

ESTIMATION OF A HIGH-DIMENSIONAL COUNTING PROCESS WITHOUT PENALTY FOR HIGH-FREQUENCY EVENTS

LUCA MUCCIANTE

Royal Holloway University of London

ALESSIO SANCETTA 

Royal Holloway University of London

This paper introduces a counting process for event arrivals in high-frequency trading, based on high-dimensional covariates. The novelty is that, under sparsity conditions on the true model, we do not need to impose any model penalty or parameters shrinkage, unlike Lasso. The procedure allows us to derive a central limit theorem to test restrictions in a two-stage estimator. We achieve this by the use of a sign constraint on the intensity which necessarily needs to be positive. In particular, we introduce an additive model to extract the nonlinear impact of order book variables on buy and sell trade arrivals. In the empirical application, we show that the shape and dynamics of the order book are fundamental in determining the arrival of buy and sell trades in the crude oil futures market. We establish our empirical results mapping the covariates into a higher-dimensional space. Consistently with the theoretical results, the estimated models are sparse in the number of parameters. Using this approach, we are also able to compare competing model hypotheses on the basis of an out-of-sample likelihood ratio type of test.

1. INTRODUCTION

Counting processes are continuous time stochastic processes with nondecreasing trajectories, taking values in the set of positive integers. This paper is concerned with the estimation of the intensity of a counting process that depends on high-dimensional covariates. In particular, we are interested in modeling the intensity of high-frequency trading events using a possibly large number of covariates, where the impact of each of them can be nonlinear. This further increases the dimension when the unknown nonlinearity is modeled by a series expansion or similar procedures. The motivation of this study is to use information from the order book to model the intensity of buy and sell arrivals. The importance of order

We are very grateful to the Editor (Peter Phillips), the Co-Editor (Eric Renault), the Associate Editor, and the Referees for their detailed comments that have led to substantial improvements both in content and presentation. Address correspondence to Alessio Sancetta, Department of Economics, Royal Holloway University of London, Egham TW20 0EX, UK; e-mail: asancetta@gmail.com.

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

book variables has already been shown by various authors with varying degrees of complexity (Hall and Hautsch, 2007; Cont, Kukanov, and Stoikov, 2014; Kercheval and Zhang, 2015; Sancetta, 2018). The intensity of the counting process can be used to predict buy or sell trade arrivals at infinitesimal timescales, using relevant microstructure variables such as order book volume imbalances, order flow, and spread as covariates.

The use of counting processes in high-frequency financial modeling was pioneered by Engle and Russell (1998). Since then, they have acquired an increasing popularity in the literature (Bauwens and Hautsch, 2009, for a survey). A possible way to characterize a counting process is via its intensity. Intuitively speaking, the intensity is the instantaneous rate of occurrence of events conditional on the past history.

For definiteness, let $N := (N(t))_{t \geq 0}$ be a counting process starting at zero, and $\lambda^* := (\lambda^*(t))_{t \geq 0}$ a predictable process, both adapted to a filtration $\mathcal{F} = (\mathcal{F}_t)_{t \geq 0}$ and such that $M := (M(t))_{t \geq 0}$, where $M(t) := (N(t) - \int_0^t \lambda^*(r) dr)$, is an \mathcal{F}_t -martingale. The process λ^* is an \mathcal{F}_t -intensity of N . We assume that there is a predictable process $X = (X(t))_{t \geq 0}$ that takes values in $[0, 1]^K$ where K is large relatively to the sample period and

$$\lambda^*(t) = X(t)' b^*, \tag{1}$$

where $b^* \in [0, \infty)^K$ is an unknown sparse parameter. All vector-valued quantities are column vectors, and the prime symbol $'$ stands for transpose. By sparse, we mean that b^* has a small number of nonzero coefficients, relatively to K . We assume that the covariates are positive stochastic processes. Together with the assumption that the true coefficient vector b^* is nonnegative, this ensures that the intensity is a positive stochastic process. We shall show that the nonnegativity restriction naturally arises in some parameterization. Then, this positivity constraint inherits a regularization property similar to Lasso. The goal of the paper is to find an estimator for b^* when we observe a single trajectory/sample from $(N(t), X(t))_{t \geq 0}$ over a time interval $[0, T]$ with $T \rightarrow \infty$.

The empirical application in Section 3 further motivates the model, and Section 3.2 provides further discussion on the scope and limitations of the modeling strategy.

1.1. Remarks on the Model Restrictions

If the covariates are bounded, the restriction to $[0, 1]^K$ is a mere linear transformation. Most covariates obtained from the order book satisfy this condition by construction. Examples include order book volume imbalances (see (10) in Section 3.1). On the other hand, variables such as durations are not bounded. However, we can always find transformations of the data to map variables into a bounded range. The type of transformation is a function of the modeling objectives. We provide more concrete remarks in Section 3.2.

The assumptions that we shall use essentially imply that the covariates process X is ergodic. We do not assume any stationarity. This is the weakest possible assumption for econometric inference. Although we do not do so in the empirical application, one example of the flexibility of the framework is to consider nonlinear Hawkes processes as one of the covariates, as long as they satisfy some stability conditions (Brémaud and Massoulié, 1996). In this case, one covariate could be set equal to $\varphi\left(\int_{(0,t)} h(t-s) dN(s)\right)$ where $\varphi(\cdot)$ is a Lipschitz function with range in $[0, 1]$, and $h(\cdot)$ is positive and integrable. If h is unknown, we could suppose a finite set of such functions $\{h^{(l)} : l = 1, 2, \dots, L\}$ and generate as many covariates as L to capture any self-exciting nature of the intensity. Alternatively, the kernel function h would have to be estimated using a sample antecedent to the one used in the estimation of the model.

1.2. Contributions and Relation to Other Work

This paper contributes to the characterization of the impact of order book variables in the intensity of buy and sell trade arrivals. In order to do so in the most robust and simple to interpret way, it relies on a methodology that is elementary from an econometric point of view, but powerful.

The literature on modeling the arrival of high-frequency events using order book variables is growing (inter alia Hall and Hautsch, 2007; Cont et al., 2014; Sancetta, 2018; Morariu-Patrichi and Pakkanen, 2022). This is an important practical problem. Algorithmic traders do track the order book (MacKenzie, 2017). The order book not only contains information about liquidity, but also helps to identify informed traders. Nowadays, unlike traditional models of informed trading (Glosten and Milgrom, 1985; Kyle, 1985), sophisticated informed investors prefer to rely on passive execution. This means that they would avoid buying at the ask and paying the spread. Instead, they would break their orders into small ones and patiently fill most of them joining the bid price. There are two reasons for this. First, informed traders want to reduce their cost. Second, they want to reduce the amount of signaling. It turns out that uninformed market participants are smart and do look at quotes to infer information beyond liquidity. In fact, our empirical results show that the intensity of buy arrivals is a nonlinear increasing function of quoted buy orders over quoted sell sizes. The empirical application focuses on crude oil futures as main instrument together with information from other auxiliary instruments.

The model in (1) is similar to Aalen (1980) multiplicative intensity model. There, we observe i.i.d. copies in the presence of censoring. For each copy, the covariates are determined at time zero, and so they are usually referred to as marks. Here, we observe a single trajectory of the data. The covariates are continuous-time stochastic processes and change as new information becomes available during trading. The absence of censoring leads to a simpler estimation procedure that does not require the use of U-statistics. U-statistics would be impractical for the sample sizes considered in this paper. The high-dimensional version of Aalen

multiplicative intensity model has been considered by several authors (e.g., Gaïffas and Guillaux, 2012, and the references therein).

From a technical point of view, our results are derived minimizing a least-squares criterion for count processes. In our problem, we impose a nonnegativity constraint without a penalty term, unlike what is usually done in the literature (Gaïffas and Guillaux, 2012; Alaya, Gaïffas, and Guillaux, 2015, and the references therein). Then, we can directly estimate the model using standard quadratic programming with no need to use a link function that ensures positivity. Hence, this work is related to results for nonnegative least squares. Under a sparsity assumption on the vector of regression coefficients, the sign constraint imposes restrictions that lead to a regularization as effective as Lasso with no need to tune a penalty parameter. We obtain results equivalent to nonnegative least squares with i.i.d. Gaussian errors (Meinshausen, 2013; Slawsky and Hein, 2013). However, in our problem, we need to control a dependent continuous-time process.

The use of quadratic programming is important in our context. In practice, models can be calibrated on months of data and the size of high-frequency data be so large that cannot be held in ready access memory. However, the problem is written in terms of sufficient statistics whose dimension does not grow with the longitude of the sample, and hence they can be held in memory making estimation very simple.

The plan for the paper is as follows. The next section describes the estimation procedure and the assumptions for validity of the statistical procedure. Then, we show that we can obtain consistency for a high-dimensional counting process with no need to use a penalty or additional tuning parameters. We conclude the section with a central limit theorem for a two-stage estimator for the parameters. Section 3 applies the results to the estimation of a point process for buy and sell arrivals where the intensity depends on order book variables in a nonlinear but additive way. In Section 3.2, we discuss the scope and limitation of the specific model we choose. Section 4 contains some final remarks. Proofs are in Section A.1 of the Supplementary Material. Additional details, which we may refer to, can also be found in the Supplementary Material. This includes a finite-sample study using simulations.

2. ASSUMPTIONS AND RESULTS

2.1. The Estimation Problem

Given (1), our estimator \hat{b} for b^* is the solution of the constrained problem

$$\min_{b \geq 0} \left\{ -2 \int_0^T X(t)' b dN(t) + \int_0^T (X(t)' b)^2 dt \right\}. \quad (2)$$

This is a standard quadratic programming problem. Throughout, all vector inequalities are meant elementwise. Replacing the constraint with a penalty, this objective

function has been used for high-dimensional problems by several authors (e.g., Gaïffas and Guilloux, 2012; Alaya et al., 2015).

Note that the population version of the objective function is proportional to

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[-2 \frac{1}{T} \int_0^T X(t)' b dN(t) + \frac{1}{T} \int_0^T (X(t)' b)^2 dt \right],$$

and using the definition of the intensity, this is equal to

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[-2 \frac{1}{T} \int_0^T X(t)' b \lambda^*(t) dt + \frac{1}{T} \int_0^T (X(t)' b)^2 dt \right]. \quad (3)$$

When $\lambda^*(t) = X(t)' b^*$, it is easy to see that the constrained minimizer of the above display is $b = b^*$, when the constraint $b \geq 0$ holds for b^* as well. In some situations, a positivity constraint on b is meaningful. Then, we do not need to use a penalty to find consistent estimators in high dimensions. We shall give examples in Section 2.3.

2.2. Assumptions

We introduce some additional notation and terminology. Fix an arbitrary positive T . We denote the first $n = N(T)$ jump times of N by $T_0 < T_1 < \dots < T_n \leq T$, where $T_0 = 0$. The integral of (1) is the compensator of $N(t)$ and $M(t) := (N(t) - \int_0^t \lambda(r) dr)$ is an \mathcal{F}_t -martingale (see the discussion around (1)). Throughout, we use index subscripts to denote the relevant entry in either vectors or matrices. Let $S = \{i \leq K : b_i^* > 0\}$ be the set of nonzero entries of b^* , $S^c = \{i \leq K : b_i^* = 0\}$ be its complement, and $s = |S|$ be the cardinality of S . For an arbitrary vector $a \in \mathbb{R}^K$ and $U \subset \{1, 2, \dots, K\}$ we denote by $a_U \in \mathbb{R}^{|U|}$ the $|U|$ -dimensional subvector of a obtained by removing all the entries with index not in the set U . Define $\hat{\Sigma} := \frac{1}{T} \int_0^T X(t) X(t)' dt$ and $\hat{\Sigma}_S := \frac{1}{T} \int_0^T X_S(t) X_S(t)' dt$. We use $\|\cdot\|_p$ to denote the ℓ_p norm $p \in (0, \infty]$. Finally, we write w.p.1. to mean with probability going to 1.

Assumption 1 (Model assumption). The point process admits the intensity (1) which is supposed to be uniformly bounded by a constant $\bar{\lambda}$ (possibly going to infinity), $X := (X(t))_{t \geq 0}$ is a predictable process with values in $[0, 1]^K$ for each $t \geq 0$, and $b^* \in [0, \infty)^K$.

One main restriction is that the intensity is bounded by a constant $\bar{\lambda}$. It could be relaxed to a moment condition, but at the cost of additional technical complications. In this case, it would not be possible to obtain an error bound that is logarithmic in K .

Assumption 2 (Eigenvalues condition). There is a constant $\phi_{\min} > 0$ such that the eigenvalues of $\hat{\Sigma}_S$ are all greater than ϕ_{\min} , w.p.1.

Let $L > 0$ and $\mathcal{R}(L, S) := \{b : \|b_{S^c}\|_1 \leq L\|b_S\|_1\}$, where S is the index set of active variables. The (L, S) restricted ℓ_1 -eigenvalue of a matrix A is defined as

$$\phi_{comp}^2(A, L, S) := \min \left\{ s \frac{b'Ab}{\|b\|_1^2} : b \in \mathcal{R}(L, S) \right\}. \tag{4}$$

A lower bound on (4) is the weakest assumption used to derive oracle inequalities for Lasso (van de Geer and Bühlmann, 2009). We shall use the following.

Assumption 3 (Compatibility condition). There is a strictly positive constant ϕ such that $\phi_{comp}^2(\hat{\Sigma}, \frac{3}{\sqrt{\nu}}, S) \geq \phi$ w.p.1, where ν is as in Assumption 4.

The positively constrained minimal ℓ_1 -eigenvalue of a matrix A is defined as

$$\phi_{pos}^2(A) := \min \left\{ \frac{b'Ab}{\|b\|_1^2} : \min_i b_i \geq 0 \right\}. \tag{5}$$

A lower bound on (5) has been used by Meinshausen (2013) in the context of nonnegative least squares.

Assumption 4 (Positive eigenvalue condition). There is a $\nu > 0$ such that $\phi_{pos}^2(\hat{\Sigma}) \geq \nu$ w.p.1.

We shall refer to Assumptions 1–4 simply as the Assumptions.

As we shall discuss next, there are reasons to make $\bar{\lambda}$, ϕ_{min} , ϕ , and ν depend on T . Hence, we shall allow $\bar{\lambda} = \bar{\lambda}(T) \rightarrow \infty$, and $\phi_{min} = \phi_{min}(T)$, $\phi = \phi(T)$, $\nu = \nu(T)$ to go to zero as $T \rightarrow \infty$, if needed. For ease of notation, we drop the dependence on T in what follows.

2.3. Remarks on the Assumptions

2.3.1. *Assumption 1.* We view the sign constraint as a hypothesis, and estimation is carried out under this hypothesis. Specifically for the problem of high-frequency trading, we have a priori knowledge whether the marginal impact of a high-frequency order book covariate is increasing or decreasing. Such information can be obtained from other studies (e.g., Kercheval and Zhang, 2015; Sancetta, 2018). Hence, given such information on the direction of the impact, we show how to improve estimation in high-dimensional problems.

For the sake of clarity, we now consider two examples. Consider the intensity $\lambda(Z(t)) = a_0 + a_1Z_1(t) - a_2Z_2(t)$ that depends on the covariate $Z(t) = [Z_1(t), Z_2(t)]'$ with values in $[0, 1]^2$, where $a_i \geq 0$, $i = 0, 1, 2$, $a_0 - a_2 \geq 0$. The parameters' restriction ensures that this intensity is always nonnegative. Then, λ can be written as (1) where $X_1(t) = 1$, $X_2(t) = Z_1(t)$, $X_3(t) = 1 - Z_2(t)$, and $b_1^* = a_0 - a_2$, $b_2^* = a_1$, $b_3^* = a_2$, where $b^* \geq 0$. Hence, in our framework, we are able to control the direction of the impact by the linear transformation $x \mapsto 1 - x$. From a computational point of view, this is equivalent to changing the sign of the covariate and imposing an additional inequality constraint.

Furthermore, possible nonlinearity of the impact of a covariate can lead to high-dimensional problems. For example, suppose that $Z(t)$ takes values in $[0, 1]$ and $\lambda(Z(t)) = g(Z(t))$ where $g: [0, 1] \rightarrow [0, \infty)$ is an unknown continuous function. Define

$$\lambda(Z(t)) = \sum_{i=0}^K a_i B_{i,K}(Z(t)), \quad (6)$$

where $B_{i,K}(z) = \binom{K}{i} z^i (1-z)^{K-i}$ is the i th term in a K th-order Bernstein polynomial. By linearity, this is in the form of (1) when $X_i(t) := B_{i-1,K}(Z(t))$ and $b_i^* = a_{i-1}$, $i = 1, 2, \dots, K+1$. The coefficients of Bernstein polynomials have a clear physical interpretation for $K \rightarrow \infty$:

$$\sup_{z \in [0, 1]} \left| g(z) - \sum_{i=0}^K a_i B_{i,K}(z) \right| \rightarrow 0$$

when we define $a_i = g(i/K)$ (Lorentz, 1986, Thm. 1.1.1). Hence, the above display suggests that if $g \geq 0$, the assumption that $a_i \geq 0$ is almost necessary when K is large. The argument can be extended to a dimension greater than 1. However, we note that the assumptions on $\hat{\Sigma}$ will fail as $K \rightarrow \infty$, unless we allow ϕ and ν to go to zero slowly enough. This is allowed in our results (Theorems 1 and 2).

In the empirical section of this paper, we view the sign constraint as a hypothesis on the direction of the impact of the covariates. Relying on several hypotheses, we estimate models and compare their performance out of sample (see Section 3 for details). This remark would suggest that impact/sign misspecification would lead to a zero coefficient. Consider the first example in the above discussion. By impact misspecification, we mean using $X_3(t) = 1 - Z_2(t)$ when in fact the impact of $Z_2(t)$ on the intensity is positive. Because of the influence from other variables, the procedure may select a covariate with the wrong impact even when solving the population objective function (3). We carried out a number of numerical examples to find the solution to (3). We found that the constraint tends to be binding for a true negative coefficient when both s and the number of wrongly signed variables are small. However, as either s or the degree of misspecification increases, the constrained population estimator may have a positive sign even when the true one is negative. For the designs we considered in the Supplementary Material, selection of a misspecified variable when solving (3) was relatively infrequent. Details can be found in Section A.4 of the Supplementary Material.

We allow $\bar{\lambda}$, the constant upper bound on the intensity λ^* , to possibly grow to infinity. This is relevant in practice, as we may only have a crude upper bound for λ^* that depends on s , the number of active variables. To see this, note that by assumption, $\lambda^* = X'b^* \leq \|b_S\|_1 = O(s)$. Even if crude, with no further information on either the covariates or the coefficients, this is the best possible upper bound on λ^* . Hence, it is relevant to allow $\bar{\lambda} = O(s)$ and s to diverge to infinity with T . To see how information can be used to find a tighter bound on $\bar{\lambda}$, consider the

intensity in (6). Then, we know that the intensity is bounded by $\max_{i \leq K} a_i$, which is the largest of the coefficients to be estimated.

2.3.2. *Assumption 2.* Unlike Meinshausen (2013) to make the proofs simpler, we use Assumption 2. This is just slightly stronger than the Compatibility Condition $\phi_{comp}^2(\hat{\Sigma}, 0, S) \geq \rho$ for some $\rho > 0$ w.p.1, which is needed in the proofs. To see this, note that, by the nonnegativity of the coefficients, the latter implies that

$$s \sum_{i,j \in S} b_i b_j \hat{\Sigma}_{i,j} \geq \rho \left(\sum_{i \in S} |b_i| \right)^2$$

w.p.1, where $\hat{\Sigma}_{i,j}$ is the i, j entry in $\hat{\Sigma}$. Suppose that ρ is a lower bound on the smallest eigenvalue of $\hat{\Sigma}_S$ w.p.1. Then,

$$s \sum_{i,j \in S} b_i b_j \hat{\Sigma}_{i,j} \geq \rho s \sum_{i \in S} b_i^2$$

w.p.1. Using the fact that the ℓ_1 norm of an s -dimensional vector is bounded by \sqrt{s} times its ℓ_2 norm, this means that $\phi_{comp}^2(\hat{\Sigma}, 0, S) \geq \rho$ holds w.p.1. In practice, this means that the covariates with index in the active set S need to be linearly independent w.p.1, when $T \rightarrow \infty$. This is plausible, as s should be relatively small, i.e., b^* is sparse.

It is instructive to consider conditions on X that imply this assumption. Recall that $\hat{\Sigma}_{i,j} := T^{-1} \int_0^T X_i(t) X_j(t)$. Suppose that $(X_i(t) X_j(t))_{t \geq 0}$ is ergodic, for all $i, j \in S$, in the sense that $\hat{\Sigma}_{i,j} \rightarrow \Sigma_{i,j}$ a.s., where $\Sigma_{i,j}$ is a constant, $i, j \in S$. By boundedness of the covariates, the convergence must also hold in L_2 . Therefore, if S has bounded cardinality, $\hat{\Sigma}_S - \Sigma_S$ converges to zero in expected Frobenius norm. This implies convergence of the eigenvalues of $\hat{\Sigma}_S$ to the ones of Σ_S (Bosq, 2000, Thm. 4.4). Incidentally, boundedness implies that all the moments of $\hat{\Sigma}_{i,j}$ converge to $\Sigma_{i,j}$, $i, j \in S$.

The above definition of ergodic is closely linked to the existence of an asymptotic mean stationary measure such that Σ_S is the expectation of $\hat{\Sigma}_S$ with respect to that measure (see Gray and Kieffer, 1980, Thm. 1, and Gray, 2009, Chap. 6, for an extensive treatment). The argument does not require stationarity of $(X_i(t) X_j(t))_{t \geq 0}$. As a simple example of why this can be relevant, consider $\mathbb{E} X_i^2(t) = \sigma_i^2(t)$. It is common in high-frequency data to have time-varying (unconditional) volatility of certain quantities, due to intraday nonstationarity. However, by ergodicity and boundedness, there is a constant $\sigma_i^2 = \lim_T \frac{1}{T} \int_0^T \sigma_i^2(t) dt$. In practice, this is plausible when T spans multiple days.

Finally, we note that if s increases, the dependence among the active variables may increase, making a lower bound ϕ_{min} on the smallest eigenvalue smaller. For this reason, we allow the bounds in our results to explicitly depend on ϕ_{min} .

2.3.3. *Assumption 3.* We use the same Compatibility Condition as in Meinshausen (2013). Meinshausen (2013) also uses $\phi_{comp}^2(\hat{\Sigma}, 0, S)$. As discussed above, by Assumption 2, we do not need this additional condition. To verify the Compatibility Condition using population quantities, we can use approximations. Suppose that there is a sequence $\epsilon_T \rightarrow 0$ such that $\max_{i,j} |\hat{\Sigma}_{i,j} - \Sigma_{i,j}| = O_P(\epsilon_T)$, where $\Sigma_{i,j}$ is understood to be the limit (in probability) of $\hat{\Sigma}_{i,j}$, $i, j \in \{1, 2, \dots, K\}$. If the number of active variables satisfies $s = o_P(v\epsilon_T^{-1})$ with v as in Assumption 3, then we can deduce that the Compatibility Condition is satisfied if $v > 0$ and the smallest eigenvalue of Σ is strictly positive (van de Geer and Bühlmann, 2009, Cor. 10.1 and discussion). A bound on the minimal eigenvalue of Σ requires that the covariates are not asymptotically dependent. To establish such a bound, we will need rates of convergence. Ergodicity alone, as discussed in the previous remark, is not sufficient even when S has finite cardinality. To see this, note that when s is bounded we still need $\max_{i,j} |\hat{\Sigma}_{i,j} - \Sigma_{i,j}| = o_P(1)$ for $i, j \in \{1, 2, \dots, K\}$ and not just for $i, j \in S$. Given that K is not necessarily bounded, this requires convergence rates.

We also note that as K increases, the smallest eigenvalue of Σ may tend to zero. For this reason, our bounds are in terms of ϕ to show how quickly the convergence rate may deteriorate if $\phi \rightarrow 0$.

2.3.4. *Assumption 4.* This assumption is the same as in Meinshausen (2013). For Assumption 4 to hold, it is sufficient that $\min_{i,j} \hat{\Sigma}_{i,j} > 0$ a.s., $i, j \in \{1, 2, \dots, K\}$. This is satisfied if the covariates $X_i(t)$ do not have a disjoint support for all $t \geq 0$. Then, the average of $X_i(t)X_j(t)$ is greater than zero because $X_i(t) \in [0, 1]$, $t \geq 0$, unless some of the covariate are exactly zero. This condition is satisfied in our empirical study. However, there are parameterizations that lead to disjoint subsets. Notable examples are splines and one-hot encoding (Alaya et al., 2019). One-hot encoding essentially builds mutually disjoint bins for each covariate and constructs dummy variables for each bin. This leads to mutually disjoint covariates. As discussed in Meinshausen (2013, Exam. III), the following setup covers splines and one-hot encoding, and satisfies Assumption 4. Fix a positive integer L , and suppose that $\{P_l : l = 1, \dots, L\}$ is a partition of $\{1, 2, \dots, K\}$ such that $\hat{\Sigma}_{i,j} > \nu L$ a.s., if $i, j \in P_l$ for some l and zero if $i \in P_l, j \in P_k$ when $k \neq l$. This means that covariates that have an index in different partitions have disjoint support. In this case,

$$b' \hat{\Sigma} b = \sum_{l=1}^L \sum_{i,j \in P_l} b_i b_j \hat{\Sigma}_{i,j} \geq \nu L \sum_{l=1}^L \sum_{i,j \in P_l} b_i b_j = \nu L \sum_{l=1}^L \left(\sum_{i \in P_l} b_i \right)^2.$$

Then, using the fact that, by Jensen’s inequality, for any constants a_1, a_2, \dots, a_L , we know that $\left(\sum_{l=1}^L a_l \right)^2 \leq L \sum_{l=1}^L a_l^2$, we deduce that

$$b' \hat{\Sigma} b \geq \nu \left[\sum_{l=1}^L \left(\sum_{i \in P_l} b_i \right) \right]^2 = \nu \|b\|_1^2$$

when $\min_i b_i \geq 0$. Then, Assumption 4 is satisfied.

As mentioned above, we require $\hat{\Sigma}_{i,j} > \nu L$. Given that $\hat{\Sigma}_{i,j} \in [0, 1]$, we need $\nu \rightarrow 0$ if the number of partitions L diverges to infinity. This is of interest for certain nonparametric estimators like one-hot encoding and splines with an increasing number of knots.

We conclude this section giving an intuition on why Assumption 4 leads to results comparable to ℓ_1 penalization. By Assumption 4, $b' \hat{\Sigma} b \geq \nu \|b\|_1^2$ w.p.1 as long as $b \geq 0$. Hence, bounds on $b' \hat{\Sigma} b$ will translate into bounds for $\|b\|_1^2$. However, as soon as $\nu \rightarrow 0$, the control of $\|b\|_1^2$ using $b' \hat{\Sigma} b$ becomes loose and the procedure will underperform standard estimation with an ℓ_1 penalty. Similar remarks, but in the context of regression analysis, are made in Meinshausen (2013).

2.4. Asymptotic Results

2.4.1. *Consistency.* We shall keep track of all the constants to see how the bound is affected. This is useful, for example, if $\phi_{\min} \rightarrow 0$ slowly enough. Similarly, we can allow $\bar{\lambda} \rightarrow \infty$. This is important, because $T/\bar{\lambda}$ is one of the main quantities affecting the convergence rate. Define $c(s) := \max \left\{ \frac{s^2}{\phi^2}, \frac{1}{\nu} \right\}$ and $\mu_T := \frac{1}{T} \int_0^T \mathbb{E} \lambda^*(t) dt$. Note that $\mu_T \leq \bar{\lambda}$. We have consistency of the estimator \hat{b} for b^* , under the ℓ_1 norm. Throughout, to simplify the notation, we assume that $K \geq 2$ in all the results that follow.

THEOREM 1. *Under the Assumptions, if $\log K = O(T\bar{\lambda})$, then*

$$\|\hat{b} - b^*\|_1 = O_P \left(\sqrt{\frac{c(s) (s^2 \mu_T \phi_{\min}^{-2} + \bar{\lambda} \log K)}{T}} \right).$$

The second result is an estimation of the prediction error.

THEOREM 2. *Under the Assumptions, if $\log K = O(T\bar{\lambda})$, then*

$$\frac{1}{T} \int_0^T \left(X(t)' \hat{b} - X(t)' b^* \right)^2 dt = O_P \left(\frac{c^{1/2}(s) (s^2 \mu_T \phi_{\min}^{-2} + \bar{\lambda} \log K)}{T} \right).$$

Theorems 1 and 2 make explicit the dependence on the parameters ϕ , ϕ_{\min} , ν , and $\bar{\lambda}$ defined in the Assumptions. Hence, we can have consistency even when $\phi_{\min}, \phi, \nu \rightarrow 0$ and $\bar{\lambda} \rightarrow \infty$ at suitable rates as $T \rightarrow \infty$. When ν, ϕ, ϕ_{\min} are fixed, $c(s) = O(s^2)$. Then, we have the following corollary.

COROLLARY 1. *Under the Assumptions, if $\bar{\lambda}$ is fixed and bounded away from infinity, ϕ_{\min}, ϕ, v are fixed and bounded away from zero, and $\log K = O(T\bar{\lambda})$, then*

$$\|\hat{b} - b^*\|_1 = O_p\left(\sqrt{\frac{s^4 + s^2 \log K}{T}}\right) \quad (7)$$

and

$$\frac{1}{T} \int_0^T (X(t)' \hat{b} - X(t)' b^*)^2 dt = O_p\left(\frac{s^3 + s \log K}{T}\right) \quad (8)$$

hold true.

The convergence rate derived here is typical of high-dimensional estimation problems under some form of regularization. For example, the convergence rate in our error bounds is similar to the one derived for the nonnegative least-squares regression problem with i.i.d. Gaussian errors (Meinshausen, 2013, Thms. 1 and 2, and Slawsky and Hein, 2013, Thm. 2). We can relate to those results, assuming that the conditions of Corollary 1 hold. Then, we have ℓ_1 consistency of the estimator if $T^{-1}(s^4 + s^2 \ln K) \rightarrow 0$ (see (7)). On the other hand, Meinshausen (2013, first part of his Thm. 1) says that, in the regression case, we have ℓ_1 consistency if $n^{-1} s^4 \ln K \rightarrow 0$, where n is the sample size.

Suppose that the regressors are bounded. Under a lower bound on the smallest nonzero elements in the regressor coefficients, results in Meinshausen (2013, Thm. 2) imply (empirical) L_2 consistency of the regression function when $n^{-1} s \ln K \rightarrow 0$. With no such condition on the regression coefficients, in the present context, we can obtain L_2 consistency of the intensity estimator when $T^{-1}(s^3 + s \ln K) \rightarrow 0$ (see (8)). We can also consider the L_2 consistency of the estimated intensity in Gaïffas and Guilloux (2012) for the Aalen intensity model, using a data-driven Lasso penalty. Using a restricted eigenvalue condition, Gaïffas and Guilloux (2012, Thm. 2) achieve consistency if $n^{-1} s \ln K \rightarrow 0$, where, again, n is the sample size. Given that the model setup differs, and we cannot rely on a penalty, the method of proof in Gaïffas and Guilloux (2012) is different. It can be difficult to discuss convergence rates beyond the aforementioned remarks.

The next result is about set identification.

COROLLARY 2. *Using the notation in Theorem 1, define $\rho := \sqrt{c(s)\bar{\lambda}(s^2\phi_{\min}^{-2} + \log K)/T}$. Suppose that the conditions of Theorem 1 hold. Furthermore, suppose that $\min_{i \in S} b_i^* > \kappa$ for some κ such that $\kappa/\rho \rightarrow \infty$. Let $\hat{S} = \{i \leq K : \hat{b}_i > 0\}$. Then, $\Pr(\hat{S} \subset S) \rightarrow 0$.*

Let $\hat{S}_\epsilon = \{i \leq K : \hat{b}_i > \epsilon\}$. Then, under the above conditions, for any ϵ such that $\epsilon/\rho \rightarrow \infty$ and $\kappa/\epsilon \rightarrow c > 1$, $\Pr(\hat{S}_\epsilon \neq S) \rightarrow 0$.

Corollary 2 says that \hat{S} , the estimated support of b^* , is a superset of S , the true support, w.p.1. Using a threshold on the coefficients, we can achieve set identification, w.p.1. Note that we can have $\kappa \rightarrow 0$ slowly enough and similarly for ϵ .

2.4.2. *Convergence in Distribution.* We obtain a central limit theorem for the ordinary least squares estimator for b^* . At first, we estimate \hat{b} in (2) and obtain \hat{S}_ϵ as in Corollary 2. Then, we compute b_ϵ^{OLS} , which is the K -dimensional vector such that its entries with index in \hat{S}_ϵ are equal to

$$\left(\int_0^T X_{\hat{S}_\epsilon}(t) X_{\hat{S}_\epsilon}(t)' dt \right)^{-1} \left(\int_0^T X_{\hat{S}_\epsilon}(t) dN(t) \right),$$

while all the other entries are zero. Note that the cardinality of \hat{S}_ϵ can still grow with T . Given a fixed K -dimensional vector α satisfying $\alpha' \alpha = 1$ and $\alpha'_{\hat{S}_\epsilon} \alpha_{\hat{S}_\epsilon} > 0$, we are interested in the asymptotic distribution of $\sqrt{T} \alpha' (b_\epsilon^{OLS} - b^*)$ to conduct inference. Define $\hat{\Sigma}^N := \frac{1}{T} \int_0^T X(t) X(t)' dN(t)$ and $\hat{\Sigma}_S^N$ to be the submatrix that includes only the entries with row and columns indices in S . We state an additional assumption.

Assumption 5. We have that $\lim_T \sum_{i,j \in S} \left[\text{Var} \left(\hat{\Sigma}_{i,j} \right) + \text{Var} \left(\hat{\Sigma}_{i,j}^N \right) \right] = o(\phi_{\min})$, where ϕ_{\min} is as in Assumption 2, and $\mathbb{E} \hat{\Sigma}_S$ and $\mathbb{E} \hat{\Sigma}_S^N$ both converge to full-rank constant matrices.

Note that $\text{Var} \left(\hat{\Sigma}_{i,j} \right) = o(1)$ under ergodicity assumptions on $X_i X_j$. As already mentioned, ergodicity and boundedness imply that $\hat{\Sigma}_{i,j}$ converges in L_2 to a constant $\Sigma_{i,j}$. By boundedness, we also know that $\mathbb{E} \hat{\Sigma}_{i,j}$ converges to a constant limit $\Sigma_{i,j}$. However, we need the rate of convergence to be fast enough to ensure convergence of all the entries in $\hat{\Sigma}_S$. For example, this is trivially satisfied if s is bounded or if we assume that $\max_{i,j \in S} \text{Var} \left(\hat{\Sigma}_{i,j} \right) = O(\epsilon_T^2)$ such that $s = o_P(\epsilon_T^{-1})$. The same argument applies to $\hat{\Sigma}_S^N$. In this case, write

$$\hat{\Sigma}_{i,j}^N = \frac{1}{T} \int_0^T X_i(t) X_j(t)' (X(t)' b^*) dt + \frac{1}{T} \int_0^T X_i(t) X_j(t)' dM(t)$$

using the definition of M and intensity (see the discussion around (1)). Being a martingale, the second term on the right-hand side converges to zero in probability $i, j \in S$. Under ergodicity assumptions on $(X_i(t) X_j(t) X_k(t))_{t \geq 0}$ for $i, j, k \in S$, we also have that the first term on the right-hand side converges to a constant and the same applies to its moments because of dominated convergence. We can now state a central limit theorem for our estimator.

THEOREM 3. *Suppose that the assumptions of Corollary 2, Assumption 5, $\inf_{t>0} \lambda^*(t) > 0$, and $\lim_T \alpha'_S (\mathbb{E} \hat{\Sigma}_S)^{-1} \alpha_S > 0$ hold. If $s^2 = o(T \phi_{\min}^2 / \bar{\lambda})$, then $\sqrt{T} / \sigma_\alpha^2 \alpha' (\hat{b}_\epsilon^{OLS} - b^*) \rightarrow \mathcal{N}(0, 1)$ in distribution, where $\mathcal{N}(0, 1)$ is the standard normal distribution and*

$$\sigma_\alpha^2 = \lim_T \alpha'_S (\mathbb{E} \hat{\Sigma}_S)^{-1} (\mathbb{E} \hat{\Sigma}_S^N)^{-1} (\mathbb{E} \hat{\Sigma}_S)^{-1} \alpha_S.$$

Moreover, $\hat{\sigma}_\alpha^2 = \alpha'_{\hat{S}_\epsilon} \hat{\Sigma}_{\hat{S}_\epsilon}^{-1} \hat{\Sigma}_{\hat{S}_\epsilon}^N \hat{\Sigma}_{\hat{S}_\epsilon}^{-1} \alpha_{\hat{S}_\epsilon}$ is a consistent estimator for σ_α^2 .

Note that the conditions are for S rather than \hat{S}_ϵ . By Corollary 2, the two sets are the same, w.p.1. By assumption, $\min_{i \in S} b_i^* > \kappa$ for κ as in Corollary 2. When κ is fixed, the post-selection asymptotics in Theorem 3 are valid as they are uniform in b^* (Leeb and Pötscher, 2005). When κ is allowed to go to zero as in Corollary 2, post-selection asymptotics are still valid, but the convergence to a normal can be arbitrarily slow (mutatis mutandis; see Leeb and Pötscher, 2005, p. 29ff). Intuitively, this follows from the fact that as $\kappa \rightarrow 0$, it becomes harder to distinguish very small coefficients from zero coefficients.

Next, we shall define a nonlinear additive model for buy and sell trades based on order book and trade variables.

3. EMPIRICAL APPLICATION: ORDER BOOK DETERMINANTS OF CRUDE OIL BUY AND SELL TRADE ARRIVALS

We estimate the intensity of buy and sell trade arrivals separately. We investigate the impact of features constructed from quote and trade data on these buy and sell events. The features include order book imbalance, trade imbalance, spread, and durations. Information from quote data appears to contain much information (Hall and Hautsch, 2007; Cont et al., 2014; Kercheval and Zhang, 2015; Sancetta, 2018). For example, order book imbalances are a quantity that has been used by practitioners for more than two decades (MacKenzie, 2017). However, most of the attention in the econometric literature has been on trade data.

We use data from the front month of crude oil futures traded on the Chicago Mercantile Exchange (CME). The CME ticker is CL, and the sample period is May 1, 2013 to June 5, 2013, each day from 13:30 to 18:00 GMT. As auxiliary instruments, we use information from heating oil (HG), natural gas (NG) and the S&P500 (ES) futures, where the CME ticker is in parenthesis. The data were collected by a high-frequency proprietary trading firm collocated in the Aurora data center in Chicago. The data are at the highest level of granularity, comprising all quotes and trades, time stamped at nanosecond resolution.

3.1. The Model

We use different model specifications to capture the nonlinearities in the impact of the covariates on the intensity. We separately consider buy and sell trade arrivals.

TABLE 1. Covariates used for estimation

Variables	Short name	Smoothing
Volume imbalance level j	$\text{VolImb}_j, j = 1, 2, \dots, 5$	
Spread	Spread	
Trade imbalance	TrdImb98	$\alpha = 0.98$
Durations	Dur98	Dur90 $\alpha = 0.98$ and 0.90

Note: The column “Smoothing” reports the smoothing parameter used if an EWMA had been applied to the original variable.

We want to assess the extent to which these variables impact the intensity in a nonlinear way. We consider different hypotheses. To evaluate the hypotheses, we carry out a test for model performance. The results for the estimated models show that the impact is nonlinear as expected and allow us to characterize the shape of the impact. In order to formulate our hypotheses, we first define the covariates.

3.1.1. *The covariates.* The covariates are reported in Table 1. We apply exponential moving average (EWMA) filters to some of the covariates. The EWMA of a variable $X(t_i)$ with smoothing parameter α is

$$EWMA(X(t_i)) = \alpha EWMA(X(t_{i-1})) + (1 - \alpha)X(t_i), \quad (9)$$

where $EWMA(X(t_1)) = X(t_1)$. Here, t_1 is the time of the first update in the variable X at the start of each day, whereas t_i is the time of the i th update. EWMA's are computed for each day. Note that the covariates are updated at discrete times for each instrument and the update times are different from the trade update times T_j . We then sample the data at times that are the union of each covariate update and the times T_j for the traded instrument only, i.e., CL. This reduces the total number of updates within each day. Finally, to ensure that the covariates are predictable, we make them left continuous by lagging them after sampling.

Variables are mapped, by linear transformation, into $[0, 1]$. However, spread and durations are first capped and then scaled by the cap so that they take values in $[0, 1]$. The book volume imbalance at level $j \in \{1, 2, \dots, 5\}$, VolImb_j , is defined as

$$\text{VolImb}_j = \frac{\text{BidSize}_j - \text{AskSize}_j}{\text{BidSize}_j + \text{AskSize}_j}, \quad (10)$$

where BidSize_j is the bid size (quantity) of the j th level bid, and similarly for AskSize_j . This variable takes values in $[-1, 1]$. We map it to $[0, 1]$ by standard linear transformation: multiply by 2 and subtract by 1. Hence, a value of 0.5 corresponds to a book volume imbalance equal to zero. The trade imbalance is computed from the EWMA of the signed traded volume every time there is a trade. We then divide it by the EWMA of the unsigned volumes. The EWMA's parameter is $\alpha = 0.98$ for both denominator and numerator. Durations are in

seconds with nanosecond decimals, capped to 1 second. They are then passed to EWMA filters with parameter $\alpha = 0.98$ and 0.90 . The spread is capped to four ticks and standardized by 4. Hence, the minimum value it can take is 0.25 . Additional details regarding the data and the calculations are in Section A.2 of the Supplementary Material.

All variables are assumed to be mapped to $[0, 1]$ as described above. Then, if we hypothesize that the impact of a covariate is negative, we apply the transformation $x \mapsto 1 - x$ as discussed in Section 2.3. For example, we would apply this transformation to durations: it is natural to assume that longer past durations are associated with a lower current intensity. We then apply a set of transformations $x \mapsto g_l(x)$, $l = 1, 2, 3$, where $g_1(x) = x$, $g_2(x) = (x - 0.5)^2$, and $g_3(x) = x^3$. The functional forms can account for different types of impacts: $g_1(x)$ has constant marginal impact as opposed to $g_3(x)$ which is close to zero for most values of x not in the proximity of one. The latter transform seems more appropriate for volume imbalances and most covariates. The function $g_2(x)$ is just convex with a minimum at $x = 0.5$. Combining these functions, we can approximate different types of impacts and still be able to interpret the results. We now state the hypotheses for model restriction.

3.1.2. Model Hypothesis 1. The impact of all volume imbalances and trade imbalance is positive for buy trade intensity and negative for sell intensity. The impact of the spread and durations is always negative. Finally, we apply, separately, the following two transformations: g_1 and g_3 , i.e., linear and cubic. Hence, this doubles the final number of covariates to estimate.

3.1.3. Model Hypothesis 2. The impact of all covariates is as in model hypothesis 1. However, we finally apply, separately, the following two transformations: g_2 and g_3 , i.e., quadratic and cubic. Again, this doubles the final number of covariates to estimate.

3.1.4. Model Hypothesis 3. The impact of all covariates is as in model hypothesis 1, except for spread, for which we now suppose a positive impact. Finally, we apply, separately, the same transformations as in model hypothesis 2.

3.1.5. Remarks on the Model Hypotheses. We briefly comment on the hypotheses. First, the direction of the impact is a maintained hypothesis when estimating the model. If this is wrong, it is likely that the estimated coefficient will be zero. Second, regarding the transformations, we allow for the possibility of combining two functions (out of three) in order to derive a more flexible model. At the same time, we also rely on the estimation procedure to select the best submodel in an automated fashion. For example, assume that the direction of the impact of the covariates is as in model hypothesis 1 and that the impact is linear. Then, model hypothesis 1 is true, whereas model hypotheses 2 is false. We have included model hypothesis 3 to illustrate the point that using the wrong direction of the impact for the spread can lead to a zero coefficient.

For this problem, we have high confidence in the direction of the impact based on the existing literature (recall the discussion in Section 2.3).

After the application of EWMA's filters and the transformations, each hypothesis has 18 covariates for each of the four instruments, which means 72 parameters. Rather than making this a pure exercise in data mining, we prefer to keep the number of covariates relatively small to simplify our discussion and focus on the significance of the results.

3.2. Scope and Limitations

The goal is to present a model in the same vein as the ones used by high-frequency trading firms. The focus is on the order book as the main source of information. The EWMA's of trade durations attempt to capture the well-known autocorrelation of durations (Engle and Russell, 1998). The model is geared toward high-frequency predictions in live trading. It is then customary to restrict variables to finite ranges to avoid consuming corrupted/invalid messages from the exchange. This may sound problematic for some fat-tailed variables like durations. However, as durations increase, the intensity should shrink. Capping durations to a relatively high value is equivalent to saying that the decreasing marginal effect of durations beyond the cap value is zero. Given that for crude oil we can have hundreds of trade arrivals within a second, capping durations to 1 second seemed reasonable.

The model does not include an intraday seasonal component. However, part of this seasonality is implicitly captured by a slow moving EWMA of durations, so we did not further increase the number of variables as we want to showcase the importance of the order book.

A final comment pertains to the separate estimation of buy and sell intensities. Suppose the following feedback loop effect:

$$\begin{aligned} \lambda^{buy} &= X' b^{buy} + \rho^{buy} \lambda^{sell}, \\ \lambda^{sell} &= X' b^{sell} + \rho^{sell} \lambda^{buy}, \end{aligned}$$

where ρ^{buy} and ρ^{sell} are constants in $[0, 1)$. This system has the reduced form

$$\begin{aligned} \lambda^{buy} &= X' (b^{buy} + \rho^{buy} b^{sell}) (1 - \rho^{buy} \rho^{sell})^{-1}, \\ \lambda^{sell} &= X' (b^{sell} + \rho^{sell} b^{buy}) (1 - \rho^{buy} \rho^{sell})^{-1}. \end{aligned}$$

In consequence, separate estimation of the buy and sell intensities is equivalent to estimation of the above reduced form model.

3.3. Results

To carry out a test of performance measurement, we split the data in two halves. We use the data until May 18, 2013 for estimation and the subsequent sample for testing. We shall only report results for buy trades. The ones for sell trades are essentially identical.

TABLE 2. Estimation results for buy trades

Hypothesis 1		Hypothesis 2		Hypothesis 3	
Active cov.	\hat{b}	Active cov.	\hat{b}	Active cov.	\hat{b}
CL_VolImb1_3	2.50	CL_VolImb1_2	1.25	CL_VolImb1_2	1.28
CL_TrIdmb98_3	0.04	CL_VolImb1_3	2.04	CL_VolImb1_3	2.05
CL_Dur98_3	0.44	CL_VolImb2_2	4.61	CL_VolImb2_2	4.65
CL_Dur90_3	3.51	CL_VolImb3_2	1.48	CL_VolImb3_2	1.51
HO_Spread_3	2.48	CL_TrIdmb98_2	8.08	CL_TrIdmb98_2	8.16
HO_Dur98_1	0.16	CL_Dur98_3	0.29	CL_Dur98_3	0.29
HO_Dur98_3	0.52	CL_Dur90_3	3.12	CL_Dur90_3	3.12
HO_Dur90_3	1.87	HO_Spread_3	1.96	HO_Dur98_1	0.29
		HO_Dur98_1	0.28	HO_Dur98_3	0.35
		HO_Dur98_3	0.38	HO_Dur90_3	1.95
		HO_Dur90_3	1.91		
K	\hat{s}/K	K	\hat{s}/K	K	\hat{s}/K
73	0.11	73	0.15	73	0.14

Notes: For each of the three hypotheses, the active covariate and the estimated coefficient are reported. The suffix and prefix in each covariate name are separated by an underscore. The suffix represents the ticker and the prefix the index in the transform g_l $l = 1, 2, 3$, defined at the end of the paragraph “the covariates” in Section 3.1. The total number of estimated parameters is denoted by K and the estimated number of nonzero elements by \hat{s} .

The estimator for all models is sparse, as only 10%–15% of the coefficients are nonzero (Table 2). An interesting finding is that the estimation based on hypothesis 2 leads to an impact of order book volume imbalances that is not monotonic. Moreover, all hypotheses suggest that order book volume imbalances beyond the third level are not important.

We then test the performance of the models out of sample. As already mentioned, we split the sample into estimation and test samples of roughly the same size each. Let $\hat{\lambda}^{(k)}$ be the intensity estimator from (2) using the estimation sample for model hypothesis $k \in \{1, 2, 3\}$. The log-likelihood on the test sample is

$$L_T^{(k)} = \int_0^T \ln \hat{\lambda}^{(k)}(t) dN(t) - \int_0^T \hat{\lambda}^{(k)}(t) dt.$$

Its predictable part is

$$H_T^{(k)} := \int_0^T \ln \hat{\lambda}^{(k)}(t) \lambda^*(t) dt - \int_0^T \hat{\lambda}^{(k)}(t) dt.$$

Under the null hypothesis that model hypothesis 1 and hypothesis 2 perform the same, we have that $H_T^{(1)} - H_T^{(2)} = 0$, against the one-sided alternative that model 2 is worse, i.e., $H_T^{(1)} - H_T^{(2)} > 0$. The standardized likelihood ratio test statistic

TABLE 3. Test of model performance for buy trades

	H2–H1	H3–H1	H2–H3
<i>t</i> -stat	38.59	37.92	2.65
<i>p</i> -value	0	0	0

Notes: The standardized out of sample likelihood ratio statistic $LR_T^{(1,2)}$ between different model combinations is reported. The columns identify the null hypothesis. For example, H2–H1 is the null that the model implied by Hypothesis 1 performs as well as the one from Hypothesis 2. A large value rejects the null in favor of the alternative that the model from Hypothesis 2 performs better.

$LR_T^{(1,2)} := (L_T^{(1)} - L_T^{(2)}) / \sqrt{T\hat{\sigma}_T^2}$ is asymptotically standard normal, where

$$\hat{\sigma}_T^2 = \frac{1}{T} \sum_{j=1}^n [\ln(\hat{\lambda}^{(1)}(T_j) / \hat{\lambda}^{(2)}(T_j))]^2$$

(Sancetta, 2018, Prop. 1). When carrying out the test, we standardize each intensity by $\hat{\lambda}_0 := \int_0^T \hat{\lambda}^{(k)}(t) dt / N(T)$ to ensure that the model comparisons are not affected by the value of the expected intensity. Recall that $\hat{\lambda}^{(k)}$ is unbiased if $\hat{\lambda}_0 = 1$.

The test is performed for different pairs of model hypotheses. We find that hypothesis 2 is favored out of sample (Table 3). From the results in Tables 2 and 3, we infer that information from auxiliary instruments is relevant and instruments are interlinked. Moreover, we find that the impact of volume imbalances is not monotonic. It is also worth mentioning that the results for the sell intensity led to similar conclusions except for one case. We did not reject the hypothesis that model hypothesis 2 (H2) is not better than model hypothesis 3 (H3). Recall that the difference in these two model hypotheses is that the impact of the spread is in the wrong expected direction for H3. In this case, the spread is not selected when estimation is carried out under H3. The spread does not seem to be an important variable in this problem because all products are very liquid. Hence, their spread tends to be very tight most of the time. The only exception is heating oil (HO). In this case, inspection of the data did show some variation for the spread. However, overall, its importance is not as crucial as the other covariates. It is reassuring that under H3, the spread for heating oil, having the wrong sign is not selected.

4. CONCLUSION

This paper studies the estimation of a point process where the intensity is a function of high-dimensional variables. We rely on a nonnegativity constraint for the intensity to show that the estimation problem can be solved by standard quadratic programming with no need to include a penalty and tune additional parameters. The resulting estimator is consistent and the error only grows logarithmically in

the number of estimated parameters, as long as the true parameter is sparse. Our motivation is the estimation of a possibly nonlinear additive model for the intensity of buy and sell trades of crude oil futures prices. The covariates that affect the intensity are order book and trade variables on crude oil together with information from auxiliary instruments such as natural gas, heating oil, and the S&P500 futures. The results show that the impact of variables constructed from the order book and trades is nonlinear and that the instruments are interconnected. A test for model performance is used, and it shows that we can compare competing hypotheses for the direction of the impact of the covariates and the functional form of the intensity.

SUPPLEMENTARY MATERIAL

To view the supplementary material for this article, please visit: <https://doi.org/10.1017/S0266466622000238>.

REFERENCES

- Aalen, O. (1980) A model for nonparametric regression analysis of counting processes. In W. Klonecki, A. Kozek, & J. Rosiński (eds), *Mathematical Statistics and Probability Theory: Proceedings of the Sixth International Conference Wisla (Poland) 1978*, Lecture Notes in Statistics, pp. 1–25. Springer.
- Alaya, M.Z., S. Bussy, S. Gaïffas, & A. Guilloux (2019) Binaricity: A penalization for one-hot encoded features in linear supervised learning. *Journal of Machine Learning Research* 20, 1–34.
- Alaya, M.Z., S. Gaïffas, & A. Guilloux (2015) Learning the intensity of time events with change-points. *IEEE Transactions on Information Theory* 61, 5148–5171.
- Bauwens, L. & N. Hautsch (2009) Modelling financial high frequency data using point processes. In T.G. Andersen, R.A. Davis, J.-P. Kreiss, & T. Mikosch (eds), *Handbook of Financial Time Series*, pp. 953–982. Springer.
- Bosq, D. (2000) *Linear Processes in Function Spaces: Theory and Applications*, Lecture Notes in Statistics, vol. 149. Springer.
- Brémaud, P. & L. Massoulié (1996) Stability of nonlinear Hawkes processes. *Annals of Probability* 24, 1563–1588.
- Cont, R., A. Kukanov, & S. Stoikov (2014) The price impact of order book events. *Journal of Financial Econometrics* 12, 47–88.
- Engle, R.F. & J.R. Russell (1998) Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Gaïffas, S. & A. Guilloux (2012) High-dimensional additive hazards models and the Lasso. *Electronic Journal of Statistics* 6, 522–546.
- Glosten, L.P. & P. Milgrom (1985) Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71–100.
- Gray, R.M. (2009) *Probability, Random Processes, and Ergodic Properties*. Springer.
- Gray, R.M. & J.C. Kieffer (1980) Asymptotically mean stationary measures. *Annals of Probability* 8, 962–973.
- Hall, D. & N. Hautsch (2007) Modelling the buy and sell intensity in a limit order book market. *Journal of Financial Markets* 10, 249–286.
- Kercheval, A. & Y. Zhang (2015) Modelling high-frequency limit order book dynamics with support vector machines. *Quantitative Finance* 15, 1315–1329.
- Kyle, A.S. (1985) Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.
- Leeb, H. & B.M. Pötscher (2005) Model selection and inference: Facts and fiction. *Econometric Theory* 21, 21–59.

- Lorentz, G.G. (1986) *Bernstein Polynomials*. Chelsea.
- MacKenzie, D. (2017) A material political economy: Automated trading desk and price prediction in high-frequency trading. *Social Studies of Science* 47, 172–194.
- Meinshausen, N. (2013) Sign-constrained least squares estimation for high-dimensional regression. *Electronic Journal of Statistics* 7, 1607–1631.
- Morariu-Patrichi, M. & M.S. Pakkanen (2022) State-dependent Hawkes processes and their application to limit order book modelling. *Quantitative Finance* 22, 563–583.
- Sancetta, A. (2018) Estimation for the prediction of point processes with many covariates. *Econometric Theory* 34, 598–627.
- Slawsky, M. & M. Hein (2013) Nonnegative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization. *Electronic Journal of Statistics* 7, 3004–3056.
- van de Geer, S.A. & P. Bühlmann (2009) On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics* 3, 1360–1392.