

Unified Inference for Nonlinear Factor Models from Panels with Fixed and Large Time Span*

Torben G. Andersen[†] Nicola Fusari[‡] Viktor Todorov[§] Rasmus T. Varneskov[¶]

January 22, 2019

Abstract

We provide unifying inference theory for parametric nonlinear factor models based on a panel of noisy observations. The panel has a large cross-section and a time span that may be either small or large. Moreover, we incorporate an additional source of information, provided by noisy observations on some known functions of the factor realizations. The estimation is carried out via penalized least squares, i.e., by minimizing the L_2 distance between observations from the panel and their model-implied counterparts, augmented by a penalty for the deviation of the extracted factors from the noisy signals of them. When the time dimension is fixed, the limit distribution of the parameter vector is mixed Gaussian with conditional variance depending on the path of the factor realizations. On the other hand, when the time span is large, the convergence rate is faster and the limit distribution is Gaussian with a constant variance. In this case, however, we incur an incidental parameter problem since, at each point in time, we need to recover the concurrent factor realizations. This leads to an asymptotic bias that is absent in the setting with a fixed time span. In either scenario, the limit distribution of the estimates for the factor realizations is mixed Gaussian, but is related to the limiting distribution of the parameter vector only in the scenario with a fixed time horizon. Although the limit behavior is very different for the small versus large time span, we develop a feasible inference theory that applies, without modification, in either case. Hence, the user need not take a stand on the relative size of the time dimension of the panel. Similarly, we propose a time-varying data-driven weighting of the penalty in the objective function, which enhances efficiency by adapting to the relative quality of the signal for the factor realizations.

Keywords: Asymptotic Bias, Incidental Parameter Problem, Inference, Large Data Sets, Nonlinear Factor Model, Options, Panel Data, Stable Convergence, Stochastic Volatility.

JEL classification: C51, C52, G12.

*Andersen and Varneskov gratefully acknowledge support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The work is partially supported by NSF Grant SES-1530748. We would like to thank the guest editor (Frank Diebold), three anonymous referees, Marcelo Fernandez (our discussant), Christian Gourieroux as well as seminar participants at various conferences and seminars for many helpful comments and suggestions.

[†]Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; NBER, Cambridge, MA; and CREATES, Aarhus, Denmark; e-mail: t-andersen@northwestern.edu.

[‡]The Johns Hopkins University Carey Business School, Baltimore, MD 21202; e-mail: nicola.fusari@jhu.edu.

[§]Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: v-todorov@northwestern.edu.

[¶]Department of Finance, Copenhagen Business School, 2000 Frederiksberg, Denmark; CREATES, Aarhus, Denmark; Multi Assets at Nordea Asset Management, Copenhagen, Denmark; e-mail: rtv.fi@cbs.dk.

1 Introduction

We develop unified inference theory for parametric nonlinear factor models from panels of noisy observations, which have an asymptotically increasing cross-sectional dimension and a time span that is either fixed or increasing. In particular, we provide inference for the time-invariant parameters of the model as well as the period-by-period latent factor realizations. Importantly, although the limit theories for the fixed and increasing time span scenarios are very different, our feasible implementation is identical across the two cases, enabling empirical work that is robust to the relative length of the time series. The theory is presented in the context of an option panel, i.e., a collection of option prices observed across time and written on a given underlying asset. Our option prices are recorded over a sample period, where each observation date typically will have a different set of option strikes and tenors (times-to-maturity) available. With the rapid increase of trading in derivatives over recent years, such option panels are readily accessible and provide a unique source of information regarding the state of the economy and the pricing of volatility and tail risk in financial markets. Moreover, with suitable adaptation, our theory is applicable for analysis of economic systems with similar panel structures, such as household and firm data or term structure and large-scale macro DSGE models.

Given a specific parametric model, the theoretical value of an option is a nonlinear function of the (known) option strike and tenor as well as of the (unknown) time-invariant parameters and the latent state vector (factor) realization.¹ Under correct specification, the observed option price equals the model-implied theoretical value plus an observation error, centered at zero. Therefore, unknown parameters and factor realizations can be estimated by minimizing the L_2 distance between observed and model-implied option prices (i.e., nonlinear least squares). We further utilize nonparametric volatility measures based on high-frequency data of the underlying asset and penalize deviations of model-implied volatilities (determined by the factor realizations) from their nonparametric high-frequency estimates in the objective function. We refer to this estimation approach as penalized least squares (PLS). In alternative applications, penalization of this form may be useful if the latent factors can be associated directly with economic variables that are observed with measurement error.

Andersen et al. (2015a) develop the inference theory for PLS and a fixed time span. In this setting, both the parameters and factor realizations are estimated at rate \sqrt{N} , where N denotes the average number of options per observation date. The limiting distribution is mixed Gaussian, i.e., it is Gaussian

¹As is common, we rely on a fully parametric specification of the risk-neutral return distribution. This ensures a tractable representation for which we are able to develop comprehensive estimation and inference techniques, in spite of the highly nonlinear dependence of the option prices on the underlying latent state variables (or factors), and the strong dependence in the option prices across nearby strikes and tenors as well as in the time series dimension. In addition, without an explicit dynamic representation of the risk-neutral distribution, it is hard to impose fundamental no-arbitrage conditions on the system. Importantly, a battery of diagnostic tests is available to guard against model misspecification, following the procedures of Andersen et al. (2015a).

conditional on the observed path of the factors. In particular, this enables the approach to convey explicitly how the precision of the estimates for the option prices and factor realizations changes from period-to-period, depending on the level of volatility of the underlying asset, the option coverage over the strike ranges and tenors, the size of the measurement errors, etc.

In this paper, we develop the asymptotic theory for the case where the time span of the panel may increase at a rate equal to that of the cross-section. This scenario induces a so-called incidental parameter problem (Neyman and Scott (1948) and Lancaster (2000)). As the time span increases, the quantities to estimate (the factor realizations in our setting) also increase. For the long time span (T large), we show that the rate of convergence of the parameter vector rises to \sqrt{NT} . Unlike for the fixed time span, the limiting distribution for the parameters is now Gaussian. In addition, due to the incidental parameter problem, there is an asymptotic bias of order $O_p(1/N)$. We propose model-based de-biasing, which allows for feasible inference on the parameters. On the other hand, the rate of recovery for the factor realizations remains \sqrt{N} and the asymptotic limit is still mixed Gaussian. We derive the joint asymptotic distribution of the parameters and the factor realizations. This is critical for exploring various economic hypotheses. For example, this allows for testing model performance across different time periods of interest, e.g., calm versus turbulent market conditions.

Despite the fact that the behavior of our PLS estimator is very different in the fixed and increasing time span scenarios – both in terms of rate of convergence and limiting distribution – we show that feasible inference can be performed in an identical manner for either of the two asymptotic schemes, provided it suitably accommodates features of both limit theories. The reason is that the estimate of the asymptotic bias for the long span case is asymptotically negligible if the time span actually is fixed. Likewise, estimates for the dependence between the parameters and factor realizations in the fixed span case become asymptotically negligible in the long span setting. Finally, the estimates for the conditional variance of the parameters automatically represent long-run asymptotic variance estimates in the long span case. The above results depend critically on the fact that the limit theory is \mathcal{F} -conditionally stable (\mathcal{F} denotes the sigma-algebra of the underlying probability space). Among other things, this implies that we can apply standard normalization techniques for the construction of confidence intervals and testing, even when the asymptotic variances are random in the limit. From a practical point of view, this set of findings is important because the applied econometrician need not take a stand on whether the time span of the panel is of a length that justifies appealing to one limit theory versus the other – the procedure automatically adapts to the relevant features of the data.

Another novel aspect of our inference procedure is the use of extra information on the latent factor. This is reminiscent of the factor augmented vector autoregression (FAVAR) proposed by Bernanke et al. (2005). In our context, the signal for the factor realization takes the form of a high-frequency

nonparametric estimator of volatility. The asymptotic role of the penalization based on this volatility estimator depends on the relative size of the option cross-section versus the sampling frequency of the underlying asset price. This determines the relative (asymptotic) efficiency in recovering the volatility from options and high-frequency data, respectively. If the former is more efficient, the specification of the penalization term ensures it has no asymptotic impact while, if the latter is more efficient, the penalization will play an asymptotic role. Therefore, the user need not take a stand on whether the option or high-frequency data are more efficient for recovering volatility. As for the span of the panel, the procedure simply adapts to the situation at hand.

To further improve efficiency, we extend the estimation to weighted penalized least squares. We do so by forming a consistent estimate for the option observation error variance across each cross-section from a first-stage PLS estimate of the parameters and factor realizations. We then weight the squared (pricing) errors in the objective function by the estimated option error variance and the penalization term by a measure of its (conditional) asymptotic variance. Again, the approach applies for both the fixed and long span case, without modification. It allows the relative weight of the penalization in the objective function to shift over the days in the sample, reflecting the time-variation in the informativeness of the option versus high-frequency data.

Our main theoretical contributions can be summarized as follows. First, we analyze nonlinear panel models in the joint asymptotic setting where the time and cross-sectional dimensions can grow at the same rate in the presence of stochastic latent factors (fixed effects). Second, we augment the estimation by including noisy factor signals in the form of nonparametric high-frequency volatility estimates. Third, we derive the joint asymptotic limit of the latent factors, the parameters and the high-frequency volatility estimates in the various asymptotic setups (T fixed or increasing up to the rate of N), which is mixed Gaussian. Fourth, and perhaps most importantly, we show that feasible inference can be conducted in a way that works irrespective of whether T is fixed or increasing.

The above contributions can be compared with prior work in the following way. First, Andersen et al. (2015a) (see also Andrews (2005) and Kuersteiner and Prucha (2013)) consider only the fixed time span case. Second, Hahn and Kuersteiner (2002), Hahn and Newey (2004), Arellano and Hahn (2007), Arellano and Bonhomme (2009), Hahn and Kuersteiner (2011) and Fernandez-Val and Weidner (2016) consider only the case when T and N grow at the same rate, while Gagliardini and Gourieroux (2014) study the case of T growing at a slower rate than N . These papers do not allow for noisy factor signals in the estimation nor do they consider the joint asymptotic behavior of factor realizations and parameters. Third, Bai and Ng (2002) and Bai (2003) consider estimation of linear factor models. The major difference from that strand of work is that our factor loadings (i.e., the sensitivity of option prices to factors) are “structural,” i.e., driven by a parameter vector of fixed size, while Bai and Ng

(2002) and Bai (2003) treat them as unknown, unrestricted coefficients. Hence, in their case, the number of parameters (factor loadings) grows in proportion to the cross-sectional dimension of the panel. This implies a very different asymptotic behavior of the factor estimates depending on the relative size of T versus N . In particular, when T is fixed, Bai (2003) shows that relatively strong assumptions, including homoskedasticity of the observation error, are necessary for consistent factor extraction. Finally, our paper relates to Gagliardini et al. (2011), who consider estimation using data on options as well as the underlying asset. Their asymptotic setup is different as they treat the case of N fixed and depend importantly on (exact) nonparametric inference from the underlying asset data.²

We conduct numerical experiments to assess the practical relevance of the developed limit theory. We rely on a popular two-factor stochastic volatility model for the underlying asset price and consider strike and tenor configurations that mimic typical option data sets. The findings illustrate the advantages of our unified inference framework. For the factor extraction, we find the “stabilization” of the standard error, captured in the long-span asymptotic setting, to become effective across widely different horizons for different factors. Thus, neither the fixed nor long span asymptotics is satisfactory across all factors simultaneously. In addition, we find fairly strong dependence in the extraction error for the factor realizations. Such dependence is of higher order in the long-span asymptotic setup and is, thus, ignored if the associated asymptotic is used. Finally, the asymptotic biases, which arise only in the long-span asymptotic setting, play a nontrivial role for some parameters, even over relatively short time spans. As noted, our feasible unified inference procedure adapts by accommodating those features of the fixed or long span asymptotics that are relevant for the situation at hand.

The rest of the paper is organized as follows. In Section 2, we introduce the parametric nonlinear factor model for the option panel, the option observation scheme and the nonparametric high-frequency volatility estimator. In Section 3, we derive the asymptotic behavior of the PLS estimator in the long time span case. Section 4 develops unified feasible inference theory for panels with long and short time spans. In Section 5, we extend the analysis to weighted penalized least squares. Section 6 contains numerical experiments, illustrating the gains from the unified inference in the context of a popular parametric asset pricing model. Section 7 concludes. The Online Appendix contains the assumptions, proofs of the theoretical results and additional simulation results.

²There is a large body of earlier empirical option pricing work, see e.g., Bates (2003), Christoffersen and Jacobs (2004), Eraker (2004), Pan (2002) and the survey by Garcia et al. (2010). The emphasis of these studies, however, is on extracting information from a limited (fixed N) number of options, which differs fundamentally from our intention of using the full set of available options, leading to the setup involving an asymptotically increasing cross-section. Moreover, unlike our analysis, the fixed cross-section typically necessitates parametric assumptions for the observation error and a parametric specification for the factor evolution.

2 Setup

We start our analysis by introducing the nonlinear factor model for the panel of option prices. We then describe the observation scheme for the options as well as the high-frequency return data for the nonparametric volatility measurement.

2.1 A Parametric Nonlinear Factor Model for the Option Panel

The options are written on an underlying process, denoted by X , which is defined on a filtered probability space $(\Omega^{(0)}, \mathcal{F}^{(0)}, (\mathcal{F}_t^{(0)})_{t \geq 0}, \mathbb{P}^{(0)})$. The process X is governed by the following general dynamics,

$$\frac{dX_t}{X_{t-}} = \alpha_t dt + \sqrt{V_t} dW_t^{\mathbb{P}} + \int_{x > -1} x \tilde{\mu}^{\mathbb{P}}(dt, dx), \quad (1)$$

where α_t and V_t are càdlàg; $W_t^{\mathbb{P}}$ is a $\mathbb{P}^{(0)}$ -Brownian motion; $\mu^{\mathbb{P}}$ is a random measure counting the jumps in X , with compensator $\tilde{\nu}^{\mathbb{P}}(dt, dx) = a_t dt \otimes \nu^{\mathbb{P}}(dx)$ for some process a_t and Lévy measure $\nu^{\mathbb{P}}(dx)$, and with associated martingale measure $\tilde{\mu}^{\mathbb{P}} = \mu^{\mathbb{P}} - \tilde{\nu}^{\mathbb{P}}$. We denote expectations under $\mathbb{P}^{(0)}$ by $\mathbb{E}(\cdot)$. This formulation is very general; essentially all parametric models in the literature satisfy equation (1).

The so-called jump activity is bounded by the parameter r , i.e., we have,

$$\int_{\mathbb{R}} (|x|^r \wedge 1) \nu^{\mathbb{P}}(dx) < \infty, \quad (2)$$

for some $r \in [0, 1)$ and we define $x \wedge y = \min(x, y)$. The leading case is $r = 0$, corresponding to jumps of finite activity, which have been predominantly used in prior work, while values of $r > 0$ allow for infinitely active jumps. The parameter r appears explicitly in our asymptotic results.

Under no-arbitrage, a risk-neutral probability measure, \mathbb{Q} , is guaranteed to exist, see, e.g., Section 6.K in Duffie (2001), and is locally equivalent to $\mathbb{P}^{(0)}$. It transforms discounted asset prices into (local) martingales. In particular, for X under \mathbb{Q} , we have,

$$\frac{dX_t}{X_{t-}} = (r_t - \delta_t) dt + \sqrt{V_t} dW_t^{\mathbb{Q}} + \int_{x > -1} x \tilde{\mu}^{\mathbb{Q}}(dt, dx), \quad (3)$$

where r_t is the risk-free interest rate and δ_t is the dividend yield, $W_t^{\mathbb{Q}}$ denotes a \mathbb{Q} -Brownian motion, and the jump martingale measure is defined with respect to a risk-neutral compensator $\tilde{\nu}^{\mathbb{Q}}(dt, dx)$.

The theoretical values of derivatives written on X equal their expected discounted payoffs under the risk-neutral measure, \mathbb{Q} . In particular, if we denote the prices of European-style out-of-the-money (OTM) options on X at time t by $O_{t,k,\tau}$, we have,

$$O_{t,k,\tau} = \begin{cases} \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^{t+\tau} r_s ds} (X_{t+\tau} - K)^+ \right], & \text{if } K > F_{t,t+\tau}, \\ \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^{t+\tau} r_s ds} (K - X_{t+\tau})^+ \right], & \text{if } K \leq F_{t,t+\tau}, \end{cases} \quad (4)$$

where τ is the tenor, K is the strike price, $F_{t,t+\tau}$ is the futures price of the underlying asset, X , at time t for the future date $t + \tau$, and $k = \ln(K/F_{t,t+\tau})$ is the log-moneyness.³

We model the risk-neutral probability measure parametrically, rendering the option prices nonlinear functions of k and τ as well as a parameter vector and a latent state vector.⁴ In particular, under the risk-neutral model, the diffusive volatility and jump processes are governed by a (latent) state vector, or factors, so that $V_t = \xi_1(\mathbf{S}_t)$ and $\tilde{\nu}^{\mathbb{Q}}(dt, dx) = \xi_2(\mathbf{S}_t) \otimes \nu^{\mathbb{Q}}(dx)$, where $\nu^{\mathbb{Q}}(dx)$ is a Lévy measure; $\xi_1(\cdot)$ and $\xi_2(\cdot)$ are known functions; and \mathbf{S}_t denotes the $p \times 1$ state vector taking values in $\mathcal{S} \subset \mathbb{R}^p$. Importantly, we only require \mathbf{S}_t to be an $\mathcal{F}_t^{(0)}$ -adapted stochastic process, whose dynamics need not be specified explicitly for the purposes of our analysis.⁵ We further assume that the model-implied option prices, quoted in terms of Black-Scholes implied variance (BSIV), are given by the functions $\kappa(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, where k and τ denote the moneyness and tenor of the option, and \mathbf{Z} and $\boldsymbol{\theta}$ denote a particular value of the state and parameter vectors, respectively. The parameter vector is of dimension $q \times 1$ and takes values $\boldsymbol{\theta} \in \Theta$, with Θ being a compact subset of \mathbb{R}^q . We denote the true values of the latent factors at time t by \mathbf{S}_t and the true value of the parameter vector by $\boldsymbol{\theta}_0$. We note that the only time-dependence of the option prices is through the value of the state vector.⁶ Our primary interest in this paper is on inference for $\boldsymbol{\theta}_0$ as well as the factor realizations $\{\mathbf{S}_t\}_{t \geq 0}$.

The above framework for option pricing nests most continuous-time models used in empirical work, including the affine jump-diffusion class of Duffie et al. (2000).⁷ The only requirement for implementation of our inference procedures is knowledge of the function $\kappa(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, either in closed-form, up to numerical integration, or via simulation.⁸ We stress, in particular, that we do allow for risk-neutral specifications in which there is dependence between the latent state vector \mathbf{S}_t and either $W_t^{\mathbb{Q}}$ or the jump measure $\mu^{\mathbb{Q}}$. That is, we accommodate the so-called leverage effect, working through either the diffusive or the jump component of X_t , or both.

2.2 Observation Scheme

The time span of the panel is $[0, T]$, and we have observations from the option surface at given integer times $t = 1, \dots, T$. For each date, the options cover a significant range of strikes k and several tenors,

³Note that by including the subscript t in the notation of the option price, we allow for the option price to change from one day to another, even when the log-moneyness of the option, k , stays the same.

⁴Our focus is on the option panel, which is fully characterized through a model for the risk-neutral probability. We impose no assumptions on the \mathbb{P} dynamics beyond what is implied by our model for the \mathbb{Q} dynamics and no-arbitrage.

⁵For convenience, we assume the functions $\xi_1(\cdot)$ and $\xi_2(\cdot)$ do not depend on the parameter vector. This is almost universally satisfied in empirical applications. In addition, r_t and δ_t should be known functions of the state vector \mathbf{S}_t .

⁶Thus, we implicitly assume stationarity of the risk-neutral distribution through the state vector.

⁷We only assume that the diffusion coefficient and the jump intensity of X are governed by a finite-dimensional state vector under the risk-neutral probability measure, but make no such assumption for the statistical measure.

⁸In our numerical exercise below, it takes 0.001 seconds to evaluate an entire option cross-section on a given trading day, and 0.06 seconds to optimize over the factors on a single day, i.e., to perform the optimization in equation (12) below on a single day.

τ . Specifically, at time t , we observe the option prices $\{O_{t,k_j,\tau_j}\}_{j=1,\dots,N_t}$ for some (large) integer N_t , where the index j runs across the full set of strikes and tenors available on that day. The number of options for the maturity τ is denoted N_t^τ , so $N_t = \sum_\tau N_t^\tau$, where N_t^τ and N_t are $\mathcal{F}_t^{(0)}$ -adapted.

Given the heterogeneity of the available option data, we allow the panel to be unbalanced, so both the total number of options and the available strike-tenor combinations (k, τ) are free to vary across time. In particular, we assume,

$$\frac{N_t^\tau}{N_t} \approx \pi_t^\tau \quad \text{and} \quad \frac{N_t}{N} \approx \varsigma_t, \quad (5)$$

where π_t^τ and ς_t are positive-valued processes, and N is an unobserved non-random number, representing the ‘‘average size of the cross-section’’. We refer to Section A.1 of the Online Appendix for a formal statement regarding the nature of the approximations in equation (5).

For each pair, (t, τ) , we denote the minimum and maximum log-moneyness by $\underline{k}(t, \tau)$ and $\bar{k}(t, \tau)$, respectively. At time t and tenor τ , we define the $\mathcal{F}_t^{(0)}$ -adapted moneyness grid as,

$$\underline{k}(t, \tau) = k_{t,\tau}(1) < k_{t,\tau}(2) \dots < k_{t,\tau}(N_t^\tau) = \bar{k}(t, \tau),$$

and we further set $\Delta_{t,\tau}(i) = k_{t,\tau}(i) - k_{t,\tau}(i-1)$, for $i = 2, \dots, N_t^\tau$. Our asymptotic scheme sequentially adds new strikes to the existing ones within $[\underline{k}(t, \tau), \bar{k}(t, \tau)]$, so that $\Delta_{t,\tau}(i)$ shrinks towards zero, asymptotically, while allowing for the strikes to cluster relatively more densely within certain regions of the range, i.e.,

$$\Delta_{t,\tau}(i) \approx \frac{1}{N_t^\tau} \psi_{t,\tau}(k_{t,\tau}(i)), \quad i = 2, \dots, N_t^\tau, \quad (6)$$

where $\psi_{t,\tau}(k)$ is some positive valued process. Again, we defer a formal characterization of the approximation in equation (6) to the Online Appendix.

Finally, \mathcal{T}_t denotes the tenors available at a given point in time t , and the vectors $\underline{\mathbf{k}}_t = (\underline{k}(t, \tau))_{\tau \in \mathcal{T}_t}$ and $\bar{\mathbf{k}}_t = (\bar{k}(t, \tau))_{\tau \in \mathcal{T}_t}$ indicate the lowest and highest log-moneyness at time t across the available tenors. These quantities may vary over time and be random, thus accommodating pronounced shifts in the observed option cross-section along these dimensions over the sample.

The above setting allows for a great deal of intertemporal heterogeneity in the observation scheme. For example, the tenors need not be identical across days, and we allow for a different number of options in the panel across days, maturities and moneyness. Also, intuitively, the relative number of options at a given time will impact the precision of the inference for the state vector compared to other dates. Likewise, the relative number of options across the maturities and the local ‘‘sparseness’’ of the strikes impact the quality of inference for parameters and state variables differentially depending on their sensitivity to tenor and moneyness. The quantities ς_t , π_t^τ and $\psi_{t,\tau}(k)$ capture these facets of the panel configuration, and they appear explicitly in the asymptotic distribution theory below.

We stipulate that option prices, quoted in terms of BSIV, are observed with error, i.e., we have,

$$\tilde{\kappa}_{t,k,\tau} = \kappa_{t,k,\tau} + \epsilon_{t,k,\tau}, \quad (7)$$

where we use the shorthand $\kappa_{t,k,\tau} = \kappa(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)$ and the observation (measurement) errors, $\epsilon_{t,k,\tau}$, are defined on a space $\Omega^{(1)} = \prod_{t \in \mathbb{N}, k \in \mathbb{R}, \tau \in \Gamma} \mathcal{A}_{t,k,\tau}$, for $\mathcal{A}_{t,k,\tau} = \mathbb{R}$, with Γ denoting the set of all possible tenors. Moreover, $\Omega^{(1)}$ is equipped with the product Borel σ -field $\mathcal{F}^{(1)}$ as well as with a transition probability $\mathbb{P}^{(1)}(\omega^{(0)}, d\omega^{(1)})$ from the original probability space $\Omega^{(0)}$ – on which X is defined – to $\Omega^{(1)}$. Finally, we define the information filtration on $\Omega^{(1)}$ via $\mathcal{F}_t^{(1)} = \sigma(\epsilon_{s,k,\tau} : s \leq t)$. Then, the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ is given as follows, $\Omega = \Omega^{(0)} \times \Omega^{(1)}$, $\mathcal{F} = \mathcal{F}^{(0)} \times \mathcal{F}^{(1)}$,

$$\mathcal{F}_t = \cap_{s > t} \mathcal{F}_s^{(0)} \times \mathcal{F}_s^{(1)}, \quad \text{and} \quad \mathbb{P}(d\omega^{(0)}, d\omega^{(1)}) = \mathbb{P}^{(0)}(d\omega^{(0)})\mathbb{P}^{(1)}(\omega^{(0)}, d\omega^{(1)}).$$

We assume the extension has the property that any local martingale and semimartingale remain a local martingale and semimartingale (with the same characteristics) on the extended probability space. Formally, this is assured if the extension is “very good”.⁹

Intuitively, we motivate this representation by the following considerations. The option errors are defined on the space $\Omega^{(1)}$. We equip this space with the simple product topology as, at any point in time, only a countable number of them appear in our estimation. Since we want to allow the errors and the underlying process X to be dependent, we define the probability measure via a transition probability distribution from $\Omega^{(0)}$ to $\Omega^{(1)}$. Importantly, we assume that the observation errors are $\mathcal{F}^{(0)}$ -conditionally centered and possibly heteroskedastic of unknown form, i.e.,

$$\mathbb{E}(\epsilon_{t,k,\tau} | \mathcal{F}^{(0)}) = 0, \quad \text{and} \quad \mathbb{E}(\epsilon_{t,k,\tau}^2 | \mathcal{F}^{(0)}) = \phi_{t,k,\tau} \text{ for some } \phi_{t,k,\tau}, \quad (8)$$

and, in addition, they are $\mathcal{F}^{(0)}$ -conditionally independent:

$$\epsilon_{t,k,\tau} \perp\!\!\!\perp \epsilon_{t',k',\tau'} | \mathcal{F}^{(0)} \quad \text{as soon as } (t, k, \tau) \neq (t', k', \tau'), \quad (9)$$

where, as is common for two vectors \mathbf{a} and \mathbf{b} , $\mathbf{a} \neq \mathbf{b}$ signifies that they differ in at least one of their elements. We refer to the Online Appendix for additional assumptions on the observation error process, but note that these are satisfied in the case of “modulated” observation noise, i.e., for $\epsilon_{t,k,\tau} = \sqrt{\phi_{t,k,\tau}} \zeta_{t,k,\tau}$, where $\phi_{t,k,\tau}$ is $\mathcal{F}_t^{(0)}$ -adapted and $\zeta_{t,k,\tau}$ is defined on $\Omega^{(1)}$ and is independent from $\mathcal{F}^{(0)}$.¹⁰

⁹That is, $\omega^{(0)} \rightarrow \int 1_{\{A\}}(\omega^{(0)}, \omega^{(1)})\mathbb{P}^{(1)}(\omega^{(0)}, d\omega^{(1)})$ is $\mathcal{F}_t^{(0)}$ -measurable for all $A \in \mathcal{F}^{(1)}$ (see Definition II.7.1 of Jacod and Shiryaev (2003)).

¹⁰We also note that the quantities associated with the sampling scheme, \underline{k} , \bar{k} and N_t , are allowed to be random and depend on $\mathcal{F}^{(0)}$. Due to the assumption $\mathbb{E}(\epsilon_{t,k,\tau} | \mathcal{F}^{(0)}) = 0$, this does not cause “endogeneity” biases in the estimation.

2.3 Volatility Estimation from High-Frequency Asset Prices

We conclude this section by introducing the second source of information, consisting of high-frequency data on the underlying asset X . These data are used only for nonparametric estimation of the volatility process, which, in turn, assist in the recovery of the state vector (or at least part of it). To this end, we assume to have an equidistant high-frequency recording of X with grid size $\Delta_n = 1/n$, i.e., we observe X at times $0, 1/n, \dots, i/n, \dots, T+1$. The log-price is denoted by $x_t = \log(X_t)$ and, similarly, its increment by $\Delta_i^n x = x_{i/n} - x_{(i-1)/n}$. Using this, our nonparametric volatility estimator at time t is defined as

$$\hat{V}_t^n = \frac{n}{k_n} \sum_{i \in I_t^n} (\Delta_i^n x)^2 1_{\{|\Delta_i^n x| \leq \alpha \Delta_n^\varpi\}}, \quad I_t^n = [tn + 1, \dots, tn + k_n], \quad (10)$$

for some $\alpha > 0$, $\varpi \in (0, 1/2)$, and $k_n \in \mathbb{N}$ with $k_n < n$. This constitutes the spot counterpart of the integrated truncated variation estimator of Mancini (2001). For satisfactory finite sample performance, it is critical to determine α adaptively, based on a preliminary tuning-free estimate of volatility, such as the bipower variation of Barndorff-Nielsen and Shephard (2004, 2006) or the median realized volatility of Andersen et al. (2012).

The limit theory for the PLS estimator is directly impacted by the degree of information about spot volatility embedded in the high-frequency data relative to the option prices. Asymptotically, this is determined by the rate at which we acquire information from the high-frequency asset returns versus the option cross-section. We summarize this feature of the system through the assumption,

$$\frac{k_n}{N} \rightarrow \varrho \quad \text{for some real number } \varrho \geq 0.$$

Note that we allow $\varrho = 0$, but not $\varrho = \infty$. This is consistent with the empirical findings, where short-dated option prices are invariably found to be very useful in pinning down the level of volatility, while the high-frequency spot estimator can be quite noisy. Nonetheless, we do expect the high-frequency data to provide auxiliary information, corresponding to a limiting ratio $\varrho > 0$.

3 Inference for Panels with Long Time Span

This section provides our econometric analysis for the scenario where the large cross-sectional dimension of the panel is accompanied by a large time span. Given the observation equation (7), we estimate the system using least squares. However, we further penalize the deviation of the model-implied volatility from its nonparametric high-frequency counterpart during estimation, so we refer to our approach as penalized least squares. Section 3.1 introduces the formal estimation criterion and establishes consistency. We then provide the notation required for the associated limit theory in

Section 3.2, before presenting the limit distribution results in Section 3.3. Finally, Section 3.4 derives results for the estimation of the factor dynamics from the extracted factor realizations.

3.1 The PLS Estimator and Consistency

Our setup is in some respects similar to that of the classical dynamic nonlinear panel models, see, e.g., Hahn and Newey (2004) and Hahn and Kuersteiner (2011), with the latent factors playing the role of the fixed effects. As in this strand of work, we do not model the dynamics of the latent factors parametrically. Consequently, given our assumption for the observation error in equation (7), it is natural to estimate the parameters and the fixed effects (latent factors) via least squares. We further augment the objective function by a penalty term, which reflects information in the high-frequency volatility estimator, \widehat{V}_t^n , about the latent factors \mathbf{S}_t .¹¹

Formally, we define the $T \times p$ matrix of factor realizations by $\mathbf{Z} = \{\mathbf{Z}'_t\}_{t=1}^T$ and use this to write the objective function, for some finite constant $\lambda \geq 0$, as,¹²

$$\begin{aligned} \mathcal{L}(\mathbf{Z}, \boldsymbol{\theta}) &\equiv \sum_{t=1}^T \mathcal{L}_t(\mathbf{Z}_t, \boldsymbol{\theta}), \quad \text{with} & (11) \\ \mathcal{L}_t(\mathbf{Z}_t, \boldsymbol{\theta}) &\equiv \left\{ \sum_{j=1}^{N_t} \frac{(\tilde{\kappa}_{t,k_j,\tau_j} - \kappa(k_j, \tau_j, \mathbf{Z}_t, \boldsymbol{\theta}))^2}{N_t} + \lambda \frac{k_n}{N_t} \left(\widehat{V}_t^n - \xi_1(\mathbf{Z}_t) \right)^2 \right\}. \end{aligned}$$

The first part of the objective function is the L_2 -distance between observed and model-implied option prices. The second is a penalization term for the deviation of the model-implied spot volatility estimate from a direct, but noisy, measure of spot volatility obtained from high-frequency return data, \widehat{V}_t^n .¹³ We recall that $V_t = \xi_1(\mathbf{S}_t)$. The use of a noisy observation of the state vector (or part of it) in the estimation procedure is reminiscent of the FAVAR approach, proposed by Bernanke et al. (2005).¹⁴ Similarly to the FAVAR, the penalization in equation (11) serves to improve the identification of the latent factors, which, as we show later, may contribute to an improvement in efficiency.

To characterize the properties of our PLS estimator, it is convenient to concentrate out the factor realizations, i.e., we set,

$$\hat{\mathbf{S}}(\boldsymbol{\theta}) = \underset{\mathbf{Z} \in \mathcal{S}^T}{\operatorname{argmin}} \mathcal{L}(\mathbf{Z}, \boldsymbol{\theta}), \quad \boldsymbol{\theta} \in \boldsymbol{\Theta}, \quad \mathcal{S} \in \mathbb{R}^p, \quad (12)$$

¹¹Note that our option pricing model does contain a parametric assumption for the evolution of \mathbf{S}_t . However, this is for its behavior under the risk-neutral measure \mathbb{Q} only. We avoid specifying the \mathbb{P} dynamics of \mathbf{S}_t parametrically in this paper (this is an interesting extension that can be pursued in future work). As discussed in Andersen et al. (2015a), this has the distinct advantage of robustness with respect to the modeling of the risk premia, i.e., the $\mathbb{P} - \mathbb{Q}$ wedge.

¹²Note that, in Section 5, we will generalize the PLS estimator by allowing λ to be time-varying and random, which will lead to efficiency improvements.

¹³Recall from equation (10) that we standardized the estimator \widehat{V}_t^n by k_n . This is the reason why we multiply the penalization term in equation (11) by k_n . It is merely a matter of normalization.

¹⁴Bernanke et al. (2005) augment a VAR of economic variables with a noisy observation of a latent factor, related to the variables in the VAR system, to exploit auxiliary information regarding the factor realizations.

and the resulting estimator, then, takes the form,

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} \mathcal{L}(\hat{\mathbf{S}}(\boldsymbol{\theta}), \boldsymbol{\theta}), \quad \hat{\mathbf{S}} = \hat{\mathbf{S}}(\hat{\boldsymbol{\theta}}). \quad (13)$$

In the case when $N \rightarrow \infty$ and T is fixed, the asymptotic properties of the estimator $(\hat{\mathbf{S}}', \hat{\boldsymbol{\theta}})$ have been derived in Andersen et al. (2015a). In this section, we consider the case where both N and T increase asymptotically. We start with the consistency result.

Theorem 1. *Let Assumptions 1-6 hold. Assume $N \rightarrow \infty$, $T \rightarrow \infty$ with $T/N \rightarrow \gamma$ for some finite $\gamma \geq 0$, $\Delta_n \rightarrow 0$ and $k_n \rightarrow \infty$ with $\frac{k_n}{N} \rightarrow \varrho$ for some $\varrho \geq 0$. In addition let $\varpi \in \left(\frac{1}{4-\tau}, \frac{1}{2}\right)$ and $k_n \sqrt{\Delta_n} \rightarrow 0$. Then, we have,*

$$\sup_{t=1, \dots, T} \|\hat{\mathbf{S}}_t - \mathbf{S}_t\| \xrightarrow{\mathbb{P}} 0 \quad \text{and} \quad \hat{\boldsymbol{\theta}} \xrightarrow{\mathbb{P}} \boldsymbol{\theta}_0. \quad (14)$$

We note that consistency holds under very general conditions for the tuning parameters of the high-frequency volatility estimator, ϖ and k_n . Of course, these conditions are absent if the penalization term is left out of the objective function, i.e., when $\lambda = 0$.

3.2 Auxiliary Notation for the Limiting Distribution

We now turn towards the limit distribution of the parameter and factor estimates. As discussed in the introduction, the model parameters will be estimated at the rate \sqrt{NT} , but are subject to an asymptotic bias due to the incidental parameter problem, while the factor (or state vector) realizations converge at the slower \sqrt{N} rate. In addition, given the multiple and potentially interacting sources of unobserved heterogeneity, the PLS estimator clearly will not utilize the optimal weighting of the observations during estimation. As a consequence, the limiting covariance matrices in the asymptotic distribution take the common sandwich form, albeit with a number of inherent features of the system impacting the representation. To state the result in an interpretable format, we require additional notation, which we is introduced in the following.

In what follows, all partial derivatives of κ that appear below are assumed to exist, as implied by our assumptions in Section A.1 of the Online Appendix. The first-order derivatives with respect to the parameter and state vectors are given by the $q \times 1$ and $p \times 1$ vectors,

$$\mathbf{D}(k, \tau, \mathbf{Z}, \boldsymbol{\theta}) = \frac{\partial}{\partial \boldsymbol{\theta}} \kappa(k, \tau, \mathbf{Z}, \boldsymbol{\theta}), \quad \mathbf{G}(k, \tau, \mathbf{Z}, \boldsymbol{\theta}) = \frac{\partial}{\partial \mathbf{Z}} \kappa(k, \tau, \mathbf{Z}, \boldsymbol{\theta}),$$

respectively, and, similarly, their first and second derivatives with respect to the state vector \mathbf{Z} are denoted by $\mathbf{D}^{(S)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, $\mathbf{G}^{(S)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, $\mathbf{D}^{(SS)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ and $\mathbf{G}^{(SS)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ (of dimensions $q \times p$, $p \times p$, $q \times p^2$ and $p \times p^2$, respectively), and the corresponding derivatives with respect to $\boldsymbol{\theta}$ are denoted by $\mathbf{D}^{(\theta)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ and $\mathbf{G}^{(\theta)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ (of dimension $q \times q$ and $p \times q$).

To simplify notation, for each pair (k_j, τ_j) , we set $\tilde{\kappa}_{j,t} \equiv \tilde{\kappa}_{t,k_j,\tau_j}$, $\kappa_j(\mathbf{Z}, \boldsymbol{\theta}) \equiv \kappa(k_j, \tau_j, \mathbf{Z}, \boldsymbol{\theta})$, and simplify the notation for the partial derivatives of the latter analogously. Furthermore, when we evaluate $\kappa_j(\mathbf{Z}, \boldsymbol{\theta})$ at the true vector $(\mathbf{S}'_t, \boldsymbol{\theta}'_0)$, we denote it by $\kappa_{j,t,0}$, and the notation for the partial derivatives of $\kappa_j(\mathbf{Z}, \boldsymbol{\theta})$ is simplified in the same manner, if evaluated at $(\mathbf{S}'_t, \boldsymbol{\theta}'_0)$. Similarly, for the penalty term,

$$\mathbf{C}(\mathbf{Z}) \equiv \frac{\partial}{\partial \mathbf{Z}} \xi_1(\mathbf{Z}) \quad \text{and} \quad \mathbf{C}^{(S)}(\mathbf{Z}) \equiv \frac{\partial^2}{\partial \mathbf{Z} \mathbf{Z}'} \xi_1(\mathbf{Z})$$

denote the $p \times 1$ gradient vector and $p \times p$ higher-order gradient matrix, respectively, and we will use the shorthand notation $\mathbf{C}_{t,0} = \mathbf{C}(\mathbf{S}_t)$ and $\mathbf{C}_{t,0}^{(S)} = \mathbf{C}^{(S)}(\mathbf{S}_t)$.

The differential rate of convergence implies that the parameters can be treated as fixed and known, when deriving the limit distribution for the estimates of the factor realizations. As a result, the quantities determining the limiting conditional covariance matrix for the latter take a fairly conventional form. The main complication arises from the fact that the usual component stemming from the least squares criterion is supplemented by a contribution from the penalization term. The relevant matrices for the factor realizations are given in the following equations. First, define the $p \times p$ matrix,

$$\mathcal{I}_{S,t} = \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \quad (15)$$

which represents an outer product of gradient type term for the least squares criterion. Notice that the relative number of options at each tenor (π_t^{τ}) and the relative denseness, or clustering, of options across certain segments of the strike range ($\psi_{t,\tau}(k)$) enter explicitly in the expression. It is important to note that $\mathcal{I}_{S,t}$ is simply the in-fill asymptotic limit of an average of outer products of gradients, computed using the available values of the pair (k, τ) at time t , i.e., we have,

$$\mathcal{I}_{S,t} \approx \frac{1}{N_t} \sum_{j=1}^{N_t} \mathbf{G}(k_j, \tau_j, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{G}(k_j, \tau_j, \mathbf{S}_t, \boldsymbol{\theta}_0)', \quad (16)$$

and analogous intuition applies to the rest of the integrals (over k) appearing in this section.

A corresponding $p \times p$ matrix, reflecting the outer product of the gradient for the penalization term, is given by,

$$\mathcal{P}_{S,t} \equiv \frac{\lambda \varrho}{\varsigma_t} \times \mathbf{C}_{t,0} \mathbf{C}'_{t,0}, \quad (17)$$

where the relative number of options available on that date (ς_t) and the relative informativeness of the high-frequency volatility estimator relative to the option data (ϱ) enter the expression.

Finally, we define the $p \times p$ asymptotic covariance matrices $\boldsymbol{\Omega}_{S,t}$ and $\boldsymbol{\Psi}_{S,t}$, which further include the variance of the observation error for the option BSIV ($\phi_{t,k,\tau}$), respectively, the variance of the nonparametric volatility estimator ($2V_t^2$),

$$\boldsymbol{\Omega}_{S,t} = \frac{1}{\varsigma_t} \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \phi_{t,k,\tau} \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \quad \boldsymbol{\Psi}_{S,t} = 2 \frac{\lambda^2 \varrho}{\varsigma_t^2} V_t^2 \mathbf{C}_{t,0} \mathbf{C}'_{t,0}.$$

We turn next to the expressions required to capture the asymptotic behavior of the estimates for the model parameters, $\boldsymbol{\theta}_0$. They are more complex and involve interaction terms among the factors and parameters as well as the features of the system determining the magnitude of the asymptotic bias. We start by defining the $q \times q$ information matrix for the parameter vector,

$$\mathcal{I}_{\boldsymbol{\theta},t} = \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \mathbf{D}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{D}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \quad (18)$$

and the $p \times q$ cross information matrix,

$$\mathcal{I}_{S,\boldsymbol{\theta},t} = \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{D}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk. \quad (19)$$

In turn, these are used to define,

$$\mathbf{R}_{t,0} \equiv -(\mathcal{I}_{S,t} + \mathcal{P}_{S,t})^{-1} \mathcal{I}_{S,\boldsymbol{\theta},t} \quad \text{and} \quad \mathbf{U}_t(k, \tau, \mathbf{Z}, \boldsymbol{\theta}) \equiv \mathbf{D}(k, \tau, \mathbf{Z}, \boldsymbol{\theta}) + \mathbf{R}'_{t,0} \mathbf{G}(k, \tau, \mathbf{Z}, \boldsymbol{\theta}), \quad (20)$$

the first of which is the influence function of the gradient with respect $\boldsymbol{\theta}$, and $\mathbf{U}_t(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ partials out the effect of the state variables on $\mathbf{D}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$. This corresponds to demeaning the score of $\boldsymbol{\theta}$, in analogy to eliminating time and individual effects in large linear panel data models. As above, we denote the first and second derivative with respect to \mathbf{Z} as well as the first derivative with respect to $\boldsymbol{\theta}$ of $\mathbf{U}_t(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$ by $\mathbf{U}_t^{(S)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, $\mathbf{U}_t^{(SS)}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, and $\mathbf{U}_t^{(\boldsymbol{\theta})}(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$, respectively. In analogy with our simplified notation for a pair (k_j, τ_j) , we set $\mathbf{U}_{j,t}(\mathbf{Z}, \boldsymbol{\theta}) = \mathbf{U}_t(k_j, \tau_j, \mathbf{Z}, \boldsymbol{\theta})$, and, when evaluated at the true vector $(\mathbf{S}'_t, \boldsymbol{\theta}'_0)$, we further write $\mathbf{U}_{j,t,0}$. Similar simplifications are made to the notation for the partial derivatives of $\mathbf{U}_t(k, \tau, \mathbf{Z}, \boldsymbol{\theta})$. With this notation in hand, we define the $q \times q$ information matrix,

$$\mathcal{I}_{U,t} = \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \mathbf{U}_t(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{U}_t(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \quad (21)$$

the $q \times q$ matrices $\mathcal{P}_{\boldsymbol{\theta},t} = \mathbf{R}'_{t,0} \mathcal{P}_{S,t} \mathbf{R}_{t,0}$, $\boldsymbol{\Psi}_{\boldsymbol{\theta},t} = \mathbf{R}'_{t,0} \boldsymbol{\Psi}_{S,t} \mathbf{R}_{t,0}$, and,

$$\begin{aligned} \boldsymbol{\Omega}_{U,t} &= \frac{1}{s_t} \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \phi_{t,k,\tau} \mathbf{U}_t(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{U}_t(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \\ \boldsymbol{\Omega}_{S,U,t} &= \frac{1}{s_t} \sum_{\tau} \pi_t^{\tau} \int_{\underline{k}(t,\tau)}^{\bar{k}(t,\tau)} \frac{1}{\psi_{t,\tau}(k)} \phi_{t,k,\tau} \mathbf{G}(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0) \mathbf{U}_t(k, \tau, \mathbf{S}_t, \boldsymbol{\theta}_0)' dk, \end{aligned}$$

which all appear, directly or indirectly, in the limiting distributions developed subsequently. Finally, we let $\boldsymbol{\mathcal{X}}$ be generic notation for any of the matrices \mathcal{I}_U , $\mathcal{P}_{\boldsymbol{\theta}}$, $\boldsymbol{\Omega}_U$, and $\boldsymbol{\Psi}_{\boldsymbol{\theta}}$, which we define by,

$$\boldsymbol{\mathcal{X}} = \begin{cases} \text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \boldsymbol{\mathcal{X}}_t, & \text{if } T \rightarrow \infty, \\ \frac{1}{T} \sum_{t=1}^T \boldsymbol{\mathcal{X}}_t, & \text{if } T \text{ is fixed.} \end{cases} \quad (22)$$

These matrices are guaranteed to exist by our assumptions.¹⁵ Note that the convergence in probability result (22) may be proven by invoking an ergodic theorem. For example, if ς_t , π_t^\top , $\psi_{t,\tau}(k)$, $\underline{k}(t, \tau)$ and $\bar{k}(t, \tau)$ are all functions of \mathbf{S}_t , do not have any other source of randomness or dependence on time, and if the process \mathbf{S}_t is stationary and ergodic with $\mathbb{E}|\mathcal{X}_t| < \infty$, then equation (22) holds.

3.3 Limiting Distribution of the PLS Estimator

The representation of the asymptotic distribution for the parameters and factor realizations involves the notion of stable convergence (denoted by $\xrightarrow{\mathcal{L}-s}$). It means that the convergence holds jointly with any bounded random variable defined on the original probability space. This feature is fundamental in establishing the following limiting distribution result.

Theorem 2. *Let Assumptions 1-8 hold. Assume $N \rightarrow \infty$, $T \rightarrow \infty$ with $T/N \rightarrow \gamma$ for some finite $\gamma \geq 0$, $\Delta_n \rightarrow 0$ and $k_n \rightarrow \infty$ with $\frac{k_n}{N} \rightarrow \varrho$ for some finite $\varrho \geq 0$. In addition, let,*

$$\varpi \in \left(\frac{1}{2(2-r)}, \frac{1}{2} \right), \quad k_n \sqrt{\Delta_n} \rightarrow 0, \quad \sqrt{T k_n} \left(k_n \Delta_n \vee \Delta_n^{(2-r)\varpi} \right) \rightarrow 0, \quad \frac{\Delta_n^{(2-\frac{r}{2})\varpi-1}}{k_n} \rightarrow 0. \quad (23)$$

Then, we have,

$$\sqrt{TN} \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 - \mathbf{B}_{T,N} \right) \xrightarrow{\mathcal{L}-s} (\mathcal{I}_U + \mathcal{P}_\theta)^{-1} (\boldsymbol{\Omega}_U + \boldsymbol{\Psi}_\theta)^{1/2} \times \mathbf{E}_\theta, \quad (24)$$

where \mathbf{E}_θ is a $q \times 1$ standard Gaussian vector, defined on an extension of the original probability space and independent from \mathcal{F} , and the asymptotic bias is given by,

$$\mathbf{B}_{T,N} = (\mathcal{I}_U + \mathcal{P}_\theta)^{-1} \times \sum_{j=1}^2 \left(\mathbf{B}_{T,N,1}^{(j)} - \mathbf{B}_{T,N,2}^{(j)} - \mathbf{B}_{T,N,3}^{(j)} \right),$$

whose components are defined in Section A.3 of the Online Appendix. In addition, for arbitrary $t_1, \dots, t_h \in \mathbb{N}$, we have,

$$\sqrt{N} \begin{pmatrix} \hat{\mathbf{S}}_{t_1} - \mathbf{S}_{t_1} \\ \vdots \\ \hat{\mathbf{S}}_{t_h} - \mathbf{S}_{t_h} \end{pmatrix} \xrightarrow{\mathcal{L}-s} \text{diag} \begin{pmatrix} (\mathcal{I}_{S,t_1} + \mathcal{P}_{S,t_1})^{-1} (\boldsymbol{\Omega}_{S,t_1} + \boldsymbol{\Psi}_{S,t_1})^{1/2} \\ \vdots \\ (\mathcal{I}_{S,t_h} + \mathcal{P}_{S,t_h})^{-1} (\boldsymbol{\Omega}_{S,t_h} + \boldsymbol{\Psi}_{S,t_h})^{1/2} \end{pmatrix} \times \begin{pmatrix} \mathbf{E}_{S_{t_1}} \\ \vdots \\ \mathbf{E}_{S_{t_h}} \end{pmatrix}, \quad (25)$$

where $\mathbf{E}_{S_{t_1}}, \dots, \mathbf{E}_{S_{t_h}}$ is a sequence of standard Gaussian vectors, defined on an extension of the original probability space, with each of them independent of each other, of \mathbf{E}_θ , as well as of \mathcal{F} , and the convergence in equation (25) holds jointly with that in equation (24).

¹⁵Typical option data sets contain deterministic cycles in the configuration of the available tenors. This feature renders quantities such as $\mathcal{I}_{U,t}$ non-stationary. However, considering $\sum_{t \in \mathcal{T}_m} \mathcal{I}_{U,t}$, for \mathcal{T}_m denoting a period covering one calendar month, such non-stationarities vanish, and standard laws of large numbers suffice for validation of equation (22).

This limit result deserves a detailed discussion. First, the rate of convergence of the parameter vector is \sqrt{NT} , so both the increase in the cross-section and time dimension of the panel help in the recovery of the parameters. By contrast, the factor realizations are estimated at the rate \sqrt{N} . This is quite intuitive as increasing the time span of the panel does not provide extra information for the factor realization at a given point in time.¹⁶ The asymptotic limit distributions of the parameters and factor realizations are fundamentally different. For the former, it is Gaussian with a constant limiting variance, as is common in econometric applications. On the other hand, it is mixed Gaussian for the estimates of the factor realizations, i.e., it is Gaussian only when conditioned on the filtration of the original probability space, and the covariance matrix depends explicitly on the latter. That is, $\mathcal{I}_{S,t} + \mathcal{P}_{S,t}$ and $\Omega_{S,t} + \Psi_{S,t}$, which appear in the limiting distribution of the factor realizations in equation (25), are stochastic. This implies that the precision in recovering the factor realizations is time-varying and depends on unobserved random quantities associated with the cross-section of option data at that particular point in time. For example, periods of elevated volatility may be associated with a larger variance of the observation errors (captured by the function ϕ), and hence noisier recovery for the factor realizations.

We can contrast the limiting behavior above with the case of a panel with fixed time span, but asymptotically increasing cross-section. In this scenario, both the parameter vector and factor realizations have mixed Gaussian limiting distributions (see also Lemma 1 below). Furthermore, unlike here, the parameter and state vector estimates display asymptotic dependence in the fixed time span case. This dependence becomes asymptotically negligible when the time span is long, since the parameter vector is now recovered at a faster rate.

We note that, given the non-standard mixed Gaussian distribution of the factors, we need a stronger notion of convergence in equation (25) – stable convergence – than what is typically derived in econometric applications, in order to specify the joint asymptotic behavior of the parameters and factors. The stable convergence is also required to accommodate the penalization with the high-frequency volatility estimator and characterize jointly the asymptotic behavior of the option error and the error coming from the high-frequency volatility estimation. To the best of our knowledge, the joint asymptotic behavior of the factor and parameter estimates in the nontrivial case, where the \mathcal{F} -conditional limiting variance of the factors is stochastic, has not been derived in earlier work on large panels, with both N and T diverging. Such results are indispensable for deriving the properties of test statistics for joint model fit across a time period of interest, involving jointly the estimates for the factors, parameters and high-frequency volatility.¹⁷

¹⁶We note that this could occur if there were pathwise restrictions on the dynamic evolution of the factors. However, existing dynamic factor models do not feature such restrictions.

¹⁷We refer to Jacod and Shiryaev (2003) for additional details regarding the notion of stable convergence. In the case where

A further distinguishing feature of the long versus fixed span case is the presence of an asymptotic bias in the parameter estimation, the explicit form of which is given in the Online Appendix. This bias reflects the effect of uncertainty in the recovery of the factor realizations on the parameter estimation. Intuitively, although the errors in the estimation of the factor realizations average out over time, and thus do not slow down the recovery rate of the parameters, the nonlinearity of the model, both in parameters and the state vector, induces an asymptotic bias of order $1/N$ in the parameter estimates. Therefore, if $T \ll N$ ($\gamma = 0$ in Theorem 2), this bias does not affect the central limit theorem for the parameter vector, i.e., the bias term in the limit result (24) is of higher order. On the other hand, if $T \asymp N$, the asymptotic bias does have a non-negligible effect on the limit distribution of the parameter estimate. The form of the bias terms resembles those in the classical panel data literature, see e.g., Hahn and Newey (2004) and Hahn and Kuersteiner (2011), with the difference here arising from the stochastic nature of the terms, the infill asymptotic setting in the cross-section as well as the penalty term in the estimation due to the factor signal from the high-frequency data. Finally, we remark that the limit results in equations (24) and (25) are infeasible because the asymptotic variance and bias term are based on population quantities. In the next section, we show that they are straightforward to estimate consistently from the data, thus enabling feasible inference.

In terms of the asymptotic role of the penalization, we have two cases. When $k_n/N \rightarrow 0$, the penalization has no first-order asymptotic effect, while for $k_n/N \rightarrow \varrho > 0$, the penalization impacts the asymptotic distribution of the estimators for both the parameters and factor realizations. The penalty condition provides a restriction on the factor realization, namely that the model-implied estimate of volatility $\xi_1(\mathbf{S}_t)$ should be close to the nonparametric estimate of volatility V_t . The option panel allows us to recover the factor realizations at the rate \sqrt{N} , as implied by equation (25). On the other hand, the rate of convergence of the high-frequency volatility estimate is $\sqrt{k_n}$ (recall, k_n is the local window for the high-frequency data used to construct the volatility estimator). Therefore, when $k_n/N \rightarrow 0$, the option data are more efficient in recovering the volatility, rendering penalization during estimation suboptimal asymptotically.¹⁸ On the other hand, when $k_n/N \rightarrow \varrho > 0$, the nonparametric high-frequency volatility estimate and the option-based are equally efficient, in terms of rate of convergence, and hence the penalization may have an asymptotic effect. This explains why, in our objective function, the penalty term is multiplied by $\lambda \frac{k_n}{N_t}$. This ensures that the penalty

the penalization term does not have an impact on the limiting distribution, then we can strengthen the convergence results of Theorem 2 to almost sure convergence, conditional on the filtration generated by the latent factor \mathbf{S}_t . When the penalization term matters asymptotically, however, we cannot prove this stronger mode of convergence as, in general, we do not have control over the resulting conditional distribution (for example, \mathbf{X} and \mathbf{S} can be strongly dependent, accommodating the so-called “leverage effect” in asset prices). The stable convergence result in Theorem 2 is sufficient for deriving the feasible limiting theory below.

¹⁸Even though it is suboptimal (asymptotically) to use penalization within this scenario, its presence may greatly facilitate the identification of the state vector in finite samples.

becomes asymptotically negligible exactly when it is suboptimal to utilize high-frequency return data for estimation. Importantly, the user does not need to take a stand on which data source, options or high-frequency data, is more efficient for recovering spot volatility. Our PLS procedure automatically adapts to the scenario at hand. Nevertheless, the user still needs to pick the constant λ . In Section 5, we propose a data-dependent choice for this parameter, which boosts efficiency of the inference.

Above, we rule out the possibility that $k_n/N \rightarrow \infty$, i.e., that the high-frequency data are more efficient than the options (in terms of rate of convergence) in recovering spot volatility.¹⁹ This case would heighten the asymptotic role of the penalization term. In fact, we have distinct scenarios depending on the dimension of the state vector and the role of its components in determining spot volatility. For example, in the one-factor stochastic volatility model of Duffie et al. (2000), the state vector equals the level of (spot) volatility. Hence, when $k_n/N \rightarrow \infty$, the rate of convergence of the state vector will be the faster $\sqrt{k_n}$ in lieu of the \sqrt{N} implied by our limit result (25), and the asymptotic bias for the parameter vector in equation (24) vanishes. In a more general parametric model with multiple factors, one of which is spot volatility, all components of the state vector, except for spot volatility, will be estimated at the rate \sqrt{N} , while the stochastic volatility will be estimated at the faster $\sqrt{k_n}$ rate. Since $k_n/N \rightarrow \infty$ is implausible on empirical grounds, at least for assets with liquid option markets, we do not pursue this case any further.

The relative growth conditions (23) are necessary for our limiting result, although they are required only if the penalty term is present in the objective function, i.e., if $\lambda > 0$. In assessing the implications of these conditions, the following points are useful. First, to avoid unnecessary truncation, it is optimal to follow the convention in the literature and set ϖ very close to $1/2$. Second, the higher the k_n , the faster the rate of convergence of the high-frequency estimator. Therefore, we prefer the largest value of k_n for which condition (23) is satisfied. In the typical case of finite activity jumps, i.e., when $r = 0$ (the value of r is typically known from the specification of the risk-neutral model), and provided ϖ is set arbitrarily close to $1/2$, the binding constraints in condition (23) are $k_n\sqrt{\Delta_n} \rightarrow 0$ and $\sqrt{Tk_n}k_n\Delta_n \rightarrow 0$. Hence, the optimal feasible choice of k_n depends on the relative size of T and $1/\Delta_n$. If T is not very large, we can set $k_n \asymp 1/\Delta_n^{1/2-\iota}$ for some arbitrary small $\iota > 0$, guaranteeing the optimal feasible rate for estimating spot volatility from the high-frequency data. On the other hand, if T is of the same order as $1/\Delta_n$, the optimal feasible rate is $k_n \asymp 1/\Delta_n^{1/3-\iota}$. From a practical perspective, this means that for panels with a large time span relative to the sampling frequency for the returns, we need to pick a smaller sized local window for volatility estimation, slowing its rate of convergence, but reducing its bias. Given the discussion in the previous paragraph, it implies that the high-frequency volatility

¹⁹The “limiting” case of $k_n/N \rightarrow \infty$ implies that part of the state vector is observed by the econometrician. This scenario can be of interest in other applications of factor augmentation than the volatility measurement case considered here.

estimator becomes comparatively noisy (k_n/N_t is small), then the penalty term may have only an asymptotically vanishing impact on our estimation.

Finally, in cases where the observation errors become negligible, we may have $\mathbf{\Omega}_U$ as well as $\mathbf{\Omega}_{S,t}$ vanish for some or all t . Some of our limit results would then become degenerate. In fact, in the extreme case when options are observed with no error, we obtain exact error-free recovery for the parameter vector and the sequence of factor realizations. Given the large percentage bid-ask spreads observed on option markets, we deem such scenarios empirically implausible.

3.4 On the Estimation of the Factor Dynamics

The limit result (25) allows us to quantify the precision in recovering the factor realizations, which is useful in a number of contexts. For example, as discussed in Andersen et al. (2015a), asset pricing models impose pathwise restrictions on the relation between derivative securities and the underlying asset, and our limit result enables us to test these formally. Another important use of the extracted factors is for the estimation of the underlying factor dynamics. However, for that application, the limit result (25) is insufficient since we need to analyze the impact of the errors on the recovery of the factors over a long time span. The following theorem provides the requisite result.

Theorem 3. *Let Assumptions 1-8 hold. Assume $N \rightarrow \infty$, $T \rightarrow \infty$ with $T/N \rightarrow \gamma$ for some finite $\gamma \geq 0$, $\Delta_n \rightarrow 0$ and $k_n \rightarrow \infty$ with $\frac{k_n}{N} \rightarrow \varrho$ for some finite $\varrho \geq 0$ as well as condition (23) of Theorem 2. Then for any twice differentiable function $g : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ with bounded first and second derivatives, we have*

$$\frac{1}{\sqrt{T}} \sum_{t=2}^T \left(g(\hat{\mathbf{S}}_t, \hat{\mathbf{S}}_{t-1}) - g(\mathbf{S}_t, \mathbf{S}_{t-1}) \right) = o_p(1). \quad (26)$$

We note from Theorem 2 that $\hat{\mathbf{S}}_t - \mathbf{S}_t = O_p(1/\sqrt{N})$. In spite of that, the error in recovering the factor realizations on the statistic in equation (26) is only $o_p(1)$, even in the case where T and N increase at the same rate. The reason is that the error in recovering the factor realizations on different days across the sample effectively gets averaged out in the time series mean in equation (26), thus mitigating its impact on the statistic sufficiently to establish the limiting result.

Under standard stationarity conditions for \mathbf{S}_t and further assumptions guaranteeing a central limit theorem (e.g., mixing conditions), $\frac{1}{\sqrt{T}} \sum_{t=2}^T (g(\mathbf{S}_t, \mathbf{S}_{t-1}) - \mathbb{E}(g(\mathbf{S}_t, \mathbf{S}_{t-1})))$ will converge to a normal random variable. Theorem 3 ensures that the recovery of the state vector from the option panel (and high-frequency asset returns) has no asymptotic effect on the long span limit result, i.e., it is valid even after replacing the state vector realizations by their estimates.

If the dynamics for the state vector takes a parametric form and \mathbf{S}_t is a Markov process under the physical probability measure, an attractive choice of g , capable of generating asymptotically efficient

estimates for the parameters of the state vector dynamics (with a \sqrt{T} rate of convergence), is based on the conditional characteristic function, namely $g(\mathbf{S}_t, \mathbf{S}_{t-1}) = e^{i\mathbf{u}'\mathbf{S}_t} - \mathbb{E}_{t-1}(e^{i\mathbf{u}'\mathbf{S}_t})$.

The use of the extracted factors for estimation of the state vector dynamics, using the above theorem, is reminiscent of the implied state inference of Pan (2002) and Pastorello et al. (2003). These papers employ a finite set of options to invert for the state vector, and then conduct inference based on the option-implied factor realizations using method of moments or maximum likelihood techniques. The major difference between these procedures and our approach is that we allow for observation errors in the option prices, while letting the number of options increase asymptotically. Because of this difference in the asymptotic setup, we find that, in contrast to Pan (2002) and Pastorello et al. (2003), the estimation of the risk-neutral parameters has no first order asymptotic effect on the state vector extraction - the risk-neutral parameters are estimated at the rate \sqrt{NT} , while the parameters of the \mathbb{P} dynamics of \mathbf{S}_t are estimated at the slower rate of \sqrt{T} .

Second, in an asymptotic setup with $T \rightarrow \infty$ and $T/N \rightarrow 0$ (and without using noisy information for the factor realizations), Gagliardini and Gourieroux (2014) model the dynamics of the state vector parametrically. They perform joint estimation of the parameters for the nonlinear factor model ($\boldsymbol{\theta}$ in our setting) and for those driving the dynamics of the state vector (for which we make no parametric assumption) via maximum likelihood, and they establish asymptotic independence between the two sets of parameters (and rates of convergence of \sqrt{NT} and \sqrt{T} , respectively). However, their approach and associated limit theory exclude the challenging case of $T \asymp N$, which can be the pertinent asymptotic scheme for many practical applications.

Finally, Theorem 3 is related to results in the large panel data literature. In particular, Fernandez-Val and Lee (2013) derive the asymptotic distribution of functions of the parameters, data and fixed effects in large panels with random coefficients, c.f. their Corollary 1. As explained above, the fixed effects in classic panel settings can be viewed as counterparts to factor realizations in our dynamic nonlinear factor model, implying that the roles of N and T are reversed in the two asymptotic schemes.²⁰ Unlike Fernandez-Val and Lee (2013), however, we do not specify the dynamics of the factors in the model (again, fixed effects in their setting), and we only establish asymptotic negligibility of the difference $\frac{1}{\sqrt{T}} \sum_{t=2}^T (g(\hat{\mathbf{S}}_t, \hat{\mathbf{S}}_{t-1}) - g(\mathbf{S}_t, \mathbf{S}_{t-1}))$ instead of deriving the limit distribution for statistics of the form $\frac{1}{\sqrt{T}} \sum_{t=2}^T (g(\hat{\mathbf{S}}_t, \hat{\mathbf{S}}_{t-1}) - \mathbb{E}[g(\mathbf{S}_t, \mathbf{S}_{t-1})])$.²¹ Importantly, and in contrast with their framework, the present setting allows the factors to display time-series dependence.

²⁰To see this, note that the asymptotic bias of the fixed model parameters in our setup is of order $O_p(1/N)$, while the corresponding bias in Fernandez-Val and Lee (2013) is of order $O_p(1/T)$.

²¹Note that the bias correction term in Corollary 1 of Fernandez-Val and Lee (2013) is negligible when the rate of convergence of the statistic is \sqrt{N} (which corresponds to \sqrt{T} in our setup).

4 Unified Feasible Inference with Short and Long Time Span

This section proceeds by developing feasible inference techniques for the parameter vector and the factor realizations, which enable us to construct empirical confidence intervals and conduct testing. Our inference covers panels with a fixed or an increasing time span.

4.1 The Unified Inference Results

We begin with a lemma that restates the limit distribution of the PLS estimator for the fixed time span, originally derived in Andersen et al. (2015a), in a manner that facilitates comparison with our limit results for the case of a large time span.

Lemma 1. *Under the conditions of Theorem 2, but with fixed time span, i.e., T fixed, then,*

$$\sqrt{N} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \xrightarrow{\mathcal{L}^{-s}} \frac{1}{\sqrt{T}} (\mathcal{I}_U + \mathcal{P}_\theta)^{-1} (\boldsymbol{\Omega}_U + \boldsymbol{\Psi}_\theta)^{1/2} \times \mathbf{E}_\theta,$$

where \mathbf{E}_θ is a $q \times 1$ vector of standard Gaussian variables, independent of each other and of \mathcal{F} . Furthermore, we have,

$$\sqrt{N} \begin{pmatrix} \hat{\mathbf{S}}_1 - \mathbf{S}_1 \\ \vdots \\ \hat{\mathbf{S}}_T - \mathbf{S}_T \end{pmatrix} \xrightarrow{\mathcal{L}^{-s}} (\mathbf{AVAR}_S)^{1/2} \times \begin{pmatrix} \mathbf{E}_{S_1} \\ \vdots \\ \mathbf{E}_{S_T} \end{pmatrix}, \quad \mathbf{AVAR}_S = \text{diag} \begin{pmatrix} \mathbf{AVAR}_{S_1} \\ \vdots \\ \mathbf{AVAR}_{S_T} \end{pmatrix} + \mathbf{ACOV}_S,$$

where \mathbf{E}_{S_t} , $t = 1, \dots, T$ is a sequence of standard Gaussian vectors, defined on an extension of the original probability space, independent of \mathcal{F} , but dependent on each other and on \mathbf{E}_θ . Moreover, the diagonal elements of asymptotic covariance matrix decomposes as $\mathbf{AVAR}_{S_t} = \sum_{i=1}^4 \mathbf{AVAR}_{S_t,i}$ where,

$$\mathbf{AVAR}_{S_t,1} = (\mathcal{I}_{S,t} + \mathcal{P}_{S,t})^{-1} (\boldsymbol{\Omega}_{S,t} + \boldsymbol{\Psi}_{S,t}) (\mathcal{I}_{S,t} + \mathcal{P}_{S,t})^{-1},$$

$$\mathbf{AVAR}_{S_t,2} = \mathbf{R}_{t,0} \mathbf{AVAR}_\theta \mathbf{R}'_{t,0} / T, \quad \mathbf{AVAR}_\theta = (\mathcal{I}_U + \mathcal{P}_\theta)^{-1} (\boldsymbol{\Omega}_U + \boldsymbol{\Psi}_\theta) (\mathcal{I}_U + \mathcal{P}_\theta)^{-1}$$

$$\mathbf{AVAR}_{S_t,3} = (\mathcal{I}_{S,t} + \mathcal{P}_{S,t})^{-1} (\boldsymbol{\Omega}_{S,U,t} + \boldsymbol{\Psi}_{S,t} \mathbf{R}_{t,0}) (\mathcal{I}_U + \mathcal{P}_\theta)^{-1} \mathbf{R}'_{t,0} / T, \quad \mathbf{AVAR}_{S_t,4} = \mathbf{AVAR}'_{S_t,3}.$$

Finally, the matrix containing the off-diagonal elements is defined as,

$$\mathbf{ACOV}_S = \begin{pmatrix} \mathbf{0}_{p \times p} & \mathbf{R}_{1,0} \mathbf{AVAR}_\theta \mathbf{R}'_{2,0} / T & \dots & \mathbf{R}_{1,0} \mathbf{AVAR}_\theta \mathbf{R}'_{T,0} / T \\ \mathbf{R}_{2,0} \mathbf{AVAR}_\theta \mathbf{R}'_{1,0} / T & \mathbf{0}_{p \times p} & & \vdots \\ \vdots & & \ddots & \\ \mathbf{R}_{T,0} \mathbf{AVAR}_\theta \mathbf{R}'_{1,0} / T & \dots & & \mathbf{0}_{p \times p} \end{pmatrix}.$$

We note that the quantities \mathcal{I}_U , \mathcal{P}_θ , $\boldsymbol{\Omega}_U$ and $\boldsymbol{\Psi}_\theta$ are random, as T is fixed (recall equation (22)). Therefore, the limit distribution of the parameter vector is now mixed Gaussian and has no asymptotic

bias. Note also that the expression for the asymptotic variance of the parameter vector looks identical to the one in equation (24). However, when T is fixed, the asymptotic variance is a random quantity. Only in the asymptotic ($T \rightarrow \infty$) limit will the latter coincide with the one for the long time span. On the other hand, the limiting variance of the factor realizations differs from the long span asymptotic case. In particular, the \mathcal{F} -conditional limiting variance in equation (25) corresponds to $\mathbf{AVAR}_{S_t,1}$ in Lemma 1 above, while the additional terms $\{\mathbf{AVAR}_{S_t,j}\}_{j=2,3,4}$ reflect the asymptotic role of the parameter estimation on the factor extraction. Note that $\{\mathbf{AVAR}_{S_t,j}\}_{j=2,3,4}$ are all of order $O(1/T)$ and, as a result, they become asymptotically negligible when T increases asymptotically, as detailed in the previous section. Likewise, \mathbf{ACOV}_S in Lemma 1 reflects the asymptotic dependence in the recovery of the factor realizations at different points in time. This dependence is due to parameter uncertainty and, since it is of the order $O(1/T)$, negligible in the long time span scenario.

Despite the distinct differences in the limiting behavior of the PLS estimator, we are able to devise a feasible central limit theory that covers both the fixed and long time span setting. The formal result is provided by the following theorem, where $\bar{N} = \frac{1}{T} \sum_{t=1}^T N_t$ denotes the average number of options per cross-section.²²

Theorem 4. *Under the conditions of Theorem 2 but with T either fixed or $T \rightarrow \infty$ with $T/N \rightarrow \gamma$ for some finite $\gamma \geq 0$,*

$$\sqrt{T\bar{N}} \left((\hat{\mathcal{I}}_U + \hat{\mathcal{P}}_\theta)^{-1} (\hat{\Omega}_U + \hat{\Psi}_\theta) (\hat{\mathcal{I}}_U + \hat{\mathcal{P}}_\theta)^{-1} \right)^{-1/2} (\hat{\theta} - \theta_0 - \hat{\mathcal{B}}_{T,N}) \xrightarrow{\mathcal{L}} \mathbf{E}_\theta,$$

where \mathbf{E}_θ is a $q \times 1$ vector of standard Gaussian variables, independent of each other, and the estimators $\hat{\mathcal{I}}_U$, $\hat{\mathcal{P}}_\theta$, $\hat{\Omega}_U$, $\hat{\Psi}_\theta$ and $\hat{\mathcal{B}}_{T,N}$ are defined in Section A.5 of the Online Appendix. Furthermore, for the state vector,

$$\sqrt{\bar{N}} \left(\widehat{\mathbf{AVAR}}_{S_t} \right)^{-1/2} (\hat{S}_t - S_t) \xrightarrow{\mathcal{L}} \mathbf{E}_{S_t},$$

where \mathbf{E}_{S_t} is a sequence of $p \times 1$ vectors of standard Gaussian variables. The estimator of the asymptotic covariance matrix, $\widehat{\mathbf{AVAR}}_{S_t}$ is, similarly, deferred to the Online Appendix.

4.2 Discussion of the Unified Inference Result

Unlike Theorem 2 and Lemma 1, the limiting variables in the convergence results of Theorem 4 are standard normal vectors. This is because we utilize consistent estimators for the asymptotic variance and bias terms in Theorem 4. Based on this result, inference is feasible regardless of whether the time span is fixed or asymptotically increasing. This is important in practice, as the user need not

²²Thus, \bar{N} can be viewed as the feasible counterpart to N , recall equation (5). We note that one does not need \bar{N} for feasible implementation of the limit results of Theorem 4. This can be easily checked by inspection of the formulas for the asymptotic biases and variances in the Online Appendix.

take a stand on whether the time span of the panel renders Lemma 1 or Theorem 2 more suitable for the case at hand. In designing the feasible inference theory, we accommodate features of the limiting distributions of either asymptotic scheme. Several key properties of the limit results for the fixed and increasing time span case are critical in this regard. First, the stable convergence for the factor realizations and parameter vector in the fixed span case enables us to apply a Slutsky type result, even if the probability limits of quantities such as $\hat{\mathcal{I}}_U$ may be random. Second, the estimates for the asymptotic variance of the parameter vector in the long span case and the \mathcal{F} -conditional asymptotic variance in the fixed time span case are identical. Third, the estimate of the asymptotic bias is asymptotically negligible for a fixed time span or when the latter increases at a rate slower than that of the cross-sectional dimension of the panel. That is, $\hat{\mathcal{B}}_{T,N} = O_p(1/N)$, and thus becomes negligible whenever $T/N \rightarrow 0$. Finally, the estimate for the asymptotic dependence between the estimates of the parameter vector and factor realizations is of higher asymptotic order for the long time span case.

We finish this section by discussing the dependence structure imposed on the observation errors and its role in deriving the asymptotic results. In equation (9), we assume $\mathcal{F}^{(0)}$ -conditional independence of the errors, both in the cross-section and over time. Nevertheless, the errors may be dependent through the information in $\mathcal{F}^{(0)}$. Moreover, the (conditional) independence of the errors in equation (9) is much stronger than what we actually need. For example, our results, and in particular Theorem 4, will continue to hold if the observation errors are of the following form,

$$\epsilon_{t,k,\tau} = \sqrt{\phi_{t,k,\tau}} \tilde{\epsilon}_t z_{t,k,\tau}, \quad (27)$$

where $\phi_{t,k,\tau}$ are $\mathcal{F}^{(0)}$ -adapted nonnegative random variables, and the sequences $\{\tilde{\epsilon}_t\}_t$ and $\{z_{t,k,\tau}\}_{(t,k,\tau)}$ are defined on $\mathcal{F}^{(1)}$, and are i.i.d., independent of each other as well as of $\mathcal{F}^{(0)}$. In the above specification, and due to the inclusion of $\tilde{\epsilon}_t$, the errors $\epsilon_{t,k,\tau}$ are no longer $\mathcal{F}^{(0)}$ -conditionally independent in the cross-section. More generally, our results will continue to hold if we replace $\mathcal{F}^{(0)}$ -conditional independence of the errors in (9) with the following condition on (cross-)products of errors:

$$\begin{aligned} \mathbb{E} \left(\epsilon_{t_1,k_1,\tau_1} \epsilon_{t_2,k_2,\tau_2} \middle| \mathcal{F}^{(0)} \right) &= \mathbb{E} \left(\epsilon_{t_1,k_1,\tau_1}^3 \epsilon_{t_2,k_2,\tau_2} \middle| \mathcal{F}^{(0)} \right) = \mathbb{E} \left(\epsilon_{t_1,k_1,\tau_1}^2 \epsilon_{t_2,k_2,\tau_2} \epsilon_{t_3,k_3,\tau_3} \middle| \mathcal{F}^{(0)} \right) \\ &= \mathbb{E} \left(\epsilon_{t_1,k_1,\tau_1} \epsilon_{t_2,k_2,\tau_2} \epsilon_{t_3,k_3,\tau_3} \epsilon_{t_4,k_4,\tau_4} \middle| \mathcal{F}^{(0)} \right) = 0, \end{aligned} \quad (28)$$

where $(t_i, k_i, \tau_i)_{i=1}^4$ differ from each other. Condition (28) replaces $\mathcal{F}^{(0)}$ -conditional independence of the errors with a weaker one imposing $\mathcal{F}^{(0)}$ -conditional uncorrelatedness of the errors and certain products of them. In particular, this allows for dependence in higher-order $\mathcal{F}^{(0)}$ -conditional moments of the observation errors, exactly as in the example (27).

Our analysis can be further extended to cover situations in which there are $\mathcal{F}^{(0)}$ -conditional dependence in the first moments of the errors. To accommodate such extensions, however, we will need to

modify the estimates of the asymptotic variances and biases in Theorem 4 in ways that will account for this dependence. Specifically, we can extend the analysis to cover situations where the errors display time series dependence, e.g., by adopting an error specification as in equation (27), but with $\tilde{\epsilon}_t$ being allowed to fluctuate systematically over time. In this case, dependence in the limiting distribution for the estimates of the factor realizations will arise, even in the case where the time span increases to infinity. Such time series dependence in the errors may be accounted for by changing our estimators of the asymptotic variances (and of the bias terms) to long-run variance type estimators.

We can also extend our analysis to handle cross-sectional dependence in the observation errors. For this to be feasible, we need both convergence in probability of, and a CLT for, cross-sectional averages of the form $\frac{1}{N_t} \sum_{j=1}^{N_t} \epsilon_{j,t} m_{j,t}$, where $m_{j,t}$ is some $\mathcal{F}^{(0)}$ -adapted random variable. This implies that we cannot have very strong dependence in the cross-section of the observation errors, and we certainly cannot accommodate scenarios in which the observations in a certain strike domain are perfectly dependent.²³ However, this case is akin to model misspecification and, arguably, such “errors” should, in fact, be part of the option pricing model. On the other hand, we can extend the analysis to handle weak cross-sectional dependence in the errors. For example, we can accommodate a specification for the error of the form $\epsilon_{t,k,\tau} = \sqrt{\phi_{t,k,\tau}} z_{t,k,\tau}$, where $\phi_{t,k,\tau}$ are $\mathcal{F}^{(0)}$ -adapted nonnegative random variables and for each $t \in \mathbb{N}$, the sequences $\{z_{j,t}\}_{j=1,\dots,N_t}$ are drawn from a stationary mixing process, independently across time. As with time series dependence in the errors, we need to adjust the estimates of the asymptotic variances and biases to ones of long-run types in order to account for the cross-sectional dependence in the errors. We note that this asymptotic setup implies that the dependence between observation errors for fixed strikes and tenors vanishes asymptotically. Of course, in practice, we work with a given panel of option observations and the limit theory is derived simply to proxy the distributional properties of the estimators based on this option panel.

5 Weighted Penalized Least Squares

We now propose a weighted counterpart to the PLS estimator, which provides efficiency gains and enables us to choose suitable values for λ , period-by-period (recall equation (11)). In analogy to the reasoning for the classical weighted least squares estimator, we enhance efficiency by weighting the options according to the variance of their observation error. In our case, this implies weighting each option price by an estimate of the \mathcal{F} -conditional variance $\phi_{t,k,\tau}$. While this is feasible, we simplify by assigning all options on a given day the identical weight. This may result in loss of efficiency due to potential heteroskedasticity in the strike and tenor dimension of the option error variance at a given

²³By contrast, since identification and inference for the parameters and the factor realization on a given day can be achieved from a single cross-section of options, we can handle situations with perfect time-series dependence in the errors.

point in time. Nevertheless, it still provides non-trivial efficiency improvements due to the pronounced heteroskedasticity of the \mathcal{F} -conditional option error variance over time.

Similarly, for efficiency improvements, we should weight the penalization term period-by-period. Unlike the option observation error, whose \mathcal{F} -conditional variance may be estimated by a preliminary PLS estimator, the error of the nonparametric volatility estimator has an unknown \mathcal{F} -conditional variance, which cannot be estimated without further assumptions. However, we can instead weight the penalization by an estimate for the \mathcal{F} -conditional asymptotic variance of the nonparametric volatility estimator. The latter is given by $2V_t^2$, and, as a result, it is trivial to construct a consistent estimator for it. Intuitively, since the asymptotic distribution of the nonparametric variance estimator is \mathcal{F} -conditionally Gaussian with a \mathcal{F} -conditional variance of $2V_t^2$, this is asymptotically equivalent to weighting with the \mathcal{F} -conditional variance of the variance estimator.²⁴

More formally, we define $\mathcal{L}^w(\mathbf{Z}, \boldsymbol{\theta}) \equiv \sum_{t=1}^T \mathcal{L}_t^w(\mathbf{Z}_t, \boldsymbol{\theta})$ with

$$\mathcal{L}_t^w(\mathbf{Z}_t, \boldsymbol{\theta}) \equiv \left\{ \frac{1}{N_t} \sum_{j=1}^{N_t} \frac{(\tilde{\kappa}_{t,k_j,\tau_j} - \kappa(k_j, \tau_j, \mathbf{Z}_t, \boldsymbol{\theta}))^2}{w(\hat{\phi}_t)} + \frac{k_n}{N_t} \frac{(\hat{V}_t^n - \xi_1(\mathbf{Z}_t))^2}{w(2(\hat{V}_t^n)^2)} \right\}, \quad (29)$$

where

$$\hat{\phi}_t = \frac{1}{N_t} \sum_{j=1}^{N_t} \left(\tilde{\kappa}_{t,k_j,\tau_j} - \kappa(k_j, \tau_j, \hat{\mathbf{S}}_t, \hat{\boldsymbol{\theta}}) \right)^2, \quad t = 1, \dots, T, \quad (30)$$

and $w(x)$ is a twice continuously differentiable function on \mathbb{R}_+ with bounded first and second derivatives. Moreover, $w(x) \geq \epsilon$ for some $\epsilon > 0$. An example of such a function is a smooth approximation of $x \vee \epsilon$. As expected, the above estimator $\hat{\phi}_t$ is consistent for the cross-sectional average \mathcal{F} -conditional option variance at time t (generally a random number). Ideally, we would pick the weight function $w(x) = x$, but we rule this out for technical reasons.²⁵ Even so, with the class of weight functions w , described above, we may get arbitrary close to the identity function.

Given the above definitions, the resulting weighted penalized least squares (WPLS) estimator is defined as

$$\hat{\mathbf{S}}^w(\boldsymbol{\theta}) = \operatorname{argmin}_{\mathbf{Z} \in \mathcal{S}^T} \mathcal{L}^w(\mathbf{Z}, \boldsymbol{\theta}), \quad \boldsymbol{\theta} \in \boldsymbol{\Theta}, \quad (31)$$

and

$$\hat{\boldsymbol{\theta}}^w = \operatorname{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \mathcal{L}^w(\hat{\mathbf{S}}^w(\boldsymbol{\theta}), \boldsymbol{\theta}), \quad \hat{\mathbf{S}}^w = \hat{\mathbf{S}}^w(\hat{\boldsymbol{\theta}}^w). \quad (32)$$

In the next theorem, we characterize its limiting distribution in the long span case, with similar extensions for the feasible central limit theorem readily applying to this scenario as well.

²⁴Again, the stable convergence is critical for proving the limit effect from the weighting.

²⁵However, under stronger conditions on the boundedness of the variances $\phi_{t,k,\tau}$ and V_t from below, we may accommodate the case $w(x) = x$.

Theorem 5. *Suppose the conditions of Theorem 2 hold. Then, convergence results for $\hat{\boldsymbol{\theta}}^w$ and $(\hat{\mathbf{S}}_{t_1}^w, \dots, \hat{\mathbf{S}}_{t_k}^w)$ similar to those in (24) and (25) apply, as long as the quantities for their asymptotic bias and variances are replaced with their weighted analogues, defined in Section A.6 of the Online Appendix.*

We note that, for our WPLS estimator, we do not impose any assumptions regarding the type of heteroskedasticity in the option observation error. This is unlike the classical weighted least squares, see, e.g., White (1982) and Robinson (1987), where heteroskedasticity is modeled as either a known or unknown function of observables. We do assume, however, that the \mathcal{F} -conditional option error variance is a smooth function of log-moneyness k . Then, our asymptotic setup of increasingly denser observations, across a given strike range for a specific point in time and tenor, generates consistent estimates of the \mathcal{F} -conditional option error variance by local (in moneyness) cross-sectional averaging.

The theorem above is stated for a general weight function w . It is optimal to set w close to the identity function. If the option error variance $\phi_{t,k,\tau}$ is independent of k and τ , while $\phi_{t,k,\tau}$ and V_t are bounded from below and the option panel is balanced, then we obtain the optimally weighted penalized least squares. The limit results are now greatly simplified and take a familiar form, as evident from the corollary below.

Corollary 1. *Suppose the conditions of Theorem 5 hold and, in addition,*

$$\inf_{t \in \mathbb{R}_+} \inf_{\tau \in \mathcal{T}_t} \inf_{k \in [\underline{k}(t,\tau), \bar{k}(t,\tau)]} \phi_{t,k,\tau} > \epsilon, \quad \text{and} \quad \inf_{t \in \mathbb{R}_+} V_t > \epsilon \quad \text{for some finite } \epsilon > 0, \quad \phi_{t,k,\tau} = \phi_t.$$

Finally, suppose $\varsigma_t = 1$ for all t . Then, if $w(x) = x$ for $x > \epsilon$, the asymptotic distribution of $\hat{\boldsymbol{\theta}}^w$ and $\hat{\mathbf{S}}^w$ is characterized as follows.

$$\sqrt{TN} \left(\hat{\boldsymbol{\theta}}^w - \boldsymbol{\theta}_0 - \mathbf{B}_{T,N}^w \right) \xrightarrow{\mathcal{L}^{-s}} (\mathcal{I}_U^w + \mathcal{P}_\theta^w)^{-1/2} \times \mathbf{E}_\theta, \quad (33)$$

where \mathbf{E}_θ is a $q \times 1$ standard Gaussian vector, defined on an extension of the original probability space and independent from \mathcal{F} , and the asymptotic bias has the form,

$$\mathbf{B}_{T,N}^w = -(\mathcal{I}_U^w + \mathcal{P}_\theta^w)^{-1} \left(\mathbf{B}_{T,N,3}^{(1,w)} + \mathbf{B}_{T,N,3}^{(2,w)} \right).$$

In addition, for arbitrary $t_1, \dots, t_h \in \mathbb{N}$, we have,

$$\sqrt{N} \begin{pmatrix} \hat{\mathbf{S}}_{t_1}^w - \mathbf{S}_{t_1} \\ \vdots \\ \hat{\mathbf{S}}_{t_h}^w - \mathbf{S}_{t_h} \end{pmatrix} \xrightarrow{\mathcal{L}^{-s}} \text{diag} \begin{pmatrix} (\mathcal{I}_{S,t_1}^w + \mathcal{P}_{S,t_1}^w)^{-1/2} \\ \vdots \\ (\mathcal{I}_{S,t_h}^w + \mathcal{P}_{S,t_h}^w)^{-1/2} \end{pmatrix} \times \begin{pmatrix} \mathbf{E}_{\mathbf{S}_{t_1}} \\ \vdots \\ \mathbf{E}_{\mathbf{S}_{t_h}} \end{pmatrix}, \quad (34)$$

where $\mathbf{E}_{S_{t_1}}, \dots, \mathbf{E}_{S_{t_h}}$ is a sequence of standard Gaussian vectors, defined on an extension of the original probability space, with each of them independent of each other, of \mathbf{E}_θ , as well as of \mathcal{F} , and the convergence in (34) holds jointly with that in (33). Finally, all the quantities describing the two asymptotic distributions, \mathcal{I}_U^w , \mathcal{P}_θ^w , \mathcal{I}_{S,t_h}^w , \mathcal{P}_{S,t_h}^w , and $\mathcal{B}_{T,N}^w$, are defined in Section A.6 of the Online Appendix.

One convenient aspect of the WPLS estimator is that we need not choose the weight assigned to the penalty term in the objective function. In particular, the tuning parameter λ in the original PLS criterion is replaced by the time-varying ratio of the option variance and the asymptotic variance of the nonparametric variance estimator. As a result, the WPLS estimator alters the objective function period-to-period, adapting to the relative precision of the option pricing and high-frequency volatility estimator. From a practical point of view, this is attractive since the precision of these measures vary substantially over time. The WPLS procedure ensures that the weight of the two components in the objective function reflects their relative information content.

6 Numerical Experiments

We now undertake numerical experiments to assess the practical relevance of our unified inference techniques. For this purpose, we rely on a representation that nests many existing models and captures the complex dynamic features of option panels quite well. It takes the form of a two-factor jump-diffusive stochastic volatility model for the risk-neutral distribution,

$$\begin{aligned} \frac{dX_t}{X_{t-}} &= (r_t - \delta_t) dt + \sqrt{V_{1,t}} dW_{1,t}^{\mathbb{Q}} + \sqrt{V_{2,t}} dW_{2,t}^{\mathbb{Q}} + \int_{\mathbb{R}^2} (e^x - 1) \tilde{\mu}^{\mathbb{Q}}(dt, dx), \\ dV_{1,t} &= \kappa_1 (\bar{v}_1 - V_{1,t}) dt + \sigma_1 \sqrt{V_{1,t}} dB_{1,t}^{\mathbb{Q}} + \mu_v \int_{\mathbb{R}^2} x^2 1_{\{x < 0\}} \mu^{\mathbb{Q}}(dt, dx), \\ dV_{2,t} &= \kappa_2 (\bar{v}_2 - V_{2,t}) dt + \sigma_2 \sqrt{V_{2,t}} dB_{2,t}^{\mathbb{Q}}, \end{aligned} \quad (35)$$

where $(W_{1,t}^{\mathbb{Q}}, W_{2,t}^{\mathbb{Q}}, B_{1,t}^{\mathbb{Q}}, B_{2,t}^{\mathbb{Q}})$ is a four-dimensional Brownian motion with $W_{1,t}^{\mathbb{Q}} \perp\!\!\!\perp W_{2,t}^{\mathbb{Q}}$, $W_{1,t}^{\mathbb{Q}} \perp\!\!\!\perp B_{2,t}^{\mathbb{Q}}$, and $W_{2,t}^{\mathbb{Q}} \perp\!\!\!\perp B_{1,t}^{\mathbb{Q}}$, while $\text{corr}(W_{1,t}^{\mathbb{Q}}, B_{1,t}^{\mathbb{Q}}) = \rho_{d,1}$ and $\text{corr}(W_{2,t}^{\mathbb{Q}}, B_{2,t}^{\mathbb{Q}}) = \rho_{d,2}$. The risk-neutral compensator for the jump measure is given by

$$\begin{aligned} \nu_t^{\mathbb{Q}}(dx, dy) &= \left\{ \left(c^-(t) 1_{\{x < 0\}} \lambda_- e^{-\lambda_- |x|} + c^+(t) 1_{\{x > 0\}} \lambda_+ e^{-\lambda_+ x} \right) \right\} dx, \\ c^-(t) &= c_1^- V_{1,t-} + c_2^- V_{2,t-}, \quad c^+(t) = c_0^+ + c_1^+ V_{1,t-} + c_2^+ V_{2,t-}. \end{aligned}$$

We exploit the parameter estimates for the model from Andersen et al. (2015b), reproduced in Table 1. Of course, these are extracted from actual option panels, thus reflecting the risk-neutral dynamics, and not necessarily the actual statistical (or physical) probability measure. In fact, option prices are known to embody sizable risk premiums, so to generate the actual evolution of the simulated system, we

modify the above specification of the volatility processes to more closely resemble the actual observed volatility dynamics, i.e., we account for the risk premiums embedded in the risk-neutral parameter estimates. In particular, we increase κ_2 to 0.69, implying that a shock to $V_{2,t}$ has a (shorter) half-life of one year, and we limit the magnitude of the jumps in V_t to not exceed 0.40.

Table 1: Parameter Values for model (35) used in Numerical Experiments

| Parameter | Value | Parameter | Value | Parameter | Value |
|-------------|--------|------------|--------|-------------|--------|
| ρ_1 | -0.762 | κ_2 | 0.169 | c_2^+ | 52.795 |
| \bar{v}_1 | 0.004 | σ_2 | 0.129 | λ_- | 16.943 |
| κ_1 | 12.831 | c_0^+ | 2.315 | λ_+ | 51.818 |
| σ_1 | 0.247 | c_1^- | 69.208 | μ_v | 6.262 |
| ρ_2 | -0.945 | c_1^+ | 13.161 | | |
| \bar{v}_2 | 0.059 | c_2^- | 97.656 | | |

In terms of the option observation scheme, we follow the Monte Carlo setup of Andersen et al. (2015a), but explore the impact of the time dimension of the panel through experiments with different time spans. Specifically, consistent with standard practice, we sample option data every fifth day (time is measured in business days). At each observation date, we obtain option prices at four separate maturities: $\tau = 10$, $\tau = 45$, $\tau = 120$ and $\tau = 252$ days. For each tenor, we have 50 OTM option prices for an equispaced log-moneyness grid, covering the range $[-4, 1] \cdot \sigma\sqrt{\tau}$, where σ is the current at-the-money BSIV. This implies a time-varying coverage in terms of moneyness, depending on the level of market volatility, roughly mimicking the characteristics of actual option panels. For the option error, we assume $\epsilon_{t,k,\tau} = \sigma_{t,k,\tau} Z_{t,k,\tau}$, where $Z_{t,k,\tau}$ are standard normal variables, independent across time, moneyness and tenor, while $\sigma_{t,k,\tau} = 0.5 \psi_k / Q_{0.995}$. For the latter, ψ_k represents the estimate of the relative bid-ask spread from a kernel regression on actual S&P 500 index options, as reported in Andersen et al. (2015a), and $Q_{0.995}$ denotes the 0.995-quantile of the standard normal distribution. This noise structure allows for significant fluctuations in the (conditional) noise variance as a function of both the level of volatility and moneyness.²⁶

Overall, our simulations involve $N = 200$ options per day, and we experiment with a span ranging from 1 to 20 years. Since we sample once per week, this generates panels of about 50 to 1,000 cross

²⁶In the absence of parametric assumptions on the actual factor dynamics (as in our theoretical treatment), the Gaussian specification of the option errors implies that the optimally weighted least squares estimator constitutes the maximum likelihood estimator.

sections, each containing 200 option observations. Since the cost in time for numerical evaluation of the option prices is nontrivial, we do not perform estimation of the parameters and the factors on the basis of the simulated option data. Instead, in what follows we report our estimates for the asymptotic variance and the asymptotic bias of the parameters and the factors at their true values.

6.1 Estimation without Penalization

We start with a baseline case where we do not penalize with a high-frequency volatility estimator and abstain from weighting the option observations.

Let us first consider the effect of the time span on the recovery of the factor realizations. Figure 1 plots the estimate for the standard error associated with the recovery of V_1 and V_2 . We consider observations recorded in the last week of the first year of the sample, as this period is present for all time spans explored in our experiments. For ease of interpretation, we report the relative standard error, obtained by dividing the standard error estimate by the true value of the factor (which is random). From Theorem 2 and Lemma 1, we recall that, regardless of the time span, the rate of convergence of the factor extraction is \sqrt{N} . Moreover, we reiterate that we keep N fixed in our experiments (mimicking what occurs in practice) and that the limit distribution is mixed Gaussian. These features are reflected in Figure 1. Indeed, the interquantile range of the relative standard error does not shrink with the length of the span. This is due to the fact that the limiting variance of the error in the factor recovery depends on the realized (random) path of the state vector. In addition, the error in recovering the factor realizations tends to be higher for the shorter time spans. This reflects the effect of parameter uncertainty captured by the terms $\{\mathbf{AVAR}_{S_i,j}\}_{j=2,3,4}$ in Lemma 1. As the time span grows, this effect diminishes, and the quantiles for the relative standard error stabilize and become insensitive to any further lengthening of the span. Interestingly, this occurs at different horizons for the two factors. For V_1 , the “stabilization” of the standard error around its long span value is complete after about 6-7 years. For V_2 , however, this occurs only after 15-16 years. The reason is that the second factor is far more persistent than the first and, hence, it is harder to achieve a proper separation between its realization and long-run parameter-driven properties. Another noteworthy discrepancy is the relatively larger error in the recovery of V_1 and the significant heterogeneity in the precision across samples. This is largely due to the erratic behavior of V_1 , which stems from the random occurrence of large jumps implied by the dynamic specification for this factor.

We next study the intertemporal persistence of the error in the extraction of the factor realizations across the sample. In Figure 2, we plot estimates for this error serial correlation over the last two weeks of the first year in the simulation. Theorem 2 and Lemma 1 imply that this dependence stems exclusively from the finite time span. That is, the intertemporal error correlation is due to

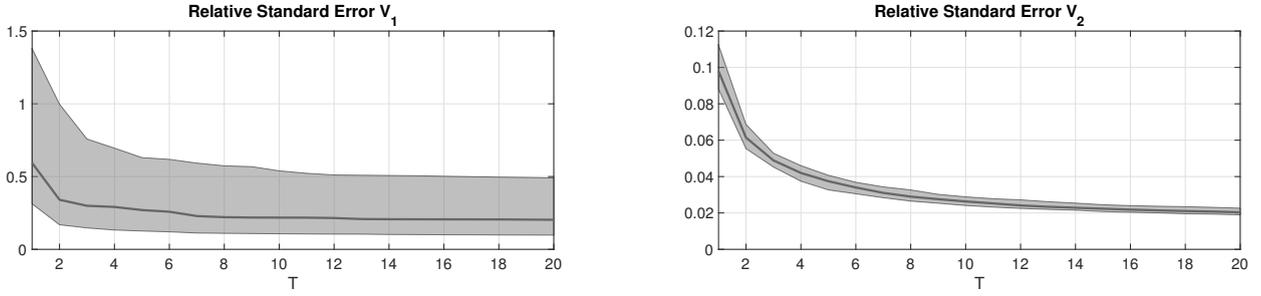


Figure 1: *Relative Standard Errors for Factors: Case of no Penalization.* Standard errors are computed for the factor realization from the last week of the first year in the sample. Standard errors are divided by the absolute value of the true value of the factor realization. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quantile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

imprecision in the estimation of the parameter vector and, as we increase the span, the effect of the latter vanishes asymptotically. Indeed, for both factors, the correlation shrinks as we increase the time span. Nonetheless, the persistence of the error in the recovery of V_2 remains fairly high, even for a panel covering 20 years. This is consistent with the evidence from Figure 1, where we noted that estimation errors for the model parameters have a strong effect on the recovery of this factor, even for panels with a long span. If the econometrician relies on our long span asymptotic result, he or she would be ignoring this nontrivial persistence in the error of the factor extraction. In contrast, our unified approach, advocated in Theorem 4, automatically accommodates this feature and more accurately conveys the degree of temporal dependence in the errors of the recovered factors.

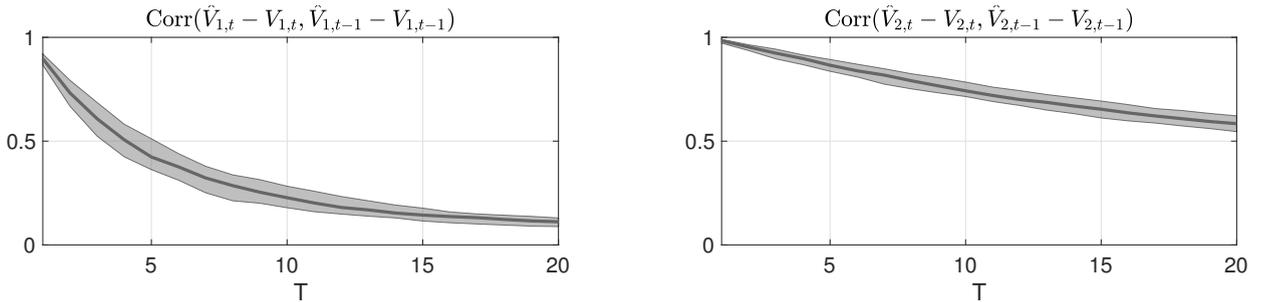


Figure 2: *Persistence of the Error in Factor Extraction: Case of no Penalization.* The plots display correlation between the errors in extracting factors over the last two weeks in the first year of the sample. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quantile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

We conclude with a discussion of the bias and standard error in parameter estimation. Figure 3 provides relative estimates for these quantities (by dividing them by their true values). To conserve

space, we present results for a subset of the parameters only, as the findings for the remaining parameters are qualitatively similar. Contrary to the short time span scenario, the long span limit distribution of the parameter vector is Gaussian, i.e., the limiting variance is independent of the sample path. We note that, for the shorter time spans, the interquantile ranges of the relative standard errors tend to be quite wide but, unlike the case for the factors, they continue to shrink as the time span grows and eventually become negligible. We also see a large reduction in the bias for some parameters as the span lengthens. However, since we keep N fixed, the estimate for the bias converges to a constant as the span grows. As for the factor recovery, the “transformation” of the features of the limiting distribution, induced by the transition from the fixed to long span asymptotics, occurs at different times across the parameters. The relative bias is larger for the parameters determining the positive jumps (c_1^+ and c_2^+) and smaller for the parameters controlling the mean reversion of the volatility factors (κ_1 and κ_2). The same ranking applies for the relative standard error of the parameters. Comparing the bias to the standard error, we also observe a nontrivial variation across the parameters. For the longest span of 20 years, the bias is approximately of the same size as the standard error for c_1^- and negligible relative to the standard error for κ_2 . Thus, ignoring the bias, i.e., performing inference according to the fixed span asymptotics, will imply poor performance for tests involving parameters such as c_1^- for time spans as short as 5 years, again corroborating the benefit of relying on the unified inference procedures.

6.2 Estimation with Penalization

We proceed by studying the impact of penalizing the estimation with a high-frequency spot volatility measure. To this end, let us set $k_n = 180$, which corresponds to using the last three hours of one-minute high frequency observations of the underlying asset. Moreover, we set λ equal to the average value of $\hat{\phi}_t/(2\hat{V}_t^2)$, computed over one realization of the options and a volatility path of length $T = 20$ years. For our parametric setup, this results in λ approximately being 0.04. Note that $\hat{\phi}_t/(2\hat{V}_t^2)$ is the optimal value of the penalization parameter, and it is generally time-varying.

Figure 4 shows that the introduction of the penalization term in the objective function has little effect on the precision of the factor extraction. The standard errors for $V_{2,t}$ are almost identical in the two cases, while the standard errors of $V_{1,t}$ are slightly higher. This may be explained by a constant λ being implemented for the whole sample. This leads to an over-weighting of the penalization term in high volatility periods, which is particularly relevant for $V_{1,t}$ as it contains frequent jumps. On the other hand, from Figure 5, we notice a nontrivial reduction in the correlation in the error of recovering the factor realizations of the more persistent volatility factor, $V_{2,t}$. This is to be expected as the high-frequency volatility estimators contain errors that are uncorrelated across days.

As evident from Figure 6, the penalization also has a nontrivial effect on the parameter recovery.

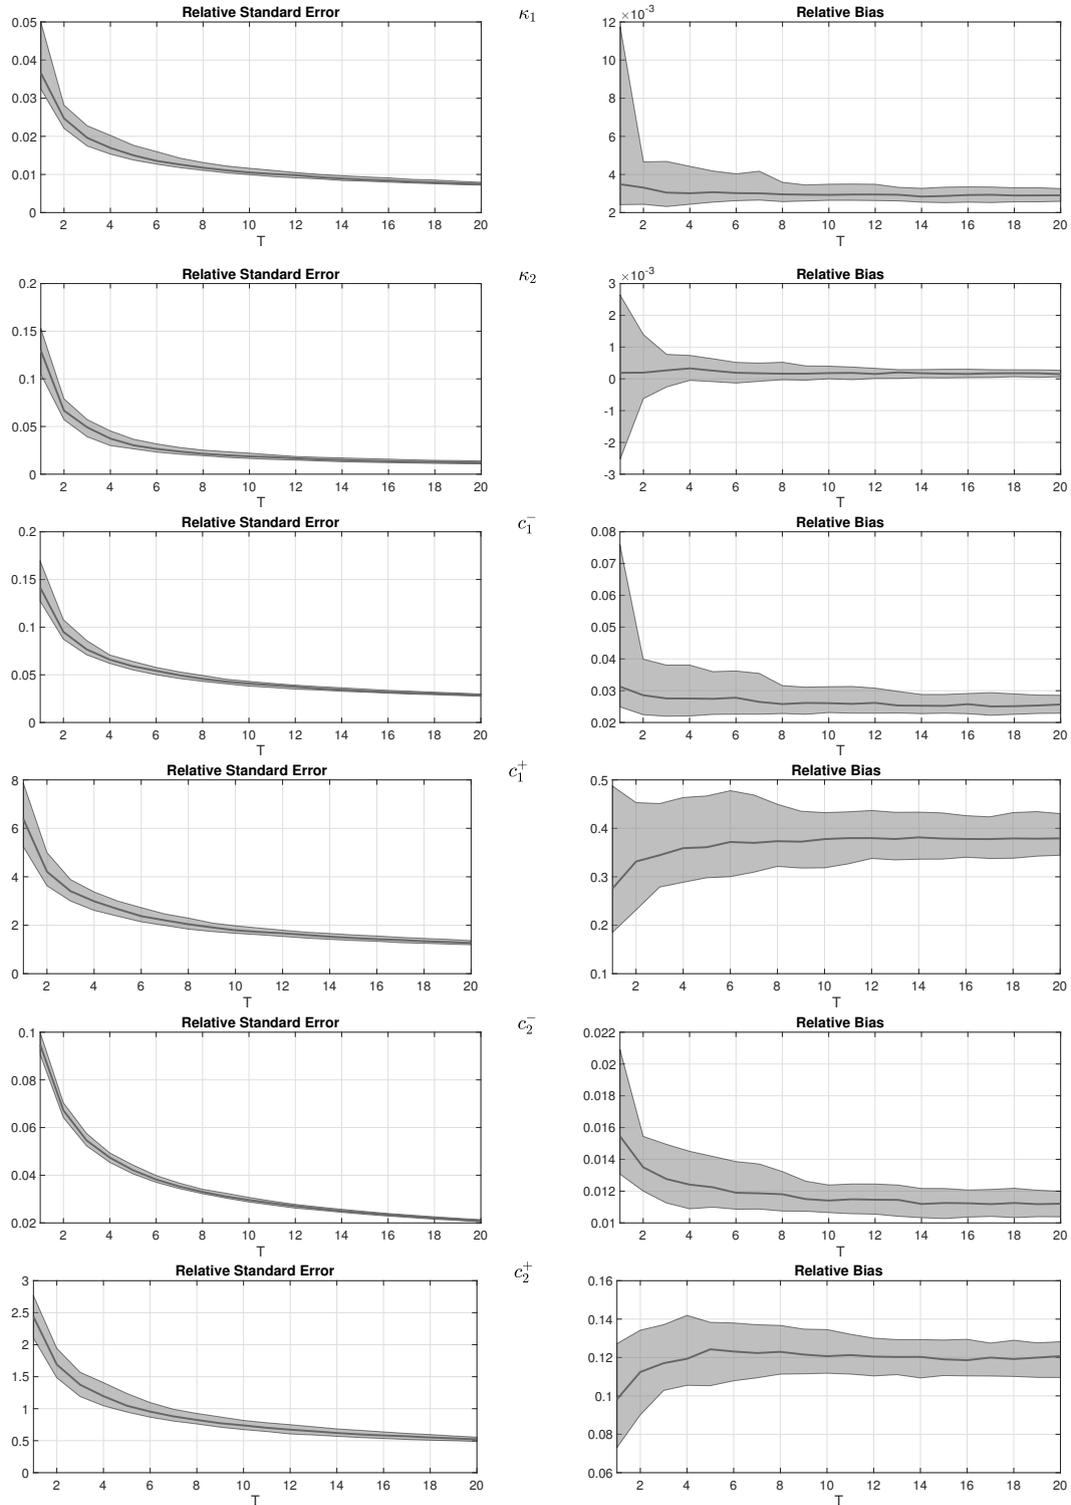


Figure 3: *Relative Standard Error and Bias for Parameters: Case of no Penalization.* Standard errors and biases are divided by the absolute value of the true value of the parameter reported in Table 1. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quartile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

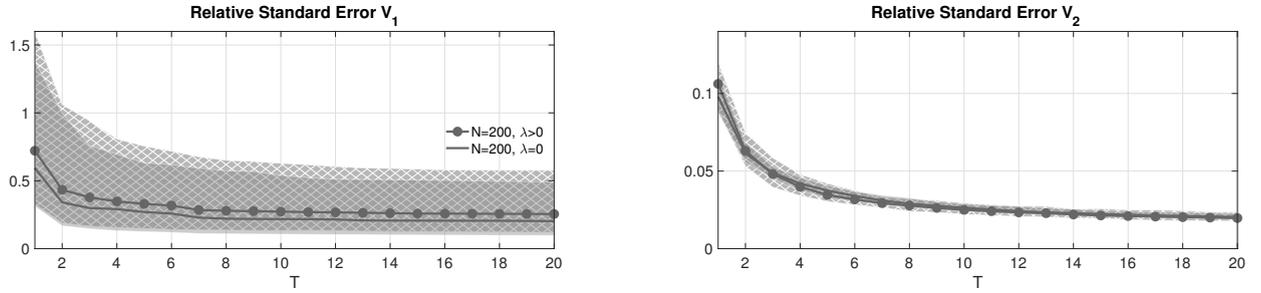


Figure 4: *Relative Standard Errors for Factors: Effect of Penalization.* Standard errors are computed for the factor realization from the last week of the first year in the sample. Standard errors are divided by the absolute value of the true value of the factor realization. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quantile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

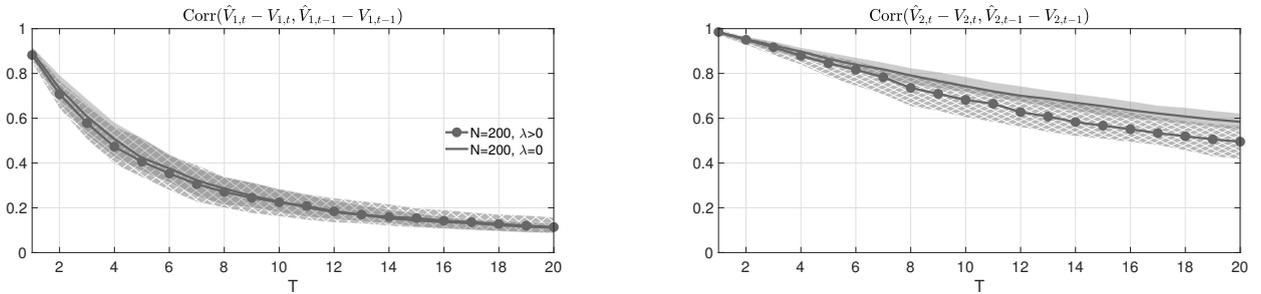


Figure 5: *Persistence of the Error in Factor Extraction: Effect of Penalization.* The plots display correlation between the errors in extracting factors over the last two weeks in the first year of the sample. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quantile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

First, the penalization affects the asymptotic bias, and its impact differs across the parameters of the model. Second, the penalization substantially reduces the standard errors of the parameters that control the jump intensity and the jump distribution. In the two-factor volatility model considered here, the jump intensity is a linear function of the two volatility states and the penalization helps to disentangle the jump intensity parameters $(c_i^\pm, \text{ for } i = 0, 1, 2)$ from the factor realizations.

7 Conclusion

In this paper, we develop unified inference theory for parametric nonlinear factor models based on a panel of noisy observations augmented with direct, albeit imperfect, measures of a known transformation of the factors. The panel has a large cross-sectional dimension while its time span is either fixed or increasing. The estimation is done via penalized least squares, with the penalty term controlling

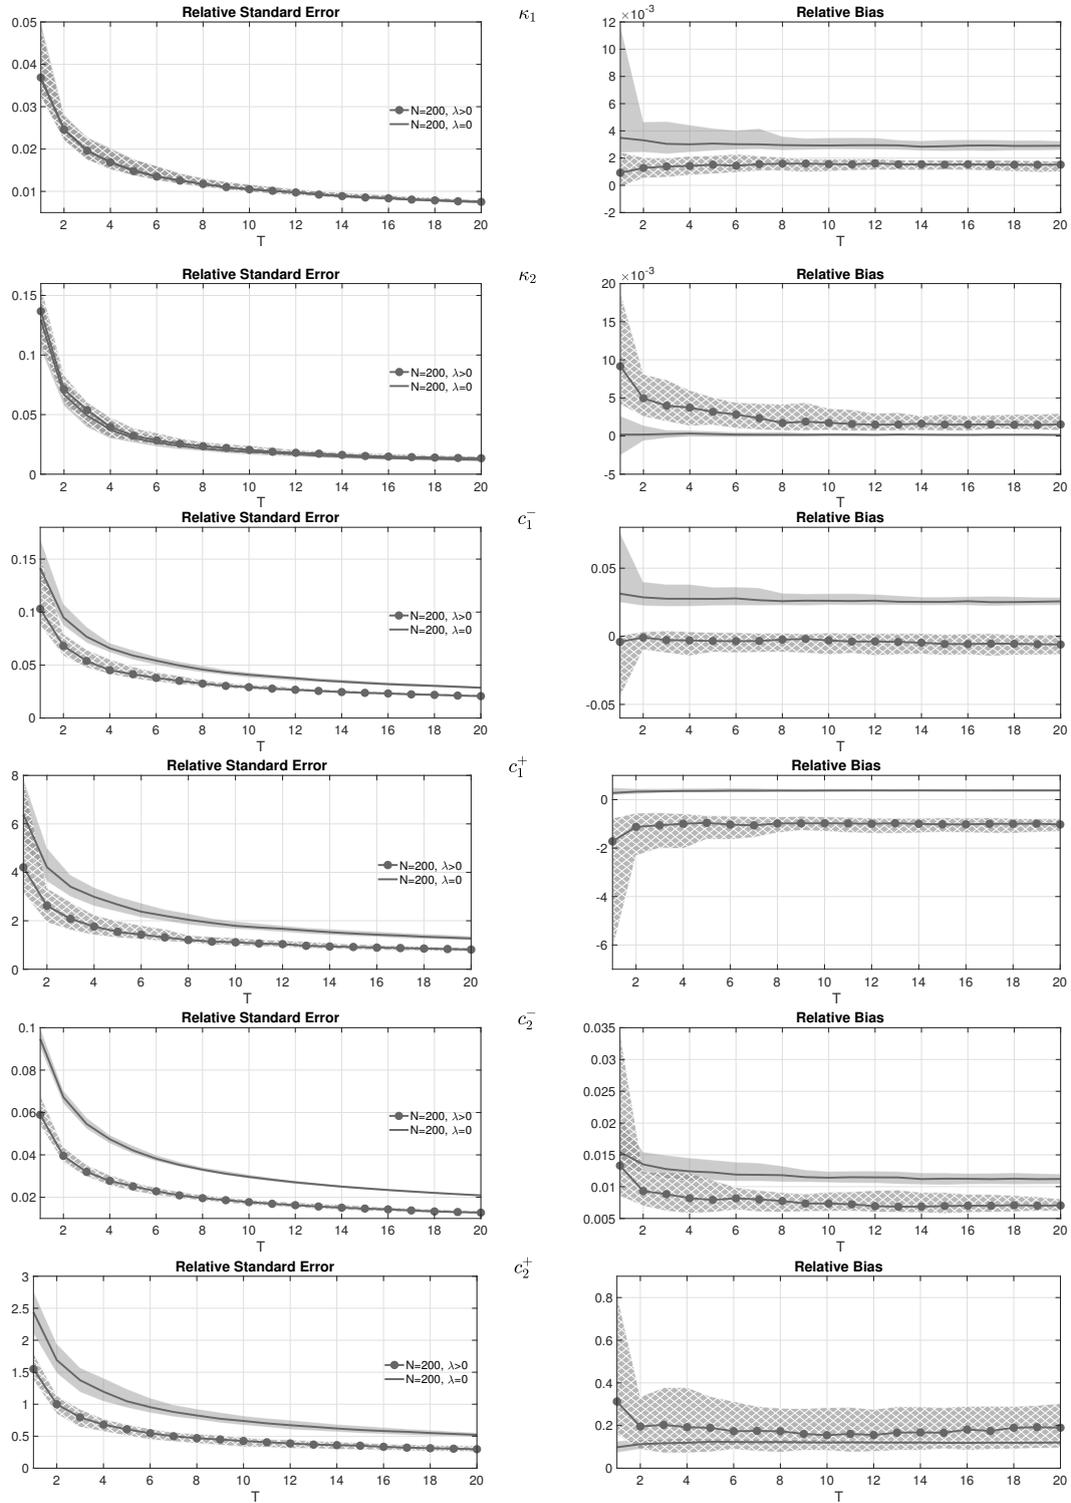


Figure 6: *Relative Standard Error and Bias for Parameters: Effect of Penalization.* Standard errors and biases are divided by the absolute value of the true value of the parameter reported in Table 1. The straight line in each of the plots corresponds to the median and the shaded area to the inter-quantile range, both computed from 100 draws of the process $(V_{1,t}, V_{2,t})$.

the size of the deviation between noisy observations and model-implied values for the known transformation of the factor realizations. Even though the limit distributions of the parameter vector and the factor realizations differ substantially in the fixed and increasing time span settings, we develop a feasible inference procedure that accommodates the critical features of the asymptotic distribution in either setup. We further propose a data-driven method for the weighting of the penalty term in the objective function, reflecting the relative information content of the noisy factor signals.

Numerical analysis demonstrates the practical relevance of our unified inference approach, as it avoids exclusive reliance on either the fixed or long time span asymptotics in any given application. In the scenarios explored, to a varying degree, we identify features of the limiting distribution that are captured best by the fixed span asymptotics (e.g., the intertemporal persistence of the error in factor recovery) and others that fall much closer to the long span asymptotics (e.g., the relative importance of bias in the parameter estimates). Our approach allows the data to determine the relative weight that is afforded the two asymptotic schemes for different aspects of the inference.

References

- Andersen, T. G., D. Dobrev, and E. Schaumburg (2012). Robust Volatility Estimation Using Nearest Neighbor Truncation. *Journal of Econometrics* 169, 75–93.
- Andersen, T. G., N. Fusari, and V. Todorov (2015a). Parametric Inference and Dynamic State Recovery from Option Panels. *Econometrica* 83, 1081–1145.
- Andersen, T. G., N. Fusari, and V. Todorov (2015b). The Risk Premia Embedded in Index Options. *Journal of Financial Economics* 117, 558–584.
- Andrews, D. (2005). Cross-Section Regression with Common Shocks. *Econometrica* 71, 1551–1585.
- Arellano, M. and S. Bonhomme (2009). Robust Priors in Nonlinear Panel Data Models. *Econometrica* 77, 489–536.
- Arellano, M. and J. Hahn (2007). Understanding Bias in Nonlinear Panel Models: Some Recent Developments. In R. Blundell, W. Newey, and T. Persson (Eds.), *Advances in Economics and Econometrics, Ninth World Congress*, Volume 3, pp. 381–409. Cambridge University Press.
- Bai, J. (2003). Inferential Theory for Factor Models of Large Dimensions. *Econometrica* 71, 135–171.
- Bai, J. and S. Ng (2002). Determining the Number of Factors in Approximate Factor Models. *Econometrica* 70, 191–221.
- Barndorff-Nielsen, O. E. and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O. E. and N. Shephard (2006). Econometrics of Testing for Jumps in Financial Economics using Bipower Variation. *Journal of Financial Econometrics* 4, 1–30.
- Bates, D. S. (2003). Empirical Option Pricing: A Retrospection. *Journal of Econometrics* 116, 387–404.
- Bernanke, B., J. Boivin, and P. Elias (2005). Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *Quarterly Journal of Economics* 120, 387–422.

- Christoffersen, P. and K. Jacobs (2004). The Importance of the Loss Function in Option Valuation. *Journal of Financial Economics* 72, 291–318.
- Duffie, D. (2001). *Dynamic Asset Pricing Theory* (3rd ed.). Princeton University Press.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.
- Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Fernandez-Val, I. and J. Lee (2013). Panel Data Models with Nonadditive Unobserved Heterogeneity: Estimation and Inference. *Quantitative Economics* 4, 453–481.
- Fernandez-Val, I. and M. Weidner (2016). Individual and Time Effects in Nonlinear Panel Models with Large N,T. *Journal of Econometrics* 196, 291–312.
- Gagliardini, P. and C. Gourieroux (2014). Efficiency in Large Dynamic Panel Models with Common Factors. *Econometric Theory* 30, 961–1020.
- Gagliardini, P., C. Gourieroux, and E. Renault (2011). Efficient Derivative Pricing by the Extended Method of Moments. *Econometrica* 79, 1181–1232.
- Garcia, R., E. Ghysels, and E. Renault (2010). The Econometrics of Option Pricing. In Y. Aït-Sahalia and L. Hansen (Eds.), *Handbook of Financial Econometrics*, Volume 1, pp. 479–552. Oxford and Amsterdam: North-Holland.
- Hahn, J. and G. Kuersteiner (2002). Asymptotically Unbiased Inference for a Dynamic Panel Model with Fixed Effects when Both N and T are Large. *Econometrica* 70, 1639–1657.
- Hahn, J. and G. Kuersteiner (2011). Bias Reduction for Dynamic Nonlinear Panel Models with Fixed Effects. *Econometric Theory* 27, 1152–1191.
- Hahn, J. and W. Newey (2004). Jackknife and Analytical Bias Reduction for Nonlinear Panel Models. *Econometrica* 72, 1295–11319.
- Jacod, J. and A. N. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.
- Kuersteiner, G. and I. Prucha (2013). Limit Theory for Panel Data Models with Cross Sectional Dependence and Sequential Exogeneity. *Journal of Econometrics* 174, 107–126.
- Lancaster, T. (2000). The Incidental Parameter Problem Since 1948. *Journal of Econometrics* 95, 391–413.
- Mancini, C. (2001). Disentangling the Jumps of the Diffusion in a Geometric Brownian Motion. *Giornale dell’Istituto Italiano degli Attuari LXIV*, 19–47.
- Neyman, J. and E. Scott (1948). Consistent Estimates Based on Partially Consistent Observations. *Econometrica* 16, 1–32.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Pastorello, S., V. Patilea, and E. Renault (2003). Iterative and Recursive Estimation in Structural Nonadaptive Models. *Journal of Business and Economic Statistics* 21, 449–482.
- Robinson, P. (1987). Asymptotically Efficient Estimation in the Presence of Heteroskedasticity of Unknown Form. *Econometrica* 55, 875–891.
- White, H. (1982). Maximum Likelihood Estimation of Misspecified Models. *Econometrica* 50, 1–25.