Observable-Validated Generator Estimation with Conditional Characteristic Functions

Johannes Bleher

June 5, 2026

Abstract

Many economic data sets record discrete adjustments, while the object of interest is the conditional distribution of an observable outcome. This paper studies an observable-relevant generator projection: the element of a transition-generator class that is distinguished by conditional characteristic restrictions and selected by held-out observable distribution scores. The estimator is a regularized conditional empirical-characteristic-function minimum-distance estimator with split-sample selection of frequency band, weighting rule, and shrinkage toward an anchor generator. The contribution is not empirical-characteristic-function estimation, conditional characteristic functions, or generator recovery alone; it is the finite-sample target and validation discipline that connect a transition law to predictive distributions of economic observables. In controlled evidence, the estimator recovers state-dependent adjustment more accurately than event-intensity benchmarks and improves characteristic-function fit, while direct quantile regression remains competitive for pinball loss. In a dairy price-adjustment application, the monthly transition generator improves whole-distribution forecasts of processor-price changes relative to Gaussian transition and direct quantile benchmarks. A historical XETRA event-history application shows how the same estimator applies when limit-order additions and removals are observed directly, with gains concentrated at low-frequency transition features.

Keywords: dynamic generators; empirical characteristic functions; conditional distributions; distributional forecasting; quantile regression

JEL codes: C14; C32; C51; C53

1 Introduction

Economic dynamics are often observed as sequences of discrete adjustments. Firms change prices, households switch choices, inventories move in jumps, orders arrive and leave, and local markets enter or exit stress states. The econometric question is rarely only which event happens next. In many applications the object of interest is the conditional distribution of an observable outcome: a price, a quantity, a duration, a transition probability, or a tail event. This paper develops a way to estimate that distribution from a model of the adjustment process.

The central difficulty is that event models and outcome models are usually estimated separately. Counting-process, duration, and point-process models describe the timing and type of adjustments (Cox, 1972; Andersen and Gill, 1982; Engle and Russell, 1998; Hawkes, 1971). Vector autoregressions and Gaussian transition models describe conditional means and second moments (Sims, 1980; Engle, 1982; Bollerslev, 1986). Quantile regressions and distributional forecasting methods describe selected features of observable predictive distributions (Koenker and Bassett, 1978; Gneiting and Raftery, 2007; Giacomini and White, 2006). Each object is useful, but the pieces do not automatically form one coherent conditional law for observables. A model that fits event likelihoods need not forecast observable tails, and a model that fits a quantile need not imply consistent probabilities at other horizons or for other functions of the state.

I ask whether a conditional adjustment generator can be estimated as an observable-relevant projection. The generator maps the current state into transition intensities. Once combined with a measurement equation, it implies a full predictive distribution for any observable function of the future state. The paper estimates this projection with a regularized conditional empirical-characteristic-function estimator. The same estimate is then evaluated by held-out characteristic losses and by proper scoring rules for observable distributions.

The contribution is methodological and specific. Characteristic-function estimation, conditional characteristic functions, and generator estimation are well established (Singleton, 2001; Jiang and Knight, 2002; Chacko and Viceira, 2003; Jiang and Knight, 2010; Chen et al., 2013; Metzner et al., 2007). The contribution here is to define and estimate the generator projection relevant for an observable distribution. Conditional characteristic moments discipline the transition law; held-out observable scores decide whether the additional transition structure is useful relative to reduced-form benchmarks.

This formulation is also close to demand-curve inversion. In demand systems, one often observes outcomes that are generated by latent shocks and structural choice rules, and the task is to recover an economically meaningful object from equilibrium observables (Berry, 1994; Berry and Haile, 2014; Albrecht and Traina, 2026). Generator inversion has the same spirit in time: observed transitions and outcomes are used to recover a dynamic law that can be pushed forward to forecast observables. The paper therefore does not require the order book to be the main empirical setting. Order books are one natural laboratory for event data, but the framework is about dynamic adjustment more generally.

The first evidence is controlled. A demand–supply simulation creates a setting in which the true generator, observable price changes, tail events, and reduced-form benchmarks are all known. It is a diagnostic test of whether the method recovers the state-dependent adjustment law and whether that recovery improves forecasts of observables. The results show a clear advantage in characteristic-function space and in several state-dependent forecast objects. Direct pinball-loss gains over quantile regression are small. The empirical sections then give two applications: dairy processor-price transitions, where the transition law is constructed from observable monthly adjustments, and a historical XETRA event-history application, where limit-order additions and removals are observed directly.

A generator that cannot improve observable forecasts is not useful for the present purpose, while a reduced-form forecast that cannot be interpreted as a dynamic adjustment law answers a different question. The paper therefore treats Doi–Peliti notation as representation language and treats the order-book application as one event-history setting, not as the definition of the framework. The dairy application demonstrates the broader transition-to-observable logic; the XETRA application demonstrates the same estimator when additions and removals are directly observed.

The remainder of the paper is organized as follows. Section 3 defines the generator and the observable target. Section 4 states the identifying content and clarifies the relation to observable models. Section 5 introduces the regularized empirical-characteristic-function estimator. Section 6 shows how generator estimates imply observable distributions and benchmark comparisons. Section 7 describes the controlled design. Section 8 reports the evidence. Section 10 concludes.

2 Relation to Existing Methods

The paper combines tools that usually appear in separate literatures. Counting-process and event-history models represent adjustment by intensities (Cox, 1972; Aalen, 1978; Andersen and Gill, 1982), with duration, transaction-time, and Hawkes-style specifications as important economic examples (Hawkes, 1971; Engle and Russell, 1998; Bowsher, 2007). Time-series models such as vector autoregressions and conditional heteroskedasticity models instead target observables directly (Sims, 1980; Engle, 1982; Bollerslev, 1986). Continuous-time Markov and diffusion methods lie between these approaches by specifying transition laws from discretely sampled data (Ethier and Kurtz, 1986; Aït-Sahalia, 2002). The present paper keeps the transition-law object but evaluates it by the predictive distribution of observables.

The estimation step follows conditional minimum-distance logic. GMM and simulated-moment methods estimate dynamic models from restrictions weaker than a full likelihood (Hansen, 1982; Hansen and Singleton, 1982; Duffie and Singleton, 1993; Gallant and Tauchen, 1996). Empirical-characteristic-function estimators use distributional restrictions and are well established for dynamic models (Feuerverger and Mureika, 1977; Carrasco and Florens, 2000; Singleton, 2001; Jiang and Knight, 2002; Chacko and Viceira, 2003; Jiang and Knight, 2010; Chen et al., 2013). Direct generator estimation for jump processes is also known (Metzner et al., 2007). The contribution is therefore not ECF estimation, conditional characteristic functions, or generator recovery by itself. It is the observable-relevant projection: estimate a conditional generator and validate it by the distribution it implies for the economic outcome.

The evaluation step follows the distributional-forecasting literature. Predictive distributions are compared by proper scores and conditional predictive ability, not only by likelihood or mean error (Diebold et al., 1998; Giacomini and White, 2006; Gneiting and Raftery, 2007). Quantile regression is a central benchmark because it targets observable quantiles directly (Koenker and Bassett, 1978). This benchmark is deliberately difficult for a generator to beat on pinball loss. The generator is useful only if the estimated transition law adds information for a broader set of observable distributional features.

The link to demand-curve inversion is conceptual. Demand estimation recovers an economically meaningful object from observed equilibrium outcomes (Berry, 1994; Berry and Haile, 2014; Albrecht and Traina, 2026). Generator inversion asks the analogous dynamic question: which conditional adjustment law rationalizes the future distribution of observables? The Doi–Peliti formalism is used only as compact representation language for jump generators and master equations (Doi, 1976; Peliti, 1985; Weber and Frey, 2017; van Kampen, 2007; Gardiner, 2009). Order books are one laboratory for this logic, as in limit-order-book intensity models (Cont et al., 2010; Bowsher, 2007), but the estimator is not tied to market microstructure.

3 Generator and Estimand

The estimand is a conditional law of adjustment. Let X_t denote the state observed just before time t. A transition of type r moves the state by ν_r. Conditional on $X_t = x$, the transition occurs with intensity $\lambda_{r,\theta}(x)$, where θ collects the unknown parameters. For any test function f, the infinitesimal generator is

$$\mathcal{L}_\theta f(x) = \sum_{r \in \mathcal{R}} \lambda_{r,\theta}(x)\{f(x+\nu_r) - f(x)\}. \tag{1}$$

The set $\mathcal{R}$ contains the economically meaningful adjustment types. In a demand–supply example, these can be demand arrivals, demand removals, supply arrivals, and supply removals. In another application, they could be firm entry and exit, product adoption, labor-market transitions, or administrative state changes.

Equation (1) separates two questions. The transition vectors ν_r define what can change. The intensities $\lambda_{r,\theta}(x)$ define how current information changes the probability of each adjustment. A linear state generator lets the log or level intensity depend on state variables. This is related to a vector autoregression in the sense that both summarize how the current state predicts future dynamics (Sims, 1980). The difference is that the generator models the law of discrete transitions before moments are taken, whereas a vector autoregression directly projects future observables on current observables.

The generator implies a local conditional characteristic function. For a frequency vector u, define

$$H_\theta(x, \mathrm{i}u) = \sum_{r \in \mathcal{R}} \lambda_{r,\theta}(x)\{\exp(\mathrm{i}u'\nu_r) - 1\}. \tag{2}$$

For a short horizon h, the transition characteristic function is approximated by

$$\phi_\theta(u \mid x, h) = \exp\{hH_\theta(x, \mathrm{i}u)\}. \tag{3}$$

This expression is the bridge from event dynamics to distributional restrictions. Low frequencies mostly identify broad location and scale features. Medium and high frequencies contain information about discreteness, tails, and non-Gaussian transition structure.

The supplement shows this bridge in a minimal birth–death generator.

The observable need not equal the state. Let $Y_{t+h} = m(X_{t+h}, \eta_{t+h})$, where m is a measurement rule and η_{t+h} is measurement noise or an auxiliary shock. Conditional on the current state, the generator and measurement rule induce

$$F_{\theta,h}(y \mid x) = \Pr_{\theta}\{Y_{t+h} \leq y \mid X_t = x\}. \tag{4}$$

This is the object used for forecasting and evaluation. Means, variances, quantiles, interval probabilities, tail probabilities, and rare-event rankings are all summaries of the same conditional distribution.

The framework applies when observables are generated by an underlying adjustment process and a transition law gives a coherent distribution for observable outcomes. Reduced-form observable models remain important benchmarks because they estimate selected summaries directly.

The distinction between the state, the adjustment, and the observable is important. The state X_t contains the information used by the adjustment law. The transition type r determines the elementary movement of the state. The observable Y_{t+h} is the object evaluated by the researcher. In some applications these objects coincide; for example, a state transition can itself be the measured outcome. In many economic applications they do not. A price can be a measurement of the balance between demand and supply states. A duration can be a measurement of the time until a transition. A market-share change can be a measurement of many latent choice adjustments. Treating these layers separately prevents the paper from confusing an event model with an outcome model.

This separation also explains why a linear state generator is not simply a vector autoregression with different notation. A vector autoregression specifies a conditional mean equation for future observables. A generator specifies the transition law before the observable is measured. If the generator is correct, the vector autoregression can be recovered as a projection of the generator-implied conditional distribution onto a linear mean function. The reverse need not hold: many transition laws can share the same conditional mean. The generator therefore contains information about higher moments, discreteness, and tail probabilities that a mean projection can discard.

The same logic applies to Gaussian transition models. A Gaussian approximation can be a useful local benchmark if only drift and variance matter. It becomes restrictive when the transition distribution is asymmetric, bounded, discrete, or driven by rare adjustment types. Characteristic functions make this contrast transparent because they evaluate the distribution directly. Low-frequency restrictions resemble moment restrictions, while higher frequencies detect distributional features that a Gaussian projection cannot represent.

The framework is intentionally modular about the measurement equation. In a purely statistical application, m can be an estimated mapping from state histories to observables. In a structural application, m can encode equilibrium or accounting restrictions. In a demand setting, this is where a demand-curve inversion enters: the observed outcome is linked to a latent or structural object, and the generator governs how that object evolves over time. The method does not solve all identification problems in the measurement equation. It provides a way to estimate and test the dynamic transition law conditional on the measurement rule used by the researcher.

4 Identification and Interpretation

Generator inversion identifies the components of the adjustment law that leave a trace in conditional transitions and observable distributions. This is a useful but limited claim. The method does not identify a structural primitive merely because it is written in generator form. With a finite frequency grid and a finite set of instruments, the estimand is a projection: the generator in the chosen class that best matches the evaluated conditional characteristic restrictions and then produces the best validated observable distribution.

To see the source of identification, compare two parameter values θ and $\tilde{\theta}$. They are observationally equivalent for the transition panel if

$$\phi_\theta(u \mid x, h) = \phi_{\tilde{\theta}}(u \mid x, h) \quad \text{for the evaluated states, horizons, and frequencies.} \tag{5}$$

They are observationally equivalent for the final forecasting problem if they also imply the same observable distribution,

$$F_{\theta,h}(y \mid x) = F_{\tilde{\theta},h}(y \mid x) \quad \text{for the evaluated outcomes and states.} \tag{6}$$

The first condition is about the event law. The second is about the observable object. A model can be well identified for one and weakly identified for the other. For example, two generators can differ in transition types that barely affect the measured outcome, while two different observable models can fit a quantile but imply different event dynamics.

The frequency grid determines which distributional features are emphasized. A low-frequency grid gives stable restrictions and is close to matching smooth moments. A medium-frequency grid adds sensitivity to non-Gaussian shape and discrete jumps. A very high-frequency grid can be informative about lattice structure, but it can also be noisy in finite samples. The paper therefore treats frequency selection as part of the estimand reported in each table, not as a hidden tuning detail.

State instruments determine which conditional variation is used. If $z(x)$ includes only a constant, the estimator matches unconditional transition distributions. If $z(x)$ includes standardized state components and interactions, the estimator also tests whether the generator tracks how the distribution changes with the state. This is the point at which generator estimation differs most sharply from unconditional empirical-characteristic-function fitting. The relevant object is not only the marginal distribution of increments; it is the conditional law of increments as the state changes.

The measurement equation determines whether the generator can be judged by observables. If Y is a known function of the future state, then the generator-implied distribution of Y is obtained by pushing the transition law through that function. If the measurement rule is estimated, then its uncertainty should be included in inference. If the measurement rule is structural, then misspecification can enter through both the adjustment law and the measurement law. In all cases, observable evaluation is essential because it prevents the generator from being judged only by internal event fit.

This perspective also clarifies the relation to quantile regression. Quantile regression directly estimates $q_\tau(x)$, the conditional τ-quantile of the observable. It does not require an event law and is therefore robust when the only target is that quantile. Generator inversion estimates a transition law and then derives $q_\tau(x)$ from the implied distribution. This is more demanding. It can lose on a single quantile score even when it gives a better description of the whole conditional distribution. The comparison is therefore not about replacing quantile regression; it is about asking whether the additional structure buys useful information.

Let $\mathcal{U}$ be the reported frequency grid and $\mathcal{Z}$ the span of reported state instruments. The population transition target is

$$\theta_{\mathcal{U},\mathcal{Z}} = \arg\min_{\theta\in\Theta} g_{\mathcal{U},\mathcal{Z}}(\theta)' W g_{\mathcal{U},\mathcal{Z}}(\theta), \tag{7}$$

where $g_{\mathcal{U},\mathcal{Z}}(\theta)$ stacks the conditional characteristic restrictions for that grid and instrument span. The observable-relevant target is the validated element of this projection class after the measurement equation is applied. Thus the object estimated in finite samples is not the full physical generator; it is the generator projection defined by $(\mathcal{U}, \mathcal{Z}, W)$, the transition dictionary, and the observable evaluation rule.

The identification claim can be summarized in three statements. First, characteristic restrictions distinguish generators only up to the features spanned by states, instruments, horizons, and frequencies. Second, observable forecast scores test whether the estimated projection is useful for the economic target. Third, disagreement between event fit and observable fit is informative rather than embarrassing: it shows whether the dynamic law captures the aspects of adjustment that matter for measured outcomes.

These statements keep the method close to the evidence. Any empirical application should state the observed state, the economically meaningful transition types, the measurement rule for the observable, and the forecast score before presenting estimates. Without those pieces, the generator is only a formal representation.

Assumption 1 (Identifying transition restrictions). *The transition dictionary is known and finite, the pre-transition state is observed on a support $\mathcal{X}$, the conditional transition characteristic function exists for all evaluated frequencies, and the instruments span the conditional variation in $\mathcal{X}$ used for estimation.*

Theorem 1 (Generator identification and projection). *Suppose Assumption 1 holds. If two parameter values imply the same conditional characteristic function for all $x \in \mathcal{X}$ and all frequencies in an open neighborhood of the origin, then they imply the same conditional transition law on $\mathcal{X}$. If the map from transition intensities to that law is injective on the transition dictionary, the conditional generator is identified on $\mathcal{X}$. With a finite grid $\mathcal{U}$ and instrument span $\mathcal{Z}$, the identified object is instead the equivalence class of generators that agree on $g_{\mathcal{U},\mathcal{Z}}(\theta)$; the reported estimator targets the minimum-distance projection in* (7).

The theorem separates the infinite-grid benchmark from the finite-grid object actually used in the evidence. This is why the paper reports losses by frequency band and treats frequency selection as part of the empirical protocol.

The supplement gives the corresponding consistency, normal approximation, and score-gap arguments.

Proposition 1 (Transition and observable equivalence). *If two generators imply the same* *conditional transition law on the support of the state and share the same measurement rule, then they imply the same conditional distribution for every observable generated by that measurement rule. The converse need not hold: two different generators can be observationally equivalent for a particular observable. Conversely, one generator can improve the finite-grid transition criterion in* (7) *without improving the observable score if the measurement rule discards the transition component on which the improvement occurs.*

The proposition motivates the estimand. The goal is not always the full physical generator. The target is the observable-relevant generator projection: the element of the generator class that is distinguishable by the chosen characteristic restrictions and useful for the chosen observable distribution. This definition is intentionally weaker than full structural recovery. It is also closer to the empirical question, because an economically irrelevant transition difference should not be treated as success merely because it improves an internal transition fit. The proposition is also the reason for reporting both transition characteristic losses and observable scores.

5 Regularized Conditional ECF Estimation

The estimator targets the transition law through characteristic restrictions. Observe transitions $\Delta X_i = X_{t_i+h_i} - X_{t_i}$, pre-transition states x_i, and horizons h_i. For each frequency u_k and instrument vector $z(x_i)$, the model implies that

$$\mathbb{E}[z(X_{t_i})\{\exp(\mathrm{i}u_k'\Delta X_i) - \phi_\theta(u_k \mid X_{t_i}, h_i)\}] = 0. \tag{8}$$

The empirical moment vector $g_n(\theta)$ stacks the real and imaginary parts of these restrictions across frequencies and instruments:

$$g_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} z(x_i) \otimes \left[\exp(\mathrm{i}u'\Delta X_i) - \phi_\theta(u \mid x_i, h_i)\right]. \tag{9}$$

The tensor notation means that each state instrument is paired with each frequency restriction. In implementation the complex moments are represented by their real and imaginary components.

The baseline estimator minimizes a weighted distance between empirical and model-implied characteristic moments:

$$\widehat{\theta} = \arg\min_{\theta\in\Theta} g_n(\theta)' W_n g_n(\theta). \tag{10}$$

The first step uses identity weighting. The second step estimates a block long-run covariance matrix $\widehat{\Omega}$ from the first-step residual moments and sets

$$W_n = (\widehat{\Omega} + \kappa I)^{-1}, \qquad \kappa > 0, \tag{11}$$

where κ is chosen on the validation block. Blocking accounts for serial dependence in the transition panel.

Regularization is useful because high-dimensional frequency and instrument grids can make the criterion noisy. The submitted estimator uses the validation-selected ridge penalty

$$\widehat{\theta}_\rho = \arg\min_{\theta\in\Theta} \left\{ g_n(\theta)' W_n g_n(\theta) + \rho\, (\theta - \theta_0)' A (\theta - \theta_0) \right\}, \tag{12}$$

where θ_0 is the anchor estimate and A is a diagonal scaling matrix based on the anchor scale. The anchor can come from likelihood, a simple transition regression, or a lower-dimensional generator. The tuning parameter ρ, the frequency band, and κ are chosen by held-out characteristic loss, not by in-sample fit.

This estimator has two roles. First, it estimates the generator. Second, it provides diagnostics. A model can be compared with an oracle generator, a Gaussian transition approximation, a Hawkes-style intensity, or a reduced-form projection by evaluating the same held-out characteristic loss. A model that improves event likelihood but fails characteristic restrictions is not a successful distributional generator.

The estimator is implemented as a split-sample procedure. Training data estimate candidate generators; validation data select the frequency band, penalty, covariance regularization, weighting rule, and anchor; test data report characteristic losses and observable forecast scores. Frequencies near zero emphasize smooth moment information, while frequencies away from zero emphasize shape, discreteness, and tail-sensitive oscillations. The penalty has a practical role: it lets the ECF criterion depart from a stable likelihood or transition-regression anchor only when held-out characteristic loss supports the departure.

The regularity requirements are those of conditional minimum-distance estimation with dependent data. The transition panel must be stationary or locally stable over the estimation window, the instruments and frequency grid must identify the relevant generator components, and the block covariance estimator must approximate the long-run covariance of the moment process. Under the usual compactness, uniform convergence, differentiability, and rank conditions, the estimator is consistent for the finite-grid minimum-distance target and has the standard GMM sandwich limit. The supplement records these arguments.

Inference is organized around loss gaps rather than only parameter standard errors. For characteristic losses and observable scores, the reported uncertainty is the sampling variation of held-out comparisons. The implementation uses block resampling of evaluation blocks, preserving local dependence in forecast errors. Unless otherwise stated, these intervals condition on the first-stage estimate and describe evaluation-sample uncertainty.

6 Observable Distribution Forecasts

The generator becomes economically useful only after it is mapped into observables. Given a current state x, an estimated generator $\widehat{\theta}$, and a forecast horizon h, draw simulated future states $X_h^{(s)}$ from the transition law implied by $\widehat{\theta}$. Applying the measurement rule gives simulated observables $Y_h^{(s)}$. The predictive distribution is

$$\widehat{F}_{G,h}(y \mid x) = \frac{1}{S}\sum_{s=1}^{S} \mathbb{1}\{Y_h^{(s)} \leq y\}, \tag{13}$$

where S is the number of simulated paths. The subscript G emphasizes that the forecast comes from the generator.

This construction answers the question that motivates the paper. Once the generator is estimated, one can forecast the entire conditional distribution of the observable, not only its mean. For any probability level τ, the generator-implied conditional quantile is the value $q_{\tau,G}(x)$ satisfying $\widehat{F}_{G,h}(q_{\tau,G}(x) \mid x) \geq \tau$. The same simulated distribution also gives prediction intervals, downside probabilities, and expected shortfall-type summaries.

The natural reduced-form benchmark is quantile regression. Let $q_{\tau,B}(x)$ be a directly estimated conditional quantile from a benchmark model B. The pinball loss for a realized outcome y is

$$\ell_\tau(q, y) = (\tau - \mathbb{1}\{y < q\})(y - q). \tag{14}$$

Quantile regression is optimized for this loss at a selected τ. The generator is not. A fair comparison therefore uses several scores: pinball loss for quantiles, an interval score for central prediction intervals, a tail-probability loss for rare outcomes, and a characteristic-function loss for the whole observable distribution.

For any scoring rule $S(\cdot,\cdot)$, the test-sample score gap is

$$\widehat{\Delta}_S = \frac{1}{n_T}\sum_{i\in T}\{S(\widehat{F}_{B,h_i}(\cdot \mid x_i), y_i) - S(\widehat{F}_{G,h_i}(\cdot \mid x_i), y_i)\}. \tag{15}$$

A positive gap means that the generator forecast has the lower loss. The comparison is observable-facing: the generator is valuable only if its implied distribution improves the forecast object or supplies coherent distributional information that a direct benchmark does not provide.

This perspective clarifies the relationship with familiar models. A vector autoregression is a linear projection for conditional observables. A Gaussian transition model matches a restricted set of moments. A Hawkes model is a structured event-intensity benchmark. Quantile regression directly estimates selected distributional summaries. Generator inversion differs by estimating one transition law and then deriving all of these summaries from that law.

The score comparison should be read with care. A proper score rewards an entire predictive distribution, but different scores emphasize different aspects of that distribution. Pinball loss emphasizes a selected quantile. Interval scores emphasize coverage and width of prediction intervals. Tail-probability losses emphasize rare events. Characteristic-function losses emphasize the shape of the distribution across selected frequencies. The paper therefore does not collapse all evidence into one ranking. It asks which method is best for which part of the forecast problem.

This is also the reason for comparing the generator with quantile regression rather than only with likelihood-based event models. Researchers often care about observables. If the generator cannot compete with a direct observable benchmark, its structural coherence may not be worth the extra modeling burden. Conversely, if the generator matches direct quantile performance and improves tail or whole-distribution diagnostics, then the additional structure has value. It supplies a single law that can be used for multiple horizons, multiple observables, and counterfactual changes in the adjustment process.

In empirical work, the same structure can be used for policy or counterfactual analysis. A reduced-form quantile regression can say how the conditional quantile changes with observed covariates. A generator can ask what happens to the observable distribution if a particular transition intensity changes. For example, one can change the arrival intensity of a state transition, simulate the new transition law, and evaluate the resulting distribution of the observable. This is not automatically causal; the intervention must be justified by the application. But the generator provides the object on which such an intervention would operate.

7 Research Design

The controlled exercise is designed to answer a methodological question: when the true adjustment law is known, does generator inversion recover the state-dependent law and does that recovery improve observable distribution forecasts? A simulation is appropriate at this stage because it separates three sources of error that are confounded in field data: misspecification of the adjustment law, misspecification of the measurement equation, and noise in the observable forecast target.

The simulated economy has demand-side and supply-side state variables. Events arrive one at a time and change one side of the state. Demand arrivals and removals shift the demand state; supply arrivals and removals shift the supply state. The event intensities depend on the current state, so the transition law is genuinely conditional. The observable price-change object is then generated from the future demand–supply state. This gives the generator an economic role: it describes how latent pressure on the two sides of the market evolves before it is measured as a price adjustment.

The estimators are compared on three levels. The first level is generator recovery. Because the true intensities are known in the simulation, the estimated intensity surface can be compared directly with the data-generating surface. This is not available in real data, but it is useful for checking whether the estimator learns the intended object. The second level is transition-distribution fit. Held-out characteristic losses compare the model-implied transition distribution with realized transition increments. The third level is observable forecast performance. Generator-implied distributions are evaluated against reduced-form quantile benchmarks using proper scores and characteristic losses for the observable.

The benchmark set covers the main competing projections. The oracle generator is included only to show the attainable loss. A CCF-only generator isolates what is gained by observable validation beyond characteristic-function fitting alone. A Gaussian truncation tests what is lost by replacing the jump law with a local Gaussian approximation. An aggregate Gaussian projection tests what is lost by ignoring state-dependent transition structure. A Hawkes-style generator tests whether a familiar event-intensity model can fit the timing and common-event likelihood while missing parts of the state-dependent adjustment law. Quantile regressions test whether directly modeling the observable quantiles is enough for the final forecasting task.

This design avoids a common ambiguity in method papers. If a generator beats all reduced-form models on every score, the exercise may simply be too favorable. If it loses everywhere, the method has little practical value. The informative case is mixed. A credible method should win where its structure is relevant and fail to dominate where a simpler benchmark is optimized for the same target. The demand–supply exercise is useful because it produces exactly this kind of diagnostic comparison.

The reported tables use reader-facing labels rather than implementation names. Losses are scaled only where needed for readability, and each table is organized by the economic object being evaluated. The evidence is kept compact: characteristic-function fit, observable distribution scores, score gaps against quantile benchmarks, and Monte Carlo stability. Trace tables, notation maps, and implementation audits are omitted because they obscure the evidence.

8 Controlled Evidence

The evidence uses a simulated demand–supply adjustment process. The state records demand-side and supply-side liquidity. Four event types move the state: demand arrival, demand removal, supply arrival, and supply removal. The observable is a price-change object generated from the state. Because the data-generating law is known, the exercise can evaluate both recovery of the generator and forecast performance for observables.

The first result is that characteristic-function diagnostics recover the relevant transition law. Table 1 compares the oracle generator, the estimated generator, a regularized characteristic-function version, Gaussian approximations, and a Hawkes-style benchmark. At low frequencies, the estimated generator has lower conditional characteristic loss than Gaussian and Hawkes-style alternatives, though the oracle remains better. At medium frequencies, the estimated generator is close to the oracle and clearly closer in oracle characteristic distance than the reduced alternatives. At high frequencies, the models are nearly indistinguishable because the evaluated transition increments contain little extra high-frequency separation.

The second result is mixed. Table 2 compares the generator-implied observable distribution with direct quantile-regression benchmarks. The generator sharply improves the observable characteristic-function loss relative to both benchmarks, which means that its full predictive distribution better matches the held-out empirical distribution in characteristic space. Average pinball loss and the continuous-ranked-probability approximation are essentially tied with the aggregate quantile regression. The quantile regression augmented with generator signals performs worse on the reported tail and characteristic scores. The result is evidence that the generator buys distributional coherence and strong whole-distribution fit, while direct quantile regression remains hard to beat on the score it targets.

Score gaps against quantile benchmarks give the same interpretation. Against aggregate quantile regression, the generator has a large positive characteristic-function gap but essentially no pinball-loss advantage. Against the quantile regression augmented with generator signals, the generator improves the tail-probability and characteristic losses. This pattern is useful because it prevents an overstatement. The generator is not winning because every score mechanically rewards it. It wins where a coherent distribution matters most.

Figure 1 shows the same point visually. The generator produces a complete predictive distribution rather than isolated conditional quantiles. This makes it possible to read

Table 1: Held-Out Characteristic-Function Fit

The table reports conditional and unconditional characteristic-function losses by frequency band. Smaller losses indicate a closer match between the model-implied transition law and held-out transition increments. The oracle row uses the known data-generating generator; other rows are estimated or reduced-form approximations.

Model	Conditional characteristic loss	Unconditional characteristic loss	Oracle characteristic distance
Low (0.25–0.75)			
Oracle jump generator	12.758	5.899	0.000
Estimated jump generator	15.784	6.381	8.898
Regularized characteristic-function jump generator	15.784	6.381	8.898
Gaussian truncation	16.217	9.575	13.549
Aggregate Gaussian projection	21.916	18.529	36.724
Hawkes-style generator	56.518	4.126	83.748
Medium (1.00–2.00)			
Oracle jump generator	20.005	31.751	0.000
Estimated jump generator	21.223	32.389	0.662
Regularized characteristic-function jump generator	21.223	32.389	0.662
Gaussian truncation	21.594	35.076	3.416
Aggregate Gaussian projection	21.221	34.290	8.663
Hawkes-style generator	24.441	31.203	9.492
High (2.50–3.10)			
Oracle jump generator	10.796	12.969	0.000
Estimated jump generator	10.799	12.974	0.000
Regularized characteristic-function jump generator	10.799	12.974	0.000
Gaussian truncation	10.809	12.989	0.000
Aggregate Gaussian projection	10.808	12.988	0.000
Hawkes-style generator	10.772	12.946	0.002

Table 2: Observable Distribution Scores

The table compares generator-implied observable forecasts with reduced-form quantile benchmarks. Pinball loss, the continuous-ranked-probability approximation, observable characteristic-function loss, tail-probability loss, and crossing share are loss or violation measures, so lower values are better for those columns. Simulation paths reports the number of simulated draws used to approximate the generator-implied distribution.

Model	Average pinball loss	Continuous ranked probability score approximation	Observable characteristic-function loss $\times 10^4$	Tail probability loss	Crossing share	Simulation paths
Generator distribution	0.0050	0.0099	24.0191	0.0696	0%	250
Aggregate quantile regression	0.0050	0.0099	161.8189	0.0695	0%	–
Quantile regression with generator signals	0.0050	0.0101	396.3519	0.0755	0%	–

central intervals, tail probabilities, and realized outcomes from one object. The figure is useful because it shows what the method ultimately delivers: not a table of transition coefficients, but a forecast distribution for an observable.

The Monte Carlo results in Table 3 summarize how stable these findings are across repeated simulated samples. The generator recovers the state-dependent intensity surface much more accurately than the Hawkes-style benchmark: the median intensity-error ratio is about 0.19, and the generator is favored in every replication. This advantage does not mechanically translate into every downstream criterion. The Hawkes-style benchmark has better event likelihood in this design, and drift and volatility errors are close. The generator performs better for price-change error and rare-liquidity probability loss, with generator-favored shares of 97% and 70%, respectively.

These results define the empirical claim. Observable-validated generator estimation is not a universal replacement for reduced-form observable models. It is a way to estimate a coherent dynamic law when the econometrician cares about several observable summaries at once, especially tail probabilities, distributional shape, and counterfactual changes in the adjustment process. If the target is a single conditional quantile, direct quantile regression is a strong benchmark and may be sufficient. If the target is the entire conditional

Table 3: Monte Carlo Evidence

The table reports median criterion ratios or differences across repeated simulated samples. Ratios below one favor the generator. Positive ranking-score differences favor the generator. The final column gives the number of finite replications used for each criterion.

Criterion	Median	Interquartile range	Generator-favored share	Replications
Generator recovery				
Intensity error ratio	0.1864	[0.1572, 0.2411]	100%	110
Event likelihood				
Log likelihood difference	-0.0270	[-0.0292, -0.0242]	0%	110
Conditional moments				
Drift error ratio	1.0152	[1.0104, 1.0193]	2%	110
Second-moment error ratio	0.9989	[0.9946, 1.0016]	64%	110
Price adjustment				
Price-change error ratio	0.8794	[0.8233, 0.9209]	98%	110
Volatility				
Volatility error ratio	1.0126	[1.0037, 1.0219]	10%	110
Rare liquidity				
Probability-loss ratio	0.9129	[0.7831, 1.0286]	70%	103
Ranking-score difference	0.0744	[-0.0069, 0.2195]	71%	103

Full Monte Carlo evidence: all reported criteria use at least 100 finite criterion-specific Monte Carlo observations.

Figure 1: Generator-Implied Distribution Forecasts

The figure plots generator-implied forecast intervals for the observable outcome against realized outcomes. The visible object is the conditional distribution delivered by the estimated generator after simulation through the measurement rule.

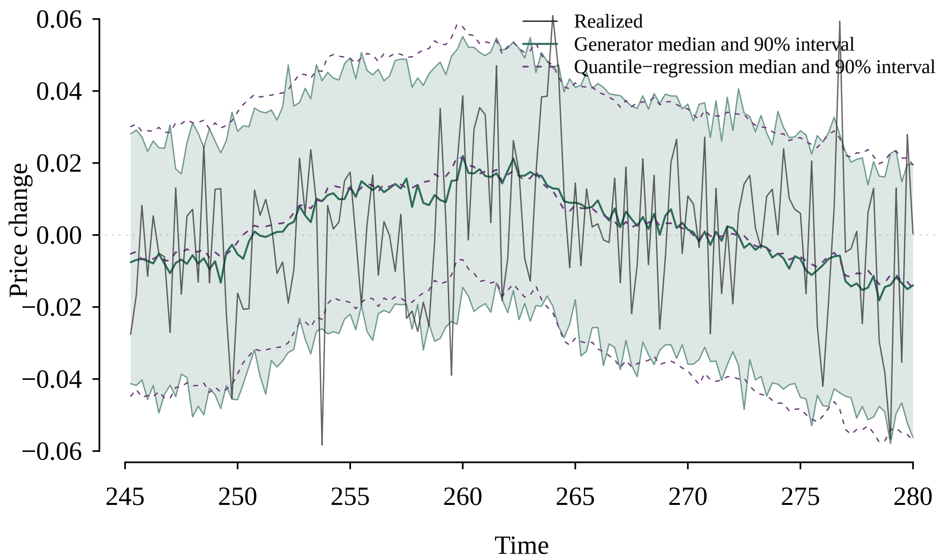

distribution induced by dynamic adjustment, the generator provides structure that direct quantile models do not.

The mixed likelihood result is especially important. The Hawkes-style benchmark does better on event likelihood in the Monte Carlo table, yet it does not recover the state-dependent intensity surface and performs worse on several observable-oriented criteria. This illustrates why the paper does not use likelihood as the sole objective. Likelihood rewards one-step event probabilities under a particular event specification. The characteristic-function and observable scores ask whether the estimated law captures the distributional features that matter for the target.

The controlled evidence also suggests where the method needs improvement. First, frequency selection should adapt to the forecast target. If the target is a tail probability, the validation loss should place more weight on frequencies and instruments that are informative about tail states. Second, the penalty should shrink different components of the generator differently. Common-event intensities and rare-event intensities have different sampling uncertainty. Third, the observable forecast step should propagate first-stage uncertainty more fully, especially when the measurement equation is estimated. These are refinements of the same framework, not changes in the estimand.

9 Empirical Applications

9.1 Dairy Price Adjustment

The dairy-market data allow a conservative empirical implementation of the framework. They do not record individual demand arrivals, supply arrivals, or order-book events. They do record monthly processor comparison prices together with spot-milk prices, butter and skimmed-milk-powder benchmark prices, futures prices, futures volumes, processor type, and processor throughput. The empirical application therefore treats the adjustment law as a law for processor-price transitions, not as a literal law for unobserved demand or supply events.

This distinction is important. The data are sufficient for a transition-to-observable design because processor payout changes are discrete, state-dependent, and economically meaningful. The data are not sufficient for a structural demand-arrival interpretation without an additional measurement model. The estimand in this application is therefore the conditional distribution of next-month processor-price adjustments, disciplined by observable benchmark states.

The transition dictionary classifies processor-price changes into upward adjustment, downward adjustment, and no material adjustment. The state contains lagged processor prices, lagged adjustment, a benchmark-price disequilibrium component, spot milk, butter and skimmed-milk-powder prices, futures prices, throughput, and processor type. The observable is the next-month processor-price change. The tail object is downside payout risk. All reported results use a time split: earlier months for estimation, middle months for validation, and later months for testing.

The usable sample starts in 2014, when the market observables needed for the state become available. The monthly-mean panel contains 923 processor transitions for 13 processors, with 195 held-out test observations. The end-of-month panel contains 845 transitions and 169 test observations. This is a compact transition panel rather than a large event-history sample. For that reason, the empirical claim is intentionally narrow.

The first empirical question is whether the regularized transition characteristic estimator improves the held-out transition law. It does. For both aggregation rules and for both frequency bands, the regularized transition estimator has lower conditional characteristic loss than the Gaussian transition anchor and the unconditional Gaussian transition. The improvement is especially large in the monthly-mean panel: at low frequencies the conditional characteristic loss falls from 124.5 under the Gaussian anchor to 6.7 under the regularized transition estimator, using the same scaling as the reported score table.

The second question is whether this transition fit improves the distribution of the observable processor-price change. Table 4 gives a mixed but informative answer. In the monthly-mean panel, the regularized transition estimator has the lowest average pinball loss, the lowest continuous-ranked-probability approximation, and by far the lowest observable characteristic-function loss. The observable characteristic loss is 26.3 after scaling, compared with 251.2 for the Gaussian transition anchor and 476.7 for the direct quantile benchmark. This is the result the method is designed to deliver: one estimated adjustment law improves the whole predictive distribution of an observable outcome.

The end-of-month panel is an aggregation robustness check. The regularized transition estimator still has the lowest observable characteristic-function loss, but the direct quantile benchmark has the best pinball and continuous-ranked-probability scores. Figure 2

Table 4: Dairy Observable Distribution Scores

The table compares predictive distributions for next-month processor-price changes. Pinball loss and the continuous-ranked-probability approximation evaluate distributional accuracy; the observable characteristic-function loss compares the whole predicted distribution with the held-out distribution; tail-probability loss evaluates downside payout risk. Crossing share measures quantile incoherence. All reported columns are losses or violation shares, so lower values are better.

Model	Average pinball loss	Continuous ranked probability score approximation	Observable characteristic-function loss $\times 10^4$	Tail probability loss	Crossing share
End-of-month aggregation					
Direct quantile benchmark	0.1325	0.2650	341.4513	0.0152	1.2%
Gaussian transition anchor	0.1415	0.2830	255.4858	0.0131	0.0%
Regularized transition characteristic estimator	0.1811	0.3623	200.2336	0.0136	0.0%
Monthly-mean aggregation					
Direct quantile benchmark	0.1521	0.3041	476.7486	0.0226	5.6%
Gaussian transition anchor	0.1504	0.3007	251.1817	0.0195	0.0%
Regularized transition characteristic estimator	0.1340	0.2681	26.3335	0.0231	0.0%

shows the same contrast as loss ratios. The generator is strongest when the target is the full conditional distribution and monthly aggregation smooths high-frequency noise in the benchmark states.

Interval diagnostics tell the same story. In the monthly-mean panel, the regularized transition estimator produces narrower intervals and lower interval scores than the Gaussian anchor and the direct quantile benchmark, with coverage close to the nominal levels. In the end-of-month panel, the transition estimator undercovers relative to the benchmarks. The evidence therefore supports the monthly-mean processor-price application more clearly than the end-of-month version.

The empirical conclusion is therefore specific. The monthly-mean version improves distributional scores relative to Gaussian transition and direct quantile benchmarks, while the end-of-month version improves characteristic-function fit but not all proper scores. The data do not justify language about literal demand arrivals or supply arrivals. They justify a price-adjustment generator whose value is assessed by the conditional distribution of

Figure 2: Dairy Benchmark Loss Ratios

The figure reports each benchmark loss divided by the regularized transition-estimator loss. The vertical line marks equal loss. Points to the right favor the transition estimator; points to the left favor the benchmark. The monthly-mean aggregation favors the transition estimator for the main distributional scores, while the end-of-month aggregation gives a more mixed comparison.

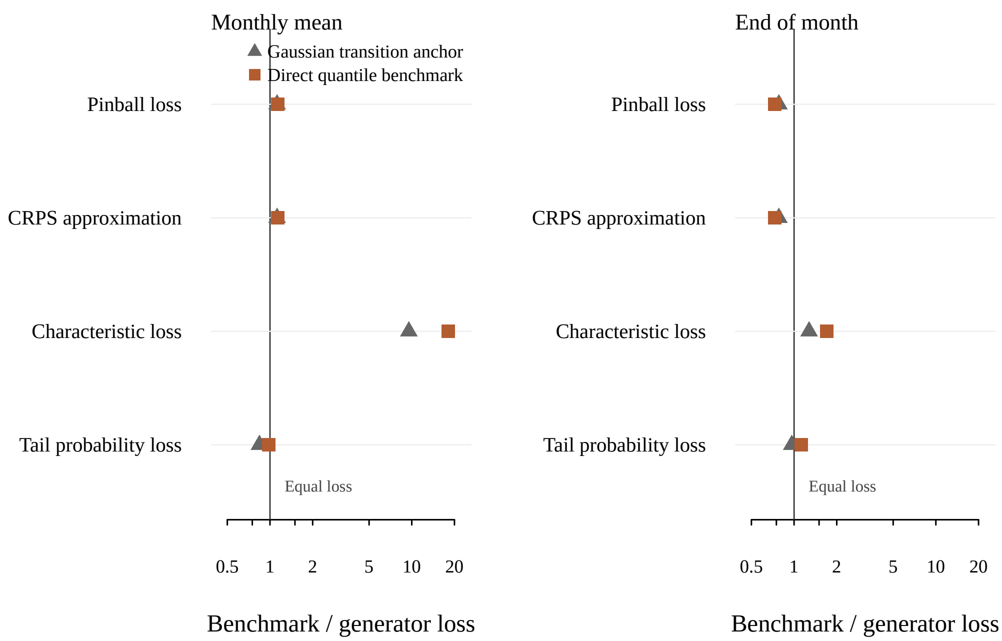

processor-price changes.

The dairy application changes the role of the order book in the paper. Order-book data remain a natural setting when the economic question is about arrivals, cancellations, executions, and liquidity provision. The dairy data show that the framework is not tied to market microstructure. It can be used whenever economically meaningful transitions can be constructed and the researcher cares about the future distribution of an observable outcome.

9.2 XETRA Order-Book Event Application

Unlike the dairy panel, the historical XETRA files record order-book events directly and come from the limit-order-book data underlying Frey and Grammig (2006). The local event histories identify limit-order additions and removals on both sides of the book, together with timestamps, side, event size, and limit price.

The application uses the self-contained raw data stored with the replication files. It takes Adidas on January 2, 2004, keeps continuous-trading additions and deletions, and maps bid-side events to demand-depth events and ask-side events to supply-depth events. The sample contains 15,808 normalized events, 3,162 held-out test events, a median waiting time of 0.451 seconds, and no invalid waiting times. The labels "demand" and "supply" mean bid-side and ask-side depth, not structural demand or supply.

The XETRA event-history application tests the estimator where the event dictionary is observed. At low frequencies, which summarize the broad balance of additions and removals, the regularized characteristic-function jump generator has the lowest scaled loss, 20.7, compared with 24.2 for the likelihood jump generator, 24.5 for the Gaussian truncation, and 1084.0 for the aggregate Gaussian projection. A six-sample robustness exercise gives the same low-frequency ranking in median loss, although the regularized estimator wins only three comparisons.

The medium-frequency comparison is less favorable. The Hawkes-style generator and Gaussian truncation perform slightly better, with scaled losses of 166.3 and 169.6 compared with 173.3 for the regularized estimator. Characteristic regularization improves broad transition fit, but the displayed-depth state specification does not pin down sharper transition shape. The supplement reports the matching transaction-price check and shows why market-order event times would be needed for a stronger price-prediction design.

10 Conclusion

This paper reframes generator estimation as an observable distribution problem. The primitive object is a conditional adjustment law. The econometric task is to estimate that law and test whether its implied distributions improve forecasts of prices, quantities, durations, choices, or other observables. Conditional empirical characteristic functions provide a practical estimator because they compare the whole transition distribution rather than a small set of moments.

The controlled evidence supports a calibrated conclusion. The generator recovers the state-dependent adjustment law and improves whole-distribution diagnostics, but reduced-form quantile regression remains competitive for direct quantile loss. The dairy application supports the same boundary in real data. With monthly-mean processor-price transitions, the regularized transition generator improves whole-distribution scores relative to Gaussian transition and direct quantile benchmarks. With end-of-month aggregation, it improves characteristic-function fit but not every proper score. The XETRA event-history application adds the event-history case: when limit-order additions and removals are observed directly, the regularized characteristic estimator improves the low-frequency held-out transition distribution, while medium-frequency fit remains competitive across benchmarks.

The empirical evidence has clear boundaries. The dairy data identify a price-adjustment transition law, not literal demand arrivals or supply arrivals. The XETRA event-history application identifies displayed depth additions and removals, but lacks a full market-order timestamp history. These boundaries clarify what the framework requires from an application. The dynamic adjustment law must be economically meaningful, and the induced observable distribution must answer a substantive question.

The contribution is therefore a finite-sample econometric target and validation discipline. Conditional ECF moments estimate the transition projection; observable scores decide whether that projection is useful for the economic target. The evidence shows where this route works and where direct reduced-form benchmarks remain strong.

Data Availability Statement

The replication materials accompanying this article are available on Zenodo at `https://doi.org/10.5281/zenodo.20514430`. They contain the code, simulated data, processed dairy transition panels, processed XETRA event file, and derived tables and figures needed to reproduce the reported results. The historical XETRA application uses the self-contained event file described in Section 9. Raw third-party source materials that are not part of the replication archive remain subject to the access conditions of their original providers.

Declaration of Interest

None.

AI-Assisted Preparation

OpenAI Codex based on GPT-5 was used during manuscript preparation for language editing, organization of revisions, coding assistance, and consistency checks. All substantive claims, calculations, interpretations, and final wording were reviewed by the author, who takes full responsibility for the content of the article. The author has checked the applicable terms of use for the tool and confirms that its use is suitable for publication.

References

Aalen, O. O. (1978). Nonparametric inference for a family of counting processes. *The Annals of Statistics*, 6(4):701–726.

Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica*, 70(1):223–262.

Albrecht, B. C. and Traina, J. (2026). Abduction and the demand curve. Working paper.

Andersen, P. K. and Gill, R. D. (1982). Cox's regression model for counting processes: A large sample study. *The Annals of Statistics*, 10(4):1100–1120.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262.

Berry, S. T. and Haile, P. A. (2014). Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bowsher, C. G. (2007). Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics*, 141(2):876–912.

Carrasco, M. and Florens, J.-P. (2000). Generalization of GMM to a continuum of moment conditions. *Econometric Theory*, 16(6):797–834.

Chacko, G. and Viceira, L. M. (2003). Spectral GMM estimation of continuous-time processes. *Journal of Econometrics*, 116(1–2):259–292.

Chen, S. X., Peng, L., and Yu, C. L. (2013). Parameter estimation and model testing for Markov processes via conditional characteristic functions. *Bernoulli*, 19(1):228–251.

Cont, R., Stoikov, S., and Talreja, R. (2010). A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B*, 34(2):187–220.

Diebold, F. X., Gunther, T. A., and Tay, A. S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39(4):863–883.

Doi, M. (1976). Second quantization representation for classical many-particle system. *Journal of Physics A: Mathematical and General*, 9(9):1465–1477.

Duffie, D. and Singleton, K. J. (1993). Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007.

Engle, R. F. and Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162.

Ethier, S. N. and Kurtz, T. G. (1986). *Markov Processes: Characterization and Convergence*. Wiley, New York.

Feuerverger, A. and Mureika, R. A. (1977). The empirical characteristic function and its applications. *The Annals of Statistics*, 5(1):88–97.

Frey, S. and Grammig, J. (2006). Liquidity supply and adverse selection in a pure limit order book market. *Empirical Economics*, 30(4):1007–1033.

Gallant, A. R. and Tauchen, G. (1996). Which moments to match? *Econometric Theory*, 12(4):657–681.

Gardiner, C. W. (2009). *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer, Berlin, 4 edition.

Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hansen, L. P. and Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, 50(5):1269–1286.

Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90.

Jiang, G. J. and Knight, J. L. (2002). Estimation of continuous-time processes via the empirical characteristic function. *Journal of Business & Economic Statistics*, 20(2):198–212.

Jiang, G. J. and Knight, J. L. (2010). ECF estimation of Markov models where the transition density is unknown. *The Econometrics Journal*, 13(2):245–270.

Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.

Metzner, P., Dittmer, E., Jahnke, T., and Schütte, C. (2007). Generator estimation of Markov jump processes. *Journal of Computational Physics*, 227(1):353–375.

Peliti, L. (1985). Path integral approach to birth-death processes on a lattice. *Journal de Physique*, 46(9):1469–1483.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1):1–48.

Singleton, K. J. (2001). Estimation of affine asset pricing models using the empirical characteristic function. *Journal of Econometrics*, 102(1):111–141.

van Kampen, N. G. (2007). *Stochastic Processes in Physics and Chemistry*. North-Holland, Amsterdam, 3 edition.

Weber, M. F. and Frey, E. (2017). Master equations and the theory of stochastic path integrals. *Reports on Progress in Physics*, 80(4):046601.