear VAR:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \mu + \mathbf{A}\begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} + \mathbf{B}x_{t-1}^2 + \epsilon_t \tag{A.1}$$

The VAR implies that $\hat{\mathbb{E}}_t^{(d)}[x_{t+1}] = \hat{\mathbf{A}}_{1,\circ}^{(d)}[x_t, y_t]' + \hat{\mu}_1^{(d)} + \hat{\mathbf{B}}_{1,1}^{(d)}x_t^2$.and that $\hat{\text{Var}}_t[x_{t+1}] = \hat{\Sigma}_{1,1}$, some constant.

Using the sequence $\hat{x}_t^{(d-1)}$ for the VAR estimation, yields the following fixed point problem (setting e_t to zero):

$$\begin{aligned}\hat{x}_t^{(d)} &= y_t - \kappa_1\hat{\mathbb{E}}_t^{(d)}[x_{t+1}] - \kappa_2\hat{\mathbb{E}}_t^{(d)}[x_{t+1}]^2 - \kappa_2\hat{\Sigma}_{1,1} \\ &= y_t - \kappa_1\hat{\mathbf{A}}_{1,\circ}^{(d)}[x_t, y_t]' - \kappa_1\hat{\mu}_1^{(d)} - \kappa_1\hat{\mathbf{B}}_{1,1}^{(d)}x_t^2 - \kappa_2(\hat{\mathbf{A}}_{1,\circ}^{(d)}[x_t, y_t]' + \hat{\mu}_1^{(d)} + \hat{\mathbf{B}}_{1,1}^{(d)}x_t^2)^2 - \kappa_2\hat{\Sigma}_{1,1}\end{aligned} \tag{A.2}$$

Note that this is a quadratic equation which, generally, has two solutions. The problem remains of which root to pick. Denote these solutions as $\hat{x}_{t,+}$ and $\hat{x}_{t,-}$.

Note that the VAR has implications for the likelihood of either solution:

$$\Pr\{x_t \leq \hat{x}_{t,+}|\hat{x}_{t-1}, y_{t-1}\} = \Phi\left(\frac{\hat{x}_{t,+} - \mu_1 - \mathbf{A}_{1,\circ}[x_{t-1}, y_{t-1}]' - \mathbf{B}_{1,1}x_{t-1}^2}{||\Sigma_{1,\circ}^{(chol)}||}\right) \tag{A.3}$$

Thus, we can choose the more likely solution.

Also, we can do more, since the forecast error in x_t is, generally, correlated with the forecast error in y_t: $\epsilon_t = \Sigma^{(chol)}u_t \sim \mathcal{N}(0, \Sigma)$. Then $\epsilon_{1,t}|\epsilon_{2,t} \sim \mathcal{N}(\Sigma_{2,2}^{-1}\Sigma_{2,1}\epsilon_{2,t}, \Sigma_{1,1} - (\Sigma_{2,2}^{-1}\Sigma_{2,1})^2)$. Defining $\hat{\epsilon}_{t,2}$ as the forecast error in y_t we then have:

$$\Pr\{x_t \leq \hat{x}_{t,+}|\hat{x}_{t-1}, y_{t-1}, y_t\} = \Phi\left(\frac{\hat{x}_{t,+} - \mu_1 - \mathbf{A}_{1,\circ}[x_{t-1}, y_{t-1}]' - \mathbf{B}_{1,1}x_{t-1}^2 - \Sigma_{2,2}^{-1}\Sigma_{2,1}\hat{\epsilon}_{2,t}}{\sqrt{\Sigma_{1,1} - (\Sigma_{2,2}^{-1}\Sigma_{2,1})^2}}\right) \tag{A.4}$$

To proceed, we could thus pick the root with the highest density, using the the pdf that corresponds to the CDF in (A.4).

A.3 Additional results

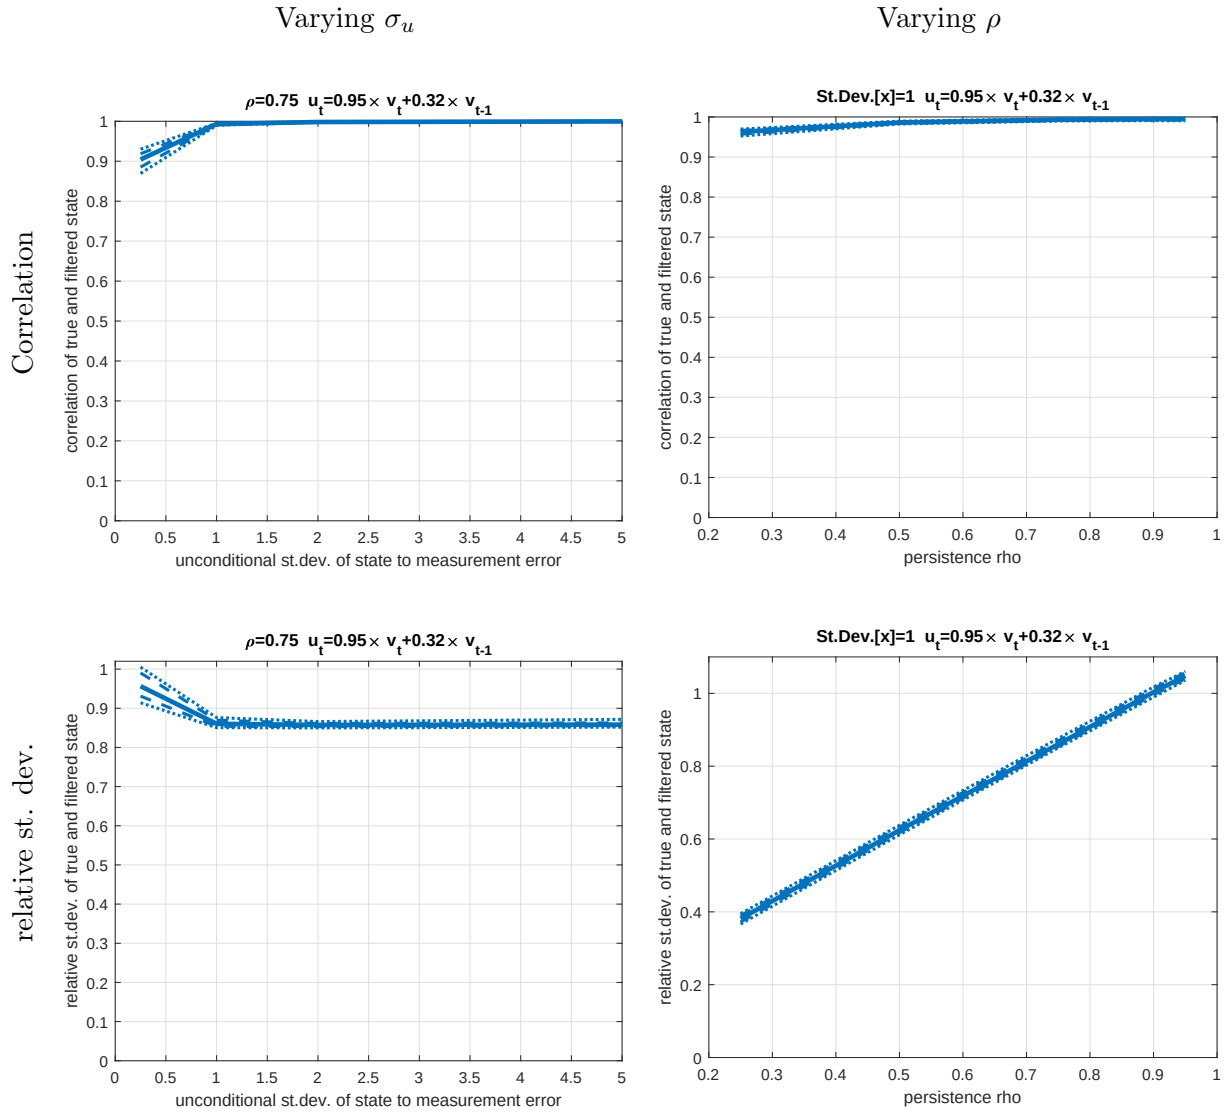

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median.

Figure A.2: Almost linear model: $\kappa_1 = 10, \kappa_2 = \frac{1}{10}$: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

(a) Varying the persistence ρ_0 of the state given $\text{Var}[e_t] = \text{Var}[x_t]$

	Linear case		Quadratic obs. eq.		Non-linear state eq.	
ρ_0	Correlation	Rel st. dev.	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.79 [0.76,0.81]	1.36 [1.29,1.48]	0.77 [0.74,0.80]	0.13 [0.12,0.14]	0.90 [0.89,0.91]	1.21 [1.19,1.25]
0.50	0.86 [0.84,0.88]	1.27 [1.20,1.35]	0.96 [0.94,0.97]	0.42 [0.40,0.46]	0.93 [0.92,0.93]	1.20 [1.17,1.23]
0.75	0.93 [0.92,0.94]	1.14 [1.08,1.21]	0.99 [0.98,0.99]	0.97 [0.94,1.00]	0.66 [0.65,0.67]	1.39 [1.37,1.40]
0.90	0.97 [0.96,0.98]	1.06 [1.02,1.11]	1.00 [1.00,1.00]	1.06 [1.04,1.12]	0.76 [0.76,0.77]	1.74 [1.72,1.76]
0.95	0.98 [0.97,0.99]	1.04 [1.01,1.07]	1.00 [1.00,1.00]	1.06 [1.05,1.08]	0.79 [0.79,0.80]	1.93 [1.91,1.95]

(b) Varying the relative standard deviation of measurement error e_t vs the state x_t given $\rho_0 = 0.75$

$\text{Var}[x_t]^{1/2}$	Linear case		Quadratic obs. eq.		Non-linear state eq.	
$/\text{Var}[e_t]^{1/2}$	Correlation	Rel st. dev.	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.53 [0.49,0.59]	2.66 [2.34,2.93]	0.80 [0.75,0.84]	1.11 [1.05,1.20]	0.65 [0.64,0.65]	1.41 [1.40,1.42]
1.00	0.93 [0.92,0.94]	1.14 [1.08,1.21]	0.99 [0.98,0.99]	0.97 [0.94,1.00]	0.66 [0.65,0.67]	1.39 [1.37,1.40]
2.00	0.98 [0.97,0.98]	1.04 [1.01,1.08]	0.99 [0.98,0.99]	0.96 [0.93,0.99]	0.88 [0.67,0.96]	1.26 [1.11,1.37]
5.00	1.00 [0.99,1.00]	1.01 [0.98,1.04]	0.99 [0.98,0.99]	0.95 [0.93,0.98]	0.97 [0.97,0.98]	1.15 [1.12,1.18]
10.00	1.00 [0.99,1.00]	1.01 [0.98,1.04]	0.99 [0.98,0.99]	0.93 [0.91,0.97]	0.98 [0.97,0.98]	1.14 [1.11,1.17]

Shown are the median correlations and relative standard deviations [68% confidence interval] across simulations.
The parameters are calibrated to target the variance of the linear x_t process.

Table A.1: Correlation and relative standard deviation of the filtered vs true state as a function of the persistence of the hidden state (top panel) and the signal to noise ratio (bottom panel). Model with measurement error and ARMA(1,1) shocks.

(a) Varying the persistence ρ_0 of the state given $\text{Var}[e_t] = \text{Var}[x_t]$

$\text{Var}[x_t]^{1/2}$	Quadratic obs. eq.: non-lin		Quadratic obs. eq.: near-lin	
ρ_0	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.80 [0.78,0.82]	0.12 [0.12,0.13]	0.96 [0.96,0.97]	0.39 [0.39,0.40]
0.50	0.97 [0.96,0.98]	0.43 [0.41,0.45]	0.99 [0.98,0.99]	0.63 [0.63,0.64]
0.75	0.99 [0.99,1.00]	0.97 [0.93,1.00]	0.99 [0.99,0.99]	0.87 [0.86,0.88]
0.90	1.00 [1.00,1.00]	1.03 [1.02,1.05]	0.99 [0.99,1.00]	1.01 [1.00,1.02]
0.95	1.00 [1.00,1.00]	1.06 [1.04,1.07]	0.99 [0.99,1.00]	1.05 [1.04,1.06]

(b) Varying the relative standard deviation of measurement error e_t vs the state x_t given $\rho_0 = 0.75$

$\text{Var}[x_t]^{1/2}$	Quadratic obs. eq.: non-lin		Quadratic obs. eq.: near-lin	
$/\text{Var}[e_t]^{1/2}$	Correlation	Rel st. dev.	Correlation	Rel st. dev.
0.25	0.83 [0.81,0.86]	1.17 [1.13,1.23]	0.90 [0.89,0.92]	0.97 [0.94,1.00]
1.00	0.99 [0.99,1.00]	0.97 [0.93,1.00]	0.99 [0.99,0.99]	0.87 [0.86,0.88]
2.00	1.00 [0.99,1.00]	0.95 [0.91,0.99]	1.00 [1.00,1.00]	0.87 [0.86,0.87]
5.00	1.00 [0.99,1.00]	0.93 [0.90,0.96]	1.00 [1.00,1.00]	0.87 [0.86,0.87]
10.00	1.00 [0.99,1.00]	0.91 [0.88,0.95]	1.00 [1.00,1.00]	0.87 [0.87,0.88]

Shown are the median correlations and relative standard deviations [68% confidence interval] across simulations.
The parameters are calibrated to target the variance of the linear x_t process.

Table A.2: Correlation and relative standard deviation of the filtered vs true state as a function of the persistence of the hidden state (top panel) and the signal to noise ratio (bottom panel). Model quadratic observation equation with very non-linear and near-linear specifications. Model with measurement error and AR(1) shocks.

A.4 Comparison with the Kalman filter

As we have shown, the partial information filter does well when the hidden state is sufficiently persistent and volatile, relative to the measurement error.

Comparing the partial filter with the Kalman filter is instructive to understand the role of measurement error. For simplicity, consider the fully linear case without expectations ($\kappa_1 = \kappa_2 = \rho_1 = \gamma = 0$). Since we treat the measurement error as zero, our partial filter gain is simply 1. In this linear environment, the Kalman filter would be optimal. Its gain would be $\widehat{\beta} = \text{Var}_{t-1}[y_t]^{-1} \text{Cov}_{t-1}[x_t] = \frac{\text{Var}_{t-1}[x_t]}{\text{Var}_{t-1}[x_t]+\sigma_e^2}$. If σ_e^2 is large, then $\widehat{\beta} \ll 1$.

Moreover, the Kalman filter is a good benchmark for the performance of the partial filter. While it is not straightforward to apply in general non-linear models, the Kalman filter is the best linear filter. We focus on the two cases when we have non-linearities. Here, we apply the Kalman filter assuming a simple AR(1) for the hidden state, as a practitioner might – ignoring the non-linearities.

Figure A.3 compares the Kalman filter for the very non-linear measurement equation, assuming a linear system with $\kappa_2 = 0$ for the Kalman filter in the very non-linear case ($\kappa_1 = \frac{1}{10}, \kappa_2 = 10$).[18] Comparing the Kalman filter in the left column with the partial filter in the right column shows that both have an excellent correlation with the true state (top row). However, the Kalman filter dramatically misses the standard deviation of the true state in its filtered estimates, unlike the partial filter (bottom row). While less dramatic, the same pattern even holds with minor non-linearities ($\kappa_1 = 10, \kappa_2 = \frac{1}{10}$) and when the hidden state follows a pure AR(1), rather than an ARMA(1,1) – see Figure A.4. Also in that case the partial filter does much better at estimating the true scale of the hidden state.

Figure A.6 shows that in the case of the non-linear state equation the partial filter outperforms a simple linear Kalman filter.[19] Both models are misspecified here, and both can attain a very high correlation of the filtered state with the truth, and almost the right scale when the signal-to-noise ratio is high. However, the performance of the Kalman filter is fairly robust, while the linear filter does poorly when the signal-to-noise ratio is low.

[18] We allow the Kalman filter to estimate κ_1 to allow it to fit the scale of x_t better. We also estimate the Kalman filter with the true κ_1 and choose the estimate with the higher likelihood for our comparison.

[19] Here, we use two different starting points of the persistence ρ_0 for the Kalman filter – both the true value, and the linear autocorrelation of the true state with its first lag in the simulated sample.

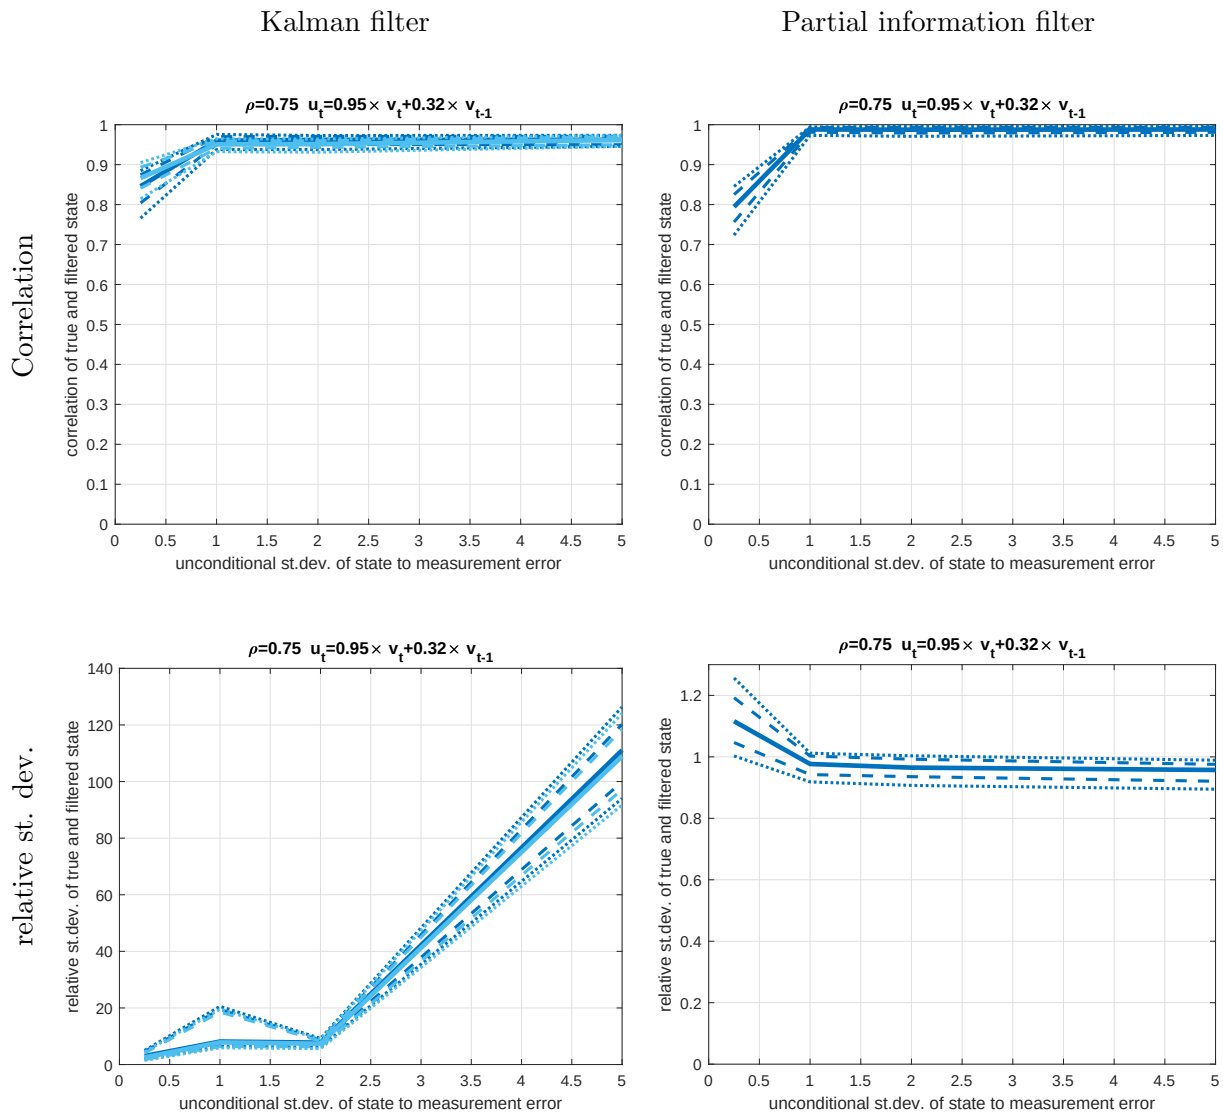

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

Figure A.3: Kalman filter vs partial information filter. Quadratic observation equation with $\kappa_1 = \frac{1}{10}, \kappa_2 = 10$ and linear state equation: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

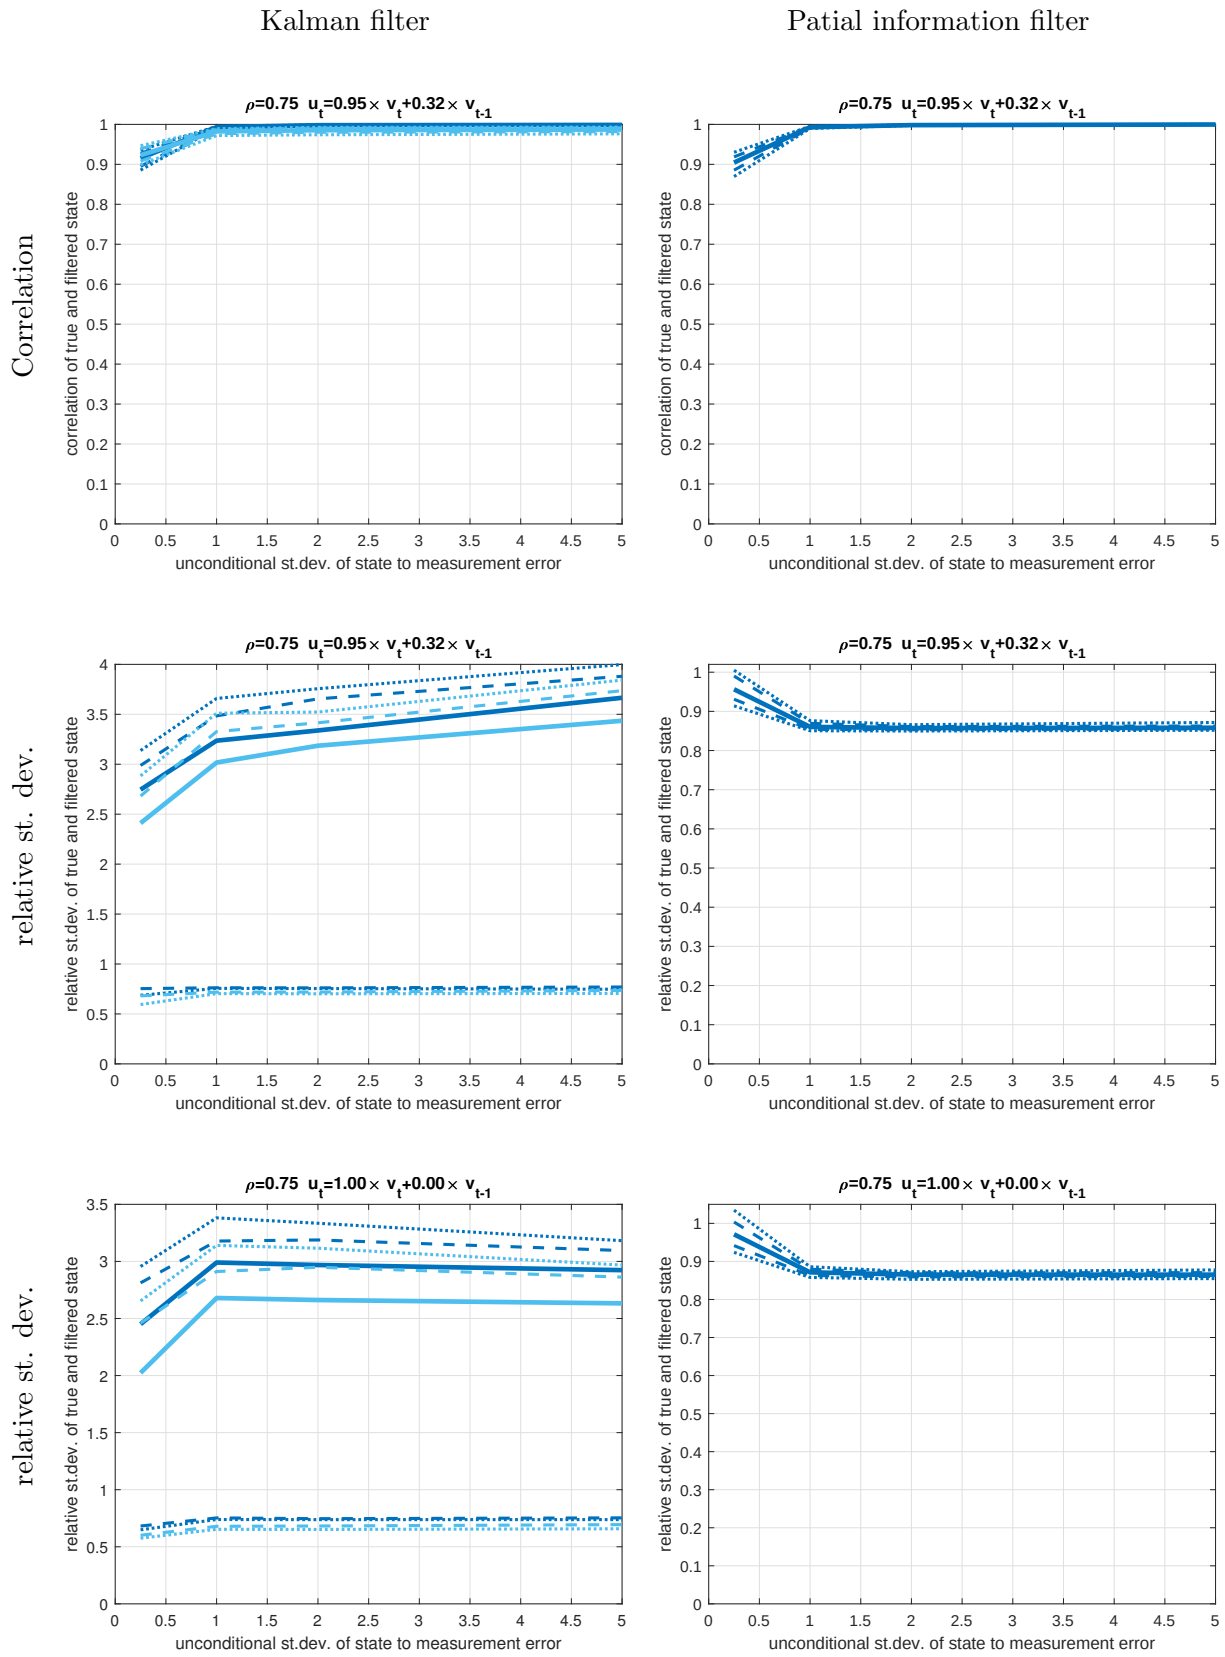

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

Figure A.4: Kalman filter vs partial information filter. Quadratic observation equation with $\kappa_1 = \frac{1}{10}, \kappa_2 = 10$ and linear state equation: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

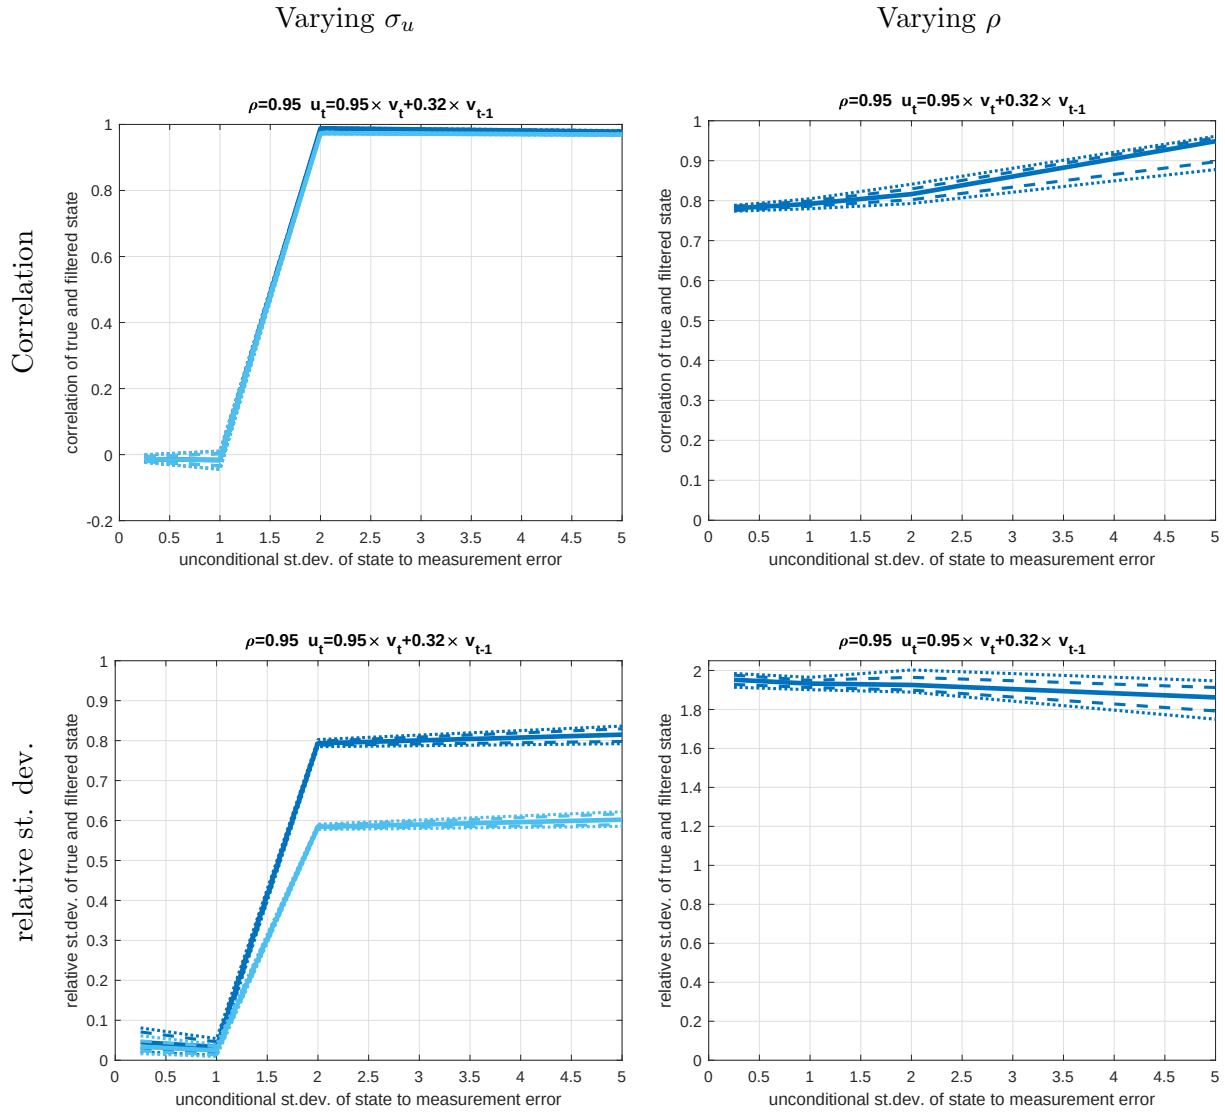

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

Figure A.5: Kalman filter vs partial information filter. Non-linear law of motion: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

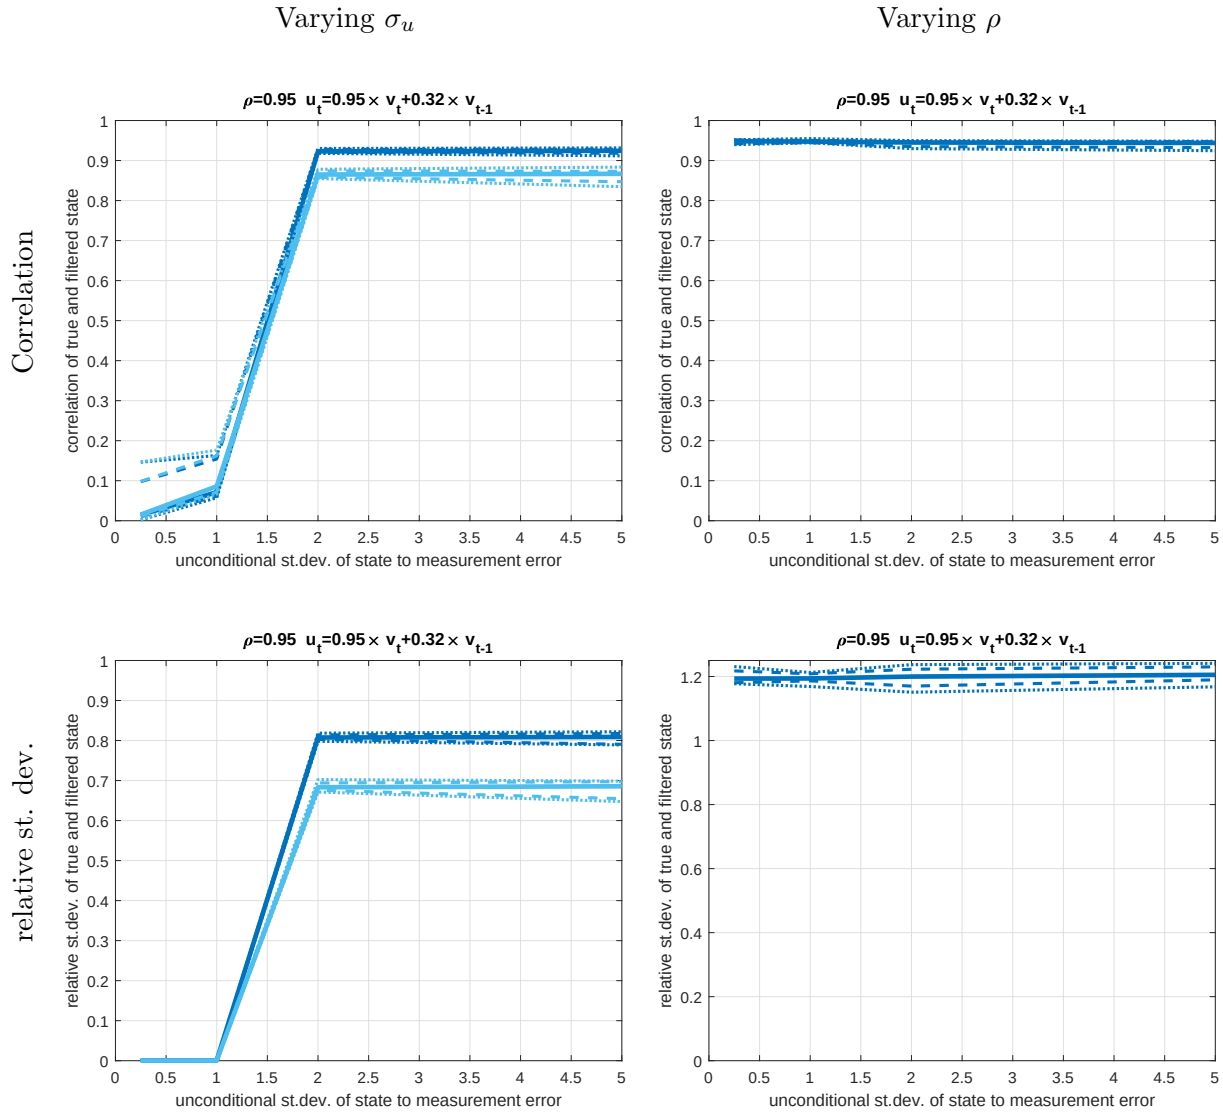

Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

Figure A.6: Kalman filter vs partial information filter. Non-linear law of motion: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

Appendix: Sudden Stop Model

This appendix presents additional results for the filtering exercise on the sudden stop model

B.1 Data

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).

Our measure for hours worked is an index of monthly hours worked in the Mexican manufacturing sector (FRED mnemonic: HOHWMN03MXQ661N). 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.[20]

With respect to the interest rate, when available, we use the LIBOR as our measure of the world interest rate (FRED mnemonic: USD3MTD156N; since 2022 this series is no longer available on FRED). Prior to the availability, we use the 3-month U.S. treasury rate (FRED mnemonic: DTB3), shifted so that during the first period of overlap the two measures agree. To convert these nominal measures into real rates, we subtract the one-year realized change in the log of the U.S. CPI (FRED mnemonic: CPIAUCNS).

We construct the dividend series using the dividends from the S&P/IFCG M MEXICO (IFGMMX) index, appended after 2008Q3 with the MSCI MEXICO (MSMEXFL) index. 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 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 INEGI.[21] Specifically, we use Annual Gross Domestic Product (GDP) by Expenditure at constant 2013 prices in National Currency, using Gross fixed capital formation and Gross domestic product (GDP). 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 according to the model up to a first-order approximation.

We obtain data for debt from Banco de Mexico.[22] We compute the quarterly debt in current prices 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.

[20] In 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). For Mexico, where we can only compare it to annual data from 1995 onward (mnemonic AVHWPEMXA065NRUG), the ratio is 1.52. While the annual data from Mexico indicate a slightly 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: $\frac{sd^q_{mfg,US}}{sd^a_{mfg,US}} sd_{mfg,MX} = \frac{2.51}{1.95} 1.52 = 1.96$.

[21] https://statistics.cepal.org/portal/cepalstat/open-data.html?lang=es.

[22] https://www.banxico.org.mx/SieInternet/consultarDirectorioInternetAction.do?sector=9&accion=consultarCuadro&idCuadro=CG7&locale=es.

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

B.2 Details of the implementation of the filter

Algorithm 1 Fixed-point algorithm for the Sudden Stop model

Step 0 Initialize the starting guesses as $q_t^{(0)} = 1+\mathcal{U}[-600^{-1}, 600^{-1}]$ and $\widetilde{\mu}_t$ 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 $[q_t^{(d-1)}/\bar{q}^{(d-1)}, \widetilde{\mu}_t^{(d-1)}]$.

Step 3 Solve for the latent variables $[q_t^{(d)}, \widetilde{\mu}_t^{(d)}]^T$ given $\hat{\mathbf{a}}^{(d)}, \hat{\mathbf{A}}^{(d)}, \hat{\mathbf{\Sigma}}^{(d)}$.

(a) Solve (4.10) and (4.11) for $[q_{*t}^{(d)}, \widetilde{\mu}_{*t}^{(d)}]^T$ 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 $\widetilde{\mu}_{*t}^{(d)} = 0 (\leq 0)$, set $\Delta_s^{(d)} = \Delta_s^{(d-)} - 4e-6$ for all s with binding constraints. If $\widetilde{\mu}_{*t}^{(d)} > 0$, but the borrowing constraint is higher than in pth percentile in sample, set $\Delta_s^{(d)} = \Delta_s^{(d-)} + 4e-6$. In the simulation, we set p equal to the target in-sample frequency plus 10pp (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 $\widetilde{\mu}_t^{(d)} = \max\{(1 - s_\mu^{(d-1)})\widetilde{\mu}_t^{(d-1)} + s_\mu^{(d-1)}\widetilde{\mu}_{*t}^{(d)}, -5e{-}4\}$. Here, $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.01$ for $d = 1, \ldots, D/10$, $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.02$ for $d = D/10 + 1, \ldots, D/5$, and $s_q^{(d-1)} = s_\mu^{(d-1)} = 0.0025$ for $d = D/5, \ldots, D$.

(d) For $d \leq D/20$ truncate $q_t^{(d)}$ at 1 ± 0.1 and 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 $\widetilde{\mu}_t$ changes within the last 100 iterations, go to Step 1. Else save and report the estimates.

Panels (a) and (b) of Figure B.1 shows that, on short-lived crisis in 1990 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.2 shows that the filtered Tobin's q closely tracks the fluctuations in measured dividends, crisis periods aside.

Figure B.3 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 $\widetilde{\mu}_t$ 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 Global Financial Crisis the 95th percentile

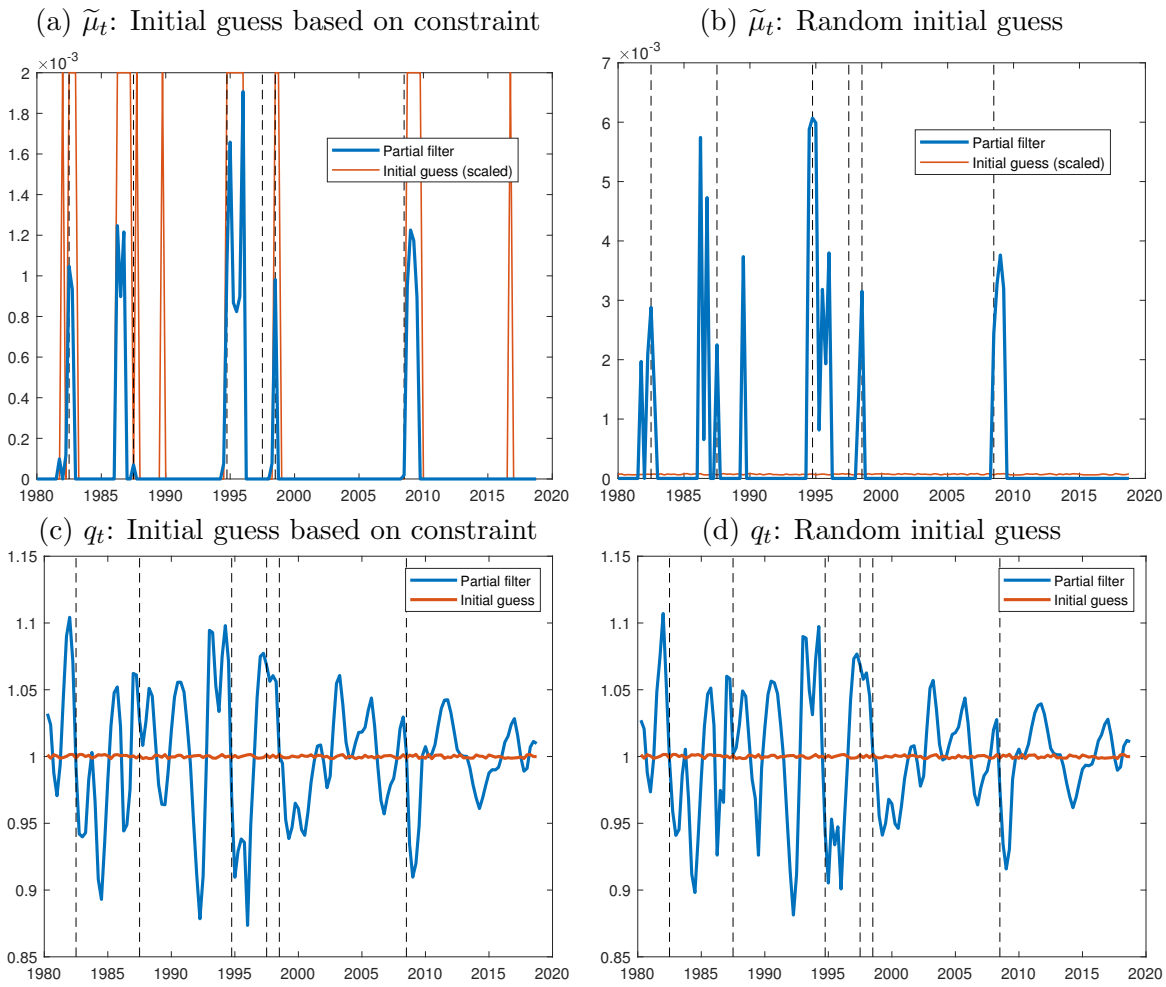

Figure B.1: Filtered values of $\widetilde{\mu}_t$ (top) and q_t (bottom) given different initial guesses.

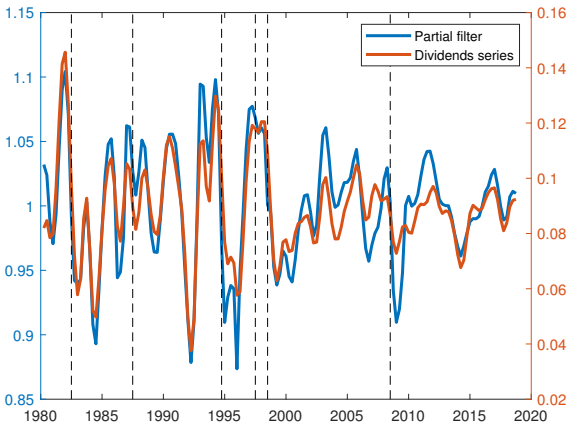

Figure B.2: Comparison of q_t in the partial filter and the dividend series.

is slightly higher in the 2nd half of the posterior than 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 $\widetilde{\mu}$.

(a) Tobin's q

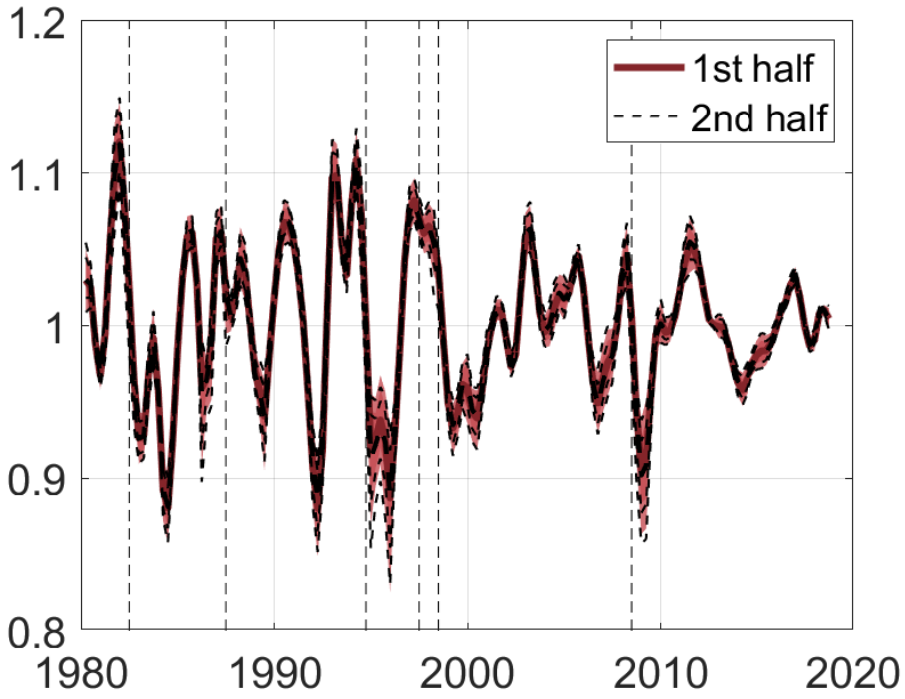

(b) Lagrange multiplier $\tilde{\mu}$

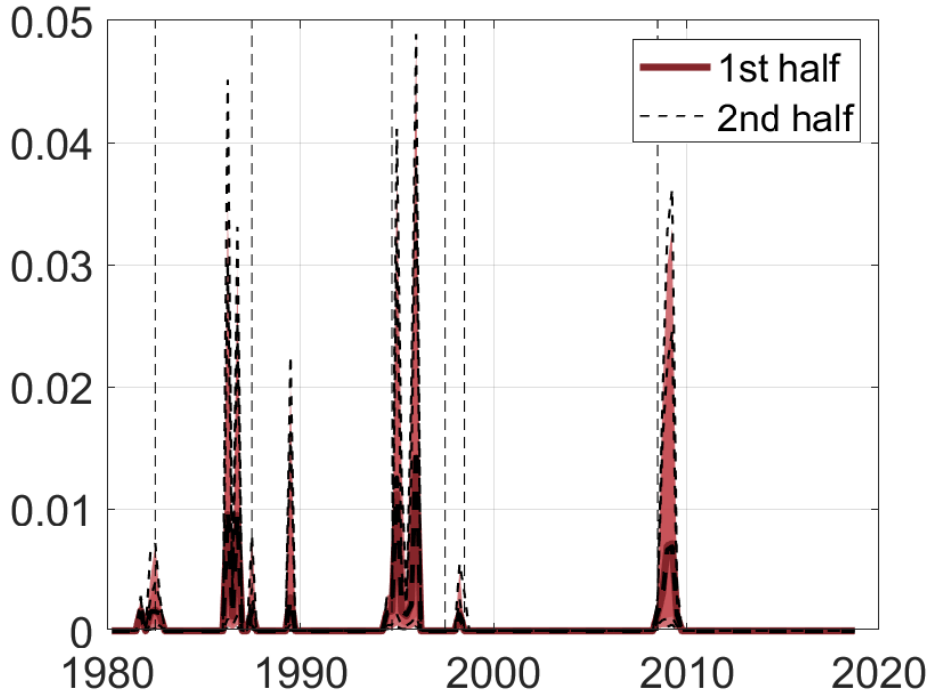

Shown are the median, 68% and 90% credible sets for the first half and the second half of the posterior draws.

Figure B.3: Comparison of the posterior based on the first half and the second half of the Gibbs sampler draws.

B.3 Filtering results without the slackness condition

Here, we compare results with and without using the slackness condition on the same simulated data. Without the slackness condition, we use the same convergence criterion for both q and $\widetilde{\mu}$ and do not use time-dummies to shift the expectation of u_{ct}. The results show an improved correlation of the filtered and the true $\tilde{\mu}$ (see Figure B.1), but a worse crisis classification (see Figure B.2).

	No slackness condition			
Sudden stop freq. [%]	corr$(\widehat{q}_t, q_t)$	corr$(\widehat{\widetilde{\mu}}_t, \widetilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\widetilde{\mu}_t$
[10.0, 15.0]	0.98 [0.94, 0.99]	0.36 [0.04, 0.59]	1.03 [0.99, 1.09]	1.06 [0.59, 1.64]
[30.0, 35.0]	0.98 [0.93, 0.99]	0.62 [0.29, 0.77]	1.05 [1.01, 1.14]	0.70 [0.50, 1.10]
[50.0, 55.0]	0.98 [0.95, 0.99]	0.75 [0.60, 0.83]	1.04 [1.01, 1.10]	0.80 [0.60, 1.00]
	With slackness condition			
Sudden stop freq. [%]	corr$(\widehat{q}_t, q_t)$	corr$(\widehat{\widetilde{\mu}}_t, \widetilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\widetilde{\mu}_t$
[10.0, 15.0]	0.84 [0.71, 0.94]	0.32 [0.12, 0.50]	1.11 [0.96, 1.37]	0.80 [0.35, 1.46]
[30.0, 35.0]	0.81 [0.65, 0.94]	0.35 [0.11, 0.53]	1.16 [0.97, 1.42]	0.68 [0.38, 1.07]
[50.0, 55.0]	0.85 [0.65, 0.95]	0.51 [0.25, 0.71]	1.03 [0.91, 1.38]	0.68 [0.50, 0.91]

Table B.1: Medians [68% CI] of correlations, relative standard deviations, and classifications errors. Comparison

	No slackness		With slackness	
	Crisis classification		Crisis classification	
Sudden stop frequency	False positives	False negatives	False positives	False negatives
[10.0, 15.0]	32.8 [16.4, 52.7]	6.8 [0.0, 27.4]	4.0 [1.6, 6.0]	13.8 [5.6, 20.8]
[30.0, 35.0]	33.6 [15.3, 48.8]	7.9 [0.0, 20.2]	9.4 [6.2, 12.9]	15.7 [8.2, 22.0]
[50.0, 55.0]	33.3 [17.0, 43.1]	2.8 [0.0, 12.9]	12.5 [10.2, 15.9]	7.1 [3.2, 15.1]

Table B.2: Medians [68% CI] of correlations, relative standard deviations, and classifications errors. Comparison