Filtering with limited information

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

August 22, 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, Rubio-Ramírez, and Schorfheide (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, Fernández-Villaverde, and Rubio-Ramírez (2017) 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

[1]For the simpler, linear Gaussian case, we can apply the Kalman filter. Some of the concerns below (i.e., curse of dimensionality) are less relevant for this situation. Others (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) still hold. As we show, a misspecified linear Kalman filter can yield misleading estimates the hidden state in our non-linear examples.

solving the model with perturbation (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, Fernández-Villaverde, and Guerrón-Quintana (2021), filters the shock process in an economy solved using the method of Andreasen, Fernández-Villaverde, and Rubio-Ramírez (2017). Gallant, Giacomini, and Ragusa (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. 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. Second, the researcher requires a consistent estimator of the model expectations.

Using VARs to approximate model expectations is not new in macroeconomics. For example, Campbell (1991), Bernanke and Kuttner (2005), and Chahrour, Cormun, Leo, Guerron-Quintana, and Valchev (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

[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.

the VAR.

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 2019. Section 5 concludes.

2 The partial information filter

In this section, we present how 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 for our procedure to work. 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 $\mathbf{f}(\circ)$ will usually stack in different rows 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 $\mathbf{g}(\circ)$ deals with the part of the model involving future states and observables. Typically, $\mathbf{x}_t$ is only a subset of the entire state vector $\mathbf{s}_t$ that determines the model expectations. The deterministic steady state of the model, $\bar{\mathbf{x}}, \bar{\mathbf{y}}$, is defined by

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

Three points about are equation (2.1) are important to highlight. First, the functions $\mathbf{f}(\circ)$ and $\mathbf{g}(\circ)$ do not necessarily need to include all equilibrium conditions required for a full solution of the model. Our goal is, precisely, to be able to get away with handling only a few of them. Second, at the cost of heavier notation, we could handle the case 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 to deal with situations where the conditional expectations depend explicitly on them. For instance, in a model with habit persistence, consumption today enters in the expected marginal utility tomorrow.

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 of the properties of $\mathbf{f}$ and $\mathbf{g}$, such as linearity, for example, or because the researcher knows that the hidden states lives 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 the vector of states other than $\mathbf{x}_t$, denoted $\mathbf{s}_{-,t}$, we have that there is a unique solution $\tilde{\mathbf{x}} \in \mathbb{X}_t$ such that $\mathbf{f}(\tilde{\mathbf{x}}, \mathbf{y}_t, \mathbb{E}[\mathbf{g}(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{y}_t)|\tilde{\mathbf{x}}, \mathbf{s}_{-,t}]) = \mathbf{0}$.*

The following scalar example is useful to fix ideas:

Example 1.

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

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

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

We consider two special cases.

(a) Linear model: $\kappa_2 = \rho_1 = 0$
Here $f(\circ) = y - \mu_y - x - \kappa_1 \mathbb{E}[x'|x]$. Under the true expectation, $\mathbb{E}[x'|x] = \mu_x + \rho_0 x$ so that $f(\circ) = y - \mu_y - (1 + \kappa_1\rho_0)x - \kappa_1\mu_x = 0$. 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:

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

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 specific equations in $\mathbf{f}(\circ)$ do not involve expectations (i.e., $\mathbf{g}(\circ)$ is a dummy argument of $\mathbf{f}(\circ)$), we can simply use functional form assumptions to 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, $\mathbb{E}_t[\mathbf{g}(\circ)]$ might come directly from survey data. We can consider the latter case as a particular application of the procedure we are about to describe.

Most commonly, however, we will deal with situations where $\mathbf{g}(\circ)$ matters and where $\mathbb{E}_t[\mathbf{g}(\circ)]$ 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 $\widehat{\mathbb{E}}_t(\circ)$ is an estimate of the conditional expectation.

To tackle the estimation of this conditional expectation, we will proceed in two steps. In the first step, we factorize the m-dimensional function $\mathbf{g}(\circ)$ 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 $g_{i,1}(\circ)$ could be a constant function, this decomposition is without loss of generality. However, in many applications to dynamic economies, 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 $\mathbf{g}(\circ)$. Choosing one or another depends on their suitability for the economic question being investigated or the data available.

Writing the terms in equation (2.5) in expected terms and using the elementary result 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 can 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 the use of equilibrium conditions, the auxiliary statistical model may help us to estimate them as well. With this evaluation (or estimation) of the expectations and the covariances estimate, we can go back to (2.4) and back out the hidden state vector.

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

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 iid shocks* ($\phi_1 = 0$) *and bounded domain:*
Given $\rho_1 = 0$, *we have that* $\mathbb{E}[g(x')|x, s_-] = \mathbb{E}[g(x')|x] = cons + (\kappa_1\rho_0 + 2\kappa_2\rho_0\mu_x)x + \kappa_2\rho_0^2 x^2$.

To bound the domain, let $\mathbb{X} = [\underline{x}, \bar{x}]$ *and* $x \in \mathbb{X}$*. Let* $\hat{\rho}$ *be a consistent estimator, such as the OLS estimator:* $\lim_{T\to\infty} \hat{\rho} = \rho_0$*. Then*

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

But because of the Continuous Mapping Theorem, we have that both $\rho_0 - \hat{\rho}$ *and* $\rho_0^2 - \hat{\rho}^2$ *converge to zero in probability. For any* δ *we can thus find a* T_δ *such that* $|\rho_0 - \hat{\rho}| < \delta$ *and* $|\rho_0^2 - \hat{\rho}^2| < \delta$*. Consequently, for any* ϵ *we can thus find an associated* δ_ϵ *and* $T_{\delta,\epsilon}$ *such that* $\sup_{x\in\mathbb{X}} |\mathbb{E}[g(x')|x, s_-] - \hat{\mathbb{E}}[g(x')|x, y]| < \epsilon,\ \forall T \geq T_{\delta,\epsilon}$*. Thus, the expectation converges uniformly.*

(b) Non-linear law of motion with iid shocks ($\phi_1 = 0$*) and linear observation equation(*$\kappa_2 = 0$*):* *Also with the non-linear law of motion, we 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* $\mathbb{E}[\mathbf{g}(\circ)|\circ]$*, and* $\mathbb{E}[\mathbf{g}(\circ)|\mathbf{x}, \circ]$ *is uniformly continuous in* $\mathbf{x}$*.*

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

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

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

To fix ideas, we model in the next subsection the covariance and the first moments 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 following notation:

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

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

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 (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 (2.4) is unique. With finite data, we can then, 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 $\widehat{\mathbb{E}}_t^{(0)}[\mathbf{g}(\circ)]$ 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. The Gibbs sampler allows us to quantify estimation uncertainty. 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 the following steps:

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 (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 from, the researcher can estimate the VAR on simulated data.[3]

2.3 Possible extensions

One could weaken the identifying assumptions above. For example, one could allow for *iid* measurement error in the filtering equation (2.1). In that case, equation (2.1) would hold only on expectation and researchers may need to adjust the estimator for the expectations to address the measurement error.[4] Below, we do 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.

Lastly, when we use a misspecified, inconsistent estimator for $\hat{\mathbb{E}}[\mathbf{g}(\circ)|\circ]$ 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]

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.

[3] In empirical applications, we typically add an initial step to either algorithm: A calibrations step, in which we calibrate first moments to be model consistent, by treating the (log) deviations of the observed series as (log) deviations from the model moments.

[4] For example, in the linear model with *iid* measurement error, used lagged filtered values as an instrument may correct for otherwise present attenuation bias.

[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.

We treat the observation equation (3.1a) as known to the researcher – given from a particular theory. Examples with such an expectation terms include a forward looking Phillips Curve or Consumption Euler equations. 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 here: x_t and v_t.

Our filtering approach involves three key steps: First, approximating the state dynamics with a (possibly non-linear) VAR. Second, using the VAR approximation to evaluate the expectations terms in the observation equation. Third, backing 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 in detail next.

Step 1. We first 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 Pth 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 $\mathbf{b}_{\ell,p}$ is statistically different from zero.

Step 2. The next step for our filter applies the well-known statistical identity that $\mathbb{E}[Z^2] = \mathbb{E}[Z]^2 + \mathrm{Var}[Z]$ to the measurement equation:

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

Step 3. The last step plugs the VAR dynamics (3.2) into this expression for the expectation terms in the observation equation and solves 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 then 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.[6]

A number of special cases of the system of equations (3.1) illustrates key 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.

In analyzing these three cases, we now 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. We then provide a comparison with the linear Kalman filter. Additionally, we consider a special case without misspecification.

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: $\{4^{-2}, 1, 2^2, 5^2\}$. We set $\phi_0 = 4 - \phi_1$, and vary the importance of the backward-looking MA component by letting $\phi_1 \in \{0, 1, 2, 3\}$. The filtering exercise also uses a burn-in in the simulations of 100 periods and then 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 Fully linear model without misspecification

Before discussing the case of misspecification, we first explain a special case with no measurement error and no MA component. This case is equivalent to increasing the signal to noise ratio to infinity in the case with measurement error.

[6] 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.

Compared to the fully linear case under misspecification, in 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:

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

The partial filter here is simple due to the perfect collinearity of x_t and y_t. We consider an AR(1), yielding $\hat{\mathbb{E}}_t[x_{t+1}^{(d)}] = \hat{A}_{1,1}[x_t^{(d)}] + \hat{a}_1$. Here, $\hat{a}_1$ is a constant term included in the AR(1), which is zero in population. While we could update $\hat{x}_t^{(d)}$ directly, we simply iterate until convergence on the following two steps:

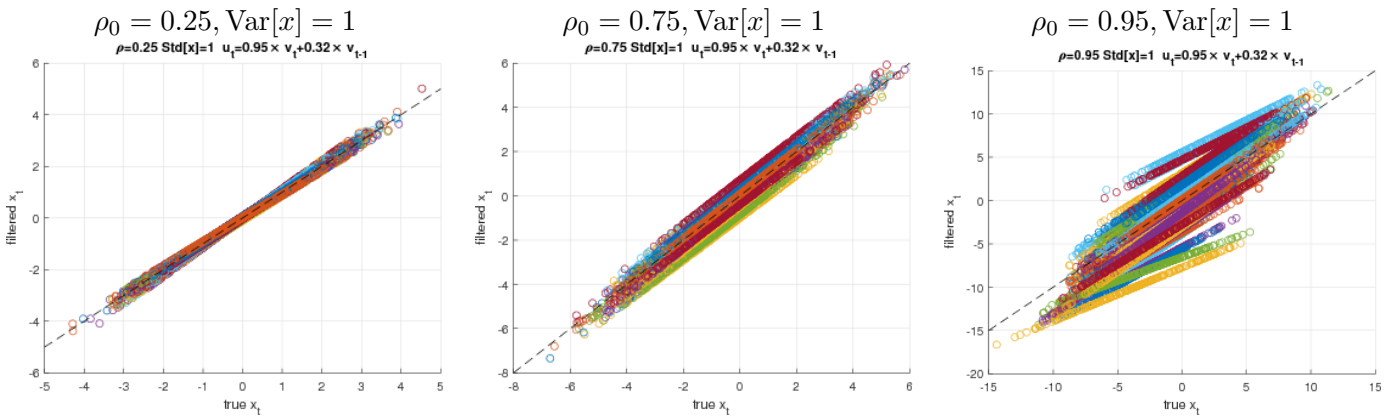
`For` $d = 1, 2, \ldots$ `do until convergence of` $\{x_t^{(d)}\}_t$

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

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

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 $\hat{\rho}_0$. We would then update $x_t = (1 + \hat{\rho}_0)^{-1} y_t$.

Figure 1 shows how the fit of the partial filter changes as a function of the persistence of the true state. Specifically, it shows scatter plots for three different levels of persistence: the left panel has a low persistence of $\rho_0 = 0.25$, the center panel of $\rho_0 = 0.75$, and the right panel has a persistence of $\rho_0 = 0.95$. Within each panel, the true state x_t is measured on the horizontal axis. Different colors correspond to the 100 different Monte Carlo samples. The dashed 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 low. 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 A for more details.



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

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

3.2 Fully 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 $\hat{\mathbb{E}}_t[x_{t+1}^{(d)}] = \hat{\mathbf{A}}_{1,\circ}[x_t^{(d)}, y_t^{(d)}] + \hat{\mathbf{a}}_1$. While we could update $\hat{x}_t^{(d)}$ 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. $x_t^{(d)} = y_t - \hat{A}_{11} x_t^{(d-1)} - \hat{A}_{12}^{(d-1)} y_t - \hat{a}_1^{(d-1)}$

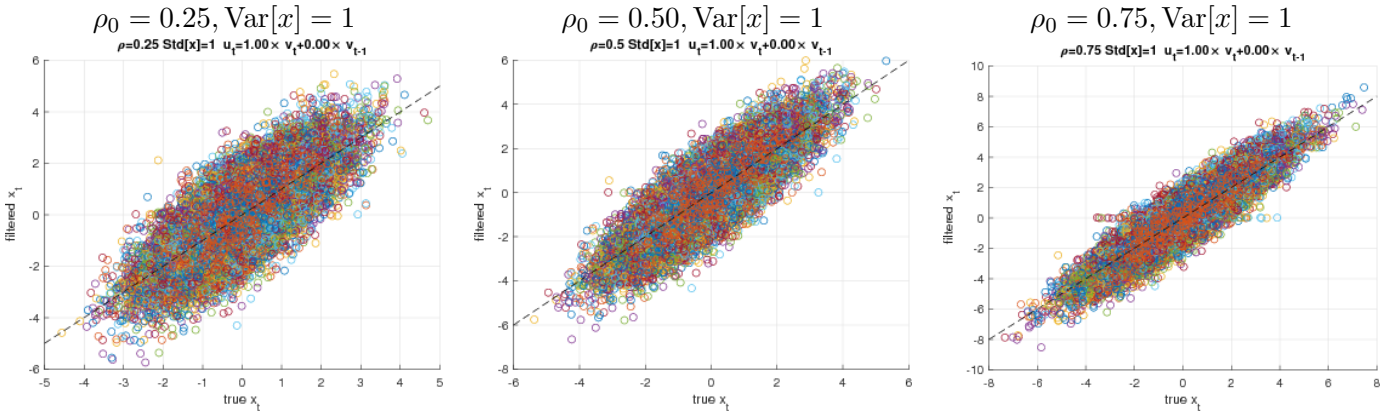
Step 2. estimate VAR coefficients $\hat{\mathbf{A}}^{(d)}, \hat{\mathbf{a}}_0^{(d)}$ regressing $\begin{bmatrix} y_1 & x_1^{(d)} \\ \vdots & \\ y_T & x_T^{(d)} \end{bmatrix}$ on $\begin{bmatrix} y_0 & x_0^{(d)} & 1 \\ \vdots & & \\ y_{T-1} & x_{T-1}^{(d)} & 1 \end{bmatrix}$.

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 $\hat{x}_t$ 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

[7] The results for non-zero ϕ_1 are similar, but the higher persistence marginally increases the correlations between the filtered and true series.

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.



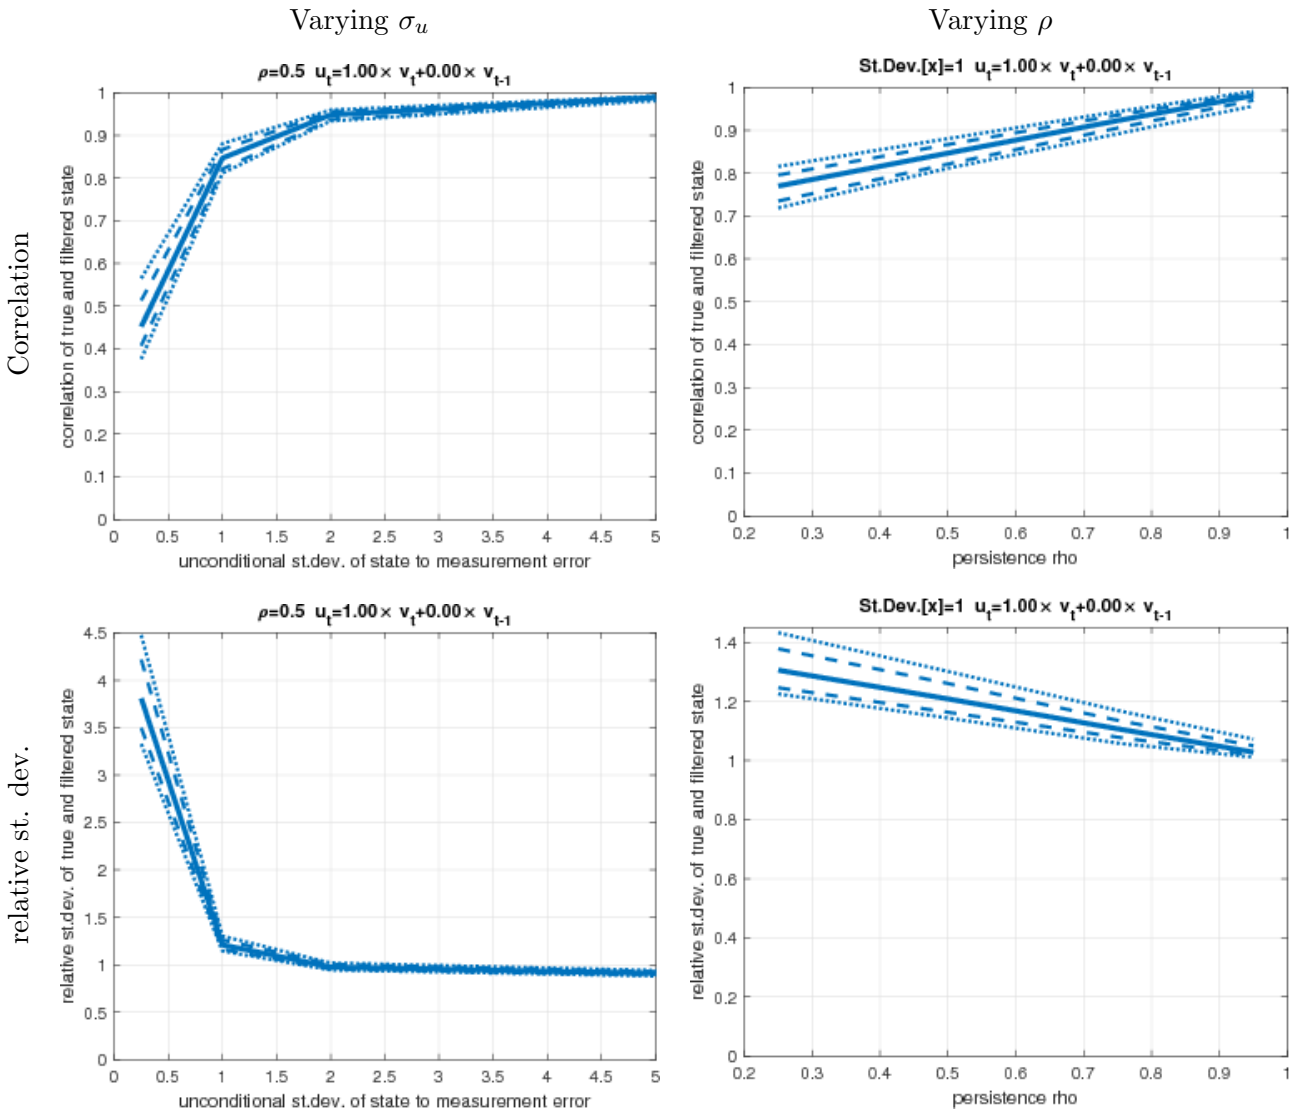
Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

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

Figure 3 evaluates the performance of the partial filter in terms of the correlation and standard deviation of the filtered and true state. The upper two panels show the correlation, while the lower row shows the relative standard deviations. The panels on the left hold the persistence ρ_0 fixed at $\rho_0 = 0.5$ and vary the signal to noise ratio, while the right panels hold the signal to noise ratio fixed and varies ρ_0. Overall, there is little variation across Monte Carlo simulations, as indicated by the narrow confidence intervals (dashed lines) relative to the median (solid line).[8] The left panel in the top row shows that the median correlations with a persistence of $\rho_0 = 0.5$ are above 0.8, unless the unconditional standard deviation of x_t is $\frac{1}{4}$, in which case it drops to 0.45. The right panel 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 0.8 to above 0.95.

The behavior of the relative standard deviation of the filtered and the true state (bottom row of Figure 3) behaves similarly 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.

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



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

Figure 3: Partial filter in the fully linear case: Correlation and relative standard deviation of the filtered state $\hat{x}_t$ vs true x_t for varying degrees of persistence ρ_0.

3.3 Quadratic observation equation with misspecification

Here, 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:

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

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 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

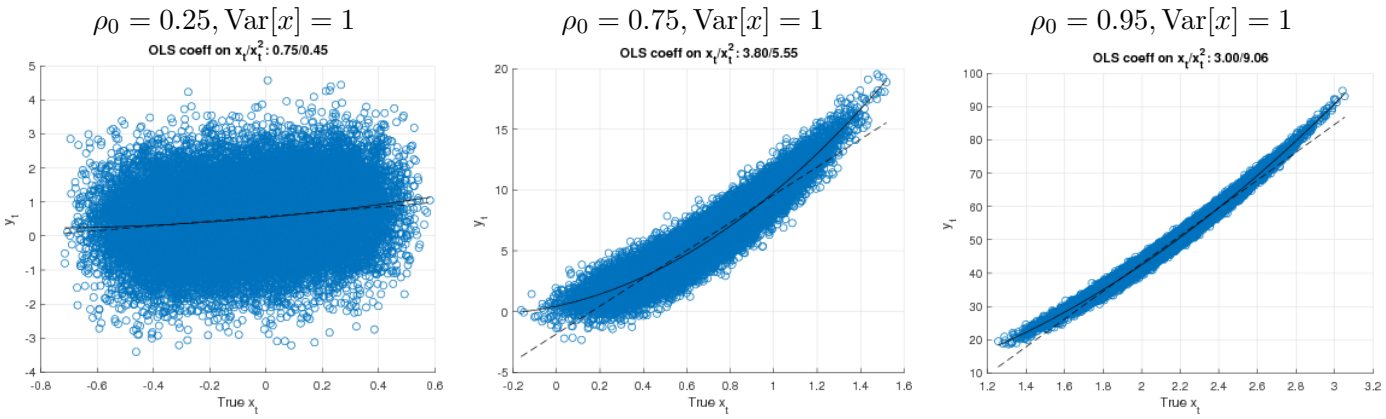
[9] We truncated the normal distribution to $[-1, 1]$ and scale it to have the desired variance σ_u^2.

with a reasonable scale. Here, we use the following guess, based only on observables and the known parameters of the observation equation:

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

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 4 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 $\rho_0 = .95$.

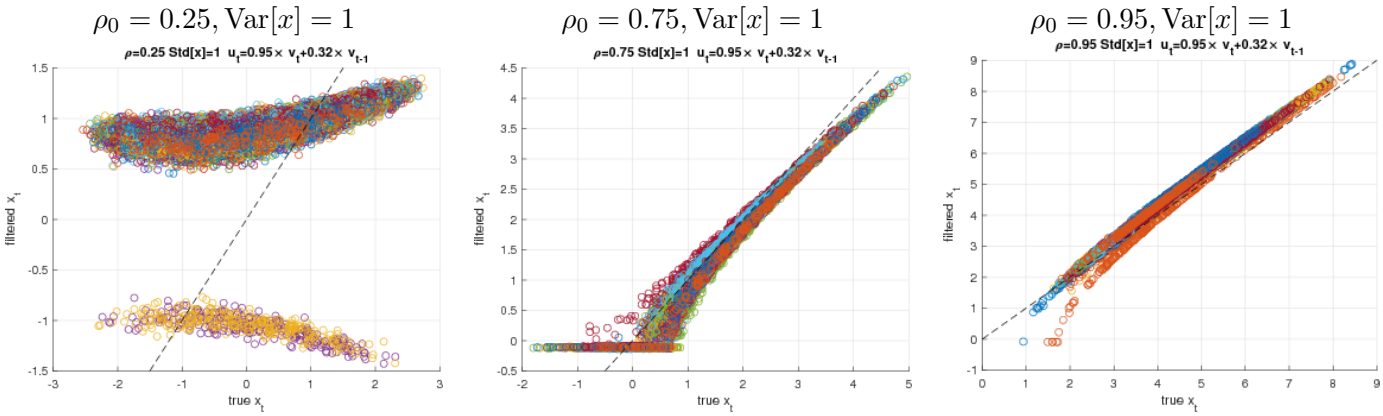


Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 4: Very-non linear case ($\kappa_1 = \frac{1}{10}, \kappa_2 = 10$) for varying degrees of persistence ρ_0. True state vs observation with T=25,000 observations.

Figure 5 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 $\hat{x}_t$ 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 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 .7 in 90% of samples. With higher persistence (center and right panel), the visual fit is remarkably high across all sample paths and the correlations exceeds .95 in 90% of sample paths.

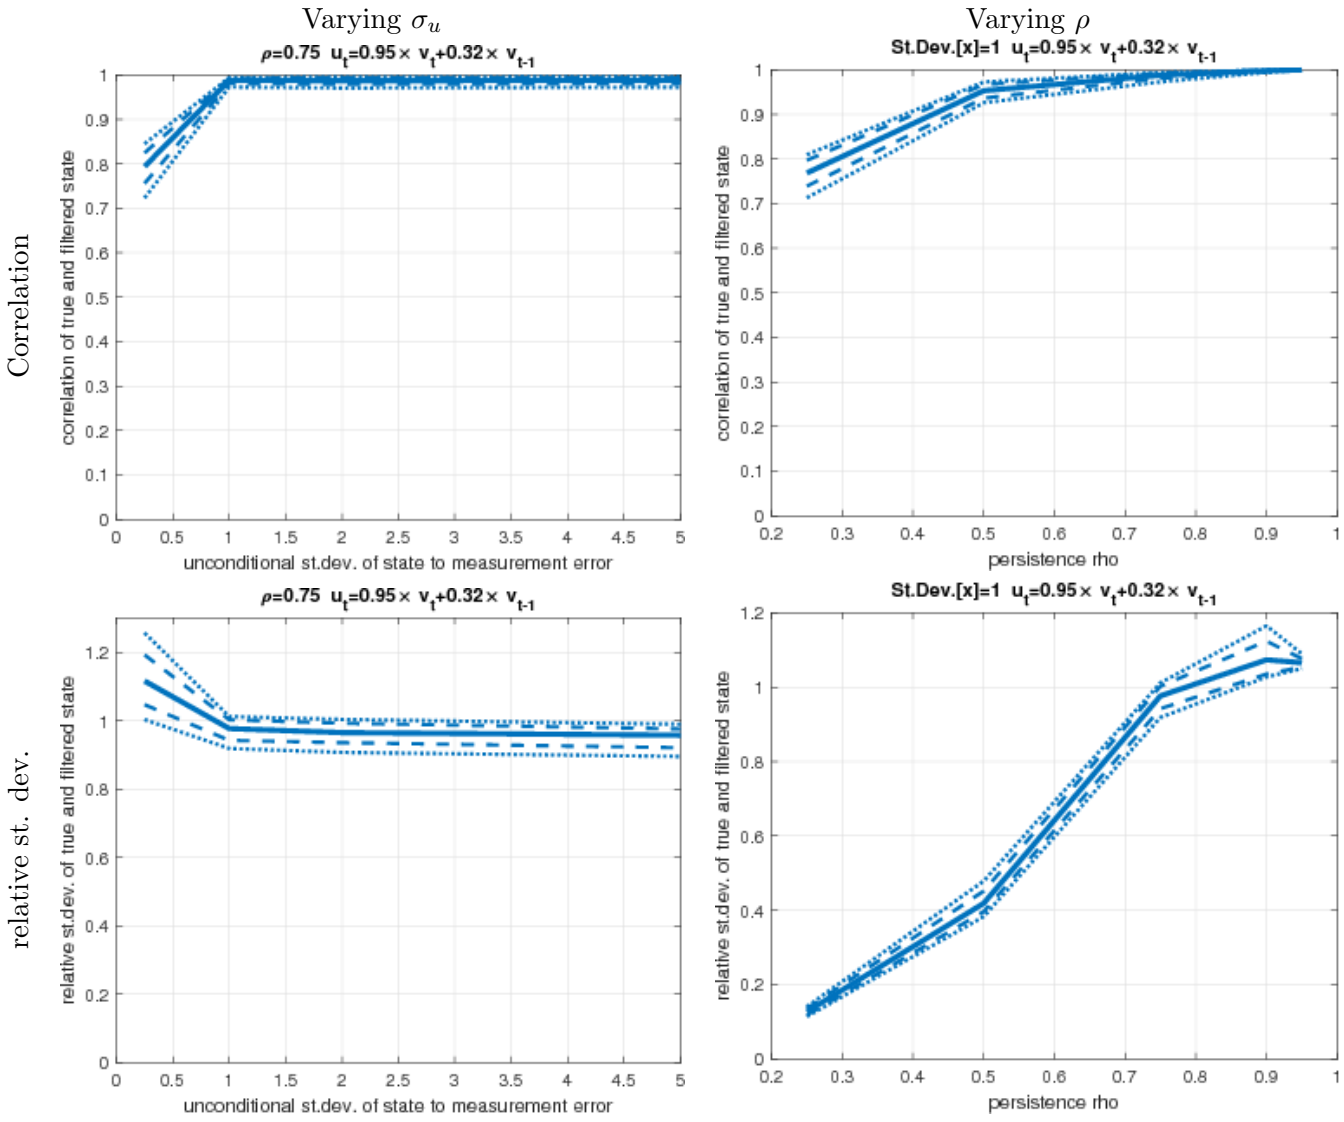


Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 5: 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 $\hat{x}_t$ vs true x_t for 100 different $T = 250$ samples.

Figure 6 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 plots the median correlation between the true and filtered states (along with 68% and 90% confidence intervals) in the top row, and the relative standard deviations of the filtered state relative to the true state in the bottom row. The two panels on the left vary σ_u, for fixed $\rho_0 = .75$. 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 above 0.75 in all cases shown, and often well above 0.9 with tight confidence intervals. The bottom row 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 5, the filter estimate is excessively smooth when the persistence is low. With $\rho_0 = 0.25$, the relative standard deviation is only about 0.15. It rises to 0.4 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. Figure A.2 in the Appendix shows the results for correlations and relative standard deviations 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 Figure 6: The correlations are now .9 or higher, as opposed to .75 or higher in the very non-linear case. The worst relative standard deviation is now .4 to .15 in the very non-linear case. In addition, the confidence intervals are tighter.



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

Figure 6: Very non-linear model: $\kappa_1 = \frac{1}{10}, \kappa_2 = 10$: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

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.

To illustrate this data-generating process, we once again turn to scatter plots. Figure 7 shows the relationship between the hidden state and the measurement in the top row. 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. The fact that this is due to the non-linear state dynamics is apparent in the bottom row. 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.

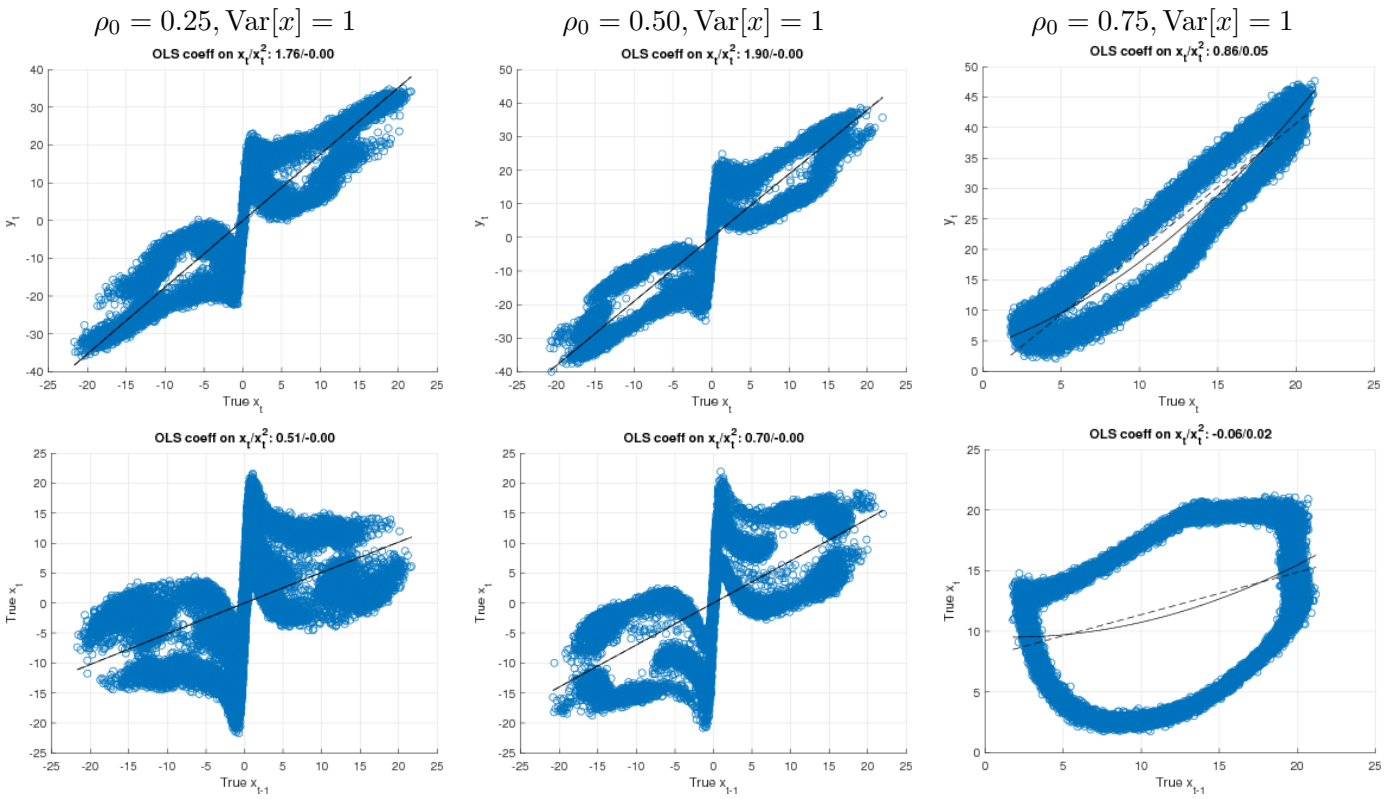
Despite the significant non-linearities, the partial filter does exceedingly well. Figure 8 shows that across all 100 sample paths, the filtered estimates are close to the 45 degree line. Even in the one sample path for $\rho_0 = 0.75$ where the state dynamics are in the negative region, a region of the state space not visited during any other simulation, the filter successfully captures this.

Figure 9 quantifies the performance of the partial filter. As before, the top row shows the correlation of the filtered estimates and the truth, with the solid line representing the median estimate. The bottom row shows the relative standard deviation of the filtered estimate. For an intermediate persistence $\rho_0 = 0.5$, the correlation is above 0.9 for all signal to noise ratios considered. 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 bottom panel is slightly high, but of the right order of magnitude.

3.5 Comparison with the Kalman Filter

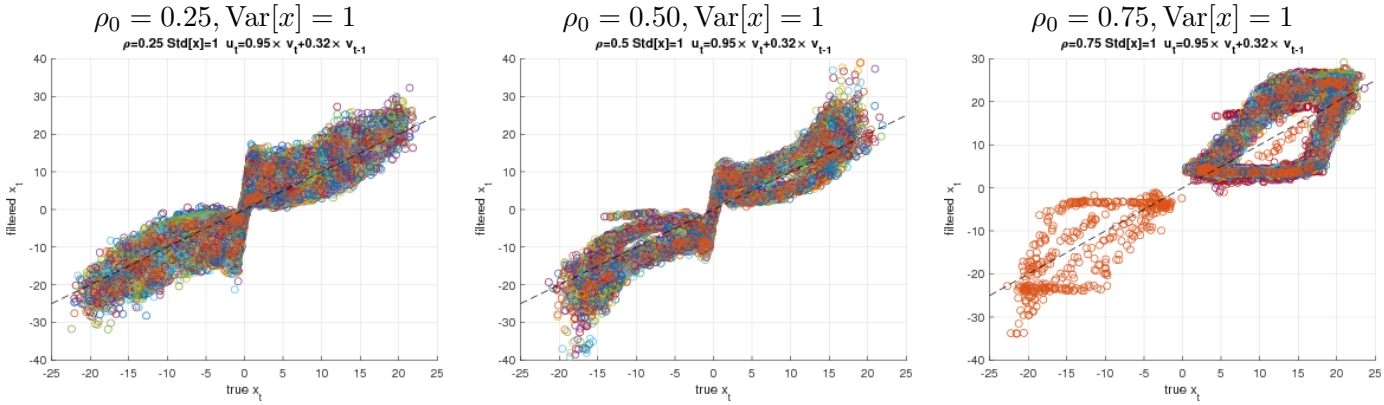
While the Kalman Filter is not applicable in all of the scalar examples that we are considering, it is still a useful benchmark. Because it is widely used for filtering dynamic economies, applied researchers may consider it as a natural starting point even when it is misspecified.

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.3.



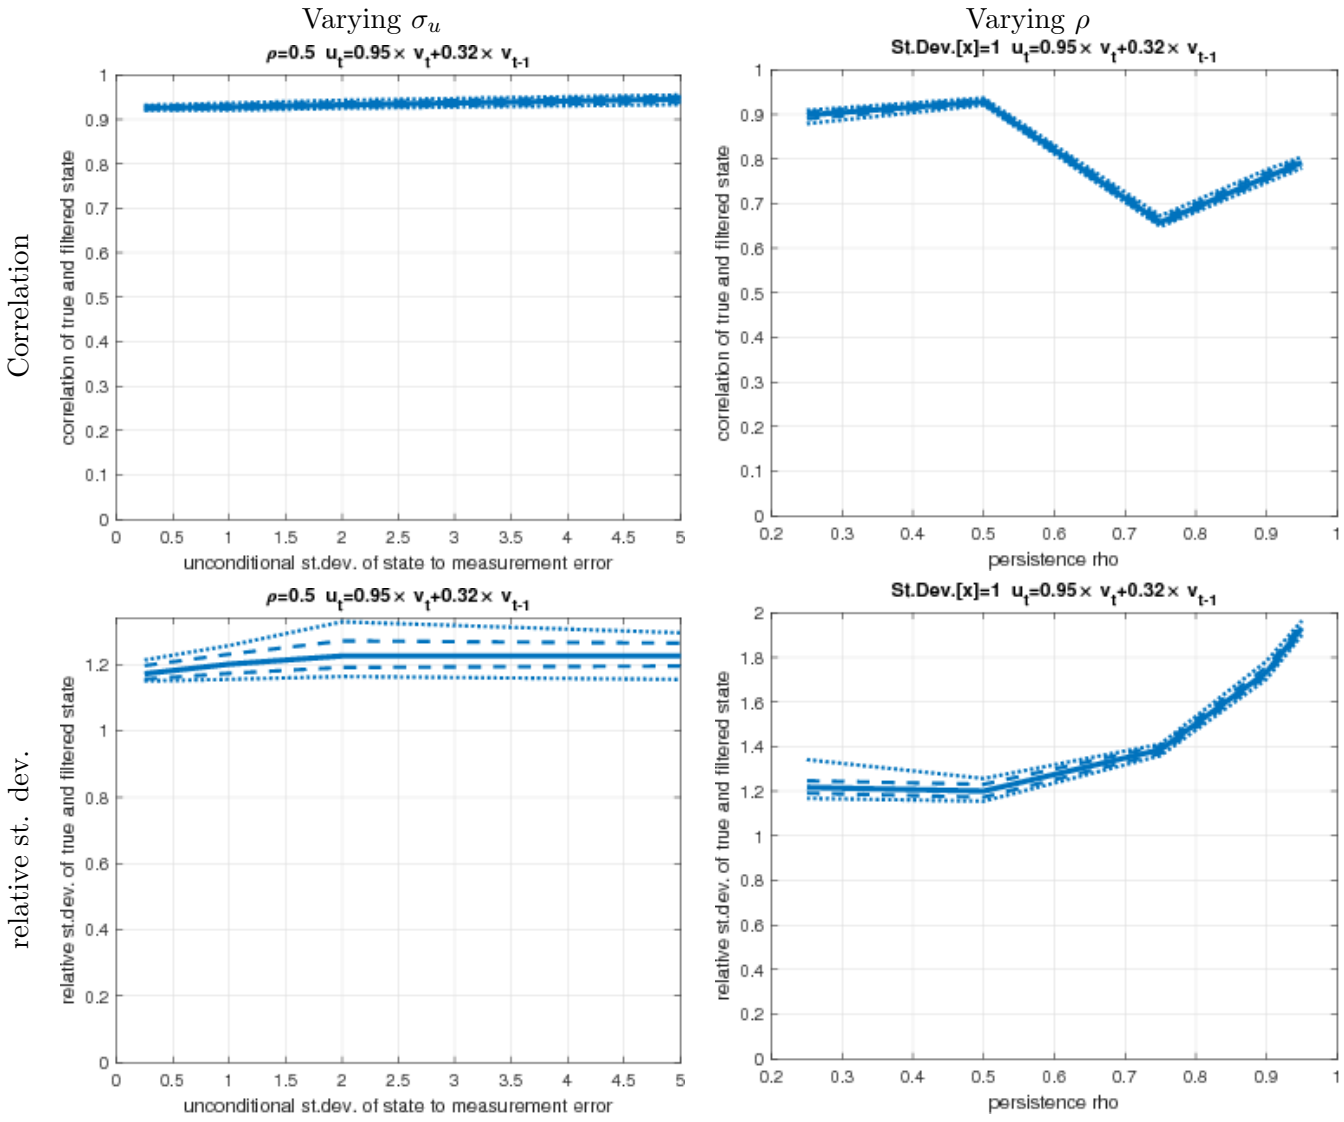
Note: Different colors correspond to different Monte Carlo samples. The dashed line is the 45 degree line.

Figure 7: Linear observation equation, non-linear state equation for varying degrees of persistence ρ_0. True state vs observation with T=25,000 observations (top row), true current state vs. lagged true state (bottom row).



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

Figure 8: Linear observation equation, non-linear state equation for varying degrees of persistence ρ_0. Filtered state $\hat{x}_t$ vs true x_t for 100 different $T = 250$ samples.



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

Figure 9: Non-linear law of motion: $\kappa_1 = \frac{1}{10}, \kappa_2 = 10$: 90 percent and 68 percent confidence intervals in Monte Carlo, along with median. $T = 250$ observations.

4 Application: Sudden Stops Model

We now apply the partial information filter to a quantitative, structural non-linear model. We choose the Sudden Stops Model of Mendoza (2010) as our application. The model introduces a loan-to-value (LTV) constraint in the workhorse RBC small-open-economy model.

4.1 Model structure and equilibrium conditions

In implementing the model, we follow Mendoza and Villalvazo (2020). The model is the same as in Mendoza (2010), except that instead of using preferences with endogenous discounting, we use standard time-separable expected utility with exogenous discounting at rate β.

The economy is inhabited by a representative firm-household with preferences defined over stochastic sequences of consumption c_t and labor supply L_t, for $t = 0, ..., \infty$, given by:

$$E_0\left[\sum_{t=0}^{\infty}\beta^t\frac{(c_t-\frac{L_t^{\omega}}{\omega})^{1-\sigma}}{1-\sigma}\right] \tag{4.1}$$

The representative agent chooses sequences of consumption, labor, investment, and holdings of real, one-period international bonds, b_{t+1} (the agent borrows when $b_{t+1} < 0$), so as to maximize the above utility function subject to the following budget and collateral constraints:

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

$$k_{t+1} = (1-\delta)k_t + \tilde{i}_t - \frac{a}{2}\frac{(k_{t+1}-k_t)^2}{k_t} \tag{4.3}$$

$$y_t = A_t F(k_t, L_t, v_t) = \exp(\epsilon_t^A) A k_t^{\gamma} L_t^{\alpha} v_t^{\eta}, \qquad \gamma = 1-\alpha-\eta, \tag{4.4}$$

$$-\kappa q_t k_{t+1} \leq q_t^b b_{t+1} - \phi R_t(w_t L_t + p_t v_t). \tag{4.5}$$

The left-hand-side of the budget constraint (4.2) is the sum of consumption expenditures, and gross investment including capital adjustment costs. Gross investment is $i_t = k_{t+1} - (1-\delta)k_t$ and gross investment inclusive of adjustment costs $\frac{a}{2}\frac{(k_{t+1}-k_t)^2}{k_t}$ is $\tilde{i}_t$. Correspondingly, capital follows the law of motion (4.3). Consumption is taxed at a constant rate τ.[10]

The right-hand-side of the budget constraint (4.2) is the sum of net profits from production, the resources generated by trading assets abroad, and government transfers. Net profits are equal to gross production y_t minus the cost of imported inputs minus the servicing of foreign working capital loans for labor and imported inputs. Gross output is represented by a constant-returns-to-scale technology, given in (4.4), that requires capital, k_t, labor, L_t and imported inputs, v_t, to produce a tradable good y_t sold at a world-determined price (normalized to unity without loss of generality). Working capital loans pay for a fraction ϕ of the cost of imported inputs and labor in

[10] Since the model does not include government expenditures, Mendoza (2010) includes a time-invariant consumption tax τ that is used to calibrate the model to match the average share of government expenditures in GDP in the data. Since the tax is constant it does not distort the savings-consumption margin and it does not create a time-varying distortion on labor supply.

advance of sales. These loans are obtained from foreign lenders at the beginning of each period and repaid at the end. Lenders charge the world gross real interest rate $R_t = R\exp(\epsilon_t^R)$ on these loans. Imported inputs are purchased at an exogenous relative price in terms of the world's numeraire $p_t = p\exp(\epsilon_t^P)$. The resources generated by trading assets abroad are given by $-q_t^b b_{t+1} + b_t$, where q_t is the price of the international bonds, which satisfies $q_t^b = R_t^{-1}$.

The collateral constraint (4.5) limits the total debt, which is equal to intertemporal debt plus working capital financing, not to exceed the fraction κ of the market value of the end-of-period capital stock. The prices q_t and w_t that appear in this constraint (and the wage in the budget constraint), are endogenous market prices taken as given by the agent when solving its optimization problem. As in Mendoza (2010), the wage rate must be on the labor supply curve (i.e. it must equal the tax-adjusted marginal disutility of labor), which requires $w_t = L^{\omega-1}(1+\tau)$, and the price of capital must satisfy the optimality condition requiring that $q_t = \frac{\partial \tilde{i}_t}{\partial k_{t+1}}$.

There are three types of shocks to the model: TFP $\log A_t = \log A + \epsilon_t^A$, where ϵ_t^A is a zero-mean random shock. Similarly, ϵ_t^P represents zero-mean shocks to the (log) world price of imported inputs (i.e., a terms-of-trade shock), centered around $\log p$. ϵ_t^R is an interest rate shock around a mean value $\log R$. The shocks ϵ_t^A, ϵ_t^R, and ϵ_t^P follow a independent first-order Markov processes.

4.2 Equilibrium Definition and Solution method

We solve for a competitive, recursive equilibrium. The representative agent maximizes the value function taking prices as given.

As shown in Appendix C, the optimality conditions can be simplified to:

$$u_c(t) = R_t\beta\mathbb{E}_t[u_c(t+1)] + \mu_t(1+\tau) \tag{4.6}$$

$$\mathbb{E}_t\left[R_{t+1}^q\right] - R_t = \frac{-cov_t(M_{t+1}, R_{t+1}^q) + \tilde{\mu}_t\left[(1+\tau) - \kappa\right]}{\mathbb{E}_t\left[M_{t+1}\right]} \tag{4.7}$$

$$L_t^{\omega-1}(1+\tau) = \frac{A_t F_{L_t}(k_t, L_t, v_t)}{(1+\phi(R_t-1) + \tilde{\mu}_t(1+\tau)\phi R_t)}, \tag{4.8}$$

where the marginal utility of consumption is given by $u_c(t) \equiv \left(c_t - \frac{L_t^\omega}{\omega}\right)^{-\sigma}$, the return on capital is given by $R_{t+1}^q \equiv (d_{t+1} + q_{t+1})/q_t$, and the stochastic discount factor is $M_{t+1} \equiv \frac{\beta u_c(t+1)}{u_c(t)}$. $\tilde{\mu}_t$ is multiplier of the budget constraint rescaled by the marginal utility of consumption: $\tilde{\mu}_t \equiv \frac{\mu_t}{u_c(t)}$.

The recursive equilibrium is solved for over a discrete state space, which requires defining discrete grids for (b, k, s). The grid for the shock triples $s \in \mathbf{S}$ comes from the discretization of the stochastic processes of the model's three shocks. We take the model calibration, the set $\mathbf{S}$, and the associated Markov transition probability matrix from Mendoza (2010), where $\mathbf{S}$ has eight triples (i.e. each shock has two realizations). For the endogenous states, we define grids with M nodes for bonds and N nodes for capital, respectively: $\mathbf{B} = \{b^1 < b^2 < ... < b^M\}$, $\mathbf{K} = \{k^1 < k^2 < ... < k^N\}$. The state space has $M \times N \times 8$ elements and is defined by all $(b, k, s) \in \mathbf{B} \otimes \mathbf{K} \otimes \mathbf{S}$.

While we can apply the partial filter to the appropriate data without solving the model, we

also use the model for a simulation study. Since the model is identical to Mendoza and Villalvazo (2020), we also use the *fixed-point iteration (FiPIt)* algorithm developed in Mendoza and Villalvazo (2020). The algorithm works by conjecturing the decision rule for bonds and the capital pricing function, deriving a set of implied decision rules that follow from these conjectures, and solving the Euler equations for bonds and capital via fixed-point iteration to find new values of the bonds decision rule and capital pricing function.

4.3 Calibration

We largely follow Mendoza and Villalvazo (2020) in the calibrating the model. They, in turn, follow Mendoza (2010) except for setting a constant discount factor β. We slightly lower the subjective discount factor β from .92 to .918 to yield an average crisis probability of 30.8%. We define a crisis as a period in which the collateral constraint is binding.. The shock processes are calibrated to match the persistence and volatility of the intermediate goods prices, the world real interest rate, and total factor productivity, as calculated by Mendoza (2010).

Parameter		c
Risk aversion coefficient	σ	2.0
Labor elasticity coefficient	ω	1.8461
Discount factor	β	0.918
Capital adjustment cost	a	2.75
Working capital parameter	ϕ	0.2579
Depreciation rate	δ	0.088
Labor share	α	0.59
Imported inputs share	η	0.10
Capital share of income	γ	0.31
Tax on consumption	τ	0.17
Average TFP	A	6.982
Collateral coefficient	κ	0.20
Intermediate goods prices: persistence		0.737
Intermediate goods prices: s.d.		3.345
World real interest rate: persistence		0.572
Intermediate goods prices: s.d.		1.958
Total factor productivity: persistence		0.537
Total factor productivity: s.d.		1.340

Table 1: Model calibration, following Mendoza and Villalvazo (2020).

The period length is one year. However, the ratio of flows to stocks scales in the period length, which allows us to use apply the model to quarterly data below.

4.4 Simulation study

Here, we use the Mendoza (2010) model to assess the performance of the partial filter in a controlled environment with a multivariate application in a non-linear model. Our univariate simulation study

has shown that the partial filter recovers the hidden state more reliably if it is more persistent. Since we want to keep the economic environment the same, we mimic this experiment by looking at sample paths with different properties. Specifically, we require sample paths in which the fraction of periods with a binding collateral constraint, $\mu_t > 0$, lies within a certain range. Implicitly, this systematically selects certain shock realizations. We retain only samples with the realized crisis probability in one of three ranges: Low crisis probability: $\sum_{t=1}^{250} \mathbf{1}\{\mu_t > 0\} \in [0.10, 0.15]$, intermediate crisis probability $\sum_{t=1}^{250} \mathbf{1}\{\mu_t > 0\} \in [0.30, 0.35]$, and high crisis probability: $\sum_{t=1}^{250} \mathbf{1}\{\mu_t > 0\} \in [0.50, 0.55]$. Since we do not change the calibration, the different selections implicitly select on different types of shock paths.[11] We take 3 samples: PCrisis ≥ 0.1, PCrisis ≥ 0.3, PCrisis ≥ 0.5.

4.4.1 Partial filter setup

Our goal is to filter the multiplier on the collateral constraint, expressed in units of consumption $\tilde{\mu}_t = \frac{\mu_t}{u_{c,t}}$, and the price of capital q_t. We thus define the hidden state (vector) $\mathbf{x}_t$ as $\mathbf{x}_t = [\tilde{\mu}_t, q_t]$.

We use (4.6) and (4.7) to filter the hidden states:

$$\mu_t(1+\tau) = u_c(t) - R_t \beta \mathbb{E}_t[u_c(t+1)] \tag{4.9}$$

$$\mathbb{E}_t\left[R^q_{t+1}\right] - R_t = \frac{-cov_t(M_{t+1}, R^q_{t+1}) + \tilde{\mu}_t\left[(1+\tau) - \kappa\right]}{\mathbb{E}_t\left[M_{t+1}\right]} \tag{4.10}$$

where $u_c(t) \equiv \left(c_t - \frac{L_t^\omega}{\omega}\right)^{-\sigma}$, $R^q_{t+1} \equiv (d_{t+1} + q_{t+1})/q_t$, $M_{t+1} \equiv \frac{\beta u_c(t+1)}{u_c(t)}$ and $\tilde{\mu}_t \equiv \frac{\mu_t}{u_c(t)}$.

Re-write using $R^q_{t+1} \equiv (d_{t+1} + q_{t+1})/q_t$:

$$\mathbb{E}_t\left[M_{t+1}\right]\mathbb{E}_t\left[d_{t+1} + q_{t+1}\right] - q_t R_t \mathbb{E}_t\left[M_{t+1}\right] = -cov_t(M_{t+1}, d_{t+1} + q_{t+1}) + q_t\tilde{\mu}_t\left[(1+\tau) - \kappa\right]$$

Solving for q_t:

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

Since we specify the VAR below in terms of marginal utility, we substitute out for the stochastic discount factor:

$$q_t = \frac{\frac{\beta}{u_{c,t}}\mathbb{E}_t\left[u_{c,t+1}\right]\mathbb{E}_t\left[d_{t+1} + q_{t+1}\right] + \frac{\beta}{u_{c,t}}\mathrm{Cov}_t(u_{c,t+1}, d_{t+1} + q_{t+1})}{R_t\frac{\beta}{u_{c,t}}\mathbb{E}_t\left[u_{c,t+1}\right] + \tilde{\mu}_t\left[(1+\tau) - \kappa\right]} \tag{4.11}$$

In addition to the two filtering conditions, we use the slackness condition:

$$\tilde{\mu}_t\left[q^b_t b_{t+1} - \phi R_t\left[w_t L_t + p_t v_t\right] + \kappa q_t k_{t+1}\right] = 0.$$

[11] Average TFP declines and the average import price increases with the frequency of crisis. TFP shocks average 1.002 when the crisis probability is in [.1, .15], but only averages .9995 when the crisis probability is in [.5, .55]. Average import price shocks average 1.0250 with a low crisis probability and 1.032 with a high crisis probability. While second moments remain stable, higher moments differ, too. For example, in the high crisis probability sample, the interest rate shock has a skewness of .36, as opposed to -.09 in the low crisis simulation.

Note that, because of the Cobb-Douglas production function, 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.

We use the information in the borrowing constraint to adjust the state estimates: When the borrowing constraint is binding but the multiplier $\tilde{m} u_t$ is not estimated to be positive, we lower our estimate of the expected future utility – and vice versa when the borrowing is slack, but the estimated multiplier is inferred to be positive. See Appendix C.4 for details.

In anticipation of our empirical specification, we assume that we can measure C_t and L_t, which gives us $u_{c,t}$. Then, we include $u_{c,t}$ and L_t in the VAR, together with R_t. These three variables are enough to be able to filter $\tilde{\mu}_t = \frac{\mu_t}{u_c}$. To be able to filter q_t, we assume we can observe d_t as well. Hence, we define the vector of observables $\mathbf{y}_t$ as $\mathbf{y}_t = [u_c, L, R_t, d_t]$. We now have a 3 dimensional g-function:

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

Then we also have for the f-function that:

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

Substituting in the VAR for the expectations and covariances, we get the following approximation. We use that: $\boldsymbol{\xi}_t^{(d)} = \mathbf{a}^{(d)} + \mathbf{A}^{(d)} \boldsymbol{\xi}_{t-1}^{(d)} + \boldsymbol{\varepsilon}_t^{(d)}$ and $\boldsymbol{\Sigma}^{(d)} = \widehat{\mathrm{Var}}[\boldsymbol{\varepsilon}_t^{(d)}]$, together with $\boldsymbol{\xi}_t^{(d)} = [u_{c,t}, L_t, R_t, d_t, \tilde{\mu}_t^{(d)}, q_t^{(d)}]$.

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

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

For given parameters, we slowly update the guesses for $\hat{q}_t$ and $\hat{\tilde{\mu}}_t$, using rules of the form $\hat{q}_t^{(d)} = (1-s)\hat{q}_t^{(d-1)} + s\mathbf{x}_{2,t}^{(d)}$, where $s = 3e-4$. Similarly, for $\tilde{\mu}$ we use $s = 1e-4$. Because in the model $\tilde{\mu}_t \geq 0$, we censor $\hat{\tilde{\mu}}_t$ at $-5e-4$ and, for numerical stability, we restrict $\hat{q}_t$ to lie within $[.9, 1.1]$ and to average 1.[12] To capture the fact that $\tilde{\mu}$ is an occasionally binding constraint, in each

[12] These bounds have minor effects in the simulated data, but improve the numerical stability in the empirical

iteration we push it down by a constant of $1.67e-7$.

4.4.2 Results

The partial filter does well in recovering the price of capital q_t, independent of the crisis probability in the sample. It is harder to recover the multiplier on the borrowing constraint, but the model produces reasonable estimates throughout.

Table B.1(a) shows the correlation and relative standard deviation for the two filtered variables for each of the three crisis probability scenarios. For Tobin's q, the median correlation across simulations is near .9 in all three scenarios. The median standard deviation of the estimated $\hat{q}_t$ relative to the true q_t rises from .93 to 1.06 as the in-sample sudden stop probability rises from $[.1, .15]$ to $[.5, .55]$. Table B.1(b) lists the false positive and false negative rates when classifying crises. The false positive rates range from a median of .07 when the true crisis probability is around .1 to .14 when the true crisis probability is between .5 and .55. The false negative rates are slightly smaller, ranging from .03 to .11.

(a) Correlations and relative standard deviations

Sudden stop probability	$\text{corr}(\hat{q}_t, q_t)$	$\text{corr}(\hat{\tilde{\mu}}_t, \tilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\tilde{\mu}_t$
[0.1,0.15]	0.92 [0.87, 0.95]	0.35 [0.17, 0.56]	.93 [.86, 1.00]	0.61 [0.39, 1.08]
[0.3, 0.35]	0.90 [0.85, 0.93]	0.30 [0.14, 0.40]	.98 [.91, 1.08]	0.50 [0.37, 0.71]
[0.5, 0.55]	0.86 [0.81, 0.91]	0.28 [0.11, 0.42]	1.06 [.95, 1.14]	0.51 [0.39, 0.66]

(b) Crisis classification errors

Sudden stop probability	Crisis classification	
	False positives	False negatives
[0.1,0.15]	0.07 [0.05, 0.09]	0.03 [0.01, 0.07]
[0.3, 0.35]	0.11 [0.09, 0.16]	0.09 [0.05, 0.15]
[0.5, 0.55]	0.14 [0.11, 0.16]	0.11 [0.07, 0.16]

Table 2: Partial filter performance across simulations as a function of the sample probability of sudden stops: Medians [68% CI] of Correlations, relative standard deviations, and classifications errors

Figure B.1 shows the corresponding data. Its three columns correspond to the three targeted in-sample crisis probabilities. The top panels show the time path of the true (blue) and estimated (orange) time series for $\hat{q}_t$. In line with the summary statistics in Table B.1, the two lines track each other well in all three cases. The amplitude is somewhat more closely aligned when the crisis probability 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. This is very clear in the bottom panel, which shows scatter plots of the relationship between the true and the filtered q_t. While centered on the 45-degree line, the filtered q_t has a bimodal distribution. This reflects the two regimes: when

application.

the borrowing constraint is slack, the price of capital reflects the standard equity premium. When the constraint binds, higher expected future excess returns on capital increase the discount rate of dividends, and lower the current asset prices. This could even set in motion the debt-deflation mechanism: when the constraint binds, agents are forced to liquidate capital. This fire sale of assets reduces the price of capital and further tightens the constraint, lowering the price even further. The scatter plots show the different simulations in different colors. This reveals that with the low crisis probability, in the bottom left panel, there is one sample path (out of 100; shown in light green) for which the partial filter has trouble recovering the true q_t, as indicated by the cluster of markers in the bottom right, away from the 45-degree line. In all other cases and sample paths, the filter performs well.

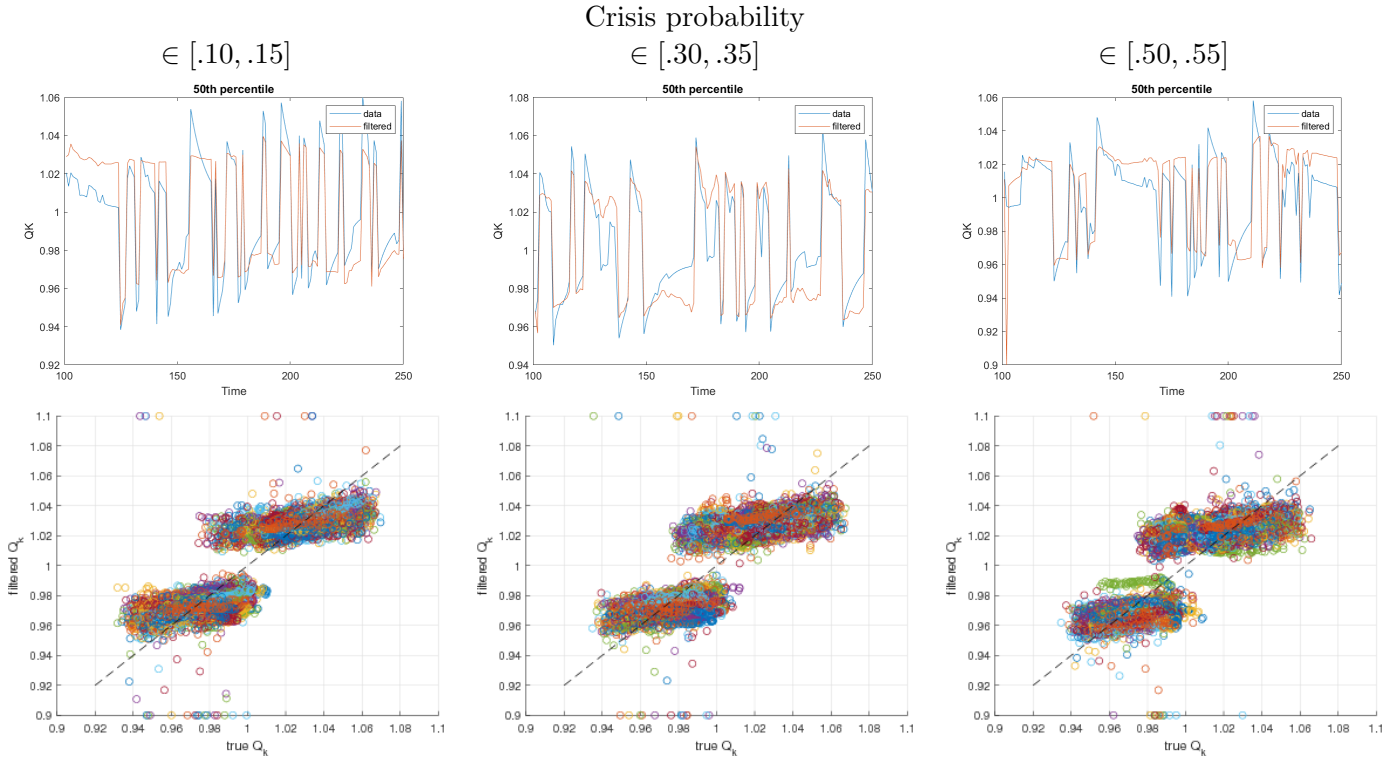


Figure 10: Sample paths for the price of capital q_t for the sample paths with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis probability.

Returning to Table 1, we see that the correlation of the filtered multiplier on the borrowing constraint is lower, and is slightly lower when the in-sample crisis probability is higher. The correlation of the filtered multiplier with the true multiplie falls from .35 to .28.[13] The relative standard deviation, however, falls from 0.61 to .51.

Figure B.2 shows representative sample paths and overall scatter plots for the multiplier on the borrowing constraint. As before, the top panels show the sample paths for the filtered and actual

[13]When we do not use the slackness condition to infer $\tilde{\mu}_t$, the increase with the crisis probability is much stronger. In Appendix B we describe the results if we do not use the slackness condition.

multipliers for the median correlation. The left panels corresponds to a low crisis probability, and the center and right panels to higher crisis probabilities. Clearly, as the crisis probability increases, the top panels are less sparse. When the crisis probability is low, the partial filter successfully identifies the most severe crisis, around period 225. It also identifies some of the smaller crisis, but largely misses the one around period 240. With the less sparsely populated panels with a higher crisis probability, the partial filter 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 axes, so there are fewer false positives and false negatives. While the mass shifts away from zero true and inferred multipliers, it is concentrated below the 45-degree line.

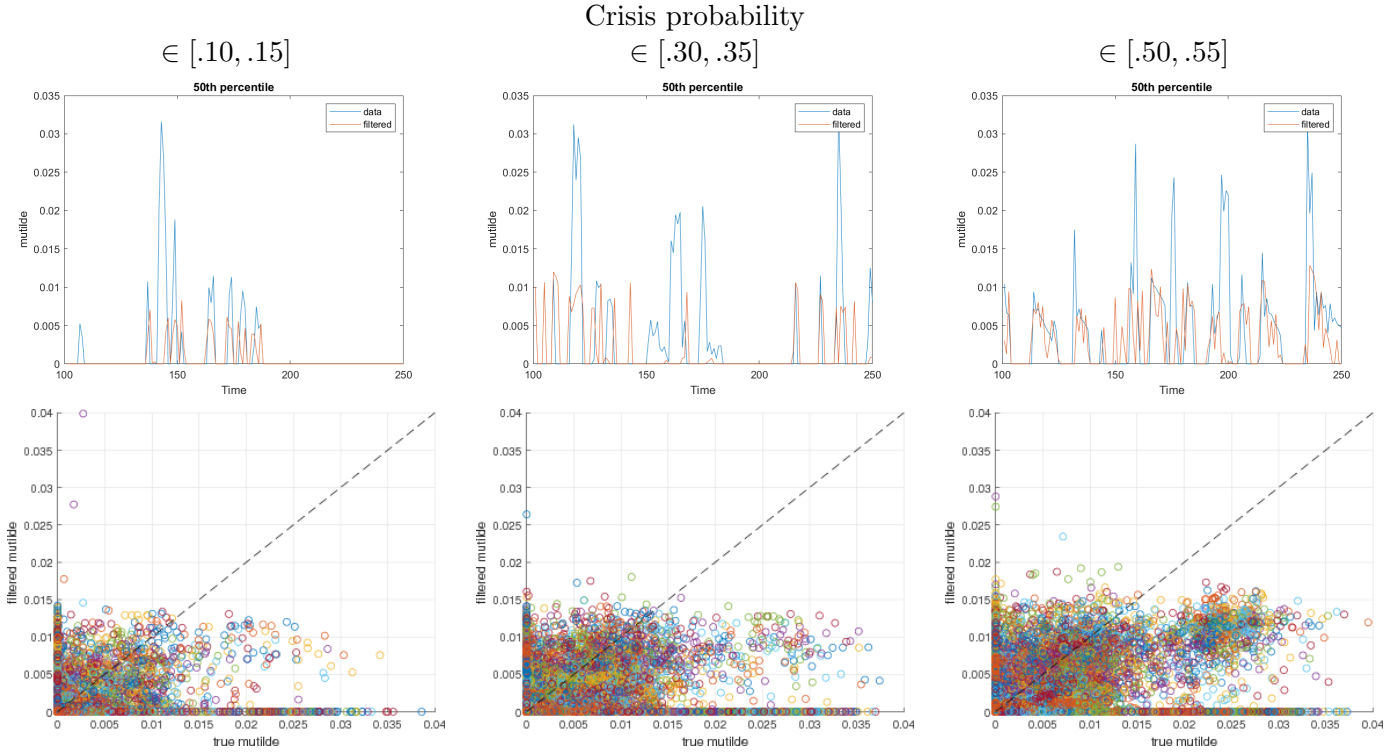


Figure 11: Sample paths for the multiplier on the borrowing constraint $\tilde{\mu}_t$ for the sample paths with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis probability.

4.5 Application to Mexico

Having established that the filter reliably recovers Tobin's q and recovers crises reasonably well, we now apply the filter to Mexican data.

4.5.1 Data and partial filter specification

In our application to Mexico, we use quarterly data. To account for trends outside the model, we use the HP-filter to detrend the variables (in logs), and then 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 evaluate the borrowing constraint using the perpetual inventory method and linearly interpolate to obtain a quarterly series. See Appendix C.5 for details.

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 push $\tilde{\mu}$ down if it is estimated to be postive but the slackness condition is above its 25th percentile. The results without using the slackness condition are summarized in Appendix C.5.

4.5.2 Results

Perhaps the most salient criterion for success is whether the partial filter successfully recovers the Mexican debt crises during our sample. To facilitate this analysis, Figure 12(a) shows the filtered multiplier on the borrowing constraint, $\tilde{\mu}_t$ and also indicates a number of economic events, that we suspect could lead to a binding borrowing constraint. These events are: (1) the debt crisis of 1982q3, (2) the collapse of oil prices in 1986 that predated Black Friday in 1987q3, (3) the so-called "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.

(a) Borrowing constraint multiplier $\tilde{\mu}_t$ (b) Tobin's q

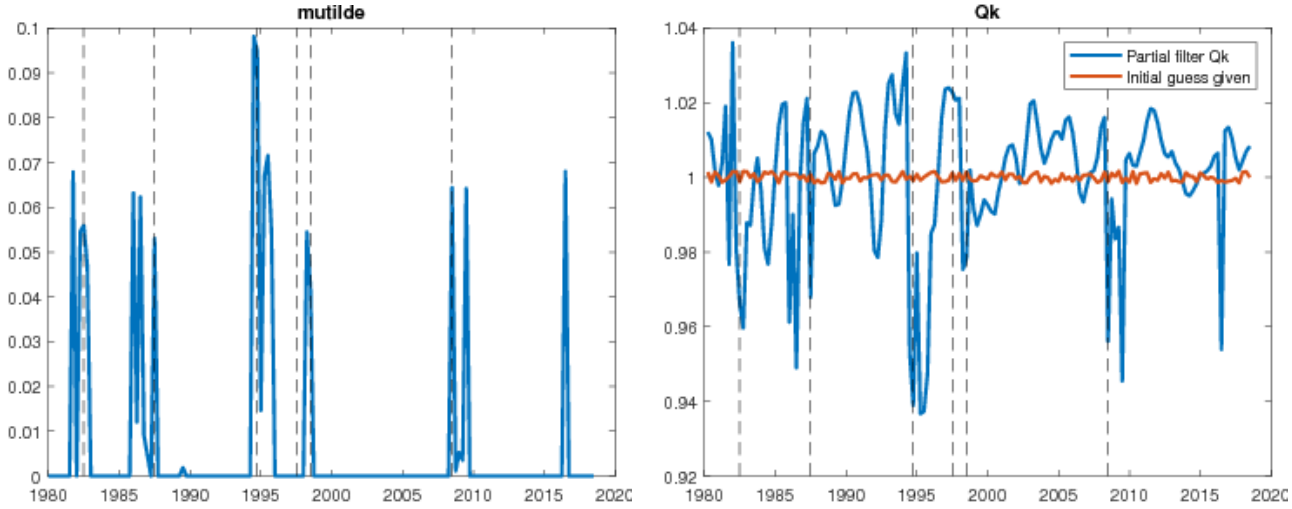


Figure 12: Filtered value of the borrowing constraint multiplier and Tobin's Q in the Mendoza (2010) model applied to Mexican data.

The partial filter is successful in identifying these crisis events. It indicates two false positive events – in 1989 and 2016, but particularly the event in 1989 is small. Small changes in the implementation of the filter eliminate both false positive events without materially affecting any

other filter implications.[14]. There is only one false negative – the Asian crisis of 1997q3. However, it is unclear whether this crisis should have an impact on Mexico and the filter identifies the subsequent financial turmoil in 1998.

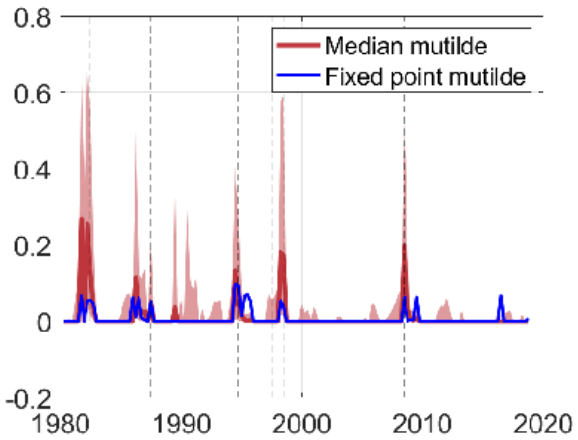
Figure 12(b) shows the corresponding movement in Tobin's q. The blue line represents the filtered value for q_t. the same crisis events that are associated with a binding borrowing constraint are also associated with a lower price of capital – which is natural because the binding borrowing constraint 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.

4.5.3 Parameter uncertainty

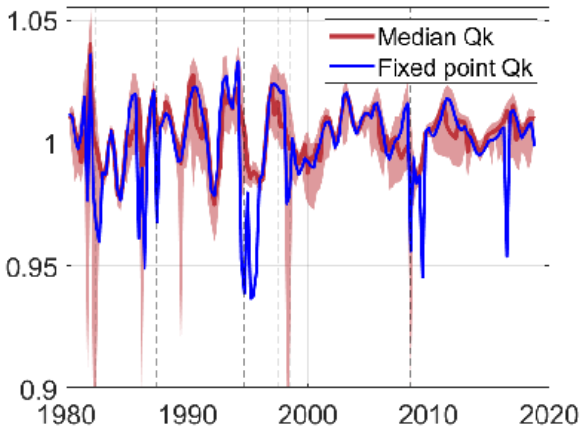
While we treated the estimated expectations as certain in the analysis so far, we now quantify the effect of parameter uncertainty. To do so, we use a Gibbs sampler. We iterate between estimating the VAR with a flat prior and filtering $\tilde{mu}_t$ and q_t, which, in turn, enter the VAR as data.

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 13 (a) shows 68% credible sets, the posterior median, and the fixed point estimates of the multiplier on the borrowing constraint. While the 68% credible sets are typically wide, the filter only identifies significantly positive multipliers in crisis times. During these times, the multiplier is positive with at least 84% posterior probability. The posterior median shows one false positive in 1989.

(a) Borrowing constraint multiplier $\tilde{\mu}_t$



(b) Tobin's q



Posterior median and 68% credible sets

Figure 13: Filtered value of the borrowing constraint multiplier and Tobin's Q in the Mendoza (2010) model applied to Mexican data. Comparing the posterior with the fixed point estimate.

Another feature of the posterior uncertainty is its asymmetry. The asymmetry extends to

[14]Either lowering the numerical threshold at which we treat the borrowing constraint as binding to a slightly negative number or not re-centering Tobin's Q at unity during the sample eliminate these false positives.

Tobin's q in panel (b). In contrast to the multiplier estimates, the estimates of Tobin's q have a larger left tail than right tail. This reflects the asymmetry of the model. Generally, the fixed point estimates are contained inside the 68% credible set, but the Gibbs sampler does not confirm the transient drop in Tobin's q in 2018 that the fixed point algorithm shows in connection with the false positive crisis event.

4.5.4 Quality check

To further assess the plausibility of our filtering application, we use the `Penn World Tables` to compare the implied price of capital based on our fixed point algorithm. To do that, we invert the first order condition to compute

$$q_t^{pwt} = \frac{\partial \tilde{i}_t}{\partial k_{t+1}} = 1 + a\left(\frac{k_{t+1} - k_t}{k_t}\right), \tag{4.12}$$

where we use the calibrated value of adjustment costs of $a = 2.75$.

Figure 14 shows the results. Qualitatively, the series agree on important movements. For example, both series agree on the two major drops, namely in 1982 and 1995, 2008, and at the end of the sample. There is disagreement in the late 1990s, when only the partial filter identifies a drop around the time of the emerging market financial crisis.

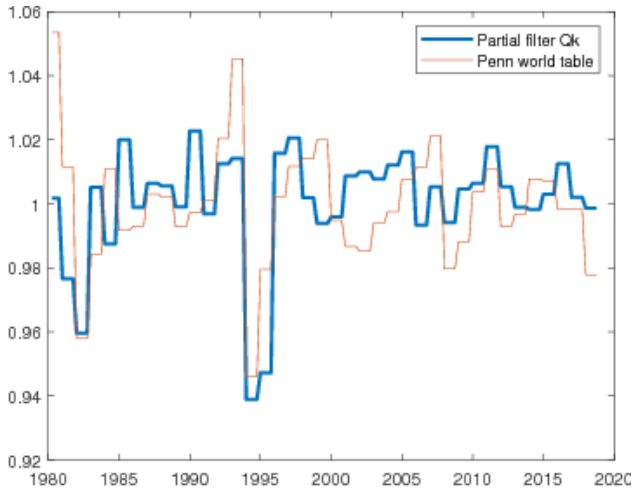


Figure 14: Comparison of Tobin's q in the partial filter (end of year) and the implied Penn World Table analogue.

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. M., J. Fernández-Villaverde, and J. F. Rubio-Ramírez (2017): "The Pruned State-Space System for Nonlinear DSGE Models: Theory and Empirical Applications," Review of Economic Studies, Forthcoming.

Bernanke, B. S., and K. N. Kuttner (2005): "What Explains the Stock Market's Reaction to Federal Reserve Policy?," The Journal of Finance, 60(3), 1221–1257.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," The Economic Journal, 101(405), 157–179.

Canova, F., and F. Ferroni (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., V. Cormun, P. D. Leo, P. Guerron-Quintana, and R. Valchev (2021): "Exchange Rate Disconnect Redux," Boston College Working Papers in Economics 1041, Boston College Department of Economics.

Chen, X., and T. M. Christensen (2015): "Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions," Journal of Econometrics, 188(2), 447–465, Heterogeneity in Panel Data and in Nonparametric Analysis in honor of Professor Cheng Hsiao.

Drautzburg, T., J. Fernández-Villaverde, and P. Guerrón-Quintana (2021): "Bargaining shocks and aggregate fluctuations," Discussion paper.

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., J. F. Rubio-Ramírez, and F. Schorfheide (2016): "Solution and Estimation Methods for {DSGE} Models," in Handbook of Macroeconomics, ed. by J. B. Taylor, and H. Uhlig, vol. 2, pp. 527–724. Elsevier.

Gabaix, X. (2020): "A Behavioral New Keynesian Model," American Economic Review, 110(8), 2271–2327.

Gallant, A. R., R. Giacomini, and G. Ragusa (2017): "Bayesian estimation of state space models using moment conditions," Journal of Econometrics, 201(2), 198 – 211, THEORETICAL AND FINANCIAL ECONOMETRICS: ESSAYS IN HONOR OF C. GOURIEROUX.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," Econometrica, 50(4), 1029–1054.

——— (2014): "Nobel Lecture: Uncertainty Outside and Inside Economic Models," Journal of Political Economy, 122(5), 945–987.

Mendoza, E., and S. Villalvazo (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.

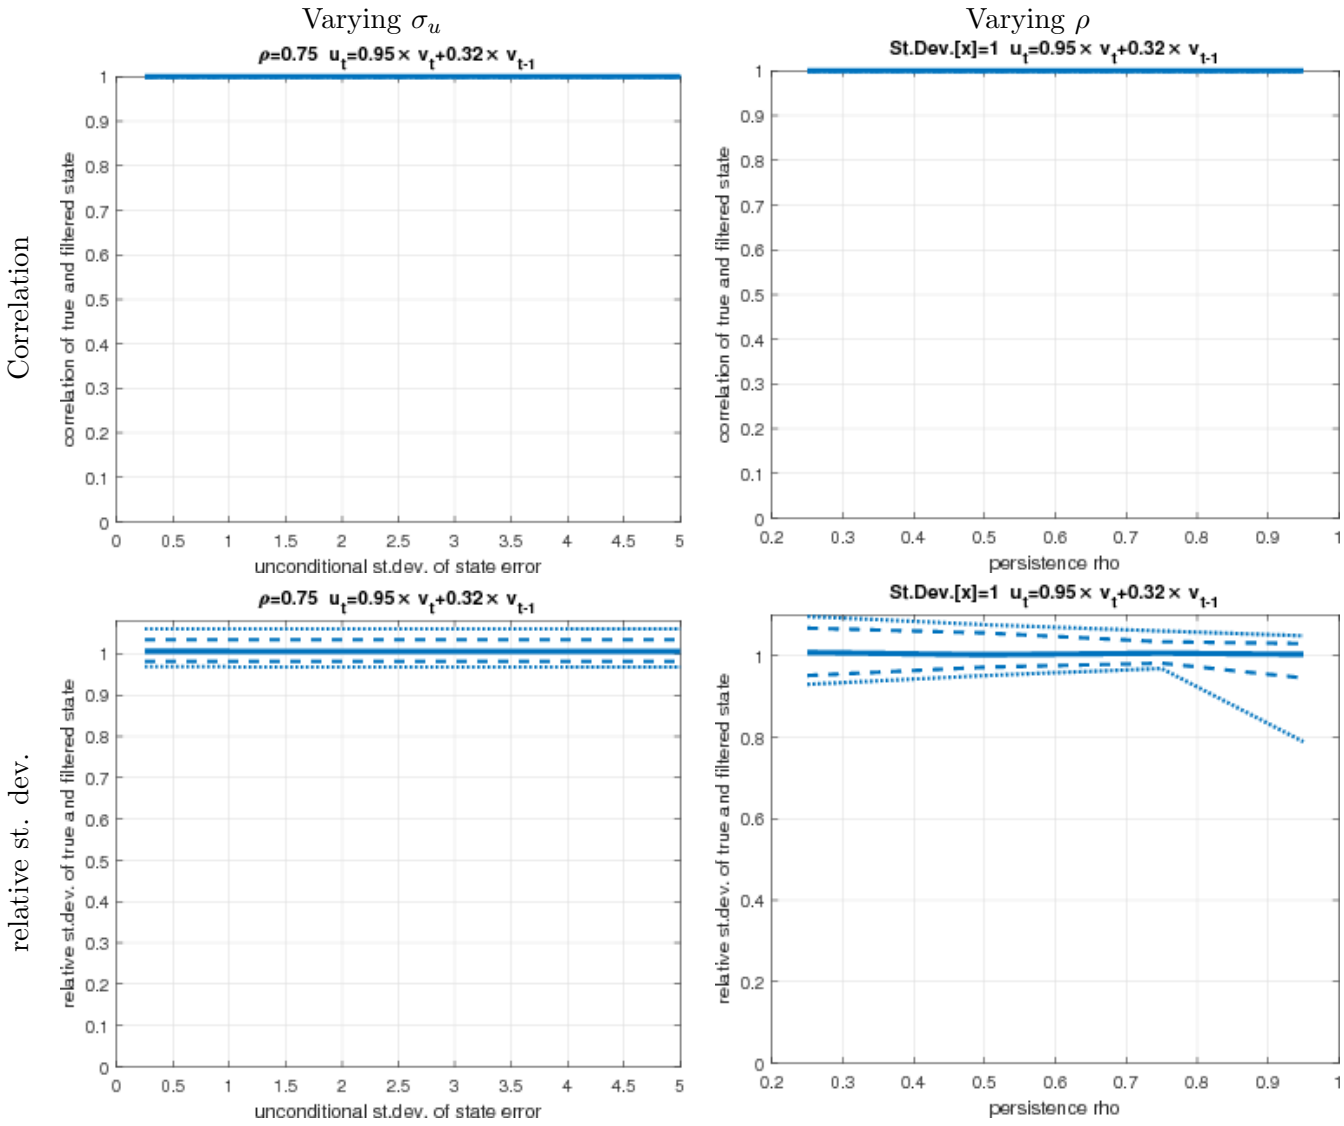
Petrova, K. (2019): "A quasi-Bayesian local likelihood approach to time varying parameter VAR models," Journal of Econometrics, 212(1), 286 – 306, Big Data in Dynamic Predictive Econometric Modeling.

Sbordone, A. M. (2002): "Prices and unit labor costs: a new test of price stickiness," Journal of Monetary Economics, 49(2), 265–292.

A Scalar model

A.1 No measurement error

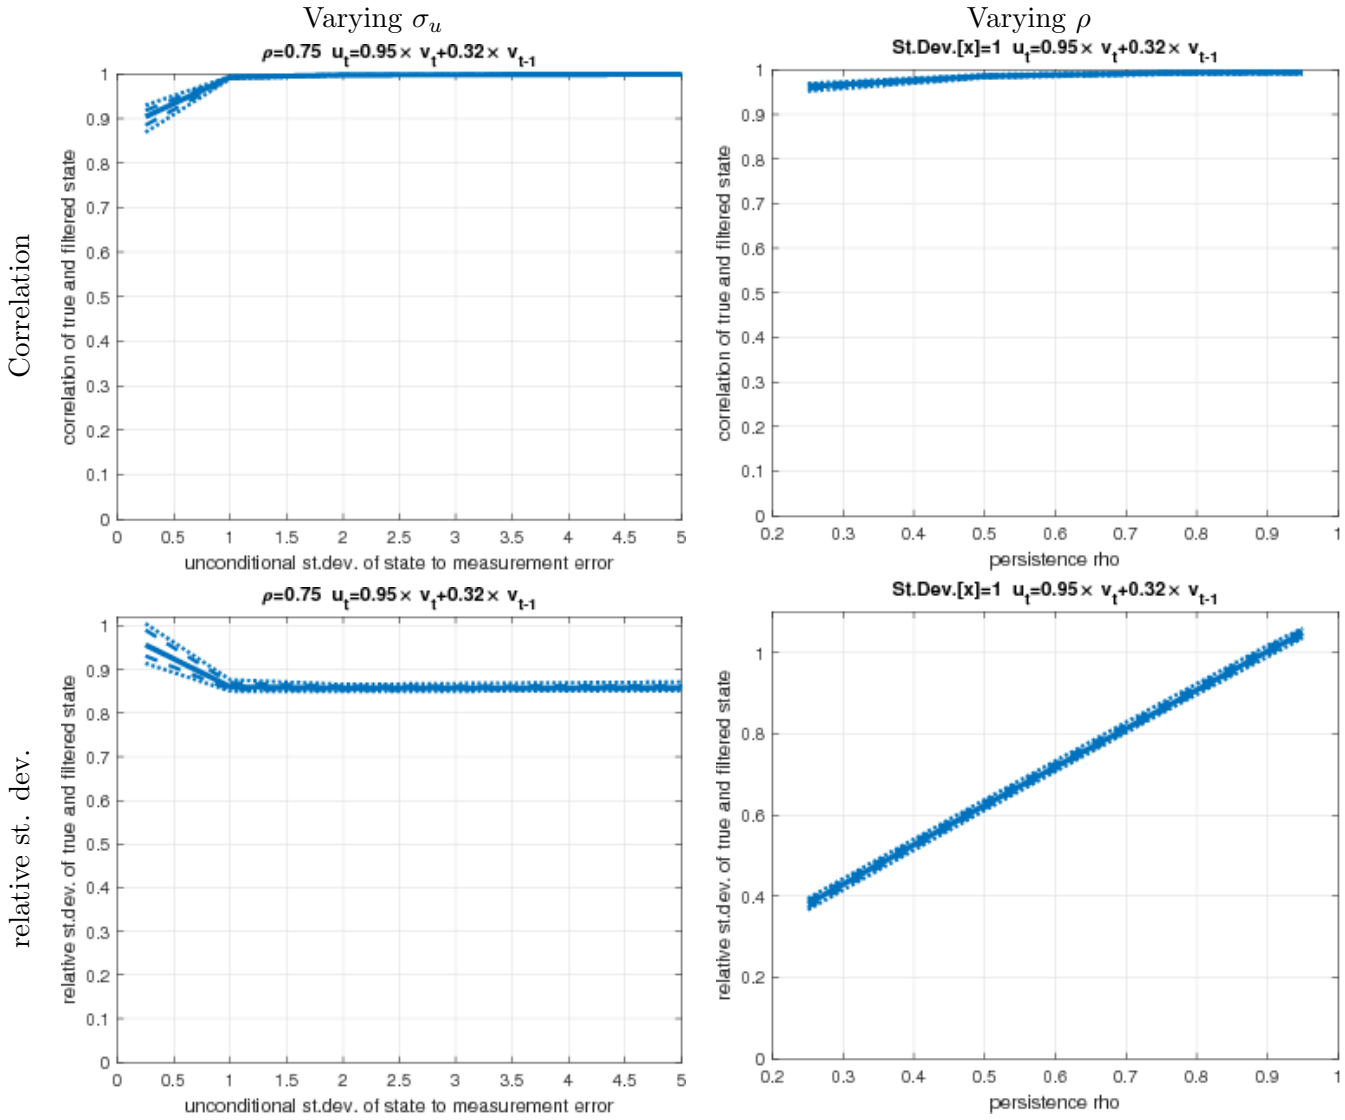
Here we provide additional correlation and standard deviation plots for the linear case without measurement error. In Figure A.1 we plot the median correlation between the true and filtered states (along with 68% and 90% confidence intervals) in the top row, and the relative standard deviations of the filtered state relative to the true state in the bottom row. 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.



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 $\hat{x}_t$ vs true x_t for varying degrees of persistence ρ_0.

A.2 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.3 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 $\hat{\beta} = \text{Var}_{t-1}[y_t]^{-1} \text{Cov}_{t-1}[x_t] = \frac{\text{Var}_{t-1}[x_t]}{\text{Var}_{t-1}[x_t]+\sigma_e^2}$. If σ_e^2 is large, then $\hat{\beta} \ll 1$.

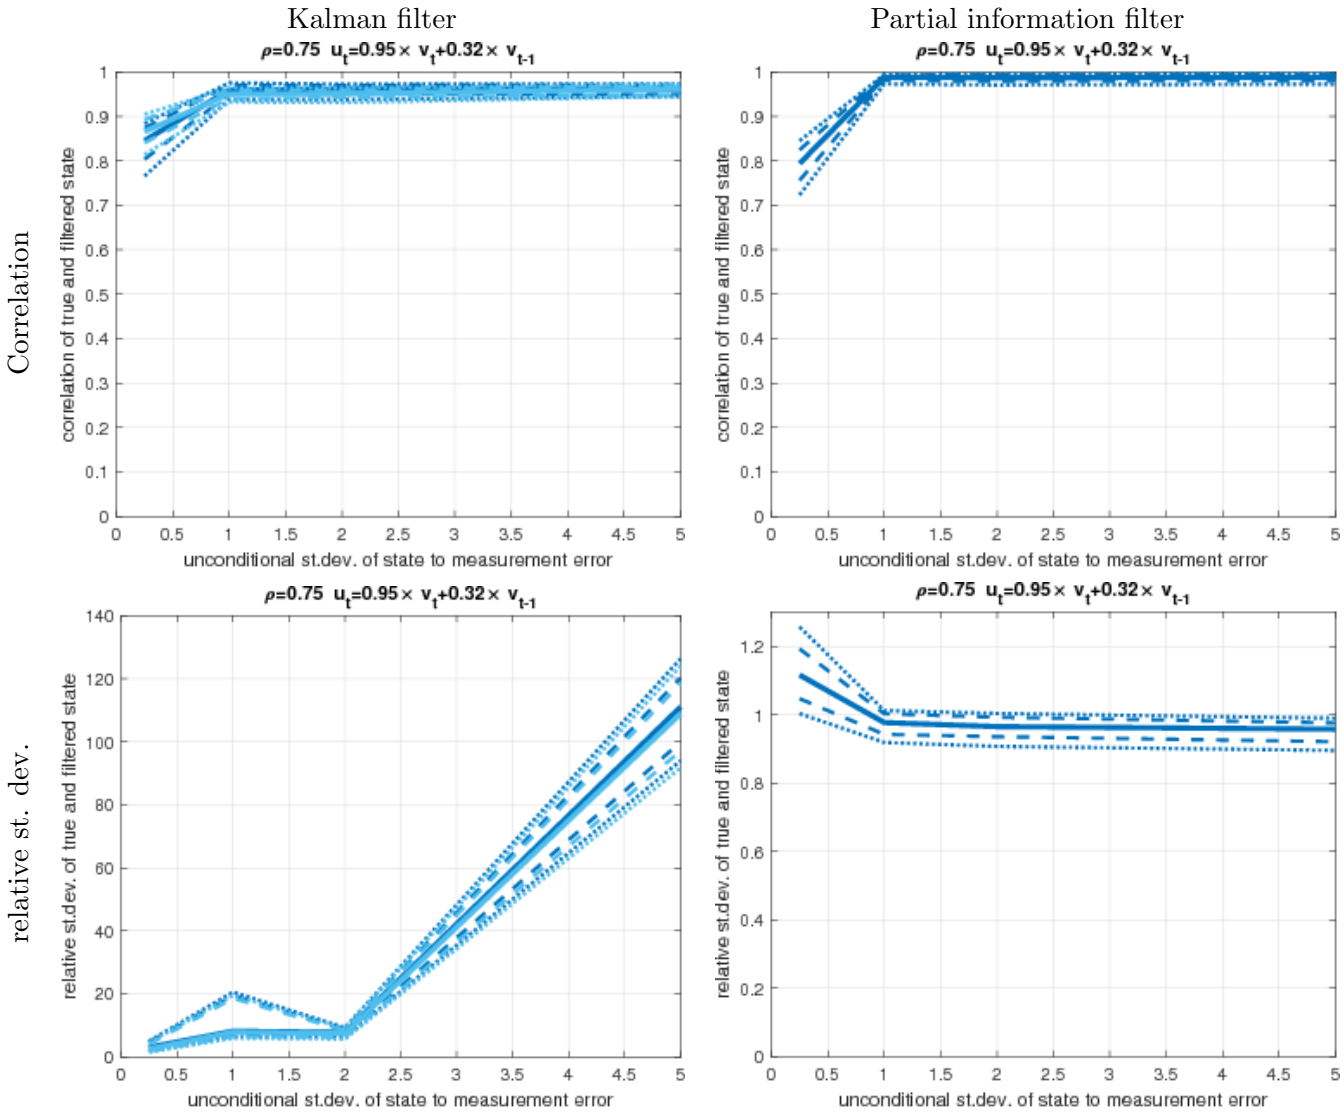
Moreover, the Kalman filter is a good benchmark for the performance of the partial filter. While it is not straightforward to apply in general non-linear models, the Kalman filter is the best linear filter. We focus on the two cases when we have non-linearities. Here, we apply the Kalman filter assuming 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 ($\kappa_1 = \frac{1}{10}, \kappa_2 = 10$).[15] 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.[16] 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.

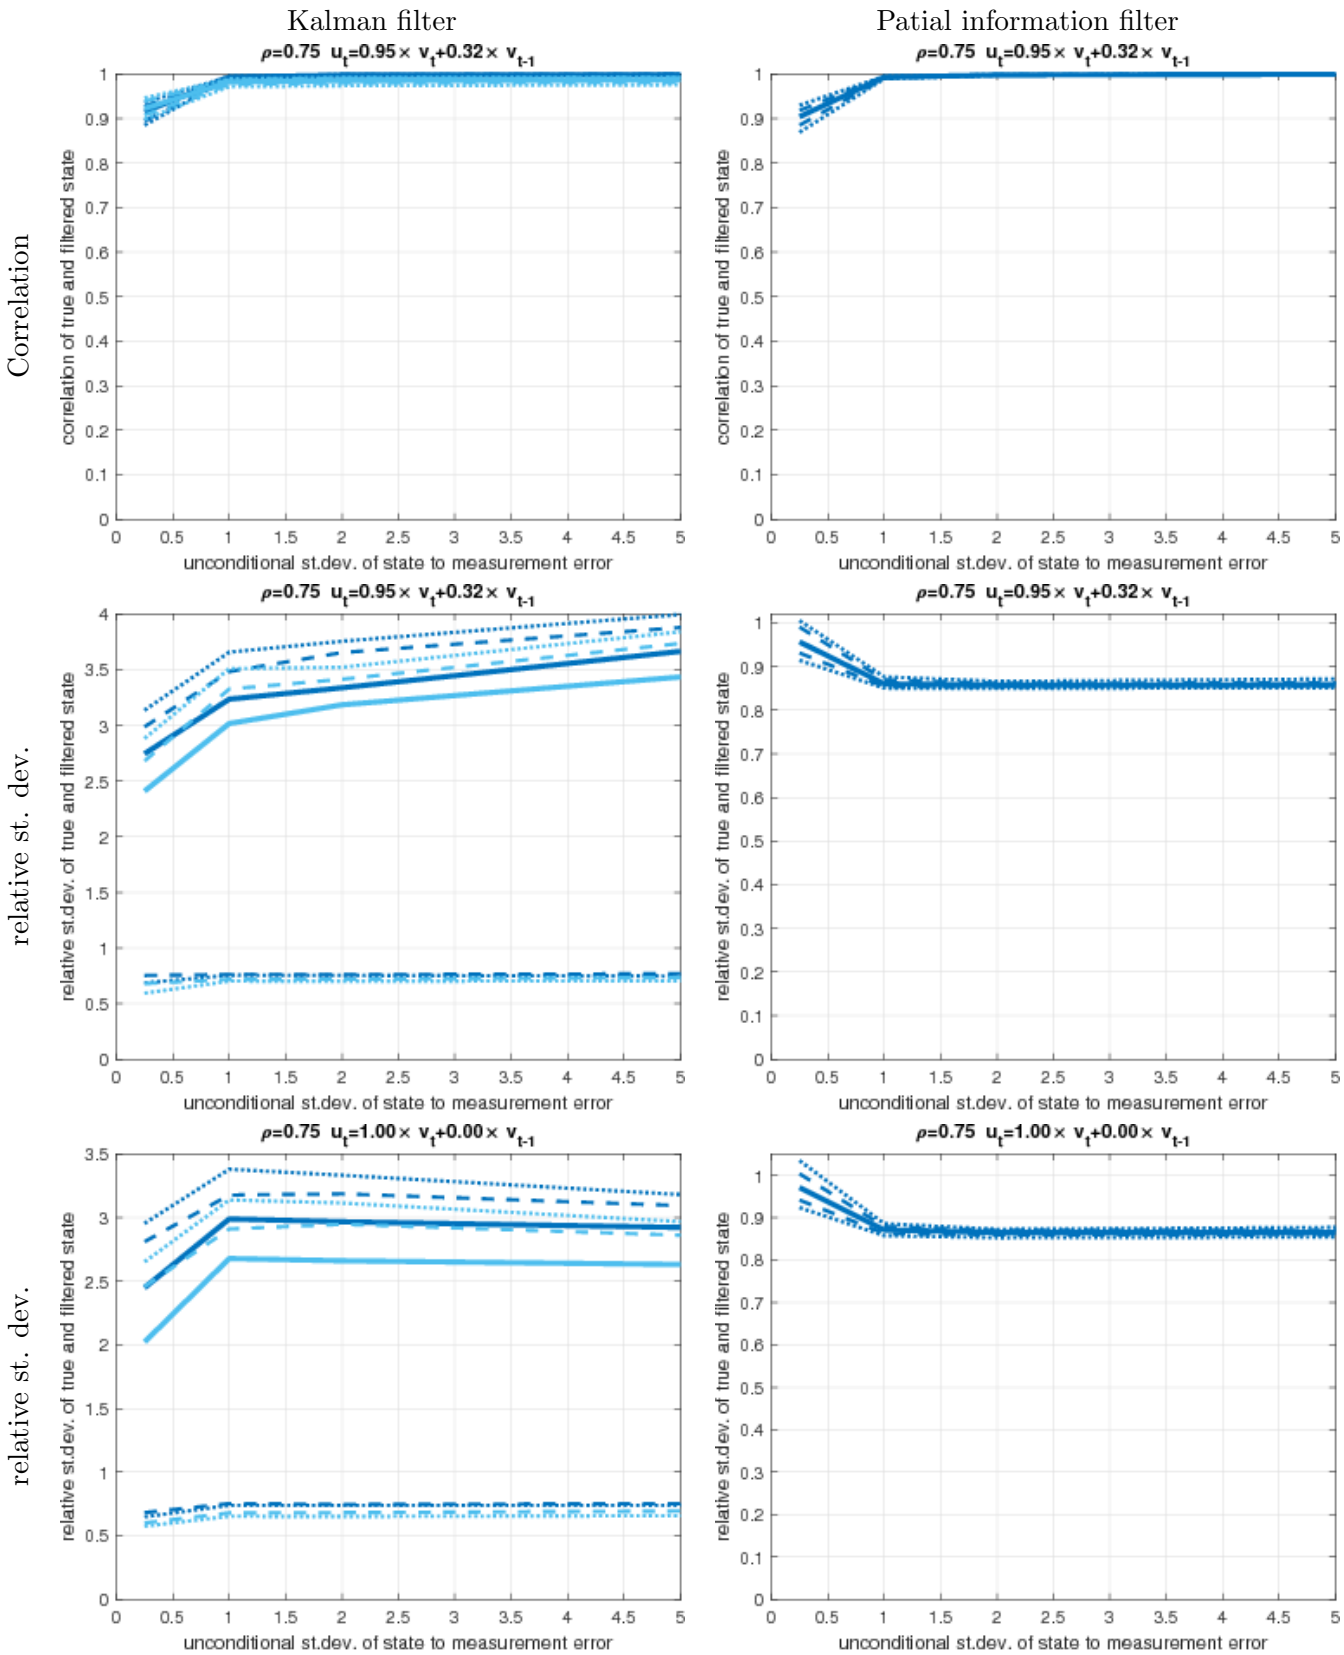
[15] We allow the Kalman filter to estimate κ_1 to allow it to fit the scale of x_t better. We also estimate the Kalman filter with the true κ_1 and choose the estimate with the higher likelihood for our comparison.

[16] 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.



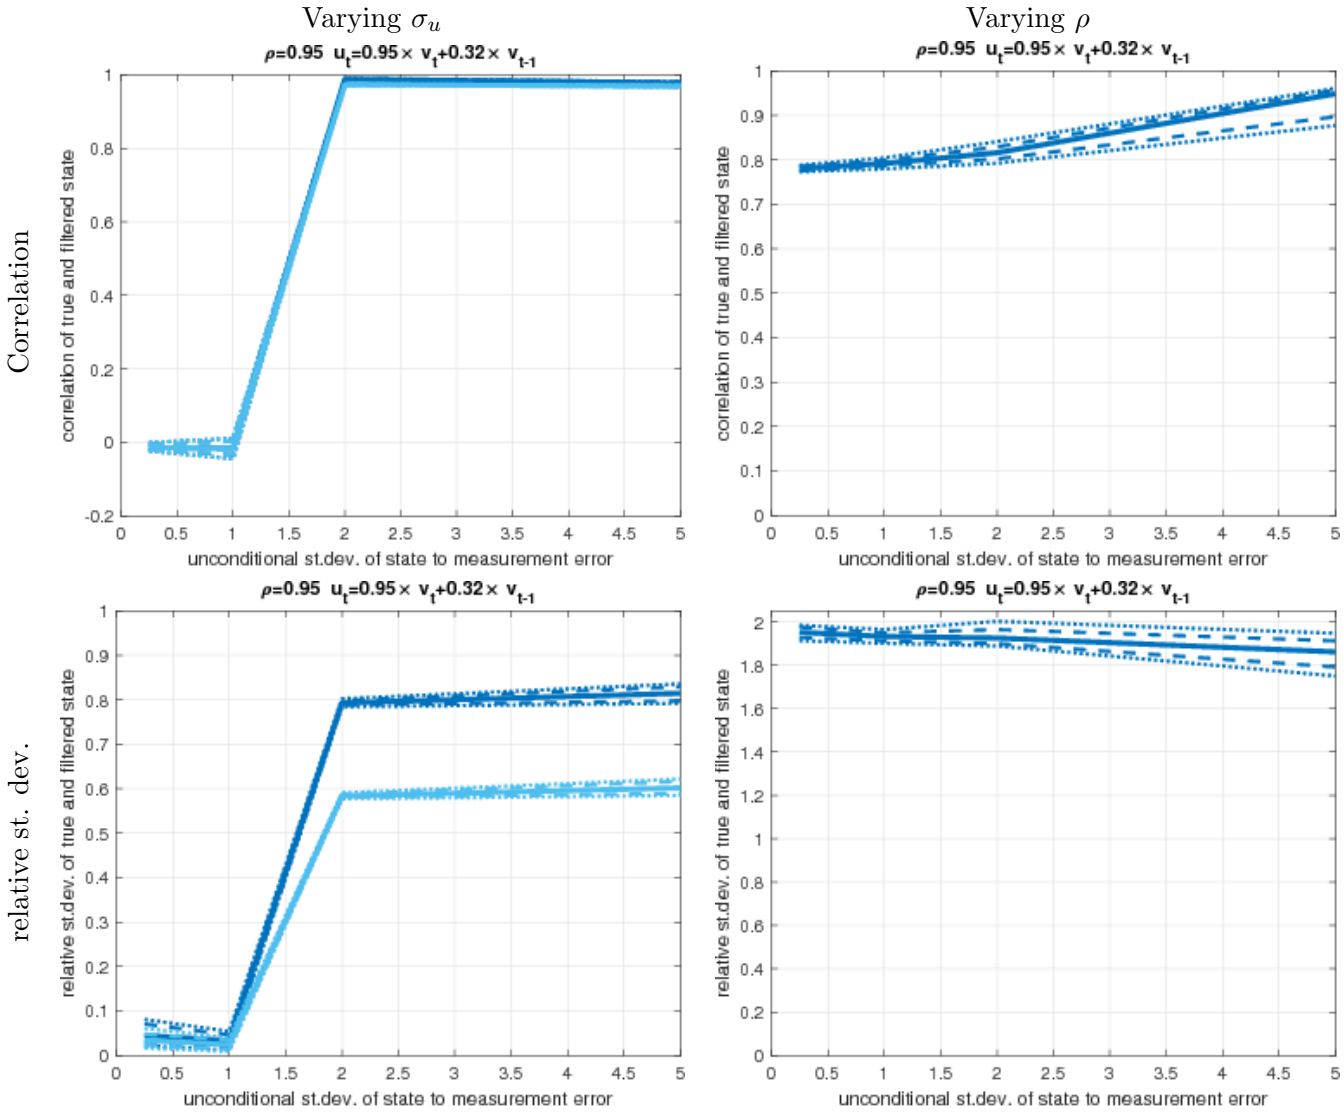
Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

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.



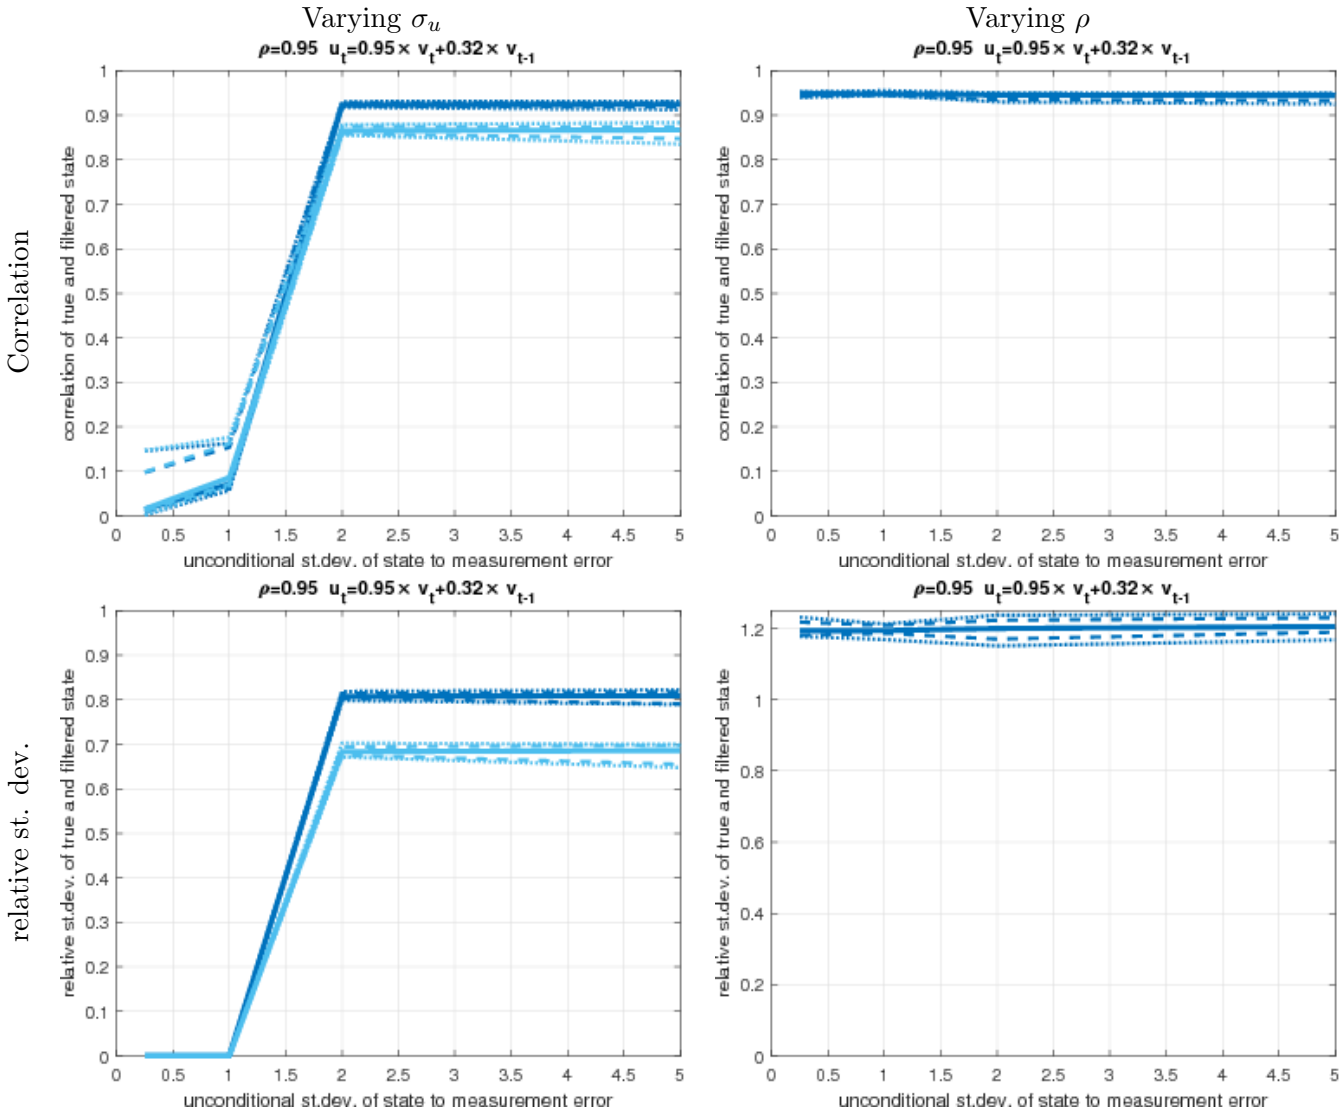
Shown are the 68% and 90% confidence intervals across Monte Carlo simulations, along with the median. Dark blue: Kalman filter. Light blue: Kalman smoother.

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.

B Sudden Stop Model Simulation Appendix

In this Appendix we describe the results for Section 4.4.2 without using the slackness condition.

(a) Correlations and relative standard deviations

Sudden stop probability	$\text{corr}(\hat{q}_t, q_t)$	$\text{corr}(\hat{\tilde{\mu}}_t, \tilde{\mu}_t)$	Rel. std. dev. q_t	Rel. std. dev. $\tilde{\mu}_t$
[0.1,0.15]	0.92 [0.86, 0.94]	0.14 [-0.02, 0.28]	.95 [.89, 1.01]	1.01 [0.69, 1.49]
[0.3, 0.35]	0.90 [0.85, 0.93]	0.23 [0.07, 0.43]	.99 [.92, 1.07]	0.50 [0.39, 0.80]
[0.5, 0.55]	0.86 [0.82, 0.91]	0.40 [0.22, 0.56]	1.04 [.94, 1.13]	0.45 [0.33, 0.59]

(b) Crisis classification errors

Sudden stop probability	Crisis classification	
	False positives	False negatives
[0.1,0.15]	0.70 [0.54, 0.79]	0.01 [0.00, 0.05]
[0.3, 0.35]	0.57 [0.50, 0.63]	0.02 [0.00, 0.08]
[0.5, 0.55]	0.42[0.36, 0.46]	0.02 [0.00, 0.06]

Table B.1: Partial filter performance across simulations as a function of the sample probability of sudden stops: Medians [68% CI] of Correlations, relative standard deviations, and classifications errors

Crisis probability

$\in [.10, .15]$ $\in [.30, .35]$ $\in [.50, .55]$

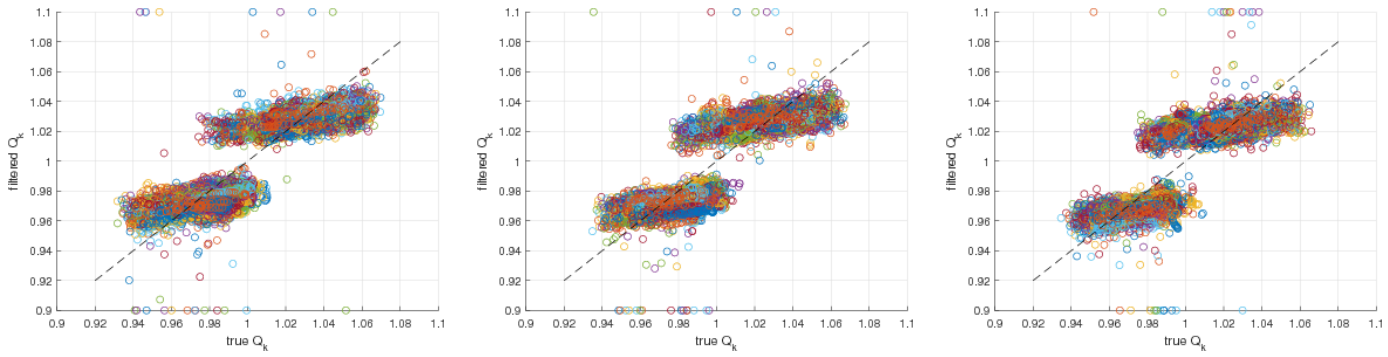


Figure B.1: Sample paths for the price of capital q_t for the sample paths with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis probability.

Crisis probability

$\in [.10, .15]$ $\in [.30, .35]$ $\in [.50, .55]$

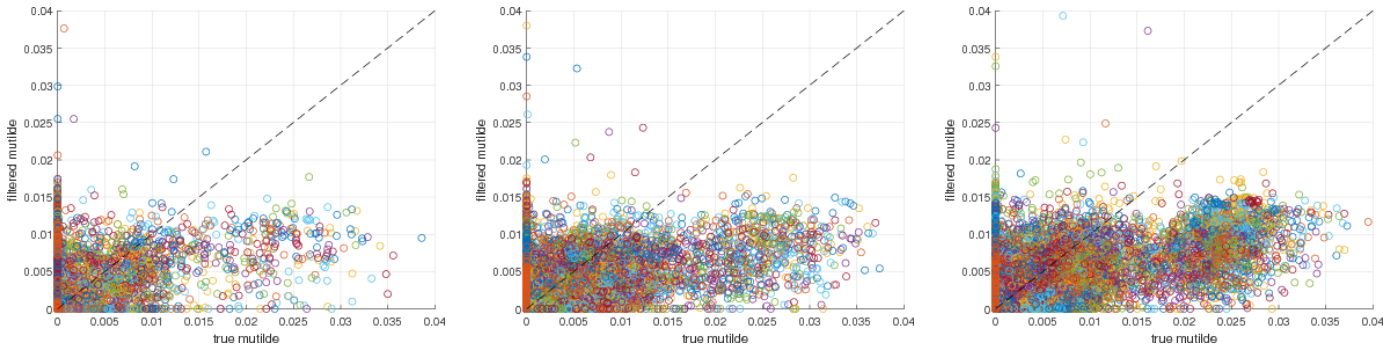


Figure B.2: Sample paths for the multiplier on the borrowing constraint $\tilde{\mu}_t$ for the sample paths with the median correlation (top panel) and scatter plot pooling all samples (bottom panel) as a function of the in-sample crisis probability.

C Sudden Stop Model Appendix

C.1 Data

When available, we use the LIBOR as our measure of the world interest rate (FRED mnemonic: USD3MTD156N). 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).

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 than overall hours, we divide the fluctuations by a factor of two.[17] Our measure for real consumption and real GDP are also taken from FRED (mnemonics NAEXKP02MXQ661S and NAEXKP01MXQ661S).

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 data from INEGI.[18] 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.[19] 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.

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

[17] 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 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 indicates a somewhat 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: $\frac{sd^q_{mfg,US}}{sd^a_{mfg,US}} sd_{mfg,MX} = \frac{2.51}{1.95} 1.52 = 1.96$.

[18] https://statistics.cepal.org/portal/cepalstat/dashboard.html?theme=2&lang=en.

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

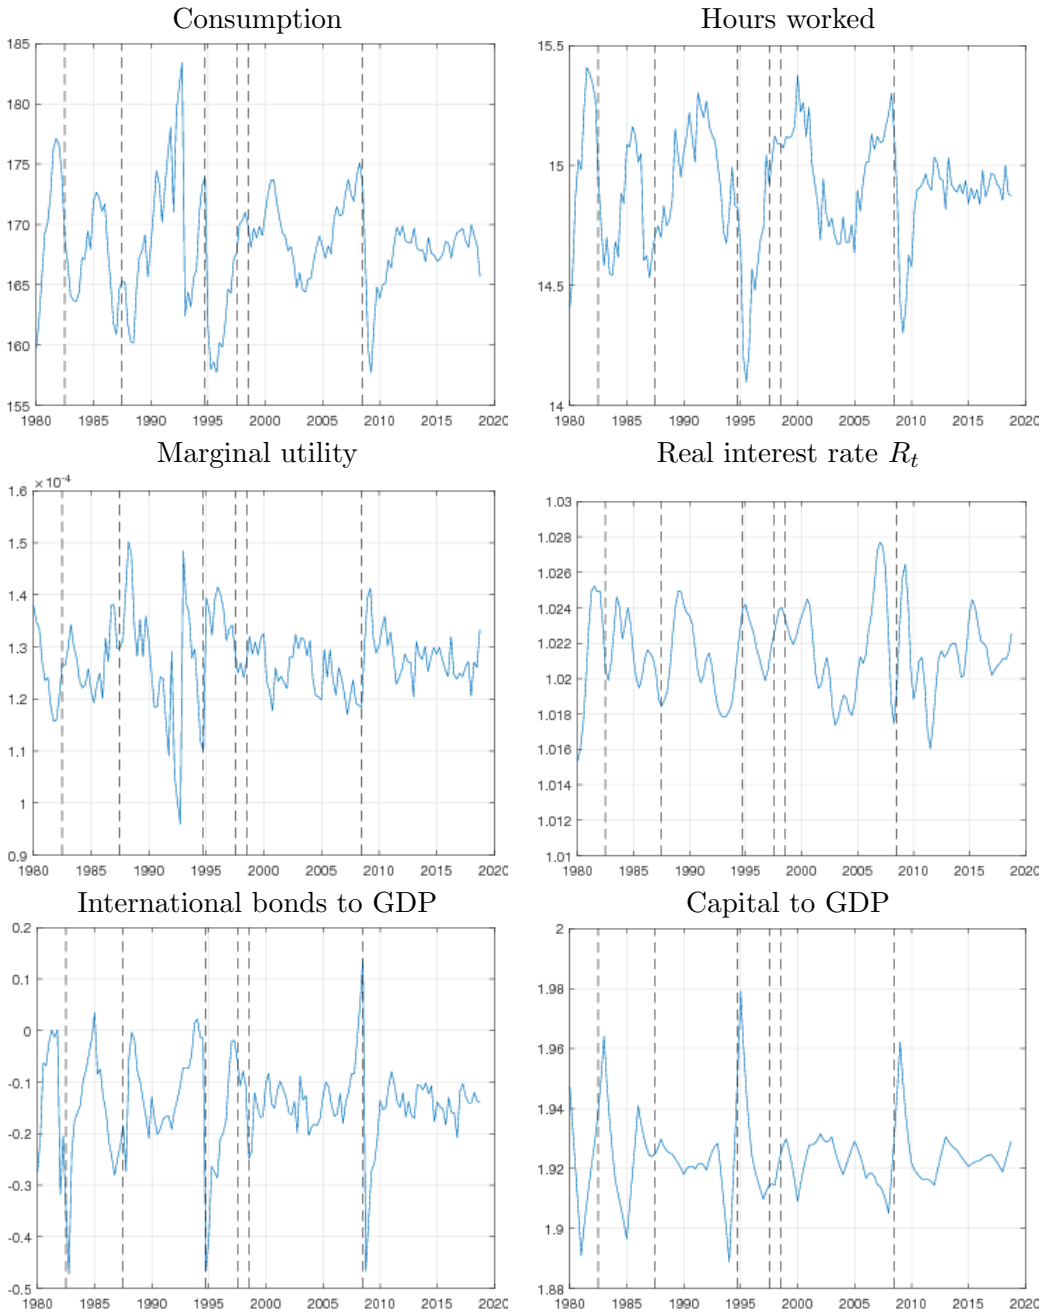


Figure C.3: Data input: Detrended data, re-centered at the theoretical means

C.2 First order conditions

Defining λ_t and μ_t as the multipliers of the budget and collateral constraints respectively, the model's equilibrium conditions in sequential form are:

$$\left(c_t - \frac{L_t^{\omega}}{\omega}\right)^{-\sigma} = \lambda_t(1+\tau) \tag{C.1}$$

$$A_t F_{L_t}(k_t, L_t, v_t) = w_t\left(1 + \phi(R_t - 1) + \frac{\mu_t}{\lambda_t}\phi R_t\right) \tag{C.2}$$

$$A_t F_{v_t}(k_t, L_t, v_t) = p_t\left(1 + \phi(R_t - 1) + \frac{\mu_t}{\lambda_t}\phi R_t\right) \tag{C.3}$$

$$\lambda_t = \frac{1}{q_t^b}\beta\mathbb{E}_t[\lambda_{t+1}] + \mu_t \tag{C.4}$$

$$\lambda_t = \frac{1}{q_t}\beta\mathbb{E}_t[\lambda_{t+1}(d_{t+1} + q_{t+1})] + \mu_t\kappa \tag{C.5}$$

$$d_t \equiv \exp(\epsilon_t^A)F_{k_t} - \delta + \frac{a}{2}\frac{(k_{t+1} - k_t)^2}{k_t^2} \tag{C.6}$$

$$q_t = \frac{\partial\tilde{i}_t}{\partial k_{t+1}} = 1 + a\left(\frac{k_{t+1} - k_t}{k_t}\right) \tag{C.7}$$

$$w_t = L_t^{\omega-1}(1+\tau) \tag{C.8}$$

$$\begin{aligned} c_t(1+\tau) + k_{t+1} - (1-\delta)k_t + \frac{a}{2}\frac{(k_{t+1} - k_t)^2}{k_t} = {} & A_t F(k_t, L_t, v_t) - p_t v_t \\ & - \phi(R_t - 1)\left(L_t^{\omega}(1+\tau) + p_t v_t\right) - q_t^b b_{t+1} + b_t \end{aligned} \tag{C.9}$$

$$\mu_t\left[q_t^b b_{t+1} - \phi R_t(w_t L_t + p_t v_t) + \kappa q_t k_{t+1}\right] = 0. \tag{C.10}$$

C.3 Recursive representation and solution method

To implement the solution method, we first re-write the equilibrium conditions in recursive form. The model has two endogenous states, b and k, and up to six exogenous states. The exogenous states consist of the three exogenous shocks, using s to denote the triple of exogenous shocks $s \equiv (A, R, p)$, which includes the shocks to TFP (A), the world interest rate (R) and the price of imported inputs (p). The recursive equilibrium is defined by a set of equilibrium recursive functions for allocations $[b'(b,k,s), k'(b,k,s), c(b,k,s), L(b,k,s), v(b,k,s)]$, prices $[w(b,k,s), q(b,k,s), d(b,k,s)]$ and multipliers $[\lambda(b,k,s), \mu(b,k,s)]$ that satisfy the following recursive representation of the equilibrium conditions:

$$\left(c(b,k,s) - \frac{L(b,k,s)^{\omega}}{\omega}\right)^{-\sigma} = \lambda(b,k,s)(1+\tau) \tag{C.11}$$

$$\alpha A k^{\gamma} L(b,k,s)^{\alpha-1} v(b,k,s)^{\eta} = w(b,k,s)\left(1 + \phi(R-1) + \frac{\mu(b,k,s)}{\lambda(b,k,s)}\phi R\right) \tag{C.12}$$

$$\eta A k^{\gamma} L(b,k,s)^{\alpha} v(b,k,s)^{\eta-1} = p\left(1+\phi(R-1)+\frac{\mu(b,k,s)}{\lambda(b,k,s)}\phi R\right) \tag{C.13}$$

$$\lambda(b,k,s) = R\beta E[\lambda(b'(b,k,s),k'(b,k,s),s')] + \mu(b,k,s) \tag{C.14}$$

$$\begin{aligned}\lambda(b,k,s) = \frac{1}{q(b,k,s)}\beta E\big[\lambda(b'(b,k,s),k'(b,k,s),s')(d(b'(b,k,s),k'(b,k,s),s') \\ + q(b'(b,k,s),k'(b,k,s),s'))\big] + \mu(b,k,s)\kappa\end{aligned} \tag{C.15}$$

$$d(b,k,s) = \gamma A k^{\gamma-1} L(b,k,s)^{\alpha} v(b,k,s)^{\eta} - \delta + \frac{a}{2}\frac{(k'(b,k,s)-k)^2}{k^2} \tag{C.16}$$

$$q(b,k,s) = 1 + a\left(\frac{k'(b,k,s)-k}{k}\right) \tag{C.17}$$

$$w(b,k,s) = L(b,k,s)^{\omega-1}(1+\tau) \tag{C.18}$$

$$\begin{aligned}c(b,k,s)(1+\tau) + k'(b,k,s) - (1-\delta)k + \frac{a}{2}\frac{(k'(b,k,s)-k)^2}{k} = Ak^{\gamma}L(b,k,s)^{\alpha}v(b,k,s)^{\eta} \\ -pv(b,k,s) - \phi(R-1)\left(L(b,k,s)^{\omega}(1+\tau) + pv(b,k,s)\right) - R^{-1}b'(b,k,s) + b\end{aligned} \tag{C.19}$$

$$\mu(b,k,s)\left[q^b b'(b,k,s) - \phi R\left[w(b,k,s)L(b,k,s) + pv(b,k,s)\right] + \kappa q(b,k,s)k'(b,k,s)\right] = 0. \tag{C.20}$$

C.4 Details of the implementation

We use the borrowing constraint as follows:

- If the borrowing constraint is binding (the term in brackets is negative), but $\tilde{\mu}$ is not positive, we shift future expected utility $u_{c,t+1}$ down by $1e-8$. This raises the current marginal utility high relative to the future, and thus increases the estimate for $\tilde{\mu}_t$ from equation (4.9).
- If the borrowing constraint hardly ever binds and $\tilde{\mu}_t$ is strictly positive, then we shift future expected utility up by $1e-8$. This lowers current marginal utility relative to the future, and thus decreases the estimate for $\tilde{\mu}_t$ from equation (4.9). We treat the borrowing constraint as hardly ever binding if the term in brackets has a value larger than 10 percentage points above the expected failure rate for the selected sample (so 0.2 for p10, 0.4 for p30, 0.6 for p50)

C.5 Results without using slackness condition

In this Appendix, we provide the results of the partial filter without using the slackness condition.

C.5.1 Simulation results

(a) Price of capital q_t

Sudden stop probability	corr$(\hat{q}_t, q_t)$ Baseline	corr$(\hat{q}_t, q_t)$ w/o slackness	Rel. std. dev. q_t Baseline	Rel. std. dev. q_t w/o slackness
[0.10, 0.15]	0.92 [0.87, 0.95]	0.92 [0.86, 0.94]	0.93 [0.86, 1.00]	0.95 [0.89, 1.01]
[0.30, 0.35]	0.90 [0.85, 0.93]	0.90 [0.85, 0.93]	0.98 [0.91, 1.08]	0.99 [0.92, 1.07]
[0.50, 0.55]	0.86 [0.81, 0.91]	0.86 [0.82, 0.91]	1.06 [0.95, 1.14]	1.04 [0.94, 1.13]

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

Sudden stop probability	corr$(\hat{\tilde{\mu}}_t, \tilde{\mu}_t)$ Baseline	corr$(\hat{\tilde{\mu}}_t, \tilde{\mu}_t)$ w/o slackness	Rel. std. dev. $\tilde{\mu}_t$ Baseline	Rel. std. dev. $\tilde{\mu}_t$ w/o slackness
[0.10, 0.15]	0.35 [0.17, 0.56]	0.14 [-0.02, 0.28]	0.61 [0.39, 1.08]	1.01 [0.69, 1.49]
[0.30, 0.35]	0.30 [0.14, 0.40]	0.23 [0.07, 0.43]	0.50 [0.37, 0.71]	0.50 [0.39, 0.80]
[0.50, 0.55]	0.28 [0.11, 0.42]	0.40 [0.22, 0.55]	0.51 [0.39, 0.66]	0.44 [0.33, 0.59]

Table C.2: Partial filter performance across simulations as a function of the sample probability of sudden stops: Medians [68% CI] of Correlations, relative standard deviations for the partial filter with and without the slackness condition.

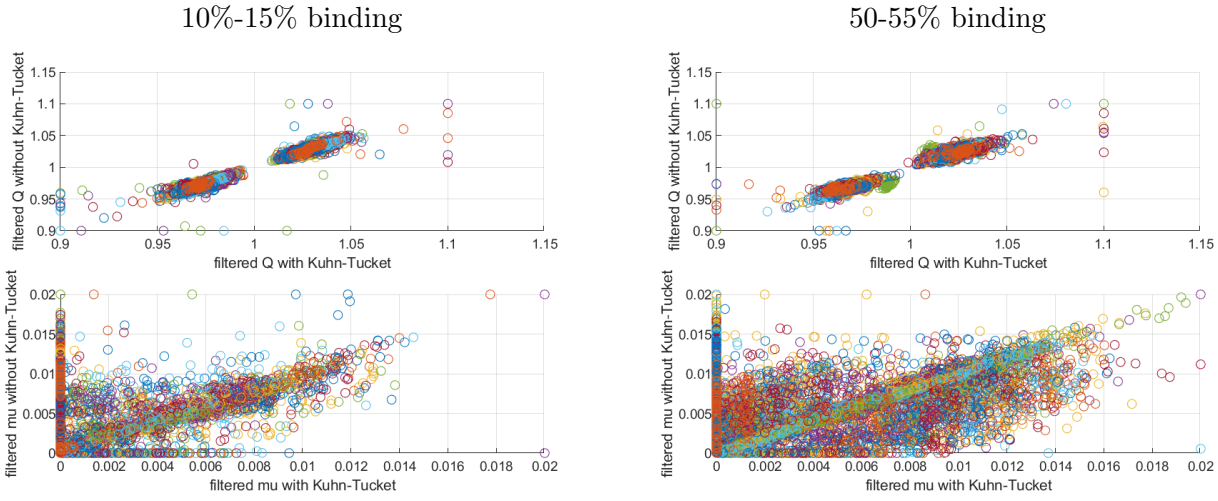


Figure C.4: Filtered value of the borrowing constraint multiplier and Tobin's Q in the Mendoza (2010) model applied to Mexican data.

C.5.2 Empirical results

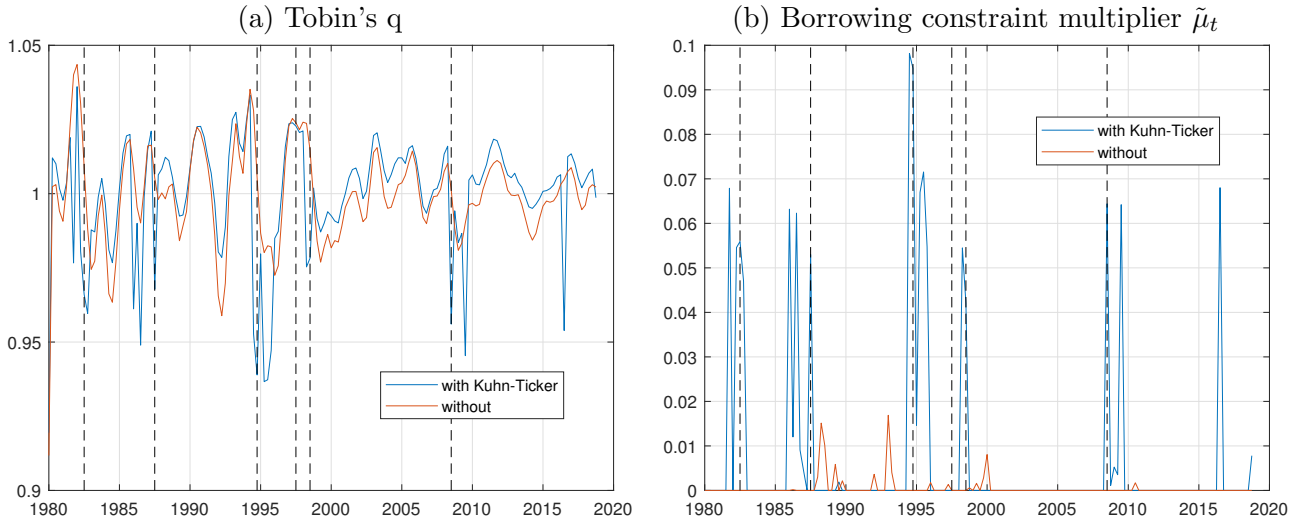


Figure C.5: Filtered value of the borrowing constraint multiplier and Tobin's Q in the Mendoza (2010) model applied to Mexican data.