

Lange, Rutger Jan

Working Paper

Bellman filtering for state-space models

Tinbergen Institute Discussion Paper, No. TI 2020-052/III

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Lange, Rutger Jan (2021) : Bellman filtering for state-space models, Tinbergen Institute Discussion Paper, No. TI 2020-052/III, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/229672>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

TI 2020-052/III
Tinbergen Institute Discussion Paper

Bellman filtering for state-space models

Revision: 8 January 2021

Rutger-Jan Lange¹

¹ Econometric Institute, Erasmus School of Economics, Rotterdam, Netherlands

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Bellman filtering for state-space models

Rutger-Jan Lange*
Econometric Institute, Erasmus School of Economics
Rotterdam, Netherlands

January 8, 2021

Abstract

This article presents a filter for state-space models based on Bellman’s dynamic programming principle applied to the mode estimator. The proposed Bellman filter (BF) generalises the Kalman filter (KF) including its extended and iterated versions, while remaining equally inexpensive computationally. The BF is also (unlike the KF) robust under heavy-tailed observation noise and applicable to a wider range of (nonlinear and non-Gaussian) models, involving e.g. count, intensity, duration, volatility and dependence. (Hyper)parameters are estimated by numerically maximising a BF-implied log-likelihood decomposition, which is an alternative to the classic prediction-error decomposition for linear Gaussian models. Simulation studies reveal that the BF performs on par with (or even outperforms) state-of-the-art importance-sampling techniques, while requiring a fraction of the computational cost, being straightforward to implement and offering full scalability to higher dimensional state spaces.

JEL Classification Codes: C32, C53, C61

Keywords: dynamic programming, curse of dimensionality, importance sampling, Laplace method, Kalman filter, maximum a posteriori (MAP) estimate, NAIS, particle filter, prediction-error decomposition, posterior mode, Viterbi algorithm

1 Introduction

1.1 State-space models

State-space models allow observations to be affected by a hidden state that changes stochastically over time. For discrete times $t = 1, 2, \dots, n$, the observation $\mathbf{y}_t \in \mathbb{R}^l$ is drawn from a conditional distribution, $p(\mathbf{y}_t|\boldsymbol{\alpha}_t)$, while the latent state $\boldsymbol{\alpha}_t \in \mathbb{R}^m$ follows a first-order Markov process with a known state-transition density, $p(\boldsymbol{\alpha}_t|\boldsymbol{\alpha}_{t-1})$, and some given initial condition, $p(\boldsymbol{\alpha}_1)$, i.e.

$$\mathbf{y}_t \sim p(\mathbf{y}_t|\boldsymbol{\alpha}_t), \quad \boldsymbol{\alpha}_t \sim p(\boldsymbol{\alpha}_t|\boldsymbol{\alpha}_{t-1}), \quad \boldsymbol{\alpha}_1 \sim p(\boldsymbol{\alpha}_1). \quad (1)$$

Here $p(\cdot|\cdot)$ and $p(\cdot)$ denote *generic* conditional and marginal densities; i.e. any two p ’s need not denote the same probability density function (e.g. Polson et al., 2008, p. 415). For a given model, the functional form of all p ’s is considered known. These densities may further depend on a fixed (hyper)parameter $\boldsymbol{\psi}$, which for notational simplicity is suppressed. Both the observation and state-transition densities may

*I thank Maksim Anisimov, Francisco Blasques, Dick van Dijk, Dennis Fok, Maria Grith, Andrew Harvey, Siem Jan Koopman, Erik Kole, Rutger Lit, Rasmus Lonn, André Lucas, Robin Lumsdaine, Andrea Naghi, Jochem Oorschot, Richard Paap, Andreas Pick, Krzysztof Postek, Rogier Quaadvlieg, Daniel Ralph, Marcel Scharth, Annika Schnücker, Stephen Thiele, Nando Vermeer, Sebastiaan Vermeulen, Michel van der Wel, Martina Zaharieva, Mikhail Zhelonkin and Chen Zhou for their valuable comments.

be non-Gaussian and involve nonlinearity. Sequences of observations are denoted $\mathbf{y}_{1:t} := (\mathbf{y}_1, \dots, \mathbf{y}_t) \in \mathbb{R}^{l \times t}$; likewise for the states, i.e. $\boldsymbol{\alpha}_{1:t} := (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_t) \in \mathbb{R}^{m \times t}$. In this article, the states are assumed to take *continuous* values in \mathbb{R}^m ; hence, the state space can be viewed as being ‘infinite dimensional’ even as m remains finite. This is in contrast with Markov-switching models (also known as hidden Markov models, see e.g. Künsch, 2001, p. 109 and Fuh, 2006, p. 2026), in which the state takes a finite number of (discrete) values.¹

Myriad examples of model (1) can be found in engineering, biology, geological physics, economics and mathematical finance (for a comprehensive overview, see Künsch, 2001). Examples in (financial) statistics with continuous state spaces include models for count data (Singh and Roberts, 1992; Frühwirth-Schnatter and Wagner, 2006), intensity (Bauwens and Hautsch, 2006), duration (Bauwens and Veredas, 2004), volatility (Tauchen and Pitts, 1983; Harvey et al., 1994; Ghysels et al., 1996; Jacquier et al., 2002; Taylor, 2008) and dependence structure (Hafner and Manner, 2012).

Model (1) presents researchers and practitioners with two important problems, known as the filtering and estimation problems. The filtering problem concerns the estimation of the latent states $\boldsymbol{\alpha}_{1:t}$ conditional on $\mathbf{y}_{1:t}$, while the constant (hyper)parameter $\boldsymbol{\psi}$ is considered known. The estimation problem entails determining the parameter $\boldsymbol{\psi}$, where both this parameter and the latent states $\boldsymbol{\alpha}_{1:t}$ are assumed to be unknown. For models with continuously varying latent states, the filtering problem can be solved in closed form only when model (1) is linear and Gaussian; Kalman’s (1960) filter then recursively computes the expectation of the state (i.e. the mean) and the most likely state (i.e. the mode), which are identical for these models. See Table 1 for a concise (by no means exhaustive) overview of well-known filtering methods.

1.2 Mode estimation

For the majority of state-space models, no exact filtering methods are available. Further, the mean is typically distinct from the mode. Most approximate filters shown in Table 1 are based on the mean (see section 1.5 for a brief literature review). The proposed Bellman filter, also included in Table 1, is based on the mode. The mode is also known as the *maximum a posteriori* (MAP, e.g. Godsill et al., 2001; Sardy and Tseng, 2004, p. 191) estimate or the *posterior mode*² (e.g. Fahrmeir and Kaufmann, 1991; Fahrmeir, 1992; Shephard and Pitt, 1997; Durbin and Koopman, 1997; So, 2003; Jungbacker and Koopman, 2007). Use of the mode is appealing due to its ‘optimality property analogous to that of maximum likelihood estimates of fixed parameters in finite samples’ (Durbin and Koopman, 2012, pp. 252-3). It is also the natural choice when considering zero-one loss functions, as in e.g. target-tracking applications (Godsill et al., 2001, 2007).

The standard method for computing the mode (i.e. by using numerical optimisation procedures) is computationally cumbersome, as it involves re-estimating the entire sequence of states for each time step. Computing times per time step typically scale as $O(t^3)$, implying a cumulative computing effort, up to time t , of $O(t^4)$. Moreover, the mode estimator per se does not address the estimation problem, making its application infeasible in practice unless supplemented with other methods. These drawbacks may explain why the mode estimator has to date received little attention as a ‘filtering’ technique for state-space models.

1.3 Filtering using dynamic programming

In this article, we circumvent both drawbacks of the mode estimator, yielding an algorithm that is both fast and feasible, while performing as well as (more computationally intensive) importance-sampling

¹Formulation (1) can accommodate Markov-switching models if $p(\boldsymbol{\alpha}_1)$ and $p(\boldsymbol{\alpha}_t | \boldsymbol{\alpha}_{t-1})$ are interpreted as probabilities rather than probability density functions (although this is beyond the scope of the present paper).

²The label ‘posterior’ does not reflect a Bayesian approach; it indicates only that the mode is computed after the data are received.

Table 1: (Non-exhaustive) overview of filtering methods.

	Discrete states	Continuously varying states	
	Exact filters	Linear & Gaussian model	Nonlinear and/or non-Gaussian model
		Exact filters	Approximate filters
Mean	Hamilton (1989)	Kalman (1960)	Iterated extended KF (e.g. Anderson and Moore, 2012) Unscented KF (Julier and Uhlmann, 1997) Masreliez (1975) filter Laplace Gaussian filter (Koyama et al., 2010)
Mode	Viterbi (1967)	Kalman (1960)	Bellman filter (BF, this article) Special case of BF: Fahrmeir’s (1992) mode estimator

Note: For brevity, the table excludes simulation-based approaches. KF= Kalman filter. BF = Bellman filter.

methods. For simplicity, we assume the mode exists (and is unique), and we take its optimality properties for granted. To address the first drawback of the mode estimator — computational complexity — we employ Bellman’s (1957) principle of dynamic programming, thus avoiding the need to solve optimisation problems involving an ever-increasing number of states. Instead, Bellman’s equation involves the maximisation over a single state vector of length m for each time step. The required computing cost per time step now remains constant over time, reducing the cumulative computational burden from $O(t^4)$ to linear in time.

If the state takes a finite number of (discrete) values, Bellman’s equation can be solved exactly for all time steps, yielding Viterbi’s (1967) algorithm. This algorithm can be used to track, in real time, the most likely state of a finite set of states (see Table 1) and has proven successful in digital signal processing applications (Forney, 1973; Viterbi, 2006). Exact solubility of Bellman’s equation is lost, however, when the states take continuous values in \mathbb{R}^m , as in this article. To address this problem, several authors have considered discretising the state space such that Viterbi’s algorithm can be used (e.g. Künsch, 2001, p. 125, Godsill et al., 2001). This approach can lead to prohibitive computational requirements (Künsch, 2001, p. 125): the discretised solution of Bellman’s equation requires the computation and storage of N^m values for each time step, where N is the number of grid points in each of m spatial directions (e.g. billions of values if $N = 100$ and $m = 5$). This is known as the ‘curse of dimensionality’ (Bellman, 1957, p. ix). Another difficulty, more so in statistics than in engineering, is that Viterbi’s algorithm leaves the estimation problem unaddressed.

When the latent states take values in a continuum, the solution to Bellman’s equation is a function, known as the *value function*, which maps the (continuous) state space \mathbb{R}^m to values in \mathbb{R} . While the value function cannot generally be found exactly, there is one exception to this rule. If model (1) is linear and Gaussian, Bellman’s equation can be solved in closed form for all time steps, yielding Kalman’s (1960) filter. In this case, the value function turns out to be multivariate quadratic with a unique global maximum for every time step; the argmax equals the Kalman-filtered state. That the Kalman filter corresponds to an exact *function-space* solution to Bellman’s equation does not appear to be widely known.

This exact solution corresponding to Kalman’s filter suggests that using quadratic approximations — which is different in principle from using discretisation methods — may be accurate in a broader context. In the early days of dynamic programming, Bellman considered polynomial approximations of value functions with the specific aim of avoiding ‘dimensionality difficulties’ associated with discretisation schemes (Bellman and Dreyfus, 1959, p. 247). Value functions in the context of filtering applications are special in that they possess a global maximum for each time step; the argmax determines the filtered state. For this reason, we consider a particularly simple polynomial: the multivariate quadratic function. While generally inexact, quadratic functions can accurately approximate smooth

value functions around their global maxima.³ The quadratic function is parametrised by its argmax and its $m \times m$ Hessian matrix, requiring $O(m^2)$ storage and $O(m^3)$ computational complexity at every time step, the latter corresponding to the inversion of an $m \times m$ matrix. The cumulative computational complexity over t time steps amounts to $O(m^3t)$, thereby offering full scalability to higher dimensional state spaces. A key contribution of this article is the insight that using function-space (rather than discrete) approximations allows us to avoid the curse of dimensionality, leading to a new class of (Bellman) filters that are computationally frugal and turn out to be remarkably accurate.

To illustrate the workings of the Bellman filter, we focus on state-space models where the observation equation may be nonlinear and/or non-Gaussian, while the state-transition equation remains linear and Gaussian; this class is still general enough for practical purposes. We compute explicit recursive formulas that constitute the Bellman filter (see Table 2), which contains as special cases (a) the Kalman filter, including its extended and iterated versions (e.g. Anderson and Moore, 2012), and (b) Fahrmeir’s (1992) approximate mode estimator (see section 4.4 for a discussion of special cases). As with the Kalman filter, the researcher keeps track of a filtered state and an associated precision matrix, which are determined by the argmax of the value function and the Hessian matrix at the peak, which are computed recursively. Like the Kalman filter, the Bellman filter is computationally inexpensive. Unlike the Kalman filter, it is driven by the *score* of the observation density rather than the prediction error. This makes the Bellman filter robust when faced with heavy-tailed observation noise, and (as we show in section 6) applicable to a wide class of nonlinear and non-Gaussian models.

1.4 Addressing the estimation problem

To circumvent the second drawback of the mode estimator — the inability to generate parameter estimates — computationally intensive (Monte Carlo) methods have been considered by many authors (Durbin and Koopman, 2000, 2002; Jungbacker and Koopman, 2007; Richard and Zhang, 2007; Koopman et al., 2015, 2016, to name a few). To achieve computational simplicity, we deviate from this strand of literature by numerically maximising the approximate log likelihood implied by the Bellman filter. We decompose the log likelihood into (a) the ‘fit’ of the Bellman-filtered states in view of the data, minus (b) the realised Kullback-Leibler (KL, see Kullback and Leibler, 1951) divergence between filtered and predicted state densities. Intuitively, we wish to maximise the congruence between Bellman-filtered states and the data, while minimising the distance between the filtered and predicted states to prevent over-fitting. All parts of the decomposition are given, or can be approximated, by the output of the Bellman filter. This means that standard gradient-based numerical optimisers can be used, making parameter estimation feasible and no more computationally demanding than ordinary estimation of the Kalman filter using maximum likelihood.

1.5 Related literature

Existing approximate filters, as shown in Table 1, typically suffer from various drawbacks. The extended Kalman filter (EKF, e.g. Anderson and Moore, 2012) and the unscented Kalman filter (UKF, Julier and Uhlmann, 1997) account for nonlinearity, but assume additive noise and make no adjustments when confronted with heavy tails. Masreliez’s (1975) filter is robust in the case of heavy-tailed observation noise, but computationally inefficient and, as the estimation problem is unaddressed, infeasible in practice. The Laplace Gaussian filter (Koyama et al., 2010) is computationally efficient, but requires tuning parameters and similarly leaves the estimation problem unaddressed. Fahrmeir’s (1992) posterior mode approximation applies to observations drawn from an exponential distribution. Despite their drawbacks, approximate filters have some advantages compared to simulation-based approaches,

³In the Bayesian literature, quadratic approximations around the mode are known as Laplace approximations (e.g. Tierney and Kadane, 1986); we avoid this term, as our approach is not Bayesian and no integrals are approximated.

such as particle-filtering methods, which have recently gained traction (e.g. Fearnhead and Clifford, 2003; Godsill et al., 2004; De Valpine, 2004; Lin et al., 2005; Andrieu et al., 2010; Bunch and Godsill, 2016; Guarniero et al., 2017; Jacob et al., 2020). One state-of-the-art approach, used in our simulation study in section 6, is the numerically accelerated importance sampling (NAIS) method in Koopman et al. (2015, 2016, 2017) and Barra et al. (2017), which is used in our simulation study in section 6. Simulation-based methods are subject to the curse of dimensionality making them computationally intensive — if not infeasible — when the dimension of the state space exceeds two (the performance of the NAIS method has not been documented in such situations). Possibly for this reason, the EKF remains the standard in many practical applications. Our contribution is to develop an approximate filter that is (a) generally applicable, (b) computationally efficient even in higher dimensions and (c) feasible in practice (i.e. allows the estimation problem to be addressed), while performing equally well as, if not better than, the NAIS method for low-dimensional state spaces.

2 Filtering using the mode

The state-space model under consideration is given in equation (1). A realised path is denoted by $\mathbf{y}_{1:t}(\omega)$ for every event $\omega \in \Omega$, where Ω denotes the event space of the underlying complete probability space of interest, denoted $(\Omega, \mathcal{F}, \mathbb{P})$. We continue to use generic notation in that we write the logarithm of joint and conditional densities as $\ell(\cdot, \cdot) := \log p(\cdot, \cdot)$ and $\ell(\cdot|\cdot) := \log p(\cdot|\cdot)$, respectively, for potentially different p 's. This section considers the filtering problem; any dependence on ψ is suppressed.

The joint log-likelihood function of the states and the data is written as $\ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t})$. Here, the data $\mathbf{y}_{1:t}$ are considered fixed and known, while the states $\mathbf{a}_{1:t}$ in Roman font are considered variables to be evaluated along any path $\mathbf{a}_{1:t} \in \mathbb{R}^{m \times t}$. Naturally, the true states $\boldsymbol{\alpha}_{1:t}$ (in Greek font) remain unknown. For the state-space model (1), the joint log likelihood can be found by means of the ‘probability chain rule’ (Godsill et al., 2004, p. 156) as follows:

$$\ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}) = \sum_{i=1}^t \ell(\mathbf{y}_i|\mathbf{a}_i) + \sum_{i=2}^t \ell(\mathbf{a}_i|\mathbf{a}_{i-1}) + \ell(\mathbf{a}_1). \quad (2)$$

The joint log likelihood $\ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t})$ is, *a priori*, a random function of the observations $\mathbf{y}_{1:t}$, even though the data are considered known and fixed *ex post*. Next, the mode is defined as the sequence of states that maximise equation (2).

Definition 1 (Mode) *Assuming it exists and is unique, the mode is*

$$\tilde{\mathbf{a}}_{1:t|t} := (\tilde{\mathbf{a}}_{1|t}, \tilde{\mathbf{a}}_{2|t}, \dots, \tilde{\mathbf{a}}_{t|t}) = \arg \max_{\mathbf{a}_{1:t} \in \mathbb{R}^{m \times t}} \ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}), \quad t \leq n. \quad (3)$$

Elements of the mode at time step t are denoted by $\tilde{\mathbf{a}}_{i|t}$ for $i \leq t$, where i denotes the state that is estimated, while t denotes the information set used. The entire solution is denoted $\tilde{\mathbf{a}}_{1:t|t} \in \mathbb{R}^{m \times t}$, which is a collection of t vectors. Iterative solution methods were proposed in Shephard and Pitt (1997), Durbin (1997) and Durbin and Koopman (2000), who use Newton’s method, and So (2003), who uses quadratic hill climbing.

Computing times for solving optimisation problem (3) typically grow as $O(t^3)$. This is unfortunate because, for the purposes of online filtering, we are predominantly interested in the last column of $\tilde{\mathbf{a}}_{1:t|t}$, i.e. $\tilde{\mathbf{a}}_{t|t}$, but for all times $t \leq n$. To obtain the desired sequence of real-time filtered states $\{\tilde{\mathbf{a}}_{t|t}\}_{t=1, \dots, n}$, we must compute the mode $\tilde{\mathbf{a}}_{1:t|t}$ for all time steps, and then, for each time step, extract the right-most column as the filtered state $\tilde{\mathbf{a}}_{t|t}$.

Definition 2 (Generally infeasible filter) *Filtered states based on the mode (3) are*

$$\tilde{\mathbf{a}}_{t_0|t_0}, \dots, \tilde{\mathbf{a}}_{t|t}, \dots, \tilde{\mathbf{a}}_{n|n}, \quad (4)$$

where $1 \leq t_0 \leq n$ is large enough to ensure that $\tilde{\mathbf{a}}_{t_0|t_0}$ exists.

Estimator (4) is generally infeasible as it is computed using the true (hyper)parameter $\boldsymbol{\psi}$. Further, estimator (4) is computationally intensive as each filtered state $\tilde{\mathbf{a}}_{t|t}$ requires the (increasingly large) optimisation problem (3) to be solved, involving $m \times t$ first-order conditions. This observation raises the question whether it is possible to proceed in real time without computing a large and increasing number of ‘smoothed’ state estimates as required by the optimisation (3). As we show in the next section, this is indeed possible when we make use of Bellman’s dynamic programming principle.

3 Filtering using dynamic programming

In this section our focus remains on the filtering problem; the (hyper)parameter $\boldsymbol{\psi}$ is considered given. To understand how a recursive approach may be feasible, we start by noting that the joint log-likelihood function (2) satisfies a straightforward recursive relation for $2 \leq t \leq n$ as follows:

$$\ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}) = \ell(\mathbf{y}_t | \mathbf{a}_t) + \ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + \ell(\mathbf{a}_{1:t-1}, \mathbf{y}_{1:t-1}). \quad (5)$$

That is, in transitioning from time $t - 1$ to time t , two terms are added: one representing the state-transition density, $\ell(\mathbf{a}_t | \mathbf{a}_{t-1})$; the other representing the observation density, $\ell(\mathbf{y}_t | \mathbf{a}_t)$. Next, we define the *value function* by maximising $\ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t})$ with respect to all states apart from the most recent state $\mathbf{a}_t \in \mathbb{R}^m$.

Definition 3 (Value function) *The value function $V_t : \Omega \times \mathbb{R}^m \rightarrow \mathbb{R}$ is*

$$V_t(\mathbf{a}_t) := \max_{\mathbf{a}_{1:t-1} \in \mathbb{R}^{m \times (t-1)}} \ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}), \quad \mathbf{a}_t \in \mathbb{R}^m. \quad (6)$$

The value function $V_t(\mathbf{a}_t)$ depends on the data $\mathbf{y}_{1:t}$, as indicated by the subscript t , which are considered fixed, and on its argument \mathbf{a}_t , which is a continuous variable in \mathbb{R}^m . Recursion (5) implies that the value function (6) satisfies Bellman’s equation, as stated below.

Proposition 1 (Bellman’s equation) *Suppose $\tilde{\mathbf{a}}_{t|t}$ exists for all $t \geq t_0$, where $1 \leq t_0 \leq n$. The value function (6) satisfies Bellman’s equation:*

$$V_t(\mathbf{a}_t) = \ell(\mathbf{y}_t | \mathbf{a}_t) + \max_{\mathbf{a}_{t-1} \in \mathbb{R}^m} \left\{ \ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + V_{t-1}(\mathbf{a}_{t-1}) \right\}, \quad \mathbf{a}_t \in \mathbb{R}^m, \quad (7)$$

for all $t_0 < t \leq n$. Further, the Bellman-filtered states, defined as

$$\mathbf{a}_{t|t} := \arg \max_{\mathbf{a}_t \in \mathbb{R}^m} V_t(\mathbf{a}_t), \quad t_0 \leq t \leq n, \quad (8)$$

satisfy $\mathbf{a}_{t|t} = \tilde{\mathbf{a}}_{t|t}$ for all $t_0 \leq t \leq n$.

Bellman’s equation (7) recursively relates the value function $V_t(\mathbf{a}_t)$ to the (previous) value function $V_{t-1}(\mathbf{a}_{t-1})$ by adding one term reflecting the state transition, $\ell(\mathbf{a}_t | \mathbf{a}_{t-1})$; one term reflecting the observation density, $\ell(\mathbf{y}_t | \mathbf{a}_t)$; and a subsequent maximisation over a single state variable, $\mathbf{a}_{t-1} \in \mathbb{R}^m$. The value function $V_t(\mathbf{a}_t)$ still depends on the data $\mathbf{y}_{1:t-1}$, but only indirectly, i.e. through the previous value function $V_{t-1}(\mathbf{a}_{t-1})$. Apart from assuming the existence of the mode, no (additional) assumptions

are needed regarding the functional forms of $\ell(\mathbf{y}_t|\mathbf{a}_t)$ and $\ell(\mathbf{a}_t|\mathbf{a}_{t-1})$; see Appendix A for the proof. As such, Bellman’s equation (7) is of quite general applicability. As the researcher receives the data \mathbf{y}_1 through \mathbf{y}_t , she can iteratively compute a sequence of value functions, which in turn imply a sequence of filtered state estimates via the argmax (8).

Remark 1 *For Markov-switching models, in which the latent state takes a finite number of (discrete) values, Bellman’s equation (7) can be solved exactly for all time steps, yielding Viterbi’s (1967) algorithm. Exact solubility of (7) is lost in general when the states take continuous values.*

When latent states are permitted to take values in a continuum, as in this article, the solution to Bellman’s equation (7) is a *function*, mapping the (continuous) state space \mathbb{R}^m to values in \mathbb{R} . While the value function cannot generally be found exactly, there is an exception to this rule. If model (1) is linear and Gaussian, the value function can be found in closed form for all time steps. In this case, the observation law $\ell(\mathbf{y}_t|\mathbf{a}_t)$ is a multivariate quadratic function of $\mathbf{a}_t \in \mathbb{R}^m$, while the state-transition law $\ell(\mathbf{a}_t|\mathbf{a}_{t-1})$ is a multivariate quadratic function in terms of both variables $\mathbf{a}_t, \mathbf{a}_{t-1} \in \mathbb{R}^m$. Finally, if the researcher’s knowledge of the previous state \mathbf{a}_{t-1} is Gaussian, then the value function $V_{t-1}(\mathbf{a}_{t-1})$, being a quantity in log space, is also a multivariate quadratic function. In this case, all optimisations required in equations (7) and (8) can be performed in closed form. The new value function $V_t(\mathbf{a}_t)$ is also multivariate quadratic, but with adjusted parameters. We then obtain Kalman’s (1960) filter, as highlighted by Corollary 1 (see next section for the proof).

Corollary 1 (Function-space solution to Bellman’s equation: Kalman filter) *Take a linear Gaussian state-space model with observation equation $\mathbf{y}_t = \mathbf{d} + \mathbf{Z} \mathbf{a}_t + \varepsilon_t$, where $\varepsilon_t \sim i.i.d. \mathbf{N}(\mathbf{0}, \mathbf{H})$, and state-transition equation $\mathbf{a}_t = \mathbf{c} + \mathbf{T} \mathbf{a}_{t-1} + \boldsymbol{\eta}_t$, where $\boldsymbol{\eta}_t \sim i.i.d. \mathbf{N}(\mathbf{0}, \mathbf{Q})$, such that Kalman’s (1960) filter applies. Assume the Kalman-filtered covariance matrices, denoted $\{\mathbf{P}_{t|t}\}$, are positive definite. Then the Bellman-filtered states $\{\mathbf{a}_{t|t}\}$ in equation (8) are identical to the Kalman-filtered states, the value function is multivariate quadratic at every time step, and its negative Hessian matrix equals $\mathbf{I}_{t|t} := \mathbf{P}_{t|t}^{-1}$ at every time step.*

The fact that the Kalman filter corresponds to a *function-space* solution to Bellman’s equation does not appear to be widely known. The main point, however, is that Bellman’s equation continues to hold in function space even when it *cannot* be solved in closed form. In this case, we must derive and store some (possibly parametric) approximation of the value function $V_t(\mathbf{a}_t)$ for each time step. Here we propose a polynomial approximation of the value function as a computationally cheap alternative to discretising the state space, which may be computationally taxing in the event of many grid points (Künsch, 2001).

While it is possible to develop asymptotic theory for approximating value functions using polynomials, doing so seems disproportionate relative to our aims. Indeed, a simpler approach will be shown to be sufficiently accurate. Motivated by Corollary 1 and the fact that our value functions possess global maxima, we approximate the value function for each time step using a multivariate quadratic function, which is parametrised by the argmax and the matrix of second derivatives at the peak. While generally inexact, multivariate quadratic functions can accurately approximate smooth value functions around their global maxima. What is more, the observation density $p(\mathbf{y}_t|\mathbf{a}_t)$ and the state-transition density $p(\mathbf{a}_t|\mathbf{a}_{t-1})$, which appear in Bellman’s equation (7), are still treated exactly. The simulation results in section 6 are so compelling that considering approximation methods more sophisticated than fitting quadratic functions appears to be unnecessary, at least for applications in time-series econometrics.

4 Bellman filter for models with linear Gaussian state dynamics

This section explicitly develops the Bellman filter for models in which the state-transition equation remains linear and Gaussian. Our focus remains on the filtering problem; the estimation problem is

Table 2: Bellman filter for model (9): A generalisation of the Kalman filter.

Step	Method	Computation
Initialise	Unconditional	Set $\mathbf{a}_{0 0} = (\mathbf{1} - \mathbf{T})^{-1} \mathbf{c}$ and $\text{vec}(\mathbf{I}_{0 0}^{-1}) = (\mathbf{1} - \mathbf{T} \otimes \mathbf{T})^{-1} \text{vec}(\mathbf{Q})$. Set $t = 1$.
	Diffuse	Set $\mathbf{a}_{0 0} = \mathbf{0}$ and $\mathbf{I}_{0 0} = \mathbf{0}$. Set $t = 1$.
Predict		$\mathbf{a}_{t t-1} = \mathbf{c} + \mathbf{T} \mathbf{a}_{t-1 t-1}$ $\mathbf{I}_{t t-1} = \mathbf{Q}^{-1} - \mathbf{Q}^{-1} \mathbf{T} (\mathbf{I}_{t-1 t-1} + \mathbf{T}' \mathbf{Q}^{-1} \mathbf{T})^{-1} \mathbf{T}' \mathbf{Q}^{-1}$
Start		Set $\mathbf{a}_{t t}^{(0)} = \mathbf{a}_{t t-1}$. Set $i = 0$.
		Alternatively, set $\mathbf{a}_{t t}^{(0)} = \arg \max_{\mathbf{a}} \ell(\mathbf{y}_t \mathbf{a})$ if this quantity exists. Set $i = 0$.
Optimise	Newton	$\mathbf{a}_{t t}^{(i+1)} = \mathbf{a}_{t t}^{(i)} + \left\{ \mathbf{I}_{i t-1} - \frac{d^2 \ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \right\}^{-1} \left\{ \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}} - \mathbf{I}_{i t-1} (\mathbf{a} - \mathbf{a}_{t t-1}) \right\} \Big _{\mathbf{a}=\mathbf{a}_{t t}^{(i)}}$
	Fisher	$\mathbf{a}_{t t}^{(i+1)} = \mathbf{a}_{t t}^{(i)} + \left\{ \mathbf{I}_{i t-1} + \mathbb{E} \left[-\frac{d^2 \ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \Big \mathbf{a} \right] \right\}^{-1} \left\{ \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}} - \mathbf{I}_{i t-1} (\mathbf{a} - \mathbf{a}_{t t-1}) \right\} \Big _{\mathbf{a}=\mathbf{a}_{t t}^{(i)}}$
	BHHH	$\mathbf{a}_{t t}^{(i+1)} = \mathbf{a}_{t t}^{(i)} + \left\{ \mathbf{I}_{i t-1} + \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}} \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}'} \right\}^{-1} \left\{ \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}} - \mathbf{I}_{i t-1} (\mathbf{a} - \mathbf{a}_{t t-1}) \right\} \Big _{\mathbf{a}=\mathbf{a}_{t t}^{(i)}}$
Stop		Set $i = i + 1$ and repeat the ‘Optimise’ step.
		Stop at $i = i_{\max}$ if some convergence criterion (e.g. that we stop after a pre-specified number of iterations i_{\max}) is satisfied.
Update		$\mathbf{a}_{t t} = \mathbf{a}_{t t}^{(i_{\max})}$
	Newton	$\mathbf{I}_{t t} = \mathbf{I}_{t t-1} - \frac{d^2 \ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \Big _{\mathbf{a}=\mathbf{a}_{t t}}$
	Fisher	$\mathbf{I}_{t t} = \mathbf{I}_{t t-1} + \mathbb{E} \left[-\frac{d^2 \ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \Big \mathbf{a} \right] \Big _{\mathbf{a}=\mathbf{a}_{t t}}$
	BHHH	$\mathbf{I}_{t t} = \mathbf{I}_{t t-1} + \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}} \frac{d\ell(\mathbf{y}_t \mathbf{a})}{d\mathbf{a}'} \Big _{\mathbf{a}=\mathbf{a}_{t t}}$
Proceed		Set $t = t + 1$ and return to the step ‘Predict’.

Note: The log-likelihood function $\ell(\mathbf{y}_t | \boldsymbol{\alpha}_t)$ is known in closed form and can be read off from the data-generating process (9). The corresponding score and the realised and expected information quantities are written as $d\ell(\mathbf{y}_t | \mathbf{a})/d\mathbf{a}$, $-d^2 \ell(\mathbf{y}_t | \mathbf{a})/(d\mathbf{a} d\mathbf{a}')$ and $\mathbb{E}[-d^2 \ell(\mathbf{y}_t | \mathbf{a})/(d\mathbf{a} d\mathbf{a}')] | \mathbf{a}$, respectively, which are viewed as functions of \mathbf{a} , to be evaluated at some state estimate. Under the steps ‘Optimise’ and ‘Update’, we list three (intentionally vanilla) optimisation methods, which may but need not be identical for both steps. Users may also implement more sophisticated optimisation methods based on the optimisation of equation (14).

addressed in section 5. The aim is to approximate, in function space and for all time steps, the solution to Bellman’s equation (7). We write the model with a linear Gaussian state equation as in Koopman et al. (2015, 2016):

$$\mathbf{y}_t \sim p(\mathbf{y}_t | \boldsymbol{\alpha}_t), \quad \boldsymbol{\alpha}_t = \mathbf{c} + \mathbf{T} \boldsymbol{\alpha}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \text{i.i.d. } \mathbf{N}(\mathbf{0}, \mathbf{Q}), \quad \boldsymbol{\alpha}_1 \sim p(\boldsymbol{\alpha}_1), \quad (9)$$

where $t = 1, \dots, n$. The system vector \mathbf{c} and system matrix \mathbf{T} are assumed to be of appropriate dimensions. The covariance matrix \mathbf{Q} is assumed symmetric and positive semi-definite. The observation density $p(\mathbf{y}_t | \boldsymbol{\alpha}_t)$ may be non-Gaussian and involve nonlinearity. We may employ exponential link functions to ensure that variables such as intensity or volatility remain positive. In our notation, the link function is left implicit; the observation density $p(\mathbf{y}_t | \boldsymbol{\alpha}_t)$ may contain any desired (nonlinear) dependence on the state $\boldsymbol{\alpha}_t$.

4.1 Deriving the prediction step

To derive the proposed Bellman filter for model (9), we start with Bellman’s equation (7). In practice, the behaviour of $V_{t-1}(\mathbf{a})$ around its peak turns out to be most relevant when determining the next

value function, $V_t(\mathbf{a})$. The value function $V_{t-1}(\mathbf{a})$ could be approximated locally, around its peak, by a multivariate quadratic function as follows:

$$V_{t-1}(\mathbf{a}_{t-1}) \approx -\frac{1}{2}(\mathbf{a}_{t-1} - \mathbf{a}_{t-1|t-1})' \mathbf{I}_{t-1|t-1} (\mathbf{a}_{t-1} - \mathbf{a}_{t-1|t-1}) + \text{constants}, \quad \mathbf{a}_{t-1} \in \mathbb{R}^m, \quad (10)$$

for some state estimate $\mathbf{a}_{t-1|t-1} \in \mathbb{R}^m$ and precision matrix $\mathbf{I}_{t-1|t-1} \in \mathbb{R}^{m \times m}$, which is assumed to be symmetric and positive definite. Here, the state $\mathbf{a}_{t-1} \in \mathbb{R}^m$ is considered a variable, while $\mathbf{a}_{t-1|t-1}$ is an estimate (i.e. a vector containing numbers). Constants can be ignored, as we are interested only in the location of the maximum and the sharpness of the peak, not in the height of the value function. Approximation (10) would be justified not only locally but also globally if our knowledge at time $t-1$ of the state \mathbf{a}_{t-1} were accurately described by a normal distribution with mean $\mathbf{a}_{t-1|t-1}$ and precision matrix $\mathbf{I}_{t-1|t-1}$. If the model is such that the Kalman filter applies, approximation (10) thus happens to be exact. Moreover, approximation (10) is exact for $t=1$ if the model is stationary and $\boldsymbol{\alpha}_0$ is drawn from the states' unconditional distribution, which is also multivariate normal (e.g. Harvey, 1990, p. 121). The initialisation of the Bellman filter based on the unconditional distribution is indicated in Table 2 under the step 'Initialise'. If the unconditional distribution does not exist, the filter can be started using the 'diffuse' initialisation given in Table 2.

The state transition in model (9) is linear and Gaussian, such that $\ell(\mathbf{a}_t|\mathbf{a}_{t-1})$ is a quadratic function of both state variables as follows:

$$\ell(\mathbf{a}_t|\mathbf{a}_{t-1}) = -\frac{1}{2}(\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1})' \mathbf{Q}^{-1} (\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1}) + \text{constants}, \quad \mathbf{a}_{t-1}, \mathbf{a}_t \in \mathbb{R}^m, \quad (11)$$

where \mathbf{a}_t and \mathbf{a}_{t-1} are continuous variables in \mathbb{R}^m . (If \mathbf{Q} is only positive semi-definite, its inverse can be interpreted in a generalised sense.) Next, substituting the quadratic approximation (10) and the exact state transition (11) into Bellman's equation (7), we obtain

$$V_t(\mathbf{a}_t) = \ell(\mathbf{y}_t|\mathbf{a}_t) + \max_{\mathbf{a}_{t-1} \in \mathbb{R}^m} \left\{ -\frac{1}{2}(\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1})' \mathbf{Q}^{-1} (\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1}) - \frac{1}{2}(\mathbf{a}_{t-1} - \mathbf{a}_{t-1|t-1})' \mathbf{I}_{t-1|t-1} (\mathbf{a}_{t-1} - \mathbf{a}_{t-1|t-1}) \right\} + \text{constants}, \quad \mathbf{a}_t \in \mathbb{R}^m, \quad (12)$$

which for the purposes of simplicity we write with equality, which is unproblematic as long as we keep in mind that the resulting value function is generally inexact. Conveniently, the variable \mathbf{a}_{t-1} appears at most quadratically on the right-hand side of equation (12). As such, the corresponding maximisation can be performed in closed form. Computing the first-order condition in equation (12) and solving for \mathbf{a}_{t-1} , we obtain

$$\mathbf{a}_{t-1}^* = (\mathbf{I}_{t-1|t-1} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T})^{-1} \{ \mathbf{I}_{t-1|t-1} \mathbf{a}_{t-1|t-1} + \mathbf{T}'\mathbf{Q}^{-1}(\mathbf{a}_t - \mathbf{c}) \}. \quad (13)$$

The solution \mathbf{a}_{t-1}^* depends linearly on \mathbf{a}_t ; in dynamic programming terms, the 'policy function' is linear in the state. Substituting argmax (13) back into equation (12), which was to be optimised, and performing some algebra (for details, see Appendix B), equation (12) becomes

$$V_t(\mathbf{a}_t) = \ell(\mathbf{y}_t|\mathbf{a}_t) - \frac{1}{2}(\mathbf{a}_t - \mathbf{a}_{t|t-1})' \mathbf{I}_{t|t-1} (\mathbf{a}_t - \mathbf{a}_{t|t-1}) + \text{constants}, \quad \mathbf{a}_t \in \mathbb{R}^m, \quad (14)$$

where \mathbf{a}_t remains as the only variable on the right-hand side, and we have defined the predicted state $\mathbf{a}_{t|t-1}$ and predicted precision matrix $\mathbf{I}_{t|t-1}$ as follows:

$$\mathbf{a}_{t|t-1} := \mathbf{c} + \mathbf{T} \mathbf{a}_{t-1|t-1}, \quad (15)$$

$$\mathbf{I}_{t|t-1} := \mathbf{Q}^{-1} - \mathbf{Q}^{-1}\mathbf{T}(\mathbf{I}_{t-1|t-1} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T})^{-1}\mathbf{T}'\mathbf{Q}^{-1}. \quad (16)$$

Equations (15) and (16) are collected in Table 2 under the step ‘Predict’. As equation (14) indicates, we are left with the log likelihood of a single observation, $\ell(\mathbf{y}_t|\mathbf{a}_t)$, and a quadratic ‘penalty’ term centred at the prediction $\mathbf{a}_{t|t-1}$, and with precision $\mathbf{I}_{t|t-1}$.

While our derivation is different, equations (15) and (16) turn out to be identical to the prediction steps of the Kalman filter. For equation (15), this is obvious; see e.g. Harvey (1990, p. 106). For equation (16), the relationship with Kalman’s prediction step is somewhat obscured because it is written in the information rather than the covariance form. To clarify, suppose that the inverses $\mathbf{P}_{t-1|t-1} := \mathbf{I}_{t-1|t-1}^{-1}$ and $\mathbf{P}_{t|t-1} := \mathbf{I}_{t|t-1}^{-1}$ exist. Following from the Woodbury matrix identity (e.g. Henderson and Searle, 1981, eq. 1), equation (16) implies $\mathbf{P}_{t|t-1} = \mathbf{T} \mathbf{P}_{t-1|t-1} \mathbf{T}' + \mathbf{Q}$, which is immediately recognisable as the covariance matrix prediction step of the Kalman filter (Harvey, 1990, p. 106).

4.2 Deriving the optimisation and updating steps

While predictions (15) and (16) turned out to be identical to those of the (information form of the) Kalman filter, the updating equations, derived next, are different in general. Taking the approximate value function (14) as given, the filtered state $\mathbf{a}_{t|t}$ and filtered precision matrix $\mathbf{I}_{t|t}$ can be found as

$$\mathbf{a}_{t|t} = \operatorname{argmax}_{\mathbf{a} \in \mathbb{R}^m} V_t(\mathbf{a}), \quad \mathbf{I}_{t|t} = - \left. \frac{d^2 V_t(\mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \right|_{\mathbf{a}=\mathbf{a}_{t|t}}. \quad (17)$$

The argmax determines our filtered state estimate, while the computation of second derivatives at the peak facilitates the proposed recursive approach, where each value function is approximated quadratically around its peak. The expression for $\mathbf{I}_{t|t}$ is ‘local’ in the sense that it utilises second derivatives at a single point; global fitting methods could also be used.

For the value function $V_t(\mathbf{a})$ in equation (14) to possess a unique global optimum, it is sufficient that the matrix of negative second derivatives, i.e. $\mathbf{I}_{t|t-1} - d^2\ell(\mathbf{y}_t|\mathbf{a})/(d\mathbf{a}d\mathbf{a}')$, is positive definite for all $\mathbf{a} \in \mathbb{R}^m$, where $-d^2\ell(\mathbf{y}_t|\mathbf{a})/(d\mathbf{a}d\mathbf{a}')$ is the *realised information*. Even if the existence of a global maximum is guaranteed, however, the potentially complicated functional form $\ell(\mathbf{y}_t|\mathbf{a}_t)$ implies that the maximisation over \mathbf{a}_t in equation (14) cannot, in general, be performed analytically. Nonetheless it is straightforward to write down analytically the steps of e.g. Newton’s optimisation method (e.g. Nocedal and Wright, 2006). Indeed, a plain-vanilla application of Newton’s method to maximising $V_t(\mathbf{a})$ with respect to its argument reads

$$\mathbf{a}_{t|t}^{(i+1)} = \mathbf{a}_{t|t}^{(i)} + \left[- \frac{d^2 V_t(\mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \right]^{-1} \left. \frac{dV_t(\mathbf{a})}{d\mathbf{a}} \right|_{\mathbf{a}=\mathbf{a}_{t|t}^{(i)}}, \quad (18)$$

where elements of the resulting sequence are denoted as $\mathbf{a}_{t|t}^{(i)}$ for $i = 0, 1, \dots$. As indicated in Table 2 under the step ‘Start’, Newton’s method (18) requires an initialisation to be specified, e.g. $\mathbf{a}_{t|t}^{(0)} = \mathbf{a}_{t|t-1}$, such that the starting point for the optimisation, at every time step, is equal to the prediction made at the previous time step.

Recalling value function (14), the gradient and negative Hessian can be approximated in closed form as follows:

$$\frac{dV_t(\mathbf{a})}{d\mathbf{a}} = \frac{d\ell(\mathbf{y}_t|\mathbf{a})}{d\mathbf{a}} - \mathbf{I}_{t|t-1}(\mathbf{a} - \mathbf{a}_{t|t-1}), \quad \mathbf{a} \in \mathbb{R}^m, \quad (19)$$

$$- \frac{d^2 V_t(\mathbf{a})}{d\mathbf{a} d\mathbf{a}'} = \mathbf{I}_{t|t-1} - \frac{d^2 \ell(\mathbf{y}_t|\mathbf{a})}{d\mathbf{a} d\mathbf{a}'}, \quad \mathbf{a} \in \mathbb{R}^m. \quad (20)$$

As the observation \mathbf{y}_t is fixed, the score in equation (19) and the realised information in equation (20) are viewed as functions of the state variable $\mathbf{a} \in \mathbb{R}^m$. Simply put, Newton’s version of the Bellman

filter can be obtained by substituting the gradient (19) and Hessian (20) into Newton’s method (18). The iterative Newton method is shown in Table 2 under the step ‘Optimise’. The computational complexity of the resulting filter is $O(m^3t)$, which is driven by the need to invert $m \times m$ matrices at every time step. The Bellman filter thus avoids the curse of dimensionality and, like the classic Kalman filter, offers full scalability to higher dimensions m . The presence of the score in the optimisation step is distinctive for the Bellman filter and guarantees its robustness if the observation density is heavy tailed. As indicated under the steps ‘Stop’ and ‘Update’, we may perform a fixed number of Newton steps, or as many as are required according to some convergence criterion, after which we set the final estimate $\mathbf{a}_{t|t}$ equal to $\mathbf{a}_{t|t}^{(i_{\max})}$, where i_{\max} is the number of iterations performed. After the updating step, we set $t = t + 1$ and return to the prediction step, as indicated in Table 2 under the step ‘Proceed’.

4.3 Alternative optimisation methods

Newton’s method is applicable if $\mathbf{I}_{t|t-1} - d^2\ell(\mathbf{y}_t|\mathbf{a})/(d\mathbf{a}d\mathbf{a}')$ is positive definite, which is guaranteed if the realised information is positive semi-definite for all realisations of $\mathbf{y}_t \in \mathbb{R}^l$ and $\mathbf{a} \in \mathbb{R}^m$. If not, we can still ensure well-defined optimisation steps by using Fisher’s scoring method or the Berndt-Hall-Hall-Hausman (BHHH) algorithm, which are also given in Table 2. These methods differ from Newton’s method in their approximation of the Hessian matrix and suggest different variations for the updated precision matrix $\mathbf{I}_{t|t}$, as indicated under the step ‘Update’ in Table 2.

The BHHH algorithm may be useful if second derivatives are hard to derive, or if the state dimension m is large; the required inverse can be computed in closed form using the Sherman-Morrison matrix-inversion lemma. Additionally, the BHHH algorithm may be attractive if the score is unbounded, in which case the BHHH updating step ensures step sizes of moderate length, such that the optimisation does not stray too far from its starting point. Generalising to more sophisticated optimisation methods, e.g. the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, is straightforward but, for applications in time-series econometrics, rarely needed.

With regard to the updated information matrix $\mathbf{I}_{t|t}$, we intuitively expect $\mathbf{I}_{t|t} \geq \mathbf{I}_{t|t-1}$, where the weak inequality means that the left-hand side minus the right-hand side is positive semi-definite. The intuition derives from the fact that missing observations can be dealt with as in the Kalman filter by setting $\mathbf{a}_{t|t} = \mathbf{a}_{t|t-1}$ and $\mathbf{I}_{t|t} = \mathbf{I}_{t|t-1}$. Next, we recognise that any (existing) observation should be weakly more informative than a non-existent one, implying $\mathbf{I}_{t|t} \geq \mathbf{I}_{t|t-1}$. The lower bound may be reached in the limit for extreme observations (i.e. outliers), which are completely uninformative. While Newton’s updating step has the advantage that it explicitly utilises the observation \mathbf{y}_t , and as such may recognise that some observations carry little information, the inequality $\mathbf{I}_{t|t} \geq \mathbf{I}_{t|t-1}$ is not guaranteed unless the realised information quantity is positive semi-definite. For Fisher’s updating step, the situation is reversed, failing to utilise the realisation \mathbf{y}_t while ensuring $\mathbf{I}_{t|t} \geq \mathbf{I}_{t|t-1}$. For some models it is possible to formulate a hybrid version, e.g. by taking a weighted average of Newton’s and Fisher’s updating steps, that achieves the best of both worlds (this will be relevant for some models in section 6).

4.4 Special cases of the Bellman filter

For linear Gaussian models, the objective function (14) is multivariate quadratic, such that both Newton’s method and Fisher’s method find the optimum in a single step, which is exactly the updating step of the Kalman filter (see Appendix C for details). For nonlinear Gaussian models, the Bellman filter contains as a special case the iterated extended Kalman filter (see Appendix D), while suggesting robust extensions in the case of heavy-tailed observation noise. The Bellman filter also generalises Fahrmeir’s (1992) approximate mode estimator (see also Fahrmeir and Tutz, 2013, p. 354). Our analysis differs from that by Fahrmeir in that (a) we show that on-line mode estimation can in theory

be performed exactly by solving Bellman’s equation, (b) we consider a general (rather than exponential) observation distribution, and (c) we allow more than one optimisation step.

5 Estimation method

This section considers the estimation problem, as distinct from the filtering problem, in that we aim to estimate both the time-varying states $\boldsymbol{\alpha}_{1:t}$ and the constant (hyper)parameter $\boldsymbol{\psi}$. As before, we take model (9) with linear Gaussian state dynamics, and we continue to assume the existence of the mode. To estimate the constant parameter $\boldsymbol{\psi}$, computationally intensive methods have been considered by many authors (see section 1). We deviate from this strand of literature by decomposing the log likelihood in terms of the ‘fit’ generated by the Bellman filter, penalised by the realised Kullback-Leibler (KL, see Kullback and Leibler, 1951) divergence between filtered and predicted states. Intuitively, we wish to maximise the congruence of the Bellman-filtered states and the data, while also minimising the distance between filtered and predicted states to prevent over-fitting. The proposed decomposition has the advantage that all terms can be evaluated or approximated using the output of the Bellman filter in Table 2; no sampling techniques or numerical integration methods are required. The resulting estimation method is as straightforward and computationally inexpensive as ordinary estimation of the Kalman filter using maximum likelihood.

To introduce the proposed decomposition, we focus on the log-likelihood contribution of a single observation, $\ell(\mathbf{y}_t|\mathcal{F}_{t-1}) := \log p(\mathbf{y}_t|\mathcal{F}_{t-1})$. The next computation is straightforward and holds for all $\mathbf{y}_t \in \mathbb{R}^l$ and all $\boldsymbol{\alpha}_t \in \mathbb{R}^m$:

$$\ell(\mathbf{y}_t|\mathcal{F}_{t-1}) = \ell(\mathbf{y}_t, \boldsymbol{\alpha}_t|\mathcal{F}_{t-1}) - \ell(\boldsymbol{\alpha}_t|\mathbf{y}_t, \mathcal{F}_{t-1}) = \ell(\mathbf{y}_t|\boldsymbol{\alpha}_t) + \ell(\boldsymbol{\alpha}_t|\mathcal{F}_{t-1}) - \ell(\boldsymbol{\alpha}_t|\mathcal{F}_t). \quad (21)$$

While the above decomposition is valid for any $\boldsymbol{\alpha}_t \in \mathbb{R}^m$, the resulting expression is not a computable quantity, as $\boldsymbol{\alpha}_t$ remains unknown. It is practical to evaluate the expression at the Bellman-filtered state estimate $\mathbf{a}_{t|t}$, such that, by swapping the order of the last two terms, we obtain

$$\ell(\mathbf{y}_t|\mathcal{F}_{t-1}) = \ell(\mathbf{y}_t|\boldsymbol{\alpha}_t) \Big|_{\boldsymbol{\alpha}_t = \mathbf{a}_{t|t}} - \left\{ \ell(\boldsymbol{\alpha}_t|\mathcal{F}_t) - \ell(\boldsymbol{\alpha}_t|\mathcal{F}_{t-1}) \right\} \Big|_{\boldsymbol{\alpha}_t = \mathbf{a}_{t|t}}. \quad (22)$$

The first term on the right-hand side, $\ell(\mathbf{y}_t|\boldsymbol{\alpha}_t)$ evaluated at $\boldsymbol{\alpha}_t = \mathbf{a}_{t|t}$, quantifies the congruence (or ‘fit’) between the Bellman-filtered state $\mathbf{a}_{t|t}$ and the observation \mathbf{y}_t , which we wish to maximise. We simultaneously want to minimise the realised KL divergence between predictions and updates, as determined by the difference between the two terms in curly brackets. The trade-off between maximising the first term and minimising the second, which appears with a minus sign, gives rise to a meaningful optimisation problem.

While the decomposition above is itself exact, we do not generally have an exact expression for the KL divergence. To ensure that the log-likelihood contribution (22) is computable, we now turn to approximating the KL divergence term. In deriving the Bellman filter for model (9), we presumed that the researcher’s knowledge, as measured in log-likelihood space for each time step, could be approximated by a multivariate quadratic function. Extending this line of reasoning, we consider the following approximations of the two terms that compose the realised KL divergence:

$$\ell(\boldsymbol{\alpha}_t|\mathcal{F}_t) \approx \frac{1}{2} \log \det\{\mathbf{I}_{t|t}/(2\pi)\} - \frac{1}{2}(\boldsymbol{\alpha}_t - \mathbf{a}_{t|t})' \mathbf{I}_{t|t} (\boldsymbol{\alpha}_t - \mathbf{a}_{t|t}), \quad (23)$$

$$\ell(\boldsymbol{\alpha}_t|\mathcal{F}_{t-1}) \approx \frac{1}{2} \log \det\{\mathbf{I}_{t|t-1}/(2\pi)\} - \frac{1}{2}(\boldsymbol{\alpha}_t - \mathbf{a}_{t|t-1})' \mathbf{I}_{t|t-1} (\boldsymbol{\alpha}_t - \mathbf{a}_{t|t-1}). \quad (24)$$

Here the state $\boldsymbol{\alpha}_t$ is understood as a variable in \mathbb{R}^m , while $\mathbf{a}_{t|t-1}$, $\mathbf{a}_{t|t}$, $\mathbf{I}_{t|t-1}$ and $\mathbf{I}_{t|t}$ are known quantities determined by the Bellman filter in Table 2. If the model is linear and Gaussian, then the

Bellman filter is exact (it is, in fact, the Kalman filter), as are equations (23)—(24).

To define our proposed approximate maximum likelihood estimator (MLE), we take the usual definition $\hat{\boldsymbol{\psi}} := \arg \max_{\boldsymbol{\psi}} \sum \ell(\mathbf{y}_t | \mathcal{F}_{t-1})$. Next, we substitute the (exact) ‘fit minus KL divergence’ decomposition (22) and the approximations (23) and (24) to obtain

$$\hat{\boldsymbol{\psi}} := \arg \max_{\boldsymbol{\psi}} \sum_{t=t_0+1}^n \left\{ \ell(\mathbf{y}_t | \mathbf{a}_{t|t}) + \frac{1}{2} \log \det(\mathbf{I}_{t|t}^{-1} \mathbf{I}_{t|t-1}) - \frac{1}{2} (\mathbf{a}_{t|t} - \mathbf{a}_{t|t-1})' \mathbf{I}_{t|t-1} (\mathbf{a}_{t|t} - \mathbf{a}_{t|t-1}) \right\}, \quad (25)$$

where all terms on the right-hand side implicitly or explicitly depend on the (hyper)parameter $\boldsymbol{\psi}$. Time $t_0 \geq 0$ is large enough to ensure the mode exists at time t_0 . If model (9) is stationary and $\boldsymbol{\alpha}_0$ is drawn from the unconditional distribution, as in our simulation studies in section 6, then $t_0 = 0$. The case $t_0 > 0$ is analogous to that for the Kalman filter when the first t_0 observations are used to construct a ‘proper’ prior (see Harvey, 1990, p. 123). The first term inside curly brackets, involving the observation density, is given by model (9). The remaining terms can be computed based on the output of the Bellman filter in Table 2. Expression (25) can be viewed as an alternative to the prediction-error decomposition for linear Gaussian state-space models (see e.g. Harvey, 1990, p. 126), the advantage being that estimator (25) is applicable more generally.

Corollary 2 *Take the linear Gaussian state-space model specified in Corollary 1. Assume that the Kalman-filtered covariance matrices $\{\mathbf{P}_{t|t}\}$ are positive definite. Estimator (25) then equals the MLE.*

Estimator (25) is only slightly more computationally demanding than standard maximum likelihood estimation of the Kalman filter. The sole source of additional computational complexity derives from the fact that the Bellman filter in Table 2 may perform several optimisation steps for each time step, while the Kalman filter performs only one. However, because each optimisation step is straightforward and few steps are typically required, the additional computational burden is negligible. Models of type (9) can now be approximately estimated with the same ease as a linear Gaussian state-space model.

6 Simulation studies

6.1 Design

We conduct a Monte Carlo study to investigate the performance of the Bellman filter for a range of data-generating processes (DGPs). We consider 10 DGPs with linear Gaussian state dynamics (9) and observation densities in Table 3, which also includes link functions, scores and information quantities. To avoid selection bias on our part, Table 3 has been adapted with minor modifications from Koopman et al. (2016). In taking the DGPs chosen by these authors, we essentially test the performance of the Bellman filter on an ‘exogenous’ set of models. While the numerically accelerated importance sampling (NAIS) method in Koopman et al. (2015, 2016) has been shown to produce highly accurate results, the Bellman filter turns out to be equally (if not more) accurate at a fraction of the computational cost.

We add one DGP to the nine considered in Koopman et al. (2016): a local-level model with heavy-tailed observation noise. While a local-level model with Gaussian observation noise would be solved exactly by the Kalman filter, the latter does not adjust for heavy-tailed observation noise. Although the Kalman filter remains the best linear unbiased estimator, the results below show that the (nonlinear) Bellman filter fares better.

For each DGP in Table 3, we simulate 1,000 time series of length 5,000, where constant (hyper)parameters for the first nine DGPs are taken from Koopman et al. (2016, Table 3).⁴ We use

⁴The state-transition equation has parameters $c = 0, T = 0.98, Q = 0.025$, except for both dependence models, in

the first 2,500 observations to estimate the constant parameters. For time steps $t = 2,501$ through $t = 5,000$, we produce one-step-ahead predictions of quantities of interest, such as λ_t , β_t , σ_t , ρ_t and μ_t . We follow Koopman et al. (2016) in predicting $k\beta_t$ and $\Gamma(1 + 1/k)\beta_t$ for the models involving the Gamma and Weibull distributions, respectively, as these quantities equal the expectation of the data, where the shape parameter k is replaced by its estimate. Finally, we compute mean absolute errors (MAEs) and root mean squared errors (RMSEs) by comparing predictions against their true (simulated) counterparts. For each DGP and each method, the reported average loss is based on $2,500 \times 1,000 = 2.5$ million predictions. We consider four methods to make predictions as follows:

1. For the generally infeasible mode estimator (4), we use the true parameters and a moving window of 250 observations, so that 250 first-order conditions are solved for each time step (larger windows result in excessive computational times). For simplicity⁵, one-step-ahead predictions of quantities of interest are obtained by applying link functions, e.g. $\tilde{\lambda}_{t|t-1} = \exp(c + T\tilde{a}_{t-1|t-1})$.
2. For the Bellman filter, the algorithm in Table 2 is initialised using the unconditional distribution. Each optimisation procedure takes as its starting point the most recent prediction and uses Newton’s method if the realised information in Table 3 is nonnegative, and Fisher’s method otherwise. The stopping criterion is either $|a_{t|t}^{(i)} - a_{t|t}^{(i-1)}| < 0.0001$ or $i_{\max} = 40$ iterations, whichever occurs first (on average, ~ 5 iterations are needed). Newton’s updating step is used if the realised information in Table 3 is nonnegative. Otherwise a weighted average of Newton’s and Fisher’s updating steps is used, where the weights are chosen to guarantee $I_{t|t} \geq I_{t|t-1}$.⁶ Predictions are made using (a) the true parameters, (b) in-sample estimated parameters, and (c) out-of-sample estimated parameters. Parameter estimation is based on estimator (25) using the first (or last) 2,500 observations for out-of-sample (or in-sample) estimation. Bellman-predicted states $a_{t|t-1}$ are transformed using link functions to obtain e.g. $\lambda_{t|t-1} = \exp(a_{t|t-1})$.
3. For the numerically accelerated importance sampling (NAIS) method, we follow Koopman et al. (2016). We deviate in computing, for each time step, not only the weighted mean but also the weighted median of the (simulated) predictions, where the weights are as in Koopman et al. (2016). We refer to these methods as NAIS-mean and NAIS-median, respectively.
4. The Kalman filter is used to estimate both stochastic volatility (SV) models and the local-level model. For both SV models, we follow the common practice of squaring the observations and taking logarithms to obtain a linear state-space model, albeit with biased and non-Gaussian observation noise (for details, see Ruiz, 1994 or Harvey et al., 1994). Predicted states can now be obtained via quasi maximum likelihood estimation (QMLE) of the Kalman filter. For the local-level model with heavy-tailed observation noise, the Kalman filter is applied directly, i.e. without adjustments, and estimated by QMLE.

which case $c = 0.02, T = 0.98, Q = 0.01$. In the observation equation, Student’s t distributions have ten degrees of freedom, i.e. $\nu = 10$, except for the local-level model, in which case $\nu = 3$. The remaining shape parameters are $k = 4$ for the negative binomial distribution, $k = 1.5$ for the Gamma distribution, $k = 1.2$ for the Weibull distribution, and $\sigma = 0.45$ for the local-level model.

⁵The transformation of predictions using (monotone) link functions is exact only if the (untransformed) predictions are based on the median, not the mode, but for simplicity we ignore this difference.

⁶For the dependence model with the Gaussian distribution, the weight placed on Fisher’s updating step should weakly exceed $1/2$. For the Student’s t distribution, this generalises to $1/2 \times (\nu + 4)/(\nu + 3)$. For the local-level model with heavy-tailed noise, the weight given to Fisher’s updating step should weakly exceed $(1 + \nu/3)/(1 + 3\nu)$.

Table 3: Overview of data-generating processes in simulation studies.

DGP		Link function	Density	Score	Realised information	Information
Type	Distribution		$p(\mathbf{y}_t \alpha_t)$	$\frac{d\ell(\mathbf{y}_t \alpha_t)}{d\alpha_t}$	$-\frac{d^2\ell(\mathbf{y}_t \alpha_t)}{d\alpha_t^2}$	$\mathbb{E}\left[-\frac{d^2\ell(\mathbf{y}_t \alpha_t)}{d\alpha_t^2}\middle \alpha_t\right]$
Count	Poisson	$\lambda_t = \exp(\alpha_t)$	$\lambda_t \exp(-\lambda_t)/y_t!$	$y_t - \lambda_t$	λ_t	λ_t
Count	Negative bin.	$\lambda_t = \exp(\alpha_t)$	$\frac{\Gamma(k+y_t) \left(\frac{k}{k+\lambda_t}\right)^k \left(\frac{\lambda_t}{k+\lambda_t}\right)^{y_t}}{\Gamma(k)\Gamma(y_t+k)}$	$y_t - \frac{\lambda_t(k+y_t)}{k+\lambda_t}$	$\frac{k\lambda_t(k+y_t)}{(k+\lambda_t)^2}$	$\frac{k\lambda_t}{k+\lambda_t}$
Intensity	Exponential	$\lambda_t = \exp(\alpha_t)$	$\lambda_t \exp(-\lambda_t y_t)$	$1 - \lambda_t y_t$	$y_t \lambda_t$	1
Duration	Gamma	$\beta_t = \exp(\alpha_t)$	$\frac{y_t^{k-1} \exp(-y_t/\beta_t)}{\Gamma(k)\beta_t^k}$	$\frac{y_t}{\beta_t} - k$	$\frac{y_t}{\beta_t}$	k
Duration	Weibull	$\beta_t = \exp(\alpha_t)$	$\frac{k (y_t/\beta_t)^{k-1}}{\beta_t \exp\{(y_t/\beta_t)^k\}}$	$k \left(\frac{y_t}{\beta_t}\right)^k - k$	$k^2 \left(\frac{y_t}{\beta_t}\right)^k$	k^2
Volatility	Gaussian	$\sigma_t^2 = \exp(\alpha_t)$	$\frac{\exp\{-y_t^2/(2\sigma_t^2)\}}{\{2\pi\sigma_t^2\}^{1/2}}$	$\frac{y_t^2}{2\sigma_t^2} - \frac{1}{2}$	$\frac{y_t^2}{2\sigma_t^2}$	$\frac{1}{2}$
Volatility	Student's t	$\sigma_t^2 = \exp(\alpha_t)$	$\frac{\Gamma(\frac{\nu+1}{2}) \left(1 + \frac{y_t^2}{(\nu-2)\sigma_t^2}\right)^{-\frac{\nu+1}{2}}}{\sqrt{(\nu-2)\pi}\Gamma(\nu/2)\sigma_t}$	$\frac{\omega_t y_t^2}{2\sigma_t^2} - \frac{1}{2}$ $\omega_t := \frac{\nu+1}{\nu-2+y_t^2/\sigma_t^2}$	$\frac{\nu-2}{\nu+1} \frac{\omega_t^2 y_t^2}{2\sigma_t^2}$	$\frac{\nu}{2\nu+6}$
Dependence	Gaussian	$\rho_t = \frac{1 - \exp(-\alpha_t)}{1 + \exp(-\alpha_t)}$	$\frac{\exp\left\{-\frac{y_{1t}^2+y_{2t}^2-2\rho_t y_{1t}y_{2t}}{2(1-\rho_t^2)}\right\}}{2\pi\sqrt{1-\rho_t^2}}$	$\frac{\rho_t}{2} + \frac{1}{2} \frac{z_{1t} z_{2t}}{1-\rho_t^2}$ $z_{1t} := y_{1t} - \rho_t y_{2t}$ $z_{2t} := y_{2t} - \rho_t y_{1t}$	$0 \not\leq \frac{1}{4} \frac{z_{1t}^2 + z_{2t}^2}{1-\rho_t^2} - \frac{1-\rho_t^2}{4}$	$\frac{1+\rho_t^2}{4}$
Dependence	Student's t	$\rho_t = \frac{1 - \exp(-\alpha_t)}{1 + \exp(-\alpha_t)}$	$\frac{\nu \left(1 + \frac{y_{1t}^2+y_{2t}^2-2\rho_t y_{1t}y_{2t}}{2(\nu-2)(1-\rho_t^2)}\right)^{-\frac{\nu+2}{2}}}{2\pi(\nu-2)\sqrt{1-\rho_t^2}}$	$\frac{\rho_t}{2} + \frac{\omega_t}{2} \frac{z_{1t} z_{2t}}{1-\rho_t^2}$ $z_{1t} := y_{1t} - \rho_t y_{2t}$ $z_{2t} := y_{2t} - \rho_t y_{1t}$	$0 \not\leq \frac{\omega_t}{4} \frac{z_{1t}^2 + z_{2t}^2}{1-\rho_t^2} - \frac{1-\rho_t^2}{4} - \frac{1}{2} \frac{\omega_t^2}{\nu+2} \frac{z_{1t}^2 z_{2t}^2}{(1-\rho_t^2)^2}$ $\omega_t := \frac{\nu+2}{\nu-2 + \frac{y_{1t}^2+y_{2t}^2-2\rho_t y_{1t}y_{2t}}{2(1-\rho_t^2)}}$	$\frac{2+\nu(1+\rho_t^2)}{4(\nu+4)}$
Local level	Student's t	$\mu_t = \alpha_t$	$\frac{\Gamma(\frac{\nu+1}{2}) \left(1 + \frac{(y_t-\mu_t)^2}{(\nu-2)\sigma^2}\right)^{-\frac{\nu+1}{2}}}{\sqrt{(\nu-2)\pi}\Gamma(\frac{\nu}{2})\sigma}$	$\frac{1}{\sigma} \frac{(\nu+1)e_t}{\nu-2+e_t^2}$ $e_t := \frac{y_t - \mu_t}{\sigma}$	$0 \not\leq \frac{\nu+1}{\sigma^2} \frac{\nu-2-e_t^2}{(\nu-2+e_t^2)^2}$	$\frac{\nu(\nu+1)}{\sigma^2(\nu-2)(\nu+3)}$

Note: The table contains ten data-generating processes (DGPs) and link functions, the first nine of which are adapted from Koopman et al. (2016). For each model, the DGP is given by the linear Gaussian state equation (9) in combination with the observation density and link functions indicated in the table. The table further displays link functions, scores, realised information quantities and expected information quantities. The realised information quantities are nonnegative, except for the bottom three models as indicated by $0 \not\leq \dots$. We deviate from Koopman et al. (2016) by computing scores and information quantities with respect to the state α_t , which is subject to linear Gaussian dynamics, rather than with respect to its transformation given by $\lambda_t, \beta_t, \sigma_t^2$ or ρ_t .

Table 4: Mean absolute errors (MAEs) of one-step-ahead predictions in simulation studies.

DGP		Infeasible estimator (4)		Bellman filter			NAIS (median)	Kalman filter using QMLE
		True parameters		True parameters	Estimated parameters (in-sample)	Estimated parameters (out-of-sample)	Estimated parameters (out-of-sample)	Estimated parameters (out-of-sample)
Type	Distribution	MAE	Relative MAE	Relative MAE			Relative MAE	Relative MAE
Count	Poisson	0.3556	1.0000	1.0029	0.9948	1.0017	1.0023	n/a
Count	Negative bin.	0.3816	1.0000	1.0006	0.9965	1.0048	1.0060	n/a
Intensity	Exponential	0.3998	1.0000	1.0036	0.9943	1.0032	1.0039	n/a
Duration	Gamma	0.5374	1.0000	1.0022	0.9952	1.0023	1.0028	n/a
Duration	Weibull	0.3493	1.0000	1.0034	0.9931	1.0014	1.0017	n/a
Volatility	Gaussian	0.1860	1.0000	1.0034	0.9934	1.0050	1.0049	1.1686
Volatility	Student's t	0.1930	1.0000	1.0014	0.9976	1.0100	1.0096	1.1709
Dependence	Gaussian	0.1131	1.0000	1.0010	0.9981	1.0166	1.0149	n/a
Dependence	Student's t	0.1160	1.0000	1.0004	0.9993	1.0204	1.0207	n/a
Local level	Student's t	0.1977	1.0000	1.0006	0.9995	1.0028	n/a	1.0793

Note: We simulated 1,000 time series each of length 5,000 for 10 data-generating processes of type (9) (the observation densities are listed in Table 3). Parameter estimation is based on the first 2,500 observations (out-of-sample estimation) or the last 2,500 observations (in-sample estimation) and is performed as follows: Bellman filter: based on estimator (25); numerically accelerated importance sampling (NAIS) method: as in Koopman et al. (2015, 2016); Kalman filter: quasi maximum likelihood estimation (QMLE). To make predictions of λ_t , β_t , σ_t , ρ_t and μ_t using the Bellman filter, we plug Bellman-predicted states $a_{t|t-1}$ into the link functions in Table 3, such that e.g. $\lambda_{t|t-1} = \exp(a_{t|t-1})$. As in Koopman et al. (2016), for the Gamma and Weibull models we predict $k\beta_t$ and $\Gamma(1 + 1/k)\beta_t$, respectively, where the shape parameter k is replaced by its estimate. To make predictions using NAIS, we compute the median of the simulations. In all cases, mean absolute errors (MAEs) are computed by comparing the last 2,500 predictions with their true (simulated) counterparts, and are reported relative to the MAE of the generally infeasible estimator (4).


6.2 Results

Table 4 contains MAEs of one-step-ahead predictions (RMSEs are shown in Appendix E). When reporting RMSEs and MAEs, we display the losses obtained from the NAIS-mean and NAIS-median methods, respectively, which are optimal for these loss functions (the Bellman filter, being based on the mode, is technically suboptimal for both loss functions).

We focus on three findings. First, comparison of the performance of the Bellman filter against that of the (generally infeasible) mode estimator (4) reveals that the MAE of the Bellman filter using true parameters is at most $\sim 0.3\%$ higher for all DGPs considered. When the parameters are estimated in an in-sample setting, the Bellman filter slightly outperforms the infeasible estimator. In an out-of-sample setting, the MAE of the Bellman filter remains within $\sim 1.7\%$ of that of the infeasible mode estimator for nine out of ten DGPs, while exceeding the MAE of the mode estimator by at most $\sim 2\%$ (for the dependence model with the Student's t distribution). Only a small fraction of the additional MAE is caused by approximate filtering and estimation. Rather, most of the additional MAE is caused by the design choice that the parameter estimation uses only the first half of the data, whereas the evaluation of MAEs pertains to the second half.

Second, although the Kalman and Bellman filters are usually in close agreement, the robustness of the Bellman filter means that it compares favourably with the Kalman filter for the SV and local-level models. Focusing on the local-level model, the performance of the Bellman filter is within $\sim 0.3\%$ of the infeasible estimator, even at parameters estimated out-of-sample. In contrast, the Kalman filter, confronted with heavy-tailed observation noise, lags $\sim 8\%$ behind the infeasible estimator. This difference is not due to the choice of loss function; the relative performance of the Kalman filter deteriorates further if we report RMSEs (see Appendix E). Moreover, the maximum absolute error in the out-of-sample period, averaged across 1,000 samples, is 1.74 for the Kalman filter; considerably higher than that for the Bellman filter (0.97). This shows that the Bellman filter is more robust in the

Figure 1: Scatter plots of 1,000 out-of-sample MAEs for each DGP as produced by the Bellman filter, horizontally, and the NAIS-median method, vertically.



Note: Each plot contains one dot for each of the 1,000 simulations. The coordinates of the dots are determined by the mean absolute error (MAE) of the Bellman filter, horizontally, and the numerically accelerated importance sampling (NAIS) median method in Koopman et al. (2016), vertically. Each dot involves an average over 2,500 out-of-sample predictions by both methods. Forty-five degree lines are also shown. For the simulation setting, see the note to Table 4.

Table 5: Average computation times (in seconds per sample).

DGP		Estimation		Filtering		
Type	Distribution	NAIS (median)	Bellman filter	Infeasible estimator (4)	NAIS (median)	Bellman filter
Count	Poisson	1.07	0.96	25.46	4.03	0.0089
Count	Negative binomial	3.06	1.11	30.61	5.20	0.0057
Intensity	Exponential	1.08	0.54	22.64	3.42	0.0080
Duration	Gamma	3.79	1.20	19.73	4.81	0.0083
Duration	Weibull	8.36	1.36	20.86	9.39	0.0088
Volatility	Gaussian	1.26	0.60	31.53	3.65	0.0055
Volatility	Student's t	2.74	1.04	22.00	5.27	0.0060
Dependence	Gaussian	2.37	1.93	35.02	5.45	0.0123
Dependence	Student's t	6.38	3.12	33.24	7.06	0.0103
Local level	Student's t	n/a	5.83	7.52	n/a	0.0168

Note: Computation times are measured on a computer running a 64-bit Windows 8.1 Pro with an Intel(R) Core(TM) i7-4810MQ CPU @ 2.80GHz with 16.0 GB of RAM. The numerical optimisation for both the NAIS method and the Bellman filter uses the Matlab function `fminunc` with identical settings.

face of heavy-tailed observation noise, while having only a single additional parameter to estimate (the degrees of freedom of the observation noise, ν). While several robust filters have been constructed ad hoc (e.g. Harvey and Luati, 2014 and Calvet et al., 2015), in our case robustness follows automatically from Bellman’s equation (14) along with the fact that the location score for a Student’s t distribution is bounded.

Third, the Bellman filter performs approximately on par with the NAIS-median method, despite the latter being more computationally intensive and, in fact, theoretically optimal for the absolute loss function. Table 4 shows that the NAIS-median method outperforms the Bellman filter by a maximum of 0.17% in terms of MAE (for the dependence model with a Gaussian distribution). For six out of nine DGPs, the roles are reversed: the Bellman filter marginally outperforms the NAIS-median method. This may be due to the fact that the NAIS method also contains an approximation; for each time step, only a finite number of predictions are simulated on which the median (or mean) is based.

Zooming in on the performance of the Bellman filter and the NAIS-median methods for individual samples, Figure 1 demonstrates that both methods perform almost identically for all 1,000 samples for each DGP; in each sub-figure, the dots are highly concentrated around the 45-degree lines. Clearly, the sample drawn is far more influential in determining the MAE than the choice between both filtering methods. Digging down even deeper, to the individual predictions, we use the 2.5 million predictions made by the Bellman filter for each DGP as an ‘exogenous’ variable to ‘explain’ the corresponding 2.5 million predictions made by the NAIS-median method. The resulting coefficients of determination (essentially R^2 values without fitting a model) exceed 99% across all DGPs, meaning that the individual predictions, too, are near identical.

Table 5 shows that the Bellman filter and the NAIS-median method differ in their computation times. In solving the estimation problem, the Bellman filter is faster by a factor 1.1 (for the Poisson distribution) to a factor ~ 6 (for the Weibull distribution). In solving the filtering problem, the Bellman filter is faster by a factor between ~ 400 (for the Poisson and exponential distributions) and $\sim 1,000$ (for the Weibull distribution).

Finally, Appendix F demonstrates that predicted confidence intervals implied by the Bellman filter, i.e. with endpoints given by $a_{t|t-1} \pm 2/\sqrt{I_{t|t-1}}$ for each time step, tend to be fairly accurate, containing the true states 93% to 96% of the time across all DGPs.

7 Conclusion

The Bellman filter for state-space models as developed in this article generalises the Kalman filter and is equally computationally inexpensive, but is robust in the case of heavy-tailed observation noise and applicable to a wider range of (nonlinear and non-Gaussian) models. Unlike the mode estimator, from which it is derived, the Bellman filter can be applied in real time and remains feasible in practice when (hyper)parameters must be estimated.

We investigated the performance of the Bellman filter in extensive simulation studies involving a wide range of models. While the predictions of the Bellman filter are near identical to those of state-of-the-art importance-sampling methods, filtering speeds are improved by a factor of ~ 400 to $\sim 1,000$ even when the state space is low dimensional. As importance-sampling techniques are — unlike the Bellman filter — subject to the curse of dimensionality, the difference in computational efficiency will only become more pronounced if the dimension of the state space is increased. Ultimately, the difference is not merely one of speed but of feasibility. When applied to state spaces with dimension greater than two, running times for the numerically accelerated importance sampling (NAIS) method (section 6) are likely to blow up exponentially, limiting the applicability of the method. In contrast, the computational complexity of the Bellman filter is only $O(m^3t)$, driven by the need to invert $m \times m$ matrices at every time step, where m is the dimension of the state space. To the best of our knowledge, therefore, the Bellman filter stands alone in offering full scalability to higher dimensional state spaces.

It is worth noting that making different choices with respect to (a) the method used to (parametrically) approximate value functions and (b) the optimisation routine used to find the argmax would result in different Bellman filters. This article has explored only a small part of this space: we have found it sufficiently accurate to approximate value functions using multivariate quadratic functions, and to make use of plain-vanilla optimisation schemes. There may be situations in which a more sophisticated approach is warranted; something we intend to explore in the future.

References

- Anderson, B. D. and Moore, J. B. (2012) *Optimal Filtering*. Courier Corporation.
- Andrieu, C., Doucet, A. and Holenstein, R. (2010) Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **72**, 269–342.
- Barra, I., Hoogerheide, L., Koopman, S. J. and Lucas, A. (2017) Joint Bayesian analysis of parameters and states in nonlinear non-Gaussian state space models. *Journal of Applied Econometrics*, **32**, 1003–1026.
- Bauwens, L. and Hautsch, N. (2006) Stochastic conditional intensity processes. *Journal of Financial Econometrics*, **4**, 450–493.
- Bauwens, L. and Veredas, D. (2004) The stochastic conditional duration model: A latent variable model for the analysis of financial durations. *Journal of Econometrics*, **119**, 381–412.
- Bellman, R. and Dreyfus, S. (1959) Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 247–251.
- Bellman, R. E. (1957) *Dynamic Programming*. Courier Dover Publications.
- Bunch, P. and Godsill, S. (2016) Approximations of the optimal importance density using Gaussian particle flow importance sampling. *Journal of the American Statistical Association*, **111**, 748–762.
- Calvet, L. E., Czellar, V. and Ronchetti, E. (2015) Robust filtering. *Journal of the American Statistical Association*, **110**, 1591–1606.
- De Valpine, P. (2004) Monte Carlo state-space likelihoods by weighted posterior kernel density estimation. *Journal of the American Statistical Association*, **99**, 523–536.
- Durbin, J. (1997) Optimal estimating equations for state vectors in non-Gaussian and nonlinear state space time series models. *Lecture Notes-Monograph Series*, 285–291.
- Durbin, J. and Koopman, S. J. (1997) Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika*, **84**, 669–684.
- (2000) Time series analysis of non-Gaussian observations based on state space models from both classical and Bayesian perspectives. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **62**, 3–56.

- (2002) A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, **89**, 603–616.
- (2012) *Time Series Analysis by State Space Methods*. Oxford University Press.
- Fahrmeir, L. (1992) Posterior mode estimation by extended Kalman filtering for multivariate dynamic generalized linear models. *Journal of the American Statistical Association*, **87**, 501–509.
- Fahrmeir, L. and Kaufmann, H. (1991) On Kalman filtering, posterior mode estimation and Fisher scoring in dynamic exponential family regression. *Metrika*, **38**, 37–60.
- Fahrmeir, L. and Tutz, G. (2013) *Multivariate Statistical Modelling based on Generalized Linear Models*. Springer Science & Business Media.
- Fearnhead, P. and Clifford, P. (2003) On-line inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **65**, 887–899.
- Forney, G. D. (1973) The Viterbi algorithm. *Proceedings of the IEEE*, **61**, 268–278.
- Frühwirth-Schnatter, S. and Wagner, H. (2006) Auxiliary mixture sampling for parameter-driven models of time series of counts with applications to state space modelling. *Biometrika*, **93**, 827–841.
- Fuh, C.-D. (2006) Efficient likelihood estimation in state space models. *The Annals of Statistics*, **34**, 2026–2068.
- Ghysels, E., Harvey, A. C. and Renault, E. (1996) Stochastic volatility. In *Handbook of Statistics*, vol. 14, 119–191. Elsevier.
- Godsill, S., Doucet, A. and West, M. (2001) Maximum a posteriori sequence estimation using Monte Carlo particle filters. *Annals of the Institute of Statistical Mathematics*, **53**, 82–96.
- Godsill, S. J., Doucet, A. and West, M. (2004) Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, **99**, 156–168.
- Godsill, S. J., Vermaak, J., Ng, W. and Li, J. F. (2007) Models and algorithms for tracking of maneuvering objects using variable rate particle filters. *Proceedings of the IEEE*, **95**, 925–952.
- Guarniero, P., Johansen, A. M. and Lee, A. (2017) The iterated auxiliary particle filter. *Journal of the American Statistical Association*, **112**, 1636–1647.
- Hafner, C. M. and Manner, H. (2012) Dynamic stochastic copula models: Estimation, inference and applications. *Journal of Applied Econometrics*, **27**, 269–295.
- Hamilton, J. D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, 357–384.
- Harvey, A. and Luati, A. (2014) Filtering with heavy tails. *Journal of the American Statistical Association*, **109**, 1112–1122.
- Harvey, A., Ruiz, E. and Shephard, N. (1994) Multivariate stochastic variance models. *The Review of Economic Studies*, **61**, 247–264.
- Harvey, A. C. (1990) *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Henderson, H. V. and Searle, S. R. (1981) On deriving the inverse of a sum of matrices. *SIAM Review*, **23**, 53–60.
- Jacob, P. E., Lindsten, F. and Schön, T. B. (2020) Smoothing with couplings of conditional particle filters. *Journal of the American Statistical Association*, **115**, 721–729.
- Jacquier, E., Polson, N. G. and Rossi, P. E. (2002) Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics*, **20**, 69–87.
- Julier, S. J. and Uhlmann, J. K. (1997) New extension of the Kalman filter to nonlinear systems. In *Signal processing, sensor fusion, and target recognition VI*, vol. 3068, 182–193. International Society for Optics and Photonics.
- Jungbacker, B. and Koopman, S. J. (2007) Monte Carlo estimation for nonlinear non-Gaussian state space models. *Biometrika*, **94**, 827–839.
- Kalman, R. E. (1960) A new approach to linear filtering and prediction problems. *American Society of Mechanical Engineers: Journal of Basic Engineering*, **82(1)**, 3545.
- Koopman, S. J., Lit, R. and Lucas, A. (2017) Intraday stochastic volatility in discrete price changes: The dynamic Skellam model. *Journal of the American Statistical Association*, **112**, 1490–1503.
- Koopman, S. J., Lucas, A. and Scharth, M. (2015) Numerically accelerated importance sampling for nonlinear non-Gaussian state-space models. *Journal of Business & Economic Statistics*, **33**, 114–127.
- (2016) Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, **98**, 97–110.

- Koyama, S., Castellanos Pérez-Bolde, L., Shalizi, C. R. and Kass, R. E. (2010) Approximate methods for state-space models. *Journal of the American Statistical Association*, **105**, 170–180.
- Kullback, S. and Leibler, R. A. (1951) On information and sufficiency. *The Annals of Mathematical Statistics*, **22**, 79–86.
- Künsch, H. (2001) State space and hidden Markov models. In *Complex Stochastic Systems* (eds. O. E. Barndorff-Nielsen and C. Kluppelberg), chap. 3. Chapman & Hall/CRC.
- Lin, M. T., Zhang, J. L., Cheng, Q. and Chen, R. (2005) Independent particle filters. *Journal of the American Statistical Association*, **100**, 1412–1421.
- Masreliez, C. (1975) Approximate non-Gaussian filtering with linear state and observation relations. *IEEE Transactions on Automatic Control*, **20**, 107–110.
- Nocedal, J. and Wright, S. (2006) *Numerical Optimization*. Springer Science & Business Media.
- Polson, N. G., Stroud, J. R. and Müller, P. (2008) Practical filtering with sequential parameter learning. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **70**, 413–428.
- Richard, J.-F. and Zhang, W. (2007) Efficient high-dimensional importance sampling. *Journal of Econometrics*, **141**, 1385–1411.
- Ruiz, E. (1994) Quasi-maximum likelihood estimation of stochastic volatility models. *Journal of Econometrics*, **63**, 289–306.
- Sardy, S. and Tseng, P. (2004) On the statistical analysis of smoothing by maximizing dirty Markov random field posterior distributions. *Journal of the American Statistical Association*, **99**, 191–204.
- Shephard, N. and Pitt, M. K. (1997) Likelihood analysis of non-Gaussian measurement time series. *Biometrika*, **84**, 653–667.
- Singh, A. and Roberts, G. (1992) State space modelling of cross-classified time series of counts. *International Statistical Review*, **60**, 321–335.
- So, M. K. (2003) Posterior mode estimation for nonlinear and non-Gaussian state space models. *Statistica Sinica*, 255–274.
- Tauchen, G. E. and Pitts, M. (1983) The price variability-volume relationship on speculative markets. *Econometrica: Journal of the Econometric Society*, 485–505.
- Taylor, S. J. (2008) *Modelling Financial Time Series*. World Scientific.
- Tierney, L. and Kadane, J. B. (1986) Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, **81**, 82–86.
- Viterbi, A. (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, **13**, 260–269.
- Viterbi, A. J. (2006) A personal history of the Viterbi algorithm. *IEEE Signal Processing Magazine*, **23**, 120–142.

A Proof of Proposition 1

Standard dynamic-programming arguments imply

$$\begin{aligned}
V_t(\mathbf{a}_t) &:= \max_{\mathbf{a}_{1:t-1} \in \mathbb{R}^{m \times (t-1)}} \ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}), \text{ by definition (6),} & (A.1) \\
&= \max_{\mathbf{a}_{1:t-1} \in \mathbb{R}^{m \times (t-1)}} \{\ell(\mathbf{y}_t | \mathbf{a}_t) + \ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + \ell(\mathbf{a}_{1:t-1}, \mathbf{y}_{1:t-1})\}, \text{ by recursion (5),} \\
&= \max_{\mathbf{a}_{t-1} \in \mathbb{R}^m} \{\ell(\mathbf{y}_t | \mathbf{a}_t) + \ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + \max_{\mathbf{a}_{1:t-2} \in \mathbb{R}^{m \times (t-2)}} \ell(\mathbf{a}_{1:t-1}, \mathbf{y}_{1:t-1})\}, \\
&\quad \text{by moving all but one maximisation inside curly brackets,} \\
&= \max_{\mathbf{a}_{t-1} \in \mathbb{R}^m} \{\ell(\mathbf{y}_t | \mathbf{a}_t) + \ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + V_{t-1}(\mathbf{a}_{t-1})\}, \text{ again by definition (6),} \\
&= \ell(\mathbf{y}_t | \mathbf{a}_t) + \max_{\mathbf{a}_{t-1} \in \mathbb{R}^m} \{\ell(\mathbf{a}_t | \mathbf{a}_{t-1}) + V_{t-1}(\mathbf{a}_{t-1})\}.
\end{aligned}$$

Further, it is evident that

$$\mathbf{a}_{t|t} := \arg \max_{\mathbf{a}_t \in \mathbb{R}^m} V_t(\mathbf{a}_t) = \arg \max_{\mathbf{a}_t \in \mathbb{R}^m} \max_{\mathbf{a}_{1:t-1} \in \mathbb{R}^{m \times (t-1)}} \ell(\mathbf{a}_{1:t}, \mathbf{y}_{1:t}) = \tilde{\mathbf{a}}_{t|t}, \quad (A.2)$$

where $\tilde{\mathbf{a}}_{t|t}$ was defined in equation (3).

B Derivation of equation (14)

In principle, equation (14) in the main text can be obtained by substituting equation (13) into equation (12) and performing algebraic manipulations. The desired result can be obtained more elegantly by ‘completing the square’ as follows. First, we replace \mathbf{a}_t with \mathbf{a}_t^* in equation (12), which then contains the following terms:

$$-\frac{1}{2}(\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1}^*)' \mathbf{Q}^{-1} (\mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1}^*) - \frac{1}{2}(\mathbf{a}_{t-1}^* - \mathbf{a}_{t-1|t-1})' \mathbf{I}_{t-1|t-1} (\mathbf{a}_{t-1}^* - \mathbf{a}_{t-1|t-1}). \quad (\text{B.1})$$

Then we recall from equation (13) that \mathbf{a}_{t-1}^* is linear in \mathbf{a}_t , such that the collection of terms in equation (B.1) above is at most multivariate quadratic in \mathbf{a}_t . Hence, we should be able to rewrite equation (B.1) as a quadratic function (i.e., by completing the square) as follows:

$$-\frac{1}{2}(\mathbf{a}_t - \mathbf{a}_{t|t-1})' \mathbf{I}_{t|t-1} (\mathbf{a}_t - \mathbf{a}_{t|t-1}) + \text{constants}, \quad (\text{B.2})$$

for some vector $\mathbf{a}_{t|t-1}$ to be found and some matrix $\mathbf{I}_{t|t-1}$ to be determined.

To do this, we note that $\mathbf{a}_{t|t-1}$ represents the argmax of equation (B.2), which can most readily be found by differentiating equation (B.1) with respect to \mathbf{a}_t and setting the result to zero. Using the envelope theorem, we need not account for the fact that \mathbf{a}_{t-1}^* depends on \mathbf{a}_t (the first derivative with respect to \mathbf{a}_{t-1}^* is zero because \mathbf{a}_{t-1}^* is optimal). Thus we set the derivative of equation (B.1) with respect to \mathbf{a}_t equal to zero, which gives $\mathbf{0} = \mathbf{a}_t - \mathbf{c} - \mathbf{T}\mathbf{a}_{t-1}^*$, or, by substituting \mathbf{a}_{t-1}^* from equation (13), we obtain

$$\begin{aligned} \mathbf{0} &= \mathbf{a}_t - \mathbf{c} - \mathbf{T}[\mathbf{I}_{t-1|t-1} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]^{-1}\mathbf{I}_{t-1|t-1}\mathbf{a}_{t-1|t-1} \\ &\quad - \mathbf{T}[\mathbf{I}_{t-1|t-1} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]^{-1}\mathbf{T}'\mathbf{Q}^{-1}(\mathbf{a}_t - \mathbf{c}). \end{aligned} \quad (\text{B.3})$$

The solution, $\mathbf{a}_{t|t-1} := \mathbf{T}\mathbf{a}_{t-1|t-1} + \mathbf{c}$, confirms prediction step (15).

Next, we compute the negative second derivative of equation (B.1) with respect to \mathbf{a}_t , which should give us $\mathbf{I}_{t|t-1}$. To account for the dependence of \mathbf{a}_{t-1}^* on \mathbf{a}_t , we use the chain rule. Specifically, in equation (13), \mathbf{a}_{t-1}^* is linear in \mathbf{a}_t , with the following Jacobian matrix:

$$\mathbf{J} := \frac{d\mathbf{a}_{t-1}^*}{d\mathbf{a}_t} = [\mathbf{I}_{t|t} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]^{-1}\mathbf{T}'\mathbf{Q}^{-1}. \quad (\text{B.4})$$

Next, the chain rule tells us that

$$\frac{d^2 \cdot}{d\mathbf{a}_t d\mathbf{a}_t'} = \begin{bmatrix} \mathbf{1} \\ \mathbf{J} \end{bmatrix}' \begin{bmatrix} \frac{\partial^2 \cdot}{\partial \mathbf{a}_t \partial \mathbf{a}_t'} & \frac{\partial^2 \cdot}{\partial \mathbf{a}_t \partial \mathbf{a}_{t-1}^{*'}} \\ \frac{\partial^2 \cdot}{\partial \mathbf{a}_{t-1}^* \partial \mathbf{a}_t'} & \frac{\partial^2 \cdot}{\partial \mathbf{a}_{t-1}^* \partial \mathbf{a}_{t-1}^{*'}} \end{bmatrix} \begin{bmatrix} \mathbf{1} \\ \mathbf{J} \end{bmatrix}, \quad (\text{B.5})$$

where instances of ∂ and d denote ‘partial’ and ‘total’ derivatives, respectively, while $\mathbf{1}$ denotes an identity matrix of appropriate size. As before, the envelope theorem ensures that no *first* derivative with respect to \mathbf{a}_t^* appears. When applying equation (B.5), we find that the negative second derivative of equation (B.1) becomes

$$\begin{aligned} &\begin{bmatrix} \mathbf{1} \\ \mathbf{J} \end{bmatrix}' \begin{bmatrix} \mathbf{Q}^{-1} & -\mathbf{Q}^{-1}\mathbf{T} \\ -\mathbf{T}'\mathbf{Q}^{-1} & \mathbf{I}_{t|t} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T} \end{bmatrix} \begin{bmatrix} \mathbf{1} \\ \mathbf{J} \end{bmatrix} \\ &= \mathbf{Q}^{-1} - \underbrace{\mathbf{Q}^{-1}\mathbf{T}\mathbf{J}} - \underbrace{\mathbf{J}'\mathbf{T}'\mathbf{Q}^{-1}} + \underbrace{\mathbf{J}'[\mathbf{I}_{t|t} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]\mathbf{J}}, \\ &= \mathbf{Q}^{-1} - \mathbf{Q}^{-1}\mathbf{T}[\mathbf{I}_{t|t} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]^{-1}\mathbf{T}'\mathbf{Q}^{-1}. \end{aligned} \quad (\text{B.6})$$

In the last line, we have used the fact that all three terms with curly brackets equal $\mathbf{Q}^{-1}\mathbf{T}[\mathbf{I}_{t|t} + \mathbf{T}'\mathbf{Q}^{-1}\mathbf{T}]^{-1}\mathbf{T}'\mathbf{Q}^{-1}$, such that two terms with curly brackets and opposite signs cancel, leaving only one term with a negative sign, which confirms prediction step (16).

C Kalman filter as a special case

Consider the linear Gaussian state-space model in Corollary 1. Suppose the inverse of the Kalman-filtered covariance matrix exists, i.e. $\mathbf{P}_{t-1|t-1}^{-1} := \mathbf{I}_{t-1|t-1}$ exists, such that the value function at time $t-1$ can be written as in equation (10), which is then exact. In Table 2, take the starting point $\mathbf{a}_{t|t}^{(0)} = \mathbf{a}_{t|t-1}$, and use Newton or Fisher optimisation steps. Given that the observation density is Gaussian, the log likelihood $\ell(\mathbf{y}_t|\mathbf{a}_t)$ is multivariate quadratic in \mathbf{a}_t , such that the entire objective function (14) turns out to be multivariate quadratic in \mathbf{a}_t . The matrix of second derivatives is constant, such that Newton and Fisher optimisation steps are identical. Moreover, given the quadratic nature of the objective function, both methods find the location of the optimum in a single step. Indeed, the result is the classic Kalman filter, albeit

written in the information form.

More explicitly, take $\mathbf{y}_t = \mathbf{d} + \mathbf{Z} \mathbf{a}_t + \boldsymbol{\varepsilon}_t$ with $\boldsymbol{\varepsilon}_t \sim \text{i.i.d. } \mathbf{N}(\mathbf{0}, \mathbf{H})$. Then

$$\ell(\mathbf{y}_t | \mathbf{a}_t) = -1/2(\mathbf{y}_t - \mathbf{d} - \mathbf{Z} \mathbf{a}_t)' \mathbf{H}^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z} \mathbf{a}_t) + \text{constants.} \quad (\text{C.1})$$

The score and realised information are

$$\frac{d\ell(\mathbf{y}_t | \mathbf{a}_t)}{d\mathbf{a}_t} = \mathbf{Z}' \mathbf{H}^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z} \mathbf{a}_t), \quad -\frac{d^2\ell(\mathbf{y}_t | \mathbf{a}_t)}{d\mathbf{a}_t d\mathbf{a}_t'} = \mathbf{Z}' \mathbf{H}^{-1} \mathbf{Z}. \quad (\text{C.2})$$

As the realised information is constant, it equals the (expected) marginal information. Taking the starting point $\mathbf{a}_{t|t}^{(0)} = \mathbf{a}_{t|t-1}$ for Newton's optimisation method, the estimate after a single Newton iteration reads

$$\mathbf{a}_{t|t}^{(1)} = \mathbf{a}_{t|t-1} + (\mathbf{I}_{t|t-1} + \mathbf{Z}' \mathbf{H}^{-1} \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{H}^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z} \mathbf{a}_{t|t-1}), \quad (\text{C.3})$$

which is exactly the Kalman filter level update written in information form. To see the equivalence with the covariance form of the Kalman filter, suppose that $\mathbf{P}_{t|t-1} := \mathbf{I}_{t|t-1}^{-1}$ exists. Then, using the Woodbury matrix-inversion formula (see e.g. Henderson and Searle, 1981, eq. 1), the expression above is equivalent to

$$\mathbf{a}_{t|t}^{(1)} = \mathbf{a}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{Z}' (\mathbf{Z} \mathbf{P}_{t|t-1} \mathbf{Z}' + \mathbf{H})^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z} \mathbf{a}_{t|t-1}), \quad (\text{C.4})$$

which is exactly the Kalman filter updating step (see e.g. Harvey, 1990, p. 106). For the information matrix update we have

$$\mathbf{I}_{t|t} = \mathbf{I}_{t|t-1} - \left. \frac{d^2\ell(\mathbf{y}_t | \mathbf{a})}{d\mathbf{a} d\mathbf{a}'} \right|_{\mathbf{a}=\mathbf{a}_{t|t}} = \mathbf{I}_{t|t-1} + \mathbf{Z}' \mathbf{H}^{-1} \mathbf{Z}. \quad (\text{C.5})$$

If the inverses $\mathbf{P}_{t|t-1} := \mathbf{I}_{t|t-1}^{-1}$ and $\mathbf{P}_{t|t} := \mathbf{I}_{t|t}^{-1}$ exist, then, again using Henderson and Searle (1981, eq. 1), we find

$$\mathbf{P}_{t|t} = \mathbf{I}_{t|t}^{-1} = (\mathbf{I}_{t|t-1} + \mathbf{Z}' \mathbf{H}^{-1} \mathbf{Z})^{-1} = \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1} \mathbf{Z}' (\mathbf{Z} \mathbf{P}_{t|t-1} \mathbf{Z}' + \mathbf{H})^{-1} \mathbf{Z} \mathbf{P}_{t|t-1}, \quad (\text{C.6})$$

which is exactly the Kalman filter covariance matrix updating step (again, see Harvey, 1990, p. 106).

D Iterated extended Kalman filter as a special case

Consider the linear Gaussian state-space model in Corollary 1, except let $\mathbf{y}_t = \mathbf{d} + \mathbf{Z}(\mathbf{a}_t) + \boldsymbol{\varepsilon}_t$ for some nonlinear vector function $\mathbf{Z}(\cdot)$ and $\boldsymbol{\varepsilon}_t \sim \text{i.i.d. } \mathbf{N}(\mathbf{0}, \mathbf{H})$. In Table 2, take the starting point $\mathbf{a}_{t|t}^{(0)} = \mathbf{a}_{t|t-1}$ and perform Fisher optimisation steps, ignoring (i.e. setting to zero) all second-order derivatives of $\mathbf{Z}(\cdot)$. The iterated extended Kalman filter is then obtained as a special case.

More explicitly, take $\mathbf{y}_t = \mathbf{d} + \mathbf{Z}(\mathbf{a}_t) + \boldsymbol{\varepsilon}_t$ with $\boldsymbol{\varepsilon}_t \sim \text{i.i.d. } \mathbf{N}(\mathbf{0}, \mathbf{H})$. Here, $\mathbf{Z}_t := \mathbf{Z}(\mathbf{a}_t)$ is a column vector of the same size as \mathbf{y}_t , where each element of \mathbf{Z}_t depends on the elements of \mathbf{a}_t . Then

$$\ell(\mathbf{y}_t | \mathbf{a}_t) = -1/2(\mathbf{y}_t - \mathbf{d} - \mathbf{Z}(\mathbf{a}_t))' \mathbf{H}^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z}(\mathbf{a}_t)) + \text{constants.} \quad (\text{D.1})$$

The score and marginal information are similar to those in Appendix C, as long as \mathbf{Z} there is replaced by the Jacobian of the transformation from \mathbf{a}_t to \mathbf{Z}_t , i.e. $d\mathbf{Z}(\mathbf{a}_t)/d\mathbf{a}_t'$. Hence

$$\frac{d\ell(\mathbf{y}_t | \mathbf{a}_t)}{d\mathbf{a}_t} = \frac{d\mathbf{Z}'}{d\mathbf{a}_t} \mathbf{H}^{-1} (\mathbf{y}_t - \mathbf{d} - \mathbf{Z}(\mathbf{a}_t)), \quad (\text{D.2})$$

$$\frac{d^2\ell(\mathbf{y}_t | \mathbf{a}_t)}{d\mathbf{a}_t d\mathbf{a}_t'} = -\frac{d\mathbf{Z}'}{d\mathbf{a}_t} \mathbf{H}^{-1} \frac{d\mathbf{Z}}{d\mathbf{a}_t'} + \text{second-order derivatives.} \quad (\text{D.3})$$

The iterated extended Kalman filter (IEKF) is obtained from the Bellman filter by choosing Newton's method and by making one further simplifying approximation: namely that all second-order derivatives of elements of \mathbf{Z}_t with respect to the elements of \mathbf{a}_t are zero. It is not obvious under what circumstances this approximation is justified, but here we are interested only in showing that the IEKF is a special case of the Bellman filter. Higher-order IEKFs may be obtained by retaining the second-order derivatives. If the observation noise $\boldsymbol{\varepsilon}_t$ is heavy tailed, however, the Bellman filter in Table 2 suggests a 'robustified' version of the Kalman filter and its extensions, in which case the tail behaviour of $p(\mathbf{y}_t | \mathbf{a}_t)$ is accounted for in the optimisation step by using the score $d\ell(\mathbf{y}_t | \mathbf{a}_t)/d\mathbf{a}_t$.

E Root mean squared errors in simulation studies

Table E.1: Root mean squared errors (RMSEs) of one-step-ahead predictions in simulation studies.

DGP		Infeasible estimator (4)		Bellman filter			NAIS (mean)	Kalman filter by QMLE
		True parameters		True parameters	Estimated parameters (in-sample)	Estimated parameters (out-of-sample)	Estimated parameters (out-of-sample)	Estimated parameters (out-of-sample)
Type	Distribution	RMSE	Relative RMSE	Relative RMSE			Relative RMSE	Relative RMSE
Count	Poisson	0.5364	1.0000	0.9977	1.0065	1.0156	1.0018	n/a
Count	Negative binomial	0.5944	1.0000	0.9981	1.0011	1.0129	0.9982	n/a
Intensity	Exponential	0.6668	1.0000	0.9961	1.0043	1.0222	1.0068	n/a
Duration	Gamma	0.8998	1.0000	1.0062	0.9819	0.9949	0.9785	n/a
Duration	Weibull	0.5933	1.0000	1.0078	0.9763	0.9941	0.9753	n/a
Volatility	Gaussian	0.2531	1.0000	1.0064	0.9830	0.9968	0.9879	1.1593
Volatility	Student's t	0.2621	1.0000	1.0037	0.9900	1.0036	0.9955	1.1507
Dependence	Gaussian	0.1455	1.0000	1.0017	0.9947	1.0126	1.0025	n/a
Dependence	Student's t	0.1487	1.0000	1.0010	0.9957	1.0167	1.0101	n/a
Local level	Student's t	0.2501	1.0000	0.9994	0.9959	0.9994	n/a	1.0969

Note: See the note to Table 4 in the main text. The only difference is that here we report root mean squared errors (RMSEs), not mean absolute errors (MAEs).

F Coverage in simulation studies

Table F.1: Coverage of Bellman-filter implied confidence intervals in simulation studies.

DGP		Bellman filter
		Estimated parameters
		Out-of-sample
Type	Distribution	Coverage of confidence interval (%)
Count	Poisson	95.2
Count	Negative binomial	94.8
Intensity	Exponential	95.8
Duration	Gamma	95.7
Duration	Weibull	96.0
Volatility	Gaussian	96.1
Volatility	Student's t	95.2
Dependence	Gaussian	93.1
Dependence	Student's t	93.0
Local level	Student's t	93.9

Note: See the note under Table 4 in the main text. This table reports how often the true states α_t were found to be within the predicted confidence interval for $t = 2,501, \dots, 5000$, where the predicted confidence intervals were derived from the Bellman filter in Table 2 of the main text with estimated parameters based on the first 2,500 observations (i.e., out-of-sample parameter estimation) using estimator (25). For scalar states, the endpoints of the approximate confidence interval are given by $a_{t|t-1} \pm 2/\sqrt{I_{t|t-1}}$.