Journal of Financial Econometrics, 2011, Vol. 9, No. 4, 685–716

Trading Dynamics in the Foreign Exchange Market: A Latent Factor Panel Intensity Approach

INGMAR NOLTE
University of Warwick

VALERI VOEV
Aarhus University

ABSTRACT

We develop a panel intensity framework for the analysis of complex trading activity datasets containing detailed information on individual trading actions in different securities for a set of investors. A feature of the model is the presence of a time-varying latent factor, which captures the influence of unobserved time effects and allows for correlation across individuals. We contribute to the literature on market microstructure and behavioral finance by providing new results on the disposition effect and on the manifestation of risk aversion on the

This work has been supported in part by the European Community's Human Potential Program under contract HPRN-CT-2002-00232, Microstructure of Financial Markets in Europe; the Fritz Thyssen Foundation through the project "Dealer-Behavior and Price-Dynamics on the Foreign Exchange Market"; and the German Research Foundation (DFG) through the research group "Price, Liquidity, and Credit Risks: Measurement and Allocation." The Center for Research in Econometric Analysis of Time Series, CREATES, is funded by The Danish National Research Foundation. The authors wish to thank René Garcia (editor), the associate editor, and two anonymous referees for very helpful comments. Moreover, they wish to express their sincere gratitude to Richard Olsen and Olsen Financial Technologies for providing us with the data. The paper benefited from several discussions with Günter Franke, Sandra Nolte, and Winfried Pohlmeier. Earlier versions of the paper circulated under the name "Panel Intensity Models with Latent Factors: An Application to the Trading Dynamics on the Foreign Exchange Market." The paper has been presented at faculty seminars at the "Helsinki School of Economics," "ECARES, Université Libre de Bruxelles," and "CREATES, Aarhus University," as well as on the "62nd Meeting of the Econometric Society (ESEM), Budapest 2007," the "Jahrestagung, Verein für Socialpolitik, Munich 2007," the "Complexity in Markets Conference, University of Warwick 2008," the "International Conference on Price, Liquidity, and Credit Risks, University of Konstanz 2008," and the "Humboldt-Copenhagen Conference, Humboldt Univesität zu Berlin 2009." Address correspondence to Ingmar Nolte, Warwick Business School, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, UK, or e-mail: ingmar.nolte@wbs.ac.uk.

doi: 10.1093/jjfinec/nbq033
Advance Access publication April 20, 2011

high-frequency trading level. These novel insights are made possible by the joint characterization of not only the decision to close (exit) a position, usually considered in isolation in the literature, but also the decision to open (enter) a position, which together describe the trading process in its entirety. While the disposition effect is defined with respect to the willingness to realize profits/losses with respect to the performance of the position under consideration, we find that the performance of the total portfolio of positions is an additional factor influencing trading decisions that can reinforce or dampen the standard disposition effect. Moreover, the proposed methodology allows the investigation of the strength of these effects for different groups of investors ranging from small retail investors to professional and institutional investors. (*JEL*: C33, C41, C50)

KEYWORDS: behavioral finance, efficient importance sampling, latent factors, market microstructure, panel intensity models, stochastic conditional intensity, trading activity datasets

1 INTRODUCTION

The complexity of financial market data containing micro-information on every individual trader's action presents new challenges to financial economists and econometricians. The immense breadth of these data opens new horizons for the analysis of market microstructure and behavior of economic agents, beyond the analyses possible with standard (trades and quotes) high-frequency data.

Trading activity datasets, now becoming increasingly available, can be considered as micro-panel datasets with four dimensions: an irregularly spaced timescale, types of trading actions, trading instruments, and investors. The marvelous amount of precise information contained in these datasets creates unique possibilities to analyze individual trading behavior since we can follow the investment activity of each individual over time. Time plays a central role in high-frequency finance, market microstructure, and the behavioral finance literature. Trading activity datasets can be thought of as field data with exact information on the timing of investment decisions hence enabling in-depth investigations of prominent behavioral finance phenomena such as the disposition effect as well as the motives driving investment decisions at specific points in time. Investors are typically heterogenous with respect to their trading and risk preferences, and trading activity datasets provide a sound data fundament for the detailed investigation of trading behavior for different groups of investors from small retail to professional and institutional investors.

In this paper, we develop an intensity-based modeling framework that is suited to characterize the data-generating process of complex dynamic systems such as trading activity datasets. The model exploits the panel structure of the trading activity dataset and characterizes a multivariate panel intensity process that is specified as a function of individual-specific effects, a set of time-varying covariates describing the investors' information set, a seasonality component, and a time-varying latent factor. The latent factor is motivated by the fact that not all individual-specific as well as public information is directly observable or measurable with available explanatory variables. These omitted unobservable factors induce dependencies across individuals that are captured by the common latent factor. The intensity-based specification is chosen since it allows us to account for the impact of time-varying covariates on the trading process. We use a simulated maximum likelihood (SML) technique to estimate the proposed model by augmenting the efficient importance sampling (EIS) method of Richard and Zhang (2007). Our approach is related to the stochastic conditional intensity (SCI) model proposed by Bauwens and Hautsch (2006) which they use to characterize a system of duration processes. Koopman, Lucas, and Monteiro (2008) develop a multistate extension of the SCI model and apply it to the analysis of credit rating transitions. The model we propose here differs from both models mentioned above in terms of its complexity and adapts the SCI model to a panel framework by adding two dimensions: individuals and trading instruments. Such degree of complexity is required in order to jointly model all aspects of a trading activity dataset, providing a framework allowing for individual- and transaction-type-specific effects.

By modeling the investors' decisions to enter and exit a position jointly and addressing directly the timing of their actions at the microlevel, we are able to offer a broad picture of investment behavior. Categorizing our investors into groups according to their total trading volume allows us to draw conclusions on the strength of various effects and the differences in the trading behavior across various types of investors.

The model is applied to the analysis of the trading behavior of investors in the foreign exchange market based on a trading activity dataset from OANDA FXTrade, containing information on every action of a large set of investors in up to thirty currency pairs over the period from October 1, 2003, to October 31, 2003. OANDA FXTrade is an electronic trading platform in the foreign exchange market in which heterogenous groups of small retail investors as well as big institutional and professional traders are active.

We investigate the disposition effect and loss aversion behavior. The disposition effect (Shefrin and Statman 1985) describes the tendency to hold positions with a paper loss longer than positions with the symmetric paper profit. A theoretical foundation for the disposition effect is based on the prospect theory of Kahneman and Tversky (1979) in which the investor evaluates the outcome of a trading strategy relative to a reference point and is risk averse if the strategy is profitable with respect to that reference point and risk seeking otherwise. Barberis and Xiong (2009) show within a theoretical framework that prospect theory preferences can cause a disposition effect. The early studies of Lease, Lewellen, and Schlarbaum (1974), Schlarbaum, Lewellen, and Lease (1978a,b), Shefrin and Statman (1985) as well as the more recent contributions of Badrinath and Lewellen (1991), Locke and

Mann (2005), and Shapira and Venezia (2001) analyze the disposition effect by comparing mean round-trip durations of profitable versus nonprofitable investments. Odean (1998a) considers the proportions of profits and losses realized over a certain time horizon, and Grinblatt and Keloharju (2001) apply ordered response models for the analysis of the disposition effect.

Most of these studies invoke an implicit narrow framing argument (Thaler 1985) and examine the disposition effect isolated for trading in a specific security in a static framework for the average investor. The studies by Shapira and Venezia (2001), Dhar and Zhu (2006), Goetzmann and Massa (2008), and Chen et al. (2007) focus on investor heterogeneity and show that professional and more sophisticated investors are less prone to the disposition effect and to behavioral biases in general.

In our analysis, we find support for the existence of the standard disposition effect with respect to a single asset position and we also find that this effect is decreasing with investor size, which is in line with the literature above. However, we also show that the profitability of the total security portfolio plays an important role for the decision to close a position, which provides counterevidence to the narrow framing argument often implicit in the literature. We find clear evidence for disposition effect behavior with respect to the total portfolio profit and loss. Interestingly, larger investors, who are more likely to trade with a portfolio strategy in mind, are prone to a disposition effect in the portfolio level but not with respect to the single-position profit/loss. Moreover, we find additional support for investment behavior consistent with prospect theory preferences by investigating the effect of profits and losses on the likelihood of opening a new position.

A related strand of the literature analyzes how investors' trading performance and activity are affected by learning and experience (cf. Nicolosi, Peng, and Zhu 2009; Seru, Shumway, and Stoffman 2009). A finding in these studies is that experience leads to more trading activity. However, there is almost no literature analyzing how risk-taking behavior affects trading strategies and activity on the high-frequency level. In many market microstructure models,[1] traders do not base their order submission strategies conditional on their current inventory or their current exposure to inventory risk. Our paper offers deeper insights into this issue that has important implications for these models. In particular, we find that investors, who temporarily face a higher than normal risk exposure, trade more cautiously in the sense that they are reluctant to further extend their position in the near future and that they tend to close out the position faster. We interpret this finding as a sign of risk-averse behavior at the high-frequency level that leads to more careful managing of active risk. Nolte and Nolte (2010) also provide support for this effect. Using the same dataset, but a different methodology, they find that traders put more effort in analyzing past order flow information whenever they decide to close a position in comparison to when they decide to open it.

[1] Cf. Chakravarty and Holden (1995), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2005), and Kaniel and Liu (2006).

The paper is structured as follows: in Section 2, we provide a theoretical description of the model. Section 3 contains the empirical analysis and a discussion of the results in light of recent behavioral finance and market microstructure theories, and Section 4 concludes. An exposition of the SML estimation procedure is presented in Appendices A and B. Additional estimation results are collected in a Web Appendix, which can be downloaded from the authors' Web pages.[2]

2 PANEL INTENSITY MODEL

2.1 Theory

Let $t \in [0, T]$ denote the physical calendar time, $n = 1, \ldots, N$ denote the nth investor, and $k = 1, \ldots, K$ denote the kth currency pair in which an investor can trade. The investor can take $s = 1, \ldots, S$ trading actions, for example, as in our application $s = 1$ increasing and $s = 2$ decreasing the position in a currency pair.[3] Note that in the foreign exchange market, both selling and buying can increase or decrease exposure to exchange rate risk. Thus, our categorization of actions is not a buy/sell categorization.

We associate the counting process $N^{skn}(t)$ with the events of type s in the kth currency pair for the nth investor defined on a joint probability space $\{\Omega, \mathfrak{F}, \mathfrak{F}_t, \mathcal{P}\}$, where the filtrations of the individual processes are denoted by $\mathfrak{F}_t^{skn} \subset \mathfrak{F}_t$. We assume that these processes are orderly.[4] The pooled process $N^{kn}(t) = \sum_{s=1}^{S} N^{skn}(t)$ that counts all the actions of a given investor in a currency pair is also assumed to be orderly, that is, the investor cannot take more than one action in the same asset instantaneously. To account for unobserved time-varying heterogeneity, we introduce a latent factor process that is defined on the overall pooled process $N(t) = \widetilde{\sum}_{n=1}^{N} \widetilde{\sum}_{k=1}^{K} N^{kn}(t)$, where the $\widetilde{\sum}$ operator counts multiple events at the same time only once since one or more investors can take actions in several currency pairs simultaneously. This convention does not exclude that the latent process is orderly as well, but it is unlikely given the large number of events. We use standard notation and denote the observation times on each of these point processes by $t_i^{(\text{superscripts})}$ with $i = 1, \ldots, N^{(\text{superscripts})}(T)$. In particular, the latent factor evolves over t_i for $i = 1, \ldots, N(T)$.

The building blocks of our model are the intensities associated with the processes $N^{skn}(t)$, which are defined as

$$\theta^{skn}(t|\mathfrak{F}_{t^-}) = \lim_{h \downarrow 0} \frac{P(N^{skn}((t+h)^-) - N^{skn}(t^-) > 0|\mathfrak{F}_{t^-})}{h}, \tag{1}$$

[2] http://www.warwick.ac.uk/staff/I.Nolte/publications/Nolte&Voev(2010)-JFEC-Web-Appendix.pdf.

[3] Generally, one can consider a broader set of actions, such as submission of special orders, order cancelation, or change.

[4] That is, $P(N^{skn}(t+\delta) - N^{skn}(t) > 1|\mathfrak{F}_t) = o(\delta)$, with $o(\cdot)$ the little Landau symbol.

where we note that the information set $\mathfrak{F}_{t^-}$ implicitly includes conditioning on the process having survived until t^-. The latent process is denoted by λ_i, and its conditional density is given by $\rho\left(\lambda_i|\mathfrak{F}_{t_i^-}\right)$. Having set up this notation, we can write the likelihood of the model as[5]

$$\mathcal{L}(W;\theta) = \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \prod_{s=1}^{S} \exp\left(d^s_{N^{kn}(t_i)} \ln \theta^{skn}\left(t^{kn}_{N^{kn}(t_i)} \middle| \mathfrak{F}_{t_i^-}, \lambda_i \right) \right.$$
$$\left. - \int_{t^{kn}_{N^{kn}(t_i)-1}}^{t^{kn}_{N^{kn}(t_i)}} \theta^{skn}(u|\mathfrak{F}_{u^-}, \lambda_{N(u^-)+1}) \mathrm{d}u \right) \rho(\lambda_i|\mathfrak{F}_{t_i^-}) \mathrm{d}\Lambda, \qquad (2)$$

where W is a generic symbol for data, θ is a parameter vector, and $d^s_{N^{kn}(t_i)}$ is a dummy that takes the value 1 whenever the event $N^{kn}(t_i)$ is of type s.[6] The set $\mathcal{C}_i = \left\{(k,n)|t_i = t^{k,n}_{N^{k,n}(t_i)}\right\}$ contains the pairs (k,n) associated with the arrival time t_i since at t_i there might be several investors trading at the same time and/or one investor trading in different currency pairs. It is natural to write the likelihood in terms of the durations $[t^{kn}_{i-1}, t^{kn}_i]$ that represent the behavior of a single investor trading in a given currency pair. Given our assumptions, this is the most aggregated orderly marked point process. These marked point processes are then aggregated over investors and currency pairs. We also find this expression instructive in terms of how one would actually implement the model. Alternatively, we can specify the model on the basis of $[t_{i-1}, t_i]$ as in Koopman, Lucas, and Monteiro (2008). In this case, however, we have to keep track of all potential investors and currency pairs that are "at risk" for an event of type s at time t_i.

An attractive feature of intensity-based modeling is that it accounts for changes in the values of time-varying covariates during a duration in a very intuitive way since it is set up in continuous time. In a discrete-time duration-based approach (e.g., the stochastic conditional duration model of Bauwens and Veredas 2004), one can also account for time-varying covariates (see Lunde and Timmermann 2005), but then the likelihood function has to be additionally adjusted (effectively this again amounts to adjusting the intensity to reflect the changes in the values of the covariates). Furthermore, the intensity-based approach allows for the characterization of the dynamic behavior of each of the s subprocesses, whereas the duration approach considers the pooled process only. The intensity can be modeled in the spirit of the SCI model of Bauwens and Hautsch (2006). We parameterize $\theta^{skn}(t|\mathfrak{F}_t^-, \lambda_{N(t^-)+1})$ generally in the following way:

$$\theta^{skn}(t|\mathfrak{F}_t^-, \lambda_{N(t)}) = (b^{skn}(t) S^{skn}(t) \Psi^{skn}(t) \lambda^{\delta^{skn}}_{N(t^-)+1}) D^{skn}(t). \qquad (3)$$

[5]See, for example, Lancaster (1997). Note that for an arbitrary time u, we need to ensure that the latent factor is indexed by the time at the end of the spell, that is, $N(u^-)+1$. At event times t_i, $N(t_i^-)+1 = i$.
[6]This describes the type of action of the nth investor in the kth currency pair at the arrival time of the latent factor time t_i.

Thereby $b^{skn}(t)$ denotes a (possibly investor, currency pair, or state-dependent) baseline intensity, $S^{skn}(t)$—a deterministic seasonality function, $\Psi^{skn}(t)$—the component capturing the effect of (time-varying) covariates, and δ^{skn} is a parameter that controls for the impact of the latent component on the s-type intensity. In our application, we need to take into account that after an action that sets the exposure in a given currency pair to zero, that is, the position is closed completely, there is no possibility for a subsequent close. Hence, the intensity $\theta^{2kn}(t|\mathfrak{F}_{t^-}, \lambda_{N(t^-)+1})$ is zero in this case. We model this through the variable

$$D^{skn}(t) = \begin{cases} 1, & \text{if } s = 1, \\ 1 - d_{cc}^{kn}(t), & \text{if } s = 2, \end{cases} \tag{4}$$

where $d_{cc}^{kn}(t)$ denotes the dummy variable that takes on the value 1 if the previous arrival time is associated with a complete close of the position in the currency pair k for investor n and 0 otherwise.

2.2 Model Parameterization

In our application, we parameterize the separate intensity components parsimoniously as outlined below.

2.2.1 Baseline intensity and individual investor-specific effects.

We assume that there are different baseline intensities for the different states and the individual investors but that they are identical across currency pairs. That is, we assume that

$$b^{skn}(t) = b^{sn}(t) \quad \text{for } k = 1, \ldots, K,\ s = 1, \ldots, S, \text{ and } n = 1, \ldots, N.$$

In the application, we use a multivariate Fourier flexible form (FFF) Weibull specification in the backward recurrence times of the following type:

$$b^{sn}(t) = \exp(\omega^{sn})\tilde{S}(\tilde{v}^s, \tilde{u}^{skn}(t), \tilde{K}^s) \prod_{r=1}^{S} u^{rkn}(t)^{\alpha_r^s - 1} \quad \text{for } s = 1, \ldots, S,$$

where $\tilde{S}(\cdot)$ with parameter vector $\tilde{v}^s$ of dimension $2\tilde{K}^s$ is defined as in Equation (5) below, $u^{skn}(t) = t - t^{skn}_{N^{skn}(t^-)}$ is the backward recurrence time of process s, and $\tilde{u}^{skn}(t) = u^{skn}(t)/T$ is the backward recurrence time standardized to the interval $[0, 1]$. The location parameters ω^{sn} are investor and state specific, and the shape parameters α_r^s are assumed to be identical for all investors but different across states. Thus, the individual effect allows for investor-specific shifts in the intensity. The FFF component is chosen to allow for humps and troughs in the baseline hazard, which cannot be modeled through the Weibull specification alone.

2.2.2 Diurnal seasonality and weekend effects. The seasonality function $S^{skn}(t)$ incorporates a diurnal seasonality component $\tilde{S}^{skn}(t)$ and a weekend component $\tilde{W}^{skn}(t)$ multiplicatively as

$$S^{skn}(t) = \tilde{S}^{skn}(t)\tilde{W}^{skn}(t).$$

In order to capture the deterministic intraday seasonality pattern of the intensity processes, we assume that

$$\tilde{S}^{skn}(t) = \tilde{S}(t) \quad \text{for } k = 1, \dots, K,\ s = 1, \dots, S, \text{ and } n = 1, \dots, N,$$

where

$$\tilde{S}(t) \equiv \tilde{S}(\nu, \tau(t), K) \equiv \exp\left(\sum_{k=1}^{K} \nu_{2k-1} \sin(2\pi(2k-1)\tau(t)) + \nu_{2k} \cos(2\pi(2k)\tau(t))\right), \tag{5}$$

which is an exponentially transformed FFF, where $\tau(t)$ denotes the intraday trading time standardized to $[0,1]$ and ν is a $2K$-dimensional parameter vector. To model the lower degree of trading activity on weekends, we specify $\tilde{W}(t)$ as

$$\tilde{W}(t) = \exp(\varpi D_W(t)),$$

where ϖ denotes a scalar and $D_W(t)$ a weekend dummy, which is 1 during weekends and 0 otherwise. According to this specification, the intensity process is dampened for $\varpi < 0$, which is the effect that we expect, and amplified for $\varpi > 0$.

2.2.3 Explanatory variables. Let z_j^{skn} denote the vector of all (time-varying) possibly investor, currency pair, and state-dependent covariates, where at least one covariate is updated at time $\tilde{t}_j^{skn}$ with $j = 1, \dots, M^{skn}(T)$, where $M^{skn}(t)$ is the corresponding counting process. We model the impact of the covariates on the s-type intensity by

$$\Psi^{skn}(t) = \exp\left(\gamma^{s\prime} z^{skn}_{M^{skn}(t^-)}\right),$$

where γ^s is the coefficient vector.

2.2.4 Latent factor. We assume that the dynamics of the latent factor are defined on the timescale t_i. This means that the latent factor changes whenever there is an action of an investor in a currency pair. Note that this does not imply that the latent factor changes if there is an update of a time-varying covariate. Since each intensity θ^{skn} depends on the current value of the latent factor, we induce at every time t a contemporaneous dependence across all intensities θ^{skn}. The magnitude of this possibly investor, currency pair, or state-specific dependence is determined by the parameters δ^{skn}. The latent factor can therefore be interpreted as an unobservable time effect that affects the decisions (open, close) of all investors at every time t by influencing the intensities of the corresponding processes. We can justify the existence of such an unobservable time effect in our model in several ways: (i) (News) effects of news announcements, not modeled due to data limitations, (ii) (Order Flow) buy or sell pressure from the interbank market, which we do not observe directly since we consider an internet trading platform, (iii) (Herding) similar behavior of traders, due to similar interpretations of any kind of technical chart patterns or further information, or (iv) (Incomplete Information) lack of information on the overall asset holdings of the investors, complexity of trading strategy, financial skills and knowledge, and access to background information (news sources such as Reuters, Bloomberg, etc.) that influence their trading behavior in the particular market we analyze.

In our model, we assume that the latent factor follows, conditional on $\mathfrak{F}_{t_i^-}$, a lognormal distribution

$$\ln \lambda_i | \mathfrak{F}_{t_i}^- \overset{\text{i.i.d.}}{\sim} N(\mu_i, 1),$$

where the dynamics is modeled through an AR(1) process

$$\ln \lambda_i = a \ln \lambda_{i-1} + \epsilon_i \quad \text{for } i = 1, \ldots, I,$$

with $\epsilon_i \overset{\text{i.i.d.}}{\sim} N(0,1)$. Let l_i denote the log of the latent factor at t_i

$$l_i \equiv \ln \lambda_i,$$

and let L_i denote the history of the log latent factor up to and including t_i

$$L_i = \{l_j\}_{j=1}^{i}.$$

With this specification, the (log) latent factor depends only on its own past, so we denote its conditional distribution by $p(l_i|L_{i-1})$. From Equation (3), it follows that the influence of the log latent factor on the s-type intensity is given by $\delta^{skn} \ln \lambda_i$. As the latent factor evolves on the pooled action process of all investors in all currencies, we model the impact of the latent factor as common for all n and k, that is, $\delta^{skn} \equiv \delta^s$. Denoting $\lambda_i^s \equiv \delta^s \ln \lambda_i$, we have that

$$\lambda_i^s = a\lambda_{i-1}^s + \delta^s \epsilon_i \quad \text{for } i = 1, \ldots, N(T).$$

Therefore, the variance of ϵ_i is set to unity, so that the conditional variance of λ_i^s is equal to $(\delta^s)^2$, which eases the interpretation of the parameter.[7]

[7]Note that this does not exclude the possibility that δ^s could be negative.

2.2.5 Integrated hazard. Implementing the likelihood function (Equation (2)) requires computing an integral of the intensities θ^{skn}. In practice, this integration is carried out in a piecewise manner. A technical point that needs consideration is that the processes θ^{skn} change not only at the times at which the latent factor is updated but also whenever a time-varying covariate is updated. For this reason, we introduce yet another process $\{\check{t}_h^{skn}\}$, $h = 1, \ldots, H^{skn}(T)$, resulting from the pooling of the latent factor process $\{t_i\}$ and the covariate process $\{\tilde{t}_j^{skn}\}$, with $H^{skn}(t)$ denoting the corresponding counting process.

The integrated intensity over an arbitrary interval $[\underline{t}, \overline{t}]$ is then computed as

$$\int_{\underline{t}}^{\overline{t}} \theta^{skn}(u|\mathfrak{F}_u^-, \lambda_{N(u^-)+1})\mathrm{d}u = \sum_{h=H^{skn}(\underline{t})}^{H^{skn}(\overline{t})} \int_{\check{t}_h^{skn}}^{\check{t}_{h+1}^{skn}} \theta^{skn}(u|\mathfrak{F}_{u^-}, \lambda_{N(u^-)+1})\mathrm{d}u. \quad (6)$$

For evaluation of the likelihood, we need the integral in Equation (6) for $[\underline{t}, \overline{t}] = [t_{i-1}^{kn}, t_i^{kn}]$. For the purposes of model evaluation, the integrated intensities with $[\underline{t}, \overline{t}] = [t_{i-1}^{skn}, t_i^{skn}]$ can be considered as generalized residuals which under the correct model specification should be i.i.d unit exponentially distributed.[8] For further details on the estimation procedure, we refer to Appendix A.

3 EMPIRICAL ANALYSIS

3.1 Data Description

We analyze a trading activity dataset obtained from OANDA FXTrade. OANDA FXTrade is a foreign exchange electronic trading platform operating 24 hours, 7 days per week. It functions as a market making system that executes orders using the exchange rate prevalent in the market determined either by their own inventory book and/or by predicted prices relying on a proprietary forecasting algorithm based on an external data-feed. The legal counterparty of every transaction is always OANDA FXTrade. OANDA FXTrade offers immediate settlement of trades and tight spreads as low as two to three pips for all transaction sizes. Given various boundary conditions, such as sufficient margin requirements are satisfied, orders are always executed.

The investors can trade in up to thirty currency pairs, including the most active ones such as EUR/USD, GBP/USD, USD/CHF, EUR/JPY, USD/JPY, etc. They can submit market orders, limit orders, take-profit orders, and stop-loss orders to the system. They can cancel or change (limits on) existing orders without incurring any extra fees. Market orders (buy or sell) are executed immediately and affect existing open positions. Limit orders are maintained in the system for up to one month. A server manages the inventory book, the current exchange rates, and

[8] See Bauwens and Hautsch (2006) and Bowsher (2007) for an in-depth discussion and proof.

the current market orders to match existing limit orders. A limit order can therefore be matched either against a market order or against a bid or an ask price obtained from the external data-feed. Stop-loss orders and take-profit orders are special limit orders in the sense that they can be set for existing open positions. They can be specified directly while entering a market or a limit order, but they can also be specified later for existing open positions. Stop-loss and take-profit orders are automatically erased from the system whenever a position is closed as a result of further trading activity. In our analysis, we only consider those actions that either lead to opening a new position, changing an existing position, or closing a position. These are market orders, executed limit orders, or executed stop-loss and take-profit orders.

Our analysis focuses on one month of trading and stretches from 00:00:00 on October 1, 2003, until 23:59:59 on October 31, 2003. The raw datasets contain 2120 different investors. Many of these investors are recorded since, for instance, they submit limit orders, which are never executed, change limits on existing limit orders, or simply undertake only a very few transactions. We try to filter out these "tiny/noise" traders and restrict our attention to traders who have at least thirty transactions and have been active in at least three currency pairs during the month. Even those traders are still quite heterogeneous with respect to their trading activity and volume, and we classify them into twenty groups, each corresponding to 5% bins (vingintiles) of the cumulative distribution function of total trading volume (in USD) in the month. Thus, the first group (the 0%–5% bin) contains the traders with the smallest total trading volume, and the last group (the 95%–100% bin) contains the traders with the largest total trading volume. Altogether we end up with forty-six investors for each group. Since the estimation of our model is very computationally intensive, we choose ten investors randomly from each group for which we estimate the model.[9] Thus, we concentrate on 200 randomly chosen investors, which amounts to more than 20% of the cleaned investor trading activity dataset.

Table 1 presents the descriptive statistics for the twenty investor groups. All the figures presented are averages over the ten investors within each group. The heterogeneity of the investors on OANDA FXTrade is clearly visible. The investors in the smallest group have an average total transaction volume of around 2700 USD per month, while the average total transaction volume is with 29.69 million USD in the largest group more than 10,000 times higher. The average median (maximum) transaction volume per trade is rising accordingly from about 50 (150) USD to 125.000 (500.000) USD from the smallest to the largest group. These figures suggest that traders in the smaller groups consist of small private retail investors, while those in the larger ones consist of smaller professional and institutional investors.[10] Although we observe that the magnitude of the realized profit or loss

[9]The estimation of the model for one group takes about two to three days on a standard office computer in GAUSS. The estimation procedure does not allow for a large degree of parallelization and simply the computation of the Hessian can take up to 12 hours.

Table 1 Descriptive statistics for the twenty investor groups

Bin	1 0%–5%	2 –10%	3 –15%	4 –20%	5 –25%	6 –30%	7 –35%	8 –40%	9 –45%	10 –50%
Tot. Vol.	2707.44	9459.47	23.31k	37.95k	63.21k	93.99k	138.75k	208.07k	278.0k	376.70k
Max. Vol.	156.31	422.81	761.95	1805.63	2808.34	3029.01	6796.08	13.24k	7043.91	17.55k
Med. Vol.	49.52	114.31	292.62	400.48	814.38	1225.45	1225.48	2538.38	2546.69	4541.82
Real. P/L	1.77	1.43	0.52	−3.07	11.46	−9.71	5.45	11.89	34.96	17.33
No Tr.	130.40	83.80	197.60	90.00	147.40	120.10	178.40	253.50	150.40	107.50
No op.	90.00	41.80	100.20	44.30	90.60	61.90	105.20	135.40	76.80	53.00
No cl.	40.40	42.00	97.40	45.70	56.80	58.20	73.20	118.10	73.60	54.50
No full cl.	33.60	28.90	42.50	38.10	35.00	40.10	47.90	60.20	63.20	42.50
No ccy pairs	7.50	6.40	8.60	8.90	5.80	4.60	6.60	5.70	5.90	7.40
Bin	11 –55%	12 –60%	13 –65%	14 –70%	15 –75%	16 –80%	17 –85%	18 –90%	19 –95%	20 –100%
Tot. Vol.	482.44k	596.09k	737.34k	1.04m	1.32m	2.01m	3.10m	4.45m	8.54m	29.68m
Med. Vol.	5643.61	4480.54	5636.22	7150.68	9648.30	9533.90	36.83k	43.60k	58.17k	124.58k
Max. Vol.	14.73k	16.86k	40.73k	33.39k	27.47k	56.12k	124.76k	156.47k	262.09k	469.39k
Real. P/L	−76.31	−47.29	−311.23	−2.14	−47.22	−296.06	−11.93	−1767.18	−71.95	−2427.41
No Tr.	114.50	196.10	200.30	166.50	277.30	363.40	393.00	134.90	240.40	370.50
No op.	58.20	91.50	104.00	91.60	140.70	214.40	196.80	68.00	145.90	196.20
No cl.	56.30	104.60	96.30	74.90	136.60	149.00	196.20	66.90	94.50	174.30
No full cl.	42.30	55.40	45.30	49.50	78.90	54.10	67.00	54.20	54.80	94.70
No ccy pairs	5.60	6.40	7.50	4.40	6.00	6.10	5.90	5.70	6.60	4.20

All figures are "averages" over the ten investors within each group. All currency values have been converted to USD. Real. P/L stands for realized profit/loss, Tr.—for transaction, Vol.—for volume, op.—for open, cl.—for close, Tot.—for total, Max.—for maximum, Med.—for median, ccy—for currency, k $\triangleq$ thousand, and m $\triangleq$ million.

slightly increases over the groups, there is no evidence that bigger investors are more profitable than smaller ones. There is, however, a slightly increasing pattern in the frequency of trading within the investor groups, with a minimum of 83 trades per investor per month for Group 2 and a maximum of 393 traders for Group 17. Moreover, we observe that the number of currency pairs in which the investors trade lies on average between 4 and 9 per investor, without any particular pattern regarding group ordering.

3.2 Estimation Results

In this section, we report the estimation results of the panel intensity model, discuss their relevance in the context of market microstructure and behavioral finance theory, and evaluate the model fit. In the main body of text, we report the estimation results and figures for Groups 1, 7, 14, and 20. Our interpretation of the results, however, is related to all groups throughout. The complete estimation results and figures for all twenty investor groups are collected in Web Appendix B.[11] The estimation results for our subset of groups are presented in Table 2. We have grouped the estimates into several categories: baseline intensity, latent factor, seasonality, and covariates.

The coefficients for the baseline intensity for all groups and all investors result in a decreasing intensity with very slight humps induced by the FFF in the backward recurrence times. The baseline intensities for a period of seven days are depicted for all twenty groups in Figure 1. The solid (dashed) line represents the baseline intensity of opening (closing) trades. We have chosen the mean over the individual $\omega^{\cdot i}$ effects as the intercepts in these graphs. We observe that, *ceteris paribus*, the longer the periods of no activity, the lower the instantaneous probability for an open or close trade. In the first minutes, there is no systematic pattern on whether one intensity dominates the other. Afterward the opening baseline intensity is always higher than the closing one. This reflects to a certain extent that in our sample period there are more opening transactions than closing ones, which is an observation that is already visible in the descriptive statistics in Table 1. Moreover, it also indicates that the durations between opening events are on average also shorter than durations between closing events, a pattern that is also visible in the mean duration statistics for the "raw" series in Tables 9 and 10 in Web Appendix C.

Figure 2 depicts the diurnal seasonality patterns in the opening and closing trading activity in Eastern Standard Time for all twenty investor groups. The seasonality patterns are somewhat different across investor groups, but we can identify as common similarities three peaks, which are, depending on the specific

[10] A description of who uses OANDA FXTrade can be found on their Web page http://fxtrade.oanda.com/. We do not have further background information on the individual investors, and we resort to the total trading volume sorting to characterize the traders.

[11] http://www.warwick.ac.uk/staff/I.Nolte/publications/Nolte&Voev(2010)-JFEC-Web-Appendix.pdf.

investor group, more or less pronounced. The first peak is usually visible at about 2–3 o'clock, the second one that is also more pronounced at about 11–12 o'clock, and the third one around 20–21 o'clock. These peaks correspond to standard trading activity patterns of Europe-, America-, and Asia-based traders. In the case of the smaller investor groups, however, it is more likely that they reflect before

Table 2 Estimation results for selected investor groups 1, 7, 14, and 20

	1		7		14		20	
Par.	Est.	Std.	Est.	Std.	Est.	Std.	Est.	Std.
	Baseline intensity							
ω^{o1}	−3.6284	0.3775	−1.6789	0.2801	−2.6730	0.2820	−2.5318	0.5974
ω^{o2}	−3.8955	0.4145	−2.1429	0.3717	−1.7475	0.2511	−2.5767	0.5874
ω^{o3}	−3.6775	0.3526	−2.0492	0.2885	−2.6790	0.2865	−2.1823	0.5954
ω^{o4}	−3.8527	0.4767	−1.0109	0.1987	−2.4264	0.2853	−2.8665	0.5962
ω^{o5}	−3.1539	0.3656	−1.0690	0.2317	−2.1260	0.2584	−1.9449	0.5951
ω^{o6}	−2.1340	0.3154	−2.4441	0.2754	−2.1029	0.2617	−2.9082	0.5988
ω^{o7}	−3.6935	0.3471	−2.5269	0.3254	−2.7882	0.3174	−3.3194	0.6153
ω^{o8}	−3.6341	0.4519	−2.3018	0.3066	−2.6071	0.2838	−2.6881	0.5995
ω^{o9}	−2.8393	0.4622	−1.9981	0.3284	−2.2985	0.3024	−2.5027	0.5764
ω^{o10}	−2.9351	0.3701	−2.3509	0.2917	−0.7721	0.2620	−2.2244	0.5923
α^o_o	0.4902	0.0141	0.5511	0.0147	0.7297	0.0203	0.6075	0.0129
α^o_c	0.8225	0.0353	0.6461	0.0216	0.5381	0.0256	0.5999	0.0184
$\tilde{\nu}^o_1$	−0.4038	0.1834	−0.5245	0.1855	−0.6802	0.3122	−1.0506	0.2162
$\tilde{\nu}^o_2$	0.2328	0.1519	−0.0016	0.0217	0.0679	0.2623	1.2589	0.7575
$\tilde{\nu}^o_3$	0.4909	0.3111	0.2807	0.2133	0.1065	0.2613	1.9579	1.5345
$\tilde{\nu}^o_4$	0.2546	0.1596	0.3477	0.2104	0.7399	0.2490	1.0430	0.4442
ω^{c1}	−3.4514	0.5650	−2.3053	0.3169	−2.8735	0.3525	−1.0395	0.6898
ω^{c2}	−3.2566	0.5455	−3.1209	0.5710	−1.8362	0.2915	−0.1375	0.6761
ω^{c3}	−2.3837	0.5278	−2.6875	0.2956	−2.3686	0.3250	−1.3857	0.6704
ω^{c4}	−3.4616	0.5303	−1.6356	0.2020	−2.4985	0.3722	−2.3678	0.7254
ω^{c5}	−2.7495	0.5464	−1.7162	0.2618	−1.9347	0.3165	−1.3663	0.6672
ω^{c6}	−2.6478	0.4977	−2.8303	0.2857	−1.6775	0.3179	−0.4958	0.5575
ω^{c7}	−2.5675	0.5404	−3.0252	0.3474	−2.3422	0.3693	−1.4901	0.6939
ω^{c8}	−2.6277	0.6195	−2.9338	0.3130	−0.4709	0.3633	−0.2550	0.6364
ω^{c9}	−2.2699	0.5076	−3.0348	0.3347	−1.8842	0.3330	−0.8946	0.6575
ω^{c10}	−2.6985	0.4827	−1.9775	0.4334	−2.2842	0.3317	−1.3123	0.6825
α^c_o	0.7725	0.0347	0.6951	0.0219	0.6466	0.0237	0.9219	0.0157
α^c_c	0.6453	0.0586	0.7809	0.0293	0.7682	0.0324	0.5182	0.0161
$\tilde{\nu}^c_1$	−0.5147	0.2216	−0.2125	0.2769	−0.7168	0.1974	−1.3323	0.4488
$\tilde{\nu}^c_2$	−0.2453	0.1209	−0.4054	0.2203	−0.0260	0.1969	0.2332	0.7820
$\tilde{\nu}^c_3$	−0.3986	0.1892	−0.3038	0.1998	−0.0516	0.3613	0.6654	1.1509
$\tilde{\nu}^c_4$	0.1485	0.2104	0.3906	0.1593	0.4321	0.2229	0.4522	0.2922

(*continued*)

Table 2 (*continued*)

	1		7		14		20	
Par.	Est.	Std.	Est.	Std.	Est.	Std.	Est.	Std.
				Latent Factor				
a	0.5914	0.3138	0.3747	0.1060	0.9070	0.1611	0.9766	0.0066
δ^o	0.2319	0.0977	0.5078	0.0343	0.1201	0.1128	0.0650	0.0166
δ^c	−0.3699	0.1617	−0.7087	0.0461	−0.1415	0.1201	−0.1220	0.0229
				Seasonality				
ν_1	0.0851	0.0574	−0.2720	0.0571	−0.1680	0.0465	−0.0650	0.0366
ν_2	−0.3241	0.0556	−0.0687	0.0474	0.2926	0.0480	0.0131	0.0848
ν_3	−0.0243	0.0526	−0.0814	0.0437	−0.1541	0.0429	−0.1833	0.0515
ν_4	−0.2254	0.0520	−0.1236	0.0409	−0.1216	0.0422	−0.1445	0.0320
ν_5	0.2921	0.0453	0.4391	0.0521	0.2592	0.0487	0.3081	0.0377
ν_6	−0.1784	0.0457	0.1762	0.0496	−0.1700	0.0429	−0.1152	0.0382
ν_7	0.3130	0.0466	0.1233	0.0462	0.0959	0.0433	0.0707	0.0465
ν_8	0.1464	0.0494	−0.0443	0.0433	0.0422	0.0400	−0.0015	0.0733
ϖ	−2.4563	0.3242	−1.8763	0.2801	−2.8239	0.3130	−2.8670	0.2352
				Covariates				
$\gamma^o_{\mathrm{P/L\,1}}$	−0.0933	0.0136	−0.0216	0.0087	−0.0335	0.0091	−0.0402	0.0126
$\gamma^o_{\mathrm{P/L\,pf}}$	−0.0631	0.0325	0.0274	0.0195	−0.0699	0.0367	−0.0151	0.0326
γ^o_{vol}	−0.2865	0.0702	−0.1113	0.0470	−0.0542	0.0614	−0.1162	0.0523
$\gamma^c_{\mathrm{P/L\,1}}$	0.1241	0.0289	0.0564	0.0083	−0.0044	0.0320	0.0006	0.0450
$\gamma^c_{\mathrm{P/L\,pf}}$	−0.0218	0.0387	−0.0215	0.0260	0.2502	0.0605	0.1327	0.0476
γ^c_{vol}	0.6250	0.1866	0.1182	0.0782	0.1383	0.0838	0.1646	0.0711
No. obs.	1164		1665		1558		3029	
Mean ll	−7.63807		−5.34914		−6.52545		−6.72727	

. The $\gamma_\cdot^\cdot$ coefficients on the covariates should be interpreted as follows: superscript "*o*" for opening intensity and superscript "*c*" for closing intensity. The subscripts stand for the corresponding variable, where "P/L 1" is the paper profit/loss in the corresponding currency pair, "P/L pf" is the paper profit/loss in the total portfolio, and "vol" is the standardized excess trading volume. Par., parameter; Est., estimate; Std., standard deviation; Obs, observations. All other coefficients are detailed in the main text. "mean ll" stands for the mean log-likelihood. Quasi maximum likelihood standard errors are reported.

and after "normal" work trading preferences and opportunities. The weekend dummy is significantly negative for all groups which is in line with the lower trading activity during weekends. We do not observe any specific pattern in the size of the weekend dummy with respect to the ordering of our investors.

Table 3 presents the likelihood ratio (LR) tests of the model containing the latent factor dynamics against the restricted model without it. The detailed

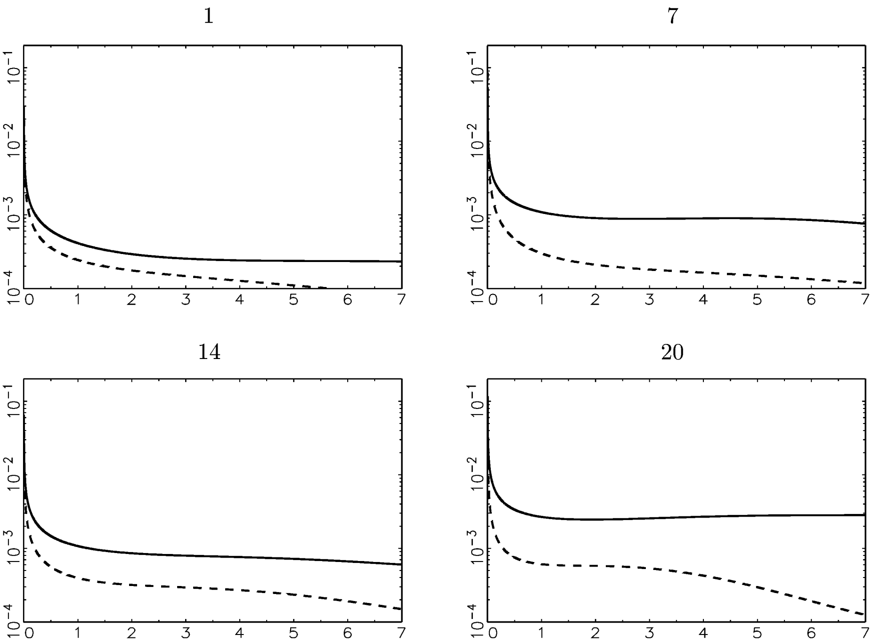

Figure 1 Opening (solid) and closing (dashed) baseline intensities for selected investor groups 1, 7, 14, and 20 over seven days (x-axis).

estimation results for the restricted model without the latent factor can be found in Web Appendix A. Generally, the sign and magnitude of the remaining parameters are relatively unaltered, which provides a robustness check regarding the stability of the market microstructure results outlined below.

The LR tests show that the parameters regarding the latent factor a, δ^o, and δ^c are jointly always significant for all twenty investor groups, which underpins the importance of the latent factor in capturing the dynamics in our model. Here, we do not observe a specific pattern over investor groups in the magnitudes of the LR test statistics which could be used as an indication that the latent factor would be more important for specific groups of investors. And, it seems that the latent factor is indeed needed for both smaller as well as bigger investors. Figure 3 depicts the parameters for every group. Whereas the impact parameters δ^o are always greater and δ^c are always smaller than 0, the autoregressive parameter a is always positive and smaller than 1, ensuring a stationary specification for the latent factor AR(1) process. It is evident that the AR coefficient is increasing over the investor groups, reflecting that the latent factor process becomes more and more persistent and thus puts increasing weights on previous observations. On the contrary, the loading coefficients become smaller (in absolute terms) the larger the investors.

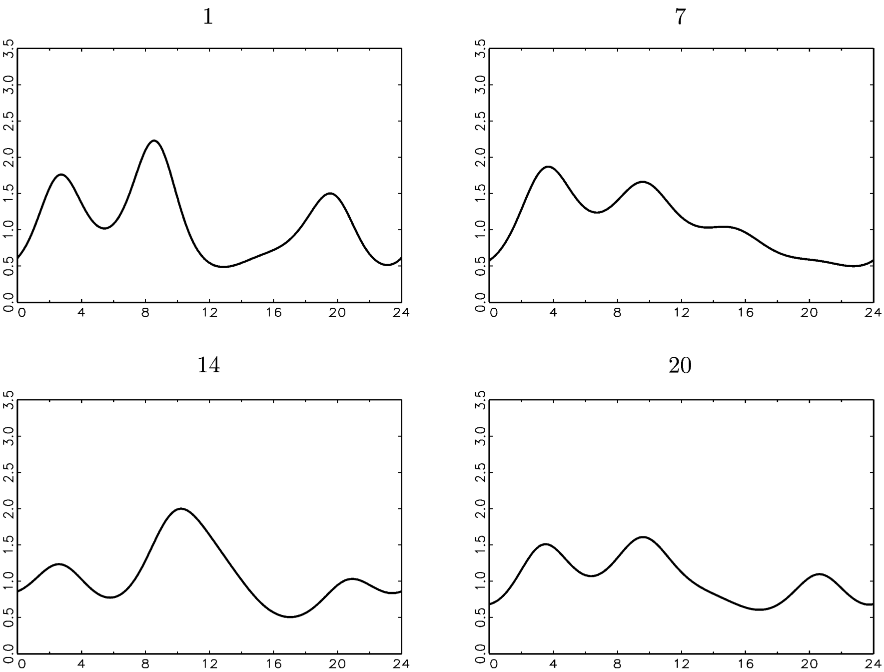

Figure 2 Diurnal seasonality patterns of trading activity for selected investor groups 1, 7, 14, and 20. The x-axes always denote time of day in Eastern Standard Time.

To shed more light on the importance of the latent factor and its behavior over time and investor groups, we compute smoothed estimates of the log latent factor following the procedure of Koopman, Lucas, and Monteiro (2008), which is detailed in Appendix B. The resulting graphs are depicted in Figure 4, in which we plot the smoothed estimates over our observation period of 31 days as a function of calendar time for selected investor groups 1, 7, 14, and 20. The figures for all investor groups are presented in Figure 3 in Web Appendix B. For all groups, the smoothed estimates show a clear time of day pattern indicating that the latent factor controls differently for unobserved heterogeneity during a trading day, with greater impact during standard trading hours. This pattern is generally consistent across all investor groups but much more pronounced and significant for the larger investors. For them it turns out that the latent factor in general seems to be much more important than for the smaller investors. One possible interpretation of this finding may be that the investment decision process of the larger investors is more complex and relies on more (unobserved) decision factors than the investment decision process of the smaller investors. Hence and despite results of the LR test, the simple model without the latent factor may already be sufficient to explain the trading dynamics of the smaller investors quite well, whereas for the larger

Table 3 Test statistics and p-values for the LR test of the full model specification against the restricted model specification without latent factor dynamics

	1	2	3	4	5	6	7	8	9	10
Observations	1164	797	1694	821	1335	1130	1665	2267	1390	1008
LR-statistics	30.8220	8.6461	45.1525	25.5950	72.0865	37.6266	145.7004	10.3997	15.5934	10.5746
p-Value	0.0000	0.0344	0.0000	0.0000	0.0000	0.0000	0.0000	0.0155	0.0014	0.0143
	11	12	13	14	15	16	17	18	19	20
Observations	1072	1761	1772	1558	2519	3179	3116	1235	2110	3029
LR-statistics	8.9694	82.4296	24.2949	31.2943	35.9178	95.3486	85.0971	24.2326	47.9862	122.5967
p-Value	0.0297	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

The test statistic is asymptotically χ^2_3 distributed.

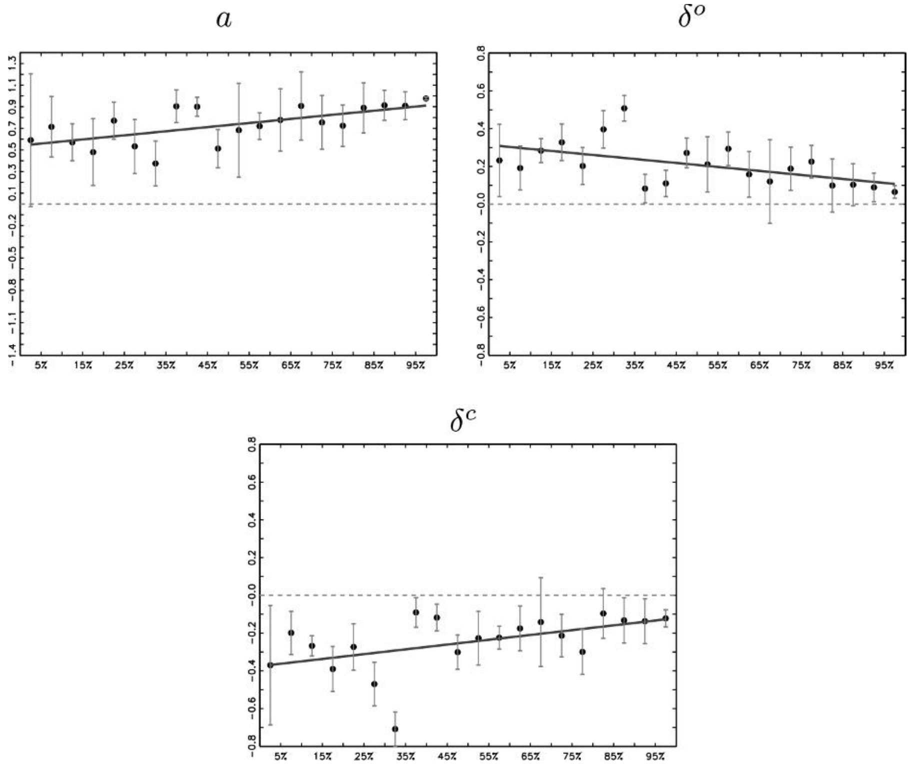

Figure 3 Parameter estimates (a, δ^o, and δ^c) for the latent factor dynamics for all twenty investor groups, ordered on the x-axis. The solid dot represents the estimated coefficient, and the bars represent the 95% confidence bounds. The solid line is the ordinary least squares (OLS) regression fit through the estimated coefficients.

investors, our results indicate the clear need to include a latent factor and to control for unobservable effects.

As covariates in our specification, we include the current (since the opening of the position) paper profit/loss in the currency pair ($\gamma_{\dot{\text{P/L 1}}}$), the paper profit/loss in the portfolio of all positions ($\gamma_{\dot{\text{P/L pf}}}$), and excess volume ($\gamma_{\dot{\text{vol}}}$). The paper profit/loss at a particular point of time is computed as the potential profit or loss (denominated in USD) that would have been obtained if the trader had decided to close his position at the prevailing market rates at that time. The portfolio paper profit/loss is the sum over the paper profit/losses of all open positions. At each event time, the excess volume variable is computed as the logarithm of the volume of the transaction divided by the average volume of the last three transactions (denominated in USD). A value larger than zero indicates that the current transaction is larger than the locally average volume a particular investor trades.

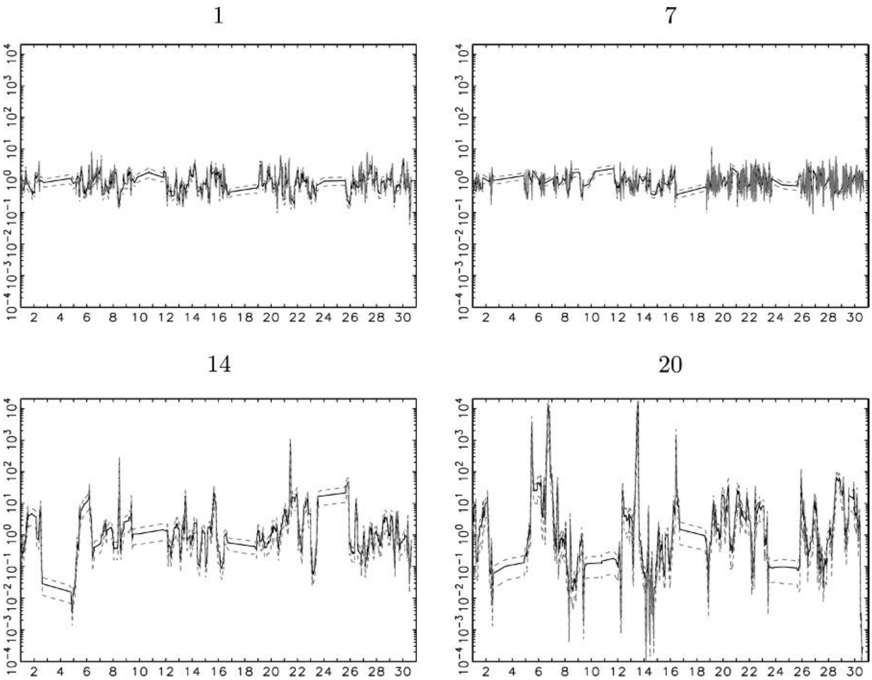

Figure 4 Smoothed estimates (solid line) of the log latent factor for selected investor groups 1, 7, 14, and 20. The x-axis always denote times measured in days, and the y-axis are presented on a log scale. The dashed lines mark 95% confidence bounds.

The proposed panel intensity model allows us to view the disposition effect and the effect of increased exposure to risk on open and close trade frequencies from a broader perspective than in the literature cited in Section 1 by focusing on the individual trader micro-level and considering the timing of trading decisions as central to the analysis. After all, the disposition effect is explicitly defined as a time effect, for the analysis of which our framework is perfectly suited. The main advantages of the model are its ability to include observable individual heterogeneity through individual fixed effects and time-varying covariates describing the investors' information set as well as to account for unobservable time-varying effects through the latent factor.

In the literature (cf. Odean 1998a; Shapira and Venezia 2001; Locke and Mann 2005; Haigh and List 2005), the disposition effect is regarded usually in isolation and with respect to a single security position. An advantage of our framework is that it allows us to analyze to what extent the portfolio profit/loss affects trading decisions in the separate currency pairs and amplifies or dampens the security-specific disposition effects. The partitioning of the population of investors into groups enables us to analyze the differences in the degree to which certain groups are prone to this behavioral effect. Recall from the descriptive analysis that there are no differences in the number of traded currency pairs from smaller to larger investors, so that differences in portfolio effects cannot be attributed to the number of traded currency pairs.

In Figure 5, we plot the parameters for the single-position paper profit/loss, the portfolio paper profit/loss, and the excess trading volume for both the open and the close intensity subprocesses across all twenty investor groups, ordered on the x-axis. The solid dots represent the estimated coefficients (with the corresponding 95% confidence intervals), and the solid line is the OLS regression fit through the estimated coefficients.

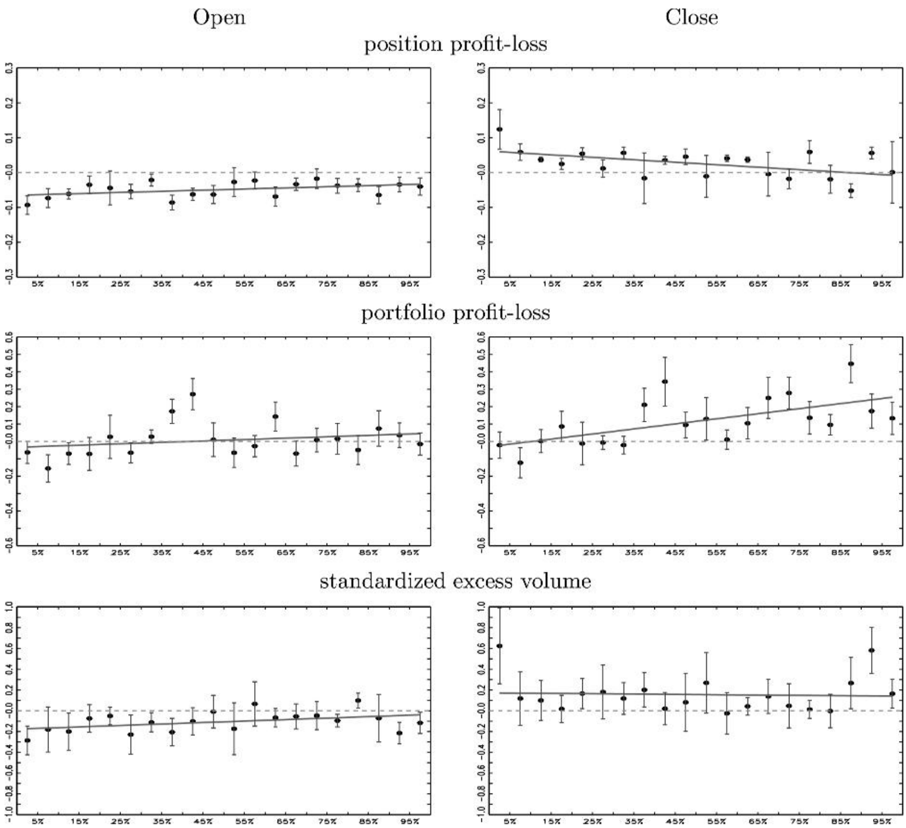

Figure 5 Parameter estimates for the single-position paper profit/loss, the portfolio paper profit/loss, and excess transaction volume for both open and close intensity subprocesses across the twenty investor groups, ordered on the x-axis. The solid dot represents the estimated coefficient, and the bars represent the 95% confidence bounds. The solid line is the OLS regression fit through the estimated coefficients.

We first focus on the upper-row panel illustrating the impact of the paper profit/loss in the single position ($\gamma_{\dot{P}/L\,1}$) on the open and close intensities. In general, we observe that the coefficients on the opening intensity are negative, while those for the closing intensity tend to be positive. Moreover, while there is only a slight upward trend for the open subprocess, the impact of the variable on the closing intensity is decreasing, the larger the investors in the group. A positive sign of the coefficient for the closing intensity implies that the larger the profit (loss), the higher (lower) the intensity to close the position. Thus, we find evidence for the presence of a disposition effect for small investors which diminishes as investors become larger. This is in line with the findings of Shapira and Venezia (2001), Dhar and Zhu (2006), Goetzmann and Massa (2008), and Chen et al. (2007) stated above. The coefficients for the open intensity cannot be interpreted in light of the narrow definition of the disposition effect, which only considers the relative propensity to realize profits compared to losses, restricting the analysis to closing events. The prospect theory of Kahneman and Tversky (1979) is a more general concept, which describes an investor as being risk averse when winning and risk seeking when losing. This implies that an investor would be less willing to increase his exposure as his position becomes more profitable. The negative coefficients for the open intensity are perfectly consistent with such loss aversion behavior and hence allow for a more general interpretation and justification of the disposition effect.

The second-row panel depicts the influence of the paper profit/loss in the portfolio of positions ($\gamma_{\dot{P}/L\,pf}$) on both intensities. This variable can be considered as a complementary decision factor, which sheds new light on traders' behavior and the disposition effect. While the impact on the open intensity is ambiguous and often insignificant, the larger investors are significantly influenced by the success of their total portfolio in their decisions to close positions. We find that these investors tend to close each of their positions faster when their portfolio is generating profits. For smaller investors, on the contrary, the portfolio profit/loss does not have an influence on the propensity to open or close a position. A possible explanation for these observations is that smaller investors may be narrow framed (Thaler 1985) and oblivious to the dependencies between their positions in the portfolio. This effect cannot be attributed to smaller investors holding only a single position, as the groups have been selected so that all traders are invested in at least three currency pairs. Also, the descriptive analysis showed that there is no systematic pattern in the number of currency pairs held. From a different perspective, these findings can be interpreted as a higher-level disposition effect to which larger investors are more susceptible.

The plots in the third-row panel illustrate the effect of a temporarily higher risk exposure (measured as excess trading volume) in a certain currency pair on the open and close intensities. Generally, a larger trade can be initiated for a variety of reasons. We argue that it signals an increase in risk-taking appetite. This may be caused on the one hand by overconfident behavior arising from spurious past trading success (self-attribution bias; Wolosin, Sherman, and Till 1973) or a wrongly perceived ability to better interpret news or an emerging chart pattern (biased perceptions of information; Alpert and Raiffa 1982). In such case, one should observe even more aggressive and frequent trading in the near future (cf. Daniel, Hirshleifer, and Subrahmanyam 1998; Odean 1998b; Wang 1998; Gervais and Odean 2001; Scheinkman and Xiong 2003). On the other hand, more active risk taking may also emerge from a rational trading strategy in which investors learn and become more experienced without misinterpretation of past information (cf. Nicolosi, Peng, and Zhu 2009; Seru, Shumway, and Stoffman 2009). We observe that in general, excess trading volume ($\gamma^{\cdot}_{\mathrm{vol}}$) has a negative influence on the open intensity and a positive impact on the close intensity. Although for some groups the coefficients turn out to be insignificant, the signs of the coefficients seem unambiguous. In the case in which an investor holds a position that is riskier than his usual positions, he seems to be reluctant to further invest into the currency pair and seems to be faster in closing out this position. Taken together, these results point at risk-averse behavior at the high-frequency level. The first observation may also simply reflect a budget constraint in the sense that there is less free investment capital available to undertake further investments. The fact that this effect seems to be also more pronounced for the small private retail than for larger investors points into the same direction. The second observation indicates that the investor seems to pay more attention to his higher risk position. For such positions, it seems reasonable that the investor spends more effort in actively monitoring the price process or sets special limit orders (stop-loss and/or take-profit) with tighter limits. In both cases, one would observe a higher closing intensity, irrespectively of the type of supervision strategy. This is consistent with the general view in the finance literature that investors like to be compensated for holding a risky position and that they tend to act more carefully when managing a larger risk. The early literature on inventory models for market makers (cf. Demsetz 1968; Madhavan and Smidt 1993) and also the related literature on asymmetric information (cf. Glosten and Milgrom 1985; Easley and O'Hara 1987) already emphasizes these effects. In a situation of increased risk and/or uncertainty, market makers tend to raise the bid–ask spread to insure themselves from being picked up and thereby implying either less or more costly trading (cf. Admati and Pfleiderer 1988). However, the focus of the market microstructure literature has moved away from inventory models to models with asymmetric information, and the importance of inventory not only for the market maker but also for the individual investors seems to be neglected. Our results and the results of Nolte and Nolte (2010) show that inventory matters.

Overall, our results are consistent with the theoretical and empirical findings on the disposition effect and its heterogeneity across different groups of investors. The novelty of the proposed modeling framework is that it highlights the importance of the time dimension, the separate open and close intensity processes, and allows for a joint analysis of complex individual trading strategies in light of behavioral finance and market microstructure theories. The proposed approach delivers new insights into the study of the disposition effect by allowing us to adopt a more general interpretation of the prospect theory. The inclusion of the portfolio profit/loss as a decision factor broadens the interpretation of the disposition effect and the analysis of trading patterns.

The goodness-of-fit of our models is evaluated by comparing properties of the "raw" inter-event durations to the model residuals. In an intensity-based framework, the integrated intensities (see Equation (6)) can be considered as generalized residuals which under the correct model specification should be i.i.d exponentially distributed with unit mean. Here, we report goodness-of-fit statistics and quantile–quantile (QQ) plots again only for investor groups 1, 7, 14, and 20. The complete statistics and figures can be found in Web Appendix B. The goodness-of-fit diagnostics are given in Table 4. While we still detect a slight overdispersion of the residuals, indicating a slight degree of misspecification, also evident in the QQ plots in Figure 6, the dynamic properties of the inter-event durations are captured well by the model. This is confirmed by the Ljung–Box and the Brock, Dechert, and Scheinkman (1987) (BDS) tests. We observe that the Ljung–Box test statistics of the generalized residual series decrease considerably in comparison to those of the raw data series in the majority of cases. The same observation also holds for the BDS test, which does not merely test for uncorrelated but for i.i.d. distributed observations.

Table 4 Diagnostics for the raw and the residual series of the open and close subprocesses for selected investor groups 1, 7, 14, and 20

	Open		Close		Open		Close	
	Raw	Res	Raw	Res	Raw	Res	Raw	Res
	1				7			
Mean	1420.94	0.97	4981.64	0.96	714.97	1.03	1994.36	1.12
Std	3496.11	1.07	7343.19	1.29	2864.14	1.10	5707.96	1.30
LB (20)	520.80	24.82	32.22	32.37	632.12	139.65	319.92	21.92
LB (50)	800.40	53.06	58.46	82.57	1012.91	177.36	333.82	53.43
BDS ($m = 2$)	10.53	1.91	1.32	0.25	9.74	5.84	14.58	1.00
BDS ($m = 3$)	12.61	2.41	0.92	−0.12	11.97	6.65	15.89	1.15
	14				20			
Mean	1063.60	1.01	1793.88	0.99	613.92	1.02	448.31	0.99
Std	3086.08	1.04	4464.85	1.13	2254.13	1.21	1657.68	1.07
LB (20)	247.64	22.02	53.17	23.62	224.26	41.60	299.85	86.60
LB (50)	309.76	63.20	88.13	44.98	314.24	66.94	433.15	129.47
BDS ($m = 2$)	6.53	1.49	4.34	0.75	7.20	2.67	7.34	2.64
BDS ($m = 3$)	8.02	2.40	5.24	1.48	8.67	3.77	8.76	2.81

The series are pooled over currency pairs and investors. LB $\triangleq$ Ljung–Box test statistic, BDS(m = embedding dimension) $\triangleq$ Brock–Dechert–Scheinkman test statistic $\sim N(0, 1)$. Std, standard deviation.

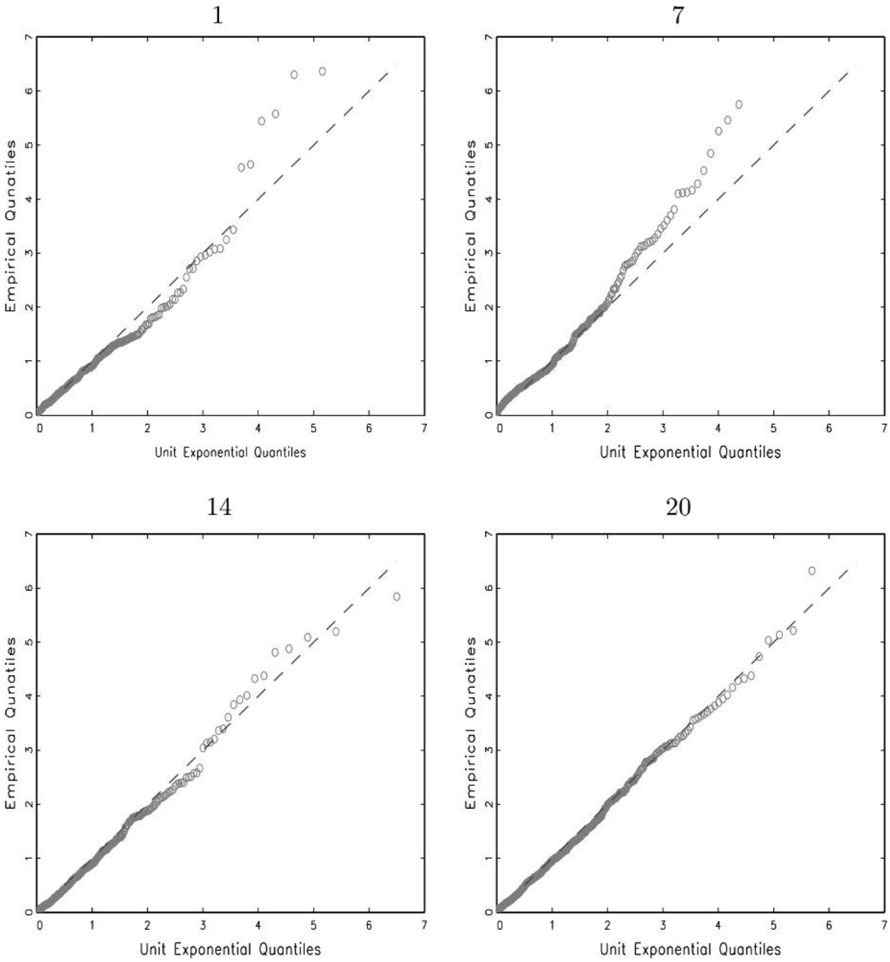

Figure 6 QQ plots of pooled open and close subprocesses residual series against the unit exponential distribution for selected investor groups 1, 7, 14, and 20.

To improve the model fit one could consider alternative baseline intensity functions of semi- or nonparametric form. In earlier versions of the paper, we used a simple multivariate Weibull baseline intensity without investor-specific effects and an FFF component, which turned out to be inferior in terms of goodness-of-fit to the one currently employed. We also considered models in which the dynamics was driven, in addition to the latent factor, by an autoregressive conditional intensity (Russell 1999) component. This complicated the interpretation of the model dynamics and did not improve the fit considerably. Further extensions of the model could allow for a broader set of individual- and currency-specific effects.

4 CONCLUSIONS

In this paper, we propose an econometric model for the analysis of complex trading activity datasets in an intensity-based framework. Such datasets contain very detailed information about the trading history of single traders and provide more insights into the market microstructure and investors' trading behavior beyond the informational content of typical high-frequency datasets. From an econometric point of view, analyzing activity datasets is rather challenging, due to their multidimensional panel structure spanning time, types of trading activity, securities, and investors, complicated by irregularly spaced observations. The model developed in the paper is suited to cope with this data structure.

An attractive feature of the intensity-based framework is its flexibility in terms of capturing the impact of observable time-varying covariates on the underlying processes. Since not all information can be observed, however, we include a latent factor in the model which is responsible for capturing hidden correlation structures. We estimate the panel intensity model adopting the EIS algorithm of Richard and Zhang (2007).

The model is applied to a trading activity dataset from OANDA FXTrade in order to analyze the trading behavior of different groups of investors, categorized according to their investment turnover. The beauty of the methodology is that time plays the central role, which allows us to draw immediate conclusions with respect to behavioral biases influencing the timing of investment decisions, such as the disposition effect. We find that the standard disposition effect is complemented by the impact of the total portfolio performance on the length of the investment periods. The joint modeling of the processes related to opening and closing financial positions allows for a broader interpretation of the disposition effect in light of the prospect theory of Kahneman and Tversky (1979). Apart from new insights into the disposition effect, the model delivers insights into the impact of risk aversion at the high-frequency level on the trading behavior.

APPENDIX A: ESTIMATION

We consider the explicit form and the estimation of the parameters in the likelihood function. Let W denote the set of data matrices W^{kn} for each currency pair $k = 1, \ldots, K$ and investor $n = 1, \ldots, N$, where the ith row of W^{kn}, w_i^{kn}, consists of the following data:

$$w_i^{kn} = (t_i^{kn}, d_i^1, \ldots, d_i^S), \quad \text{with } i = 1, \ldots, N^{kn}(T).$$

With W_i^{kn} we denote the history of w_i^{kn} up to and including t_i^{kn}, that is,

$$W_i^{kn} = \{w_j^{kn}\}_{j=1}^{i}.$$

Furthermore, let Z_i^{kn} for $k = 1, \ldots, K$ and $n = 1, \ldots, N$ denote the set consisting of the following time-varying covariate data:

$$Z_i^{kn} = \{\{z_j^{1kn} | j = 1, \ldots, M^{1kn}(t_i^{kn-})\}, \ldots, \{z_j^{Skn} | j = 1, \ldots, M^{Skn}(t_i^{kn-})\}\}.$$

Recall that the likelihood function of our model is given by

$$\begin{aligned}
\mathcal{L}(W;\theta) &= \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \prod_{s=1}^{S} \exp\left(d^s_{N^{kn}(t_i)} \ln \theta^{skn}\left(t^{kn}_{N^{kn}(t_i)} \middle| \mathfrak{F}_{t_i^-}, \lambda_i \right) \right. \\
&\quad \left. - \int_{t^{kn}_{N^{kn}(t_i)-1}}^{t^{kn}_{N^{kn}(t_i)}} \theta^{skn}(u | \mathfrak{F}_{u^-}, \lambda_{N(u^-)+1}) \mathrm{d}u \right) \rho\left(\lambda_i | \mathfrak{F}_{t_i^-} \right) \mathrm{d}\Lambda \\
&= \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \prod_{s=1}^{S} \exp\left(d^s_{N^{kn}(t_i)} \ln \theta^{skn}\left(t^{kn}_{N^{kn}(t_i)} \middle| \mathfrak{F}_{t_i^-}, l_i \right) \right. \\
&\quad \left. - \int_{t^{kn}_{N^{kn}(t_i)-1}}^{t^{kn}_{N^{kn}(t_i)}} \theta^{skn}(u | \mathfrak{F}_{u^-}, l_{N(u^-)+1}) \mathrm{d}u \right) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(l_i - \mu_i)^2}{2} \right) \mathrm{d}L,
\end{aligned}$$

where $L = \ln \Lambda$ and the second equality follows from a change of the variable λ to l. Using the datasets defined above, the likelihood function can be rewritten as

$$\begin{aligned}
\mathcal{L}(W;\theta) &= \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} g^{kn}\left(w^{kn}_{N^{kn}(t_i)} | W^{kn}_{N^{kn}(t_i)-1}, L_i, Z^{kn}_{N^{kn}(t_i)} \right) p(l_i | L_{i-1}) \mathrm{d}L \\
&= \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \varphi^{kn}\left(w^{kn}_{N^{kn}(t_i)}, l_i | W^{kn}_{N^{kn}(t_i)-1}, L_{i-1}, Z^{kn}_{N^{kn}(t_i)} \right) \mathrm{d}L, \qquad \text{(A1)}
\end{aligned}$$

where g^{kn} denotes the product of the survivor and the intensity functions, p the density of the conditional normal distribution, and φ^{kn} the resulting corresponding joint conditional density. Since this likelihood involves the computation of an $N(T)$-dimensional integral, we employ the EIS technique of Liesenfeld and Richard (2003), which has been used for estimating SCI models by Bauwens and Hautsch (2006). The EIS technique is based on simulation of the likelihood function (Equation (A1)) which can be rewritten as

$$\begin{aligned}
\mathcal{L}(W;\theta) = \int_{\mathbb{R}^{N(T)}} \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \frac{\varphi^{kn}\left(w^{kn}_{N^{kn}(t_i)}, l_i | W^{kn}_{N^{kn}(t_i)-1}, L_{i-1}, Z^{kn}_{N^{kn}(t_i)} \right)}{m(l_i | L_{i-1}, \phi_i)} \\
\times \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} m(l_i | L_{i-1}, \phi_i) \mathrm{d}L,
\end{aligned}$$

where $m(l_i | L_{i-1}, \phi_i)$ is a sequence of auxiliary importance samplers which are used to draw a trajectory of the latent factor, given some additional parameters ϕ_i of the sampler. The estimation then proceeds by generating R trajectories of the latent factor and averaging over the draws

$$\mathcal{L}_R(W;\theta) = \frac{1}{R}\sum_{r=1}^{R} \frac{\prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \varphi^{kn}\left(w^{kn}_{N^{kn}(t_i)}, l_i^{(r)} | W^{kn}_{N^{kn}(t_i)-1}, L_{i-1}^{(r)}, Z^{kn}_{N^{kn}(t_i)}\right)}{\prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} m(l_i^{(r)} | L_{i-1}^{(r)}, \phi_i)}, \quad \text{(A2)}$$

where the bracketed superscript r indicates the values of the corresponding variable or set for the rth repetition. The idea of the EIS approach is to find the values of the parameters ϕ_i for $i = 1, \ldots, N(T)$ such that the sampling variance of $\mathcal{L}_R(W;\theta)$ is minimized. We sketch the main idea and steps involved in the EIS methodology and refer the reader to Richard and Zhang (2007) for details.

The densities $m(l_i|L_{i-1}, \phi_i)$ can be decomposed as

$$m(l_i|L_{i-1}, \phi_i) = \frac{k(L_i, \phi_i)}{\chi(\phi_i, L_{i-1})}, \quad \text{(A3)}$$

which can be interpreted as writing the density of l_i conditional on its past as the ratio of the joint density and the marginal density or as decomposing the density into a kernel and an integrating constant independent of l_i. Richard and Zhang (2007) show that the optimal values of ϕ_i can be determined by solving recursively for $i = N(T), N(T) - 1, \ldots, 2, 1$ a minimization problem, which in our case can be written as

$$\hat{\phi}_i(\theta) = \text{argmin}_{\phi_i} \sum_{r=1}^{R} \left(\ln \left(\prod_{\mathcal{C}_i} \varphi^{kn} \left(w^{kn}_{N^{kn}(t_i)}, l_i^{(r)} \middle| W^{kn}_{N^{kn}(t_i)-1}, L_{i-1}^{(r)}, Z^{kn}_{N^{kn}(t_i)} \right) \right.\right.$$
$$\left.\left. \chi\left(\phi_{i+1}, L_i^{(r)}\right)\right) - \phi_{0,i} - \ln(k(L_i^{(r)}, \phi_i)) \right)^2, \quad \text{(A4)}$$

where $\phi_{0,i}$ are additional scalar parameters and $\chi(\phi_{N(T)+1}, L_{N(T)}) \equiv 1$. The intuition behind the minimization problem is that it looks for the values of ϕ_i which minimize the distance between the densities φ and m, where m has been decomposed as in Equation (A3). To make Equation (A4) operational, one needs to specify the functional form of the kernel k. We follow Liesenfeld and Richard (2003) and choose as a kernel a parametric extension to the direct samplers p given by

$$k\left(L_i, \phi_i\right) = p(l_i|L_{i-1})\zeta\left(l_i, \phi_i\right),$$

where ζ is itself a Gaussian density kernel:

$$\zeta\left(l_i, \phi_i\right) = \exp(\phi_{1,i} l_i + \phi_{2,i} l_i^2).$$

Since a product of normal kernels is a normal kernel as well, we obtain for $k\left(L_i, \phi_i\right)$

$$k(L_i, \phi_i) \propto \exp\left((\phi_{1,i} + \mu_i) l_i + \left(\phi_{2,i} - \frac{1}{2}\right) l_i^2 - \frac{1}{2}\mu_i^2 \right)$$
$$= \exp\left(-\frac{1}{2\pi_i^2}(l_i - \kappa_i)^2 \right) \exp\left(\frac{\kappa_i^2}{2\pi_i^2} - \frac{1}{2}\mu_i^2 \right),$$

where

$$\pi_i^2 = (1 - 2\phi_{2,i})^{-1} \quad \text{and} \tag{A5}$$

$$\kappa_i = (\phi_{1,i} + \mu_i)\pi_i^2. \tag{A6}$$

It follows that

$$\chi(\phi_i, L_{i-1}) = \exp\left(\frac{\kappa_i^2}{2\pi_i^2} - \frac{\mu_i^2}{2}\right). \tag{A7}$$

Under this choice of kernels class, $p(l_i|L_{i-1})$ cancels out in the minimization problem (A4), which can then be rewritten as

$$\hat{\phi}_i(\theta) = \text{argmin}_{\phi_i} \sum_{r=1}^{R} \left(\ln\left(\prod_{\mathcal{C}_i} g^{kn}\left(w^{kn}_{N^{kn}(t_i)} | W^{kn}_{N^{kn}(t_i)-1}, L_i^{(r)}, Z^{kn}_{N^{kn}(t_i)}\right)\chi(\phi_{i+1}, L_i^{(r)})\right) - \phi_{0,i} - \ln(\zeta(l_i^{(r)}, \phi_i))\right)^2. \tag{A8}$$

The implementation of the sequential ML-EIS approach can be summarized in the following steps:

STEP 1. Draw R trajectories $\{l_i^{(r)}\}_{i=1}^{N(T)}$ from $\{N(\mu_i, 1)\}_{i=1}^{N(T)}$.

STEP 2. For each i with i: $N(T) \to 1$ solve the R-dimensional OLS problem in (A8).

STEP 3. Calculate the sequences $\{\pi_i^2\}_{i=1}^{N(T)}$ and $\{\kappa_i\}_{i=1}^{N(T)}$ from equations (A5) and (A6) and draw R trajectories of $\{l_i^{(r)}\}_{i=1}^{N(T)}$ from $\{N(\kappa_i, \pi_i^2)\}_{i=1}^{N(T)}$ to compute the likelihood function given in (A2).

APPENDIX B: SMOOTHED ESTIMATES

For evaluation purposes of the latent factor, we compute smoothed estimates of it following the procedure outlined in Koopman, Lucas, and Monteiro (2008). Using the notation from Appendix A, we construct an estimate of the latent factor given the data $\hat{l}_{i|N(T)}$ by weighting each latent factor $l_i^{(r)}$ by its contribution to the likelihood function. Hence,

$$\hat{l}_{i|N(T)} = \frac{\sum_{r=1}^{R} \pi^{(r)} l_i^{(r)}}{\sum_{r=1}^{R} \pi^{(r)}},$$

with

$$\pi^{(r)} = \prod_{i=1}^{N(T)} \prod_{\mathcal{C}_i} \frac{\varphi^{kn}\left(w^{kn}_{N^{kn}(t_i)}, l_i^{(r)} | W^{kn}_{N^{kn}(t_i)-1}, L_i^{(r)}, Z^{kn}_{N^{kn}(t_i)}\right)}{m(l_i^{(r)} | L_{i-1}^{(r)}, \hat{\phi}_i)}.$$

The standard errors for $\hat{l}_{i|N(T)}$ are computed following Koopman, Lucas, and Monteiro (2008) as

$$\sqrt{\frac{\sum_{r=1}^{R} \pi^{(r)} (l_i^{(r)})^2}{\sum_{r=1}^{R} \pi^{(r)}} - (\hat{l}_{i|N(T)})^2}.$$

REFERENCES

Admati, A., and P. Pfleiderer. 1988. A Theory of Intra-Day Patterns: Volume and Price Variability. *Review of Financial Studies* 1: 3–40.

Alpert, M., and H. Raiffa. 1982. "A Progress Report on the Training of Probability Assessors." In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press, pp. 294–305.

Badrinath, S. G., and W. G. Lewellen. 1991. Evidence on Tax-Motivated Securities Trading Behavior. *Journal of Finance* 46: 369–382.

Barberis, N., and W. Xiong. 2009. What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation. *Journal of Finance* 64: 751–784.

Bauwens, L., and N. Hautsch. 2006. Stochastic Conditional Intensity Processes. *Journal of Financial Econometrics* 4: 450–493.

Bauwens, L., and D. Veredas. 2004. The Stochastic Conditional Duration Model: A Latent Factor Model for the Analysis of Financial Durations. *Journal of Econometrics* 119: 381–412.

Bowsher, C. 2007. Modelling Security Market Events in Continuous Time: Intensity Based, Multivariate Point Process Models. *Journal of Econometrics* 141: 876–912.

Brock, W. A., W. D. Dechert, and J. A. Scheinkman. 1987. "A Test for Independence Based on the Correlation Dimension." Working Paper, Department of Economics, University of Wisconsin-Madison.

Chakravarty, S., and C. Holden. 1995. An Integrated Model of Market and Limit Orders. *Journal of Financial Intermediation* 4: 213–241.

Chen, G.-M., K. A. Kim, J. R. Nofsinger, and O. M. Rui. 2007. Trading Performance, Disposition Effect, Overconfidence, Representativeness Bias, and Experience of Emerging Market Investors. *Journal of Behavioral Decision Making* 20: 425–451.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam. 1998. Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance* 53: 1839–1885.

Demsetz, H. 1968. The Cost of Transacting. *Quarterly Journal of Economics* 82:33–53.

Dhar, R., and N. Zhu. 2006. Up Close and Personal: An Individual Level Analysis of the Disposition Effect. *Management Science* 52: 726–740.

Easley, D., and M. O'Hara. 1987. Price, Trade Size, and Information in Securities Markets. *Journal of Financial Economics* 19: 69–90.

Foucault, T. 1999. Order Flow Composition and Trading Costs in a Dynamic Limit Order Market. *Journal of Financial Markets* 2: 99–134.

Foucault, T., O. Kadan, and E. Kandel. 2005. Limit Order Book as a Market for Liquidity. *Review of Financial Studies* 18: 1171–1217.

Gervais, S., and T. Odean. 2001. Learning to be Overconfident. *Review of Financial Studies* 14: 1–27.

Glosten, L. R., and P. R. Milgrom. 1985. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics* 14: 71–100.

Goetzmann, W., and M. Massa. 2008. Disposition Matters: Volume, Volatility and Price Impact of a Behavioral Bias. *Journal of Portfolio Management* 34: 103–125.

Grinblatt, M., and M. Keloharju. 2001. What Makes Investors Trade? *Journal of Finance* 56: 1053–1073.

Haigh, M. S., and J. A. List. 2005. Do Professional Trades Exhibit Myopic Loss Realization Aversion? An Experimental Analysis. *Journal of Finance* 60: 523–534.

Kahneman, D., and A. Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47: 263–292.

Kaniel, R., and H. Liu. 2006. So What Orders Do Informed Traders Use? *Journal of Business* 79: 1867–1913.

Koopman, S. J., A. Lucas, and A. Monteiro. 2008. The Multi-state Latent Factor Intensity Model for Credit Rating Transitions. *Journal of Econometrics* 142: 399–424.

Lancaster, T. 1997. *The Econometric Analysis of Transition Data*. Cambridge: Cambridge University Press.

Lease, R. C., W. G. Lewellen, and G. G. Schlarbaum. 1974. The Individual Investor: Attributes and Attitudes. *Journal of Finance* 29: 413–433.

Liesenfeld, R., and J.-F. Richard. 2003. Univariate and Multivariate Stochastic Volatility Models: Estimation and Diagnostics. *Journal of Empirical Finance* 10: 505–531.

Locke, P. R., and S. C. Mann. 2005. Professional Trader Discipline and Trade Disposition. *Journal of Financial Economics* 76: 401–444.

Lunde, A., and A. Timmermann. 2005. Completion Time Structures of Stock Price Movements. *Annals of Finance* 1: 293–326.

Madhavan, A., and S. Smidt. 1993. An Analysis of Daily Changes in Specialists' Inventories and Quotations. *Journal of Finance* 48: 1595–1628.

Nicolosi, G., L. Peng, and N. Zhu. 2009. Do Individual Investors Learn from Their Trading Experience? *Journal of Financial Markets* 12: 317–336.

Nolte, I., and S. Nolte. 2010. "How do Individual Investors Trade?" Warwick Business School Working Paper, University of Warwick.

Odean, T. 1998a. Are Investors Reluctant to Realize Their Losses? *Journal of Finance* 53: 1775–1798.

Odean, T. 1998b. Volume, Volatility, Price, and Profit When All Traders Are Above Average. *Journal of Finance* 53: 1887–1934.

Parlour, C. 1998. Price Dynamics in Limit Order Markets. *Review of Financial Studies* 11: 789–816.

Richard, J.-F., and W. Zhang. 2007. Efficient High-Dimensional Importance Sampling. *Journal of Econometrics* 141: 1385–1411.

Russell, J. R. 1999. "Econometric Modeling of Multivariate Irregularly-Spaced High-Frequency Data." Working Paper, University of Chicago.

Scheinkman, J. A., and W. Xiong. 2003. Overconfidence and Speculative Bubbles. *Journal of Political Economy* 111: 1183–1219.

Schlarbaum, G. G., W. G. Lewellen, and R. C. Lease. 1978a. The Common-Stock-Portfolio Performance Record of Individual Investors: 1964–70. *Journal of Finance* 33: 429–441.

Schlarbaum, G. G., W. G. Lewellen, and R. C. Lease. 1978b. Realized Returns on Common Stock Investments: The Experience of Individual Investors. *Journal of Business* 51: 299–325.

Seru, A., T. Shumway, and N. Stoffman. 2009. Learning by Trading. *Review of Financial Studies* 23: 705–739.

Shapira, Z., and I. Venezia. 2001. Patterns of Behavior of Professionally Managed and Independent Investors. *Journal of Banking and Finance* 25: 1573–1587.

Shefrin, H., and M. Statman. 1985. The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *Journal of Finance* 40: 770–790.

Thaler, R. 1985. Mental Accounting and Consumer Choice. *Marketing Science* 4: 199–214.

Wang, F. A. 1998. Strategic Trading, Asymmetric Information and Heterogeneous Prior Beliefs. *Journal of Financial Markets* 1: 321–352.

Wolosin, R., S. Sherman, and A. Till. 1973. Effects of Cooperation and Competition on Responsibility Attribution After Success and Failure. *Journal of Experimental and Social Psychology* 9: 220–235.