Modeling Asset Prices for Algorithmic and High Frequency Trading[*]
Forthcoming in Applied Mathematical Finance

Álvaro Cartea[†]   and   Sebastian Jaimungal[‡]

January 7, 2013

Abstract

Algorithmic Trading (AT) and High Frequency (HF) trading, which are responsible for over 70% of US stocks trading volume, have greatly changed the microstructure dynamics of tick-by-tick stock data. In this paper we employ a hidden Markov model to examine how the intra-day dynamics of the stock market have changed, and how to use this information to develop trading strategies at high frequencies. In particular, we show how to employ our model to submit limit-orders to profit from the bid-ask spread and we also provide evidence of how HF traders may profit from liquidity incentives (liquidity rebates). We use data from February 2001 and February 2008 to show that while in 2001 the intra-day states with shortest average durations (waiting time between trades) were also the ones with very few trades, in 2008 the vast majority of trades took place in the states with shortest average durations. Moreover, in 2008 the states with shortest durations have the smallest price impact as measured by the volatility of price innovations.

1   Introduction

Not too long ago the vast majority of the transactions in stock exchanges were executed by humans or required frequent human input along the trading process. This trend has changed dramatically over the last decade, and especially over the last five years, where fast computers now conduct most of the transactions. The use of computer algorithms that make trading decisions, submit orders, and manage those orders after submission, is known as algorithmic trading (AT). This technological change has taken over most exchanges and different sources report that between 50% to 77% of trading volume in the US equities markets is due to AT, SEC (2010), and Cvitanić and Kirilenko (2010).

Trading on the back of powerful computers and software that relies heavily on the ability to process and react quickly to the flux of trades and market information, has made it possible to execute large volumes of trades over short periods of time. Some of the effects of AT in stock exchanges can be gauged in disparate ways including: daily volume, speed of execution, daily

---

[†]Department of Mathematics, University College London, UK; *a.cartea@ucl.ac.uk*

[‡]Department of Statistical Sciences, University of Toronto, Canada; *sebastian.jaimungal@utoronto.ca*

trades, and average trade size. For example, the SEC reports that in the NYSE between 2005 and 2009: consolidated average daily share volume increased 181%; average speed of execution for small, immediately executable (marketable) orders shrunk from 10.1 to 0.7 seconds; consolidated average daily trades increased 662%; and consolidated average trade size decreased from 724 to 268 shares, SEC (2010). These substantial changes in the aggregate figures are the tip of the iceberg in modern electronic trading and are showing a particular aspect of how AT is changing financial markets in general and equity markets in particular.

But what are the fundamental changes in the tick-by-tick dynamics of stock prices as a consequence of AT? From the aggregate figures it is not clear if new trading patterns have emerged, and if they have, what are their key characteristics. AT has become an arms race and the profitability of these algorithms not only depends on the level of participation of other types of traders, for instance liquidity or noise traders, but also on how AT strategies coexist with other algorithmic traders.

In this paper we model stock-price dynamics and extract important information on changes in the market's behavior at a tick-by-tick level and use this information to design AT strategies. To model the tick-by-tick dynamics we start from the fact that AT has considerably changed the way in which trading is done and that historical stylized facts of tick-by-tick data might have been altered in a substantial way. In general, at this point one can only conjecture what are the principal strategies that AT deploy and how do they affect stock prices at high frequencies. However, in equilibrium, which patterns emerge or what are the new stylized facts of tick-by-tick dynamics are questions that can be answered and are key in the development of trading algorithms.

The majority of AT strategies are designed to compete for profits or manage risks whilst others are designed to execute third-party trades at best prices. Examples of types of strategies include: i) high frequency (HF) market making strategies which are designed to operate on extremely short-time scales. Currently, any strategies which are designed and/or are able to react within 100 milliseconds are considered HF, see Cartea, Jaimungal, and Ricci (2011), Cartea and Jaimungal (2012), Latza, Marsh, and Payne (2012). ii) Strategies that are designed to: minimize price impact when a large order must be executed over a fixed horizon; trigger other algorithmic traders into action; or other proprietary strategies based on speed of execution and information processing, see Cartea and Penalva (2012), Almgren (2003), Almgren (2009), and Lorenz and Almgren (2011). The complexity of these strategies and their effect on the dynamics of tick-by-tick stock prices requires a modeling approach that can describe the different states in which financial markets could be and how the market transitions between these states. Ideally, one would want to model states of the market where the presence of a type of strategy (or types of AT) is the main source that drives trading (or the lack of) activity. For instance, in situations where HF traders are active, one expects to be in a state where duration between trades is very low (very short periods of time between consecutive trades) until the market 'moves on' to another state where the underlying reasons for trading is a release of a piece of news or the market transitions to a state of more calm where less trading takes place.[1]

The overall effect of all these new trading strategies in the market at a macroscopic level might be easy to measure, but the microscopic changes are far from clear. In the era of superfast electronic trading the dynamics of prices at high frequencies will be a consequence of many economic and financial factors, but ultimately the trading decisions and the management of these orders are handled by AT. Thus, at an intraday level the market can show bursts of activity which may be accompanied by high or low volatility of price revisions (measured in transaction time), times of relatively low activity but with high volatility, and many other features very difficult to see at the aggregate level. Therefore, to model the tick-by-tick dynamics of stock prices we use a Hidden Markov Model (HMM) in order to capture the different states in which the market can be. In particular, our model determines the different states by: (i) the existence of regimes or states of

---

[1]In general, AT can refer to a wide range of computerized strategies including: technical indicators that alert traders when to enter/exit positions, computers arbitraging different exchanges or statistical patterns.

2

intra-day activity characterized by the intra-day trading intensity of market orders and how the market switches between these regimes; (ii) the state-dependent distribution of price revisions in transaction time controlling for trades that generate no change in prices and those that do; and (iii) the distribution of the duration between trades which is an important variable in intra-day AT and HF trading strategy design.

Our approach allows us to address two issues. First, from a purely financial viewpoint, how has the market changed in the recent years when AT has had an increasing role? Second, if nowadays most of what we see at the tick-by-tick stock price level is due to AT, can our model be used to design and execute high frequency trading strategies?

We summarize some of our findings as a response to these two questions. First, we employ tick-by-tick data for six stocks[2] over the two separate periods February 2001 and February 2008 to estimate the model parameters. Our empirical findings show that over the last decade the increasing presence of AT has not only changed the speed at which trades take place, but that there have been other fundamental changes in the intra-day characteristics of stock price behavior. We start by describing the characteristics that have changed little in the two periods: In 2001 and 2008 we find that i) for all but one asset, the states with shortest average durations is where the highest probability of observing zero price innovations occur; and ii) the states with longest average durations are generally the ones where the probability of observing a zero price innovation is lowest. Some of the changes between the two periods are: i) Across all stocks we study in 2008 the intra-day states with shortest average durations are also the states with lowest volatility of price revisions. The same is not true for 2001 where there is no general connection between states of high activity and volatility. ii) For all stocks in 2001 the intra-day state with the shortest durations is also the state where the least amount of trades took place. On the other hand, in 2008 we find the opposite result where, generally, the intra-day states with the longest durations have the least number of trades. Our empirical results are consistent with the theoretical predictions of Cvitanić and Kirilenko (2010) who show that the introduction of HF traders (HFTs) increases trading activity (by reducing the waiting time between trades) and modifies the distribution of price revisions by increasing mass around the center and thinning the tails.

Second, an advantage of our approach is that the HMM identifies not only the intra-day states of trading, and their persistence, but also captures the probability of trades with zero price revision and is able to capture the distribution of non-zero price revisions. This information allows us to discuss the potential profits from HF trading strategies such as rebate trading.

Moreover, the HMM allows us to develop a tick-by-tick trading strategy for an HF investor that posts immediate-or-cancel buy and sell limit-orders to profit from the bid-ask spread. An HF investor would execute this strategy over a time interval of length T which usually ranges between a couple of minutes and at most one day. The optimal strategy indicates the buy and sell quantities that the investor should post and how to update them every time a trade has occurred. These quantities depend on: the rate of arrival of trades, the intra-day-state of the market, the within state volatility of price revisions, the inventories which track the investor's accumulated stock, and finally, the proximity to the terminal investment horizon. We show that the spread posted by the HF investor is wider (tighter) when the volatility of the price innovation is high (low). Moreover, as the investor accumulates a long (short) position, the investor's bid-price (ask-price) moves away from the mid-price and the ask-price (bid-price) moves in towards it – inducing the investor to sell (buy) assets – which induces the inventories to mean-revert towards zero. Finally, all else equal, as the investment horizon approaches T, the investor submits buy and sell limit-orders which are tighter around the mid-price; a strategy that stresses the fact that the HF investor aims at holding zero inventories at time T.

As a particular example of this tick-by-tick strategy we calibrate the model to PCP data and find the profit and loss (PnL) distribution of an HF investor who posts limit-orders on PCP shares

---

[2]The six stocks are: AA, AMZN, HNZ, IBM, KO, and PCP.

based on a two-regime model and the PnL distribution of a less informed HF trader who cannot distinguish between the different regimes PCP may be in. We show that the less informed trader's PnL is almost always underperforming that of the better informed trader. This difference in PnLs can be in part attributed to adverse selection costs; the better informed trader is able to adjust her posts so that she is able to avoid losses as a consequence of being picked off by better informed traders.

The remainder of this article is organized as follows. Section 2 discusses how we jointly model durations and price revisions using an HMM. Section 3 describes the data used throughout the article and discusses some estimation issues. Section 4 presents and interprets the results. Section 5 presents a discussion of how HFTs can use the information provided by our model to execute certain trading strategies. Finally, Section 6 concludes.

2   Joint modeling of durations and price revisions

Over the last twenty years a substantial body of literature known as market microstructure has focused on the study of price formation at an intra-day level. Initially, most of the studies were at a theoretical level and particular attention was devoted to market structure and market designs and how these affect price formation – see e.g. de Jong and Rindi (2009). More recently, the availability of intra-day high-frequency data has enabled researchers to test some of the previous theories of market microstructure and to attempt to describe the stylized facts of high-frequency price dynamics.

Prior to the days when AT dominated most of the trading volume in the US equity markets, empirical studies with tick-by-tick data document some of the salient features of the intraday behavior of stock prices. For example most of the volume of transactions generally takes place at the opening and closing of the market, together with the U-shaped pattern of volatility over the day, see Engle (2000). Other studies, both theoretical and empirical, show that although traditional stock price models that assume that trades occur at every instant in time (or that they occur at equally spaced time-intervals) may be harmless at long-time scales, it is an unsuitable assumption for high-frequency data modeling. In particular, these studies show that at high frequencies, duration between trades conveys relevant information about the dynamics of tick-by-tick trades, including: the pace of the market, the presence of uninformed or informed traders, the volatility of price revisions, and implied volatility from the option markets, see Diamond and Verrechia (1987), Easley and O'Hara (1992), Engle and Russell (1998), Engle (2000), Dufour and Engle (2000), Manganelli (2005), and Cartea and Meyer-Brandis (2010).

Thus, duration is one of the features of stock price behavior that becomes highly relevant over short periods of time. This random variable is generally overlooked in most asset pricing models that have horizons of at least a few days because it is believed that any effect that durations may have are dissipated very quickly. But nowadays, when the majority of trades are executed by AT that process information on a tick-by-tick level, duration becomes an important variable to model because it conveys relevant information about the market over short-time intervals. From a statistical point of view, the calendar-time distribution of stock price dynamics (on small timescales) depends not only on the distribution of price revisions, but also on the distribution of duration. From a financial viewpoint, trading strategies are specifically designed to profit from price patterns and behavior over ever shrinking timescales.

As mentioned in the introduction, the speed of trade execution shrunk by a factor of ten in the last five years, strongly indicating that trading very quickly over short periods of time is at the heart of modern trading in general, and AT in particular. There are many factors that have contributed to the increase of AT. The introduction of limit order markets and changes in market structure have lowered the entry barriers to new participants. At the same time computer power has spectacularly increased and its costs dramatically decreased. Thus, the number of market
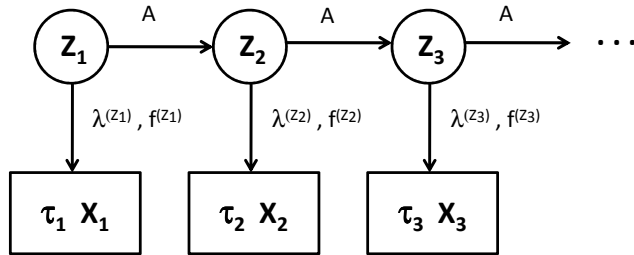
Figure 1: The intra-day-states Z_t evolve according to discrete time Markov chain with transition matrix A. Trades arrive at a rate of $\lambda^{(Z_t)}$ and have price revisions with pdf $f^{(Z_t)}$. Once a trade occurs, the world-state evolves.

participants has increased and the speed at which trading occurs has also increased.

The econometrics literature focusing on trade arrival started in earnest with the work of Engle and Russell (1998) who propose the autoregressive conditional duration (ACD) model to capture the time of arrival of financial data. Since then, most models have extended the ACD framework in different directions. See for example the logarithmic model of Bauwens and Giot (2000) and the augmented class of Fernandes and Grammig (2005) among others. Other extensions are based on regime-shifting and mixture ACD models, see for example Maheu and McCurdy (2000), Zhang, Russell, and Tsay (2001), Meitz and Terasvirta (2006), and Hujer, Vuletic, and Kokot (2002), and the recent work of Renault, van der Heijden, and Werker (2010) which proposes a structural model for durations between events and associated marks. For a comprehensive account of ACD models we refer the reader to Bauwens and Hautsch (2009).

Departing from the more traditional literature based on ACD models, we propose a finite-state HMM for the high-frequency dynamics of spot prices. We take this approach because it provides us not only with a good description of the statistical properties of the arrival of trades, but also, and more importantly, it provides us with a framework that is applicable to algorithmic and HF tick-by-tick trading design. Specifically, our model zooms in to the fine structure of price dynamics and is able to: distinguish between different trading regimes throughout the trading day and how the intra-day market switches between the different states; capture the distribution of durations between trades; and model the regime-dependent distribution of price revisions (trade and volatility clustering). The rest of this Section discusses the model we propose and Section 5 looks at tick-by-tick trading strategies.

We employ a finite state $\{1, \ldots, K\}$ discrete-time Markov chain Z_t, with transition matrix A, to modulate intra-day states. The time index in the Markov chain corresponds to the number of trades that have occurred during the trading day – in other words the time index marks the business time. Within a given intra-day state (or regime) the arrival of trades is governed by the regime-dependent hazard rate $\lambda_t = \lambda(Z_t)$, and price revisions are distributed according to a discrete-continuous mixture model. The discrete part of the distribution of price innovations models a zero price revision upon a trade occurring, while the continuous portion models non-zero price revisions, where all parameters are dependent on the intra-day-state. Specifically, we assume that the size of the log-mid-price revision X, in state $k \in \{1, \ldots, K\}$, has pdf

$$f_{X|Z_t=k}(x) \triangleq f_X^{(k)}(x) = p^{(k)}\, \delta(x) + (1 - p^{(k)})\, g^{(k)}(x)\,, \tag{1}$$

where $\delta(x)$ represents a probability mass (or Dirac measure) at $x = 0$, $g^{(k)}(x)$ represents the continuous distribution of the non-zero price revisions, and $p^{(k)}$ represents the probability of observing a

regime	A		λ	p	σ
1	0.80	0.20	1.37	0.56	2.9×10^{-4}
2	0.43	0.57	0.14	0.14	6.3×10^{-4}

Table 1: Parameters used to generate the sample price path in Figure 2. These parameters were estimated from the PCP Feb 2008 data set assuming a two-regime model.

trade with zero price innovation. In principle, conditional on a non-zero price revision, any reasonable distribution could be used to model the price innovations, for example: Gaussian, student-t, double exponential, etc. Moreover, in this framework there is ample flexibility to choose how to model durations within a given regime, for example using a hyper-exponential, Coxian class, or more generally, using phase-type distributions which uniquely describe the state-dependent hazard rate $\lambda_t = \lambda(Z_t)$. Moreover, it is also possible to introduce co-dependence between the duration and price revision within a given regime through a copula. However, we have found that having independence of duration and price revision within a fixed regime aptly captures the stylized features of the data. Figure 1 shows how the intra-day-states evolve according to the discrete-time Markov chain with transition matrix A, and where upon a trade occurring in regime i it enters regime j with probability A_{ij}.

Now, equipped with the Markov chain Z_t, the regime contingent rate of arrival function $\lambda^{(k)}$ and the regime contingent price revision distribution $F_X^{(k)}(x) = \int_{-\infty}^{x} f_X^{(k)}(z)dz$ with $k \in \{1, \dots, K\}$, we model the tick-by-tick price process of the asset as a marked point process as follows:

$$S_t = S_0 \exp \left\{ \sum_{n=1}^{N_t} \varepsilon_n^{(Z_{t_n-})} \right\}, \qquad (2)$$

where $\left\{ \varepsilon_1^{(k)}, \varepsilon_2^{(k)}, \dots \right\}$ are i.i.d. random variables with distribution function $F_X^{(k)}(x)$, and where $\{t_1, t_2, \dots\}$ are the arrival times of the trades and $N_t = \sup\{n : t_n < t\}$ is the counting process corresponding to trade arrivals.

For simplicity, we assume that the non-zero price revisions are Gaussian, that is $g^{(k)}(x) = \phi\left(x; \sigma^{(k)}\right)$ where $\phi(x; \sigma)$ denotes the pdf of a Gaussian random variable with zero mean and standard deviation σ, and that the state-dependent hazard function $\lambda_t = \lambda(Z_t)$ is a constant which implies that within the regimes the waiting times are exponentially distributed. We remark that our HMM is able to capture the long and short durations exhibited by financial data because the chain meanders through the different regimes according to the transition matrix A, we return to this point below.

In Figure 2, we use equation (2) to simulate a high-frequency sample path of stock prices using a two-state HMM with parameters given in Table 1 which have been estimated from PCP February 2008 data. Notice that in regime 1 (depicted by blue 'x's) durations are fairly short and the price innovations tend to be small; moreover, the chain persists in this regime for some time. Once the chain migrates to regime 2 (depicted by green circles), durations are longer and the price innovations have larger variance; however, the chain eventually switches back to regime 1 at a faster rate than the rate at which it originally switched into regime 2 with. This simple example shows some of the characteristics of prices on a tick-by-tick level. There are times when the market experiences bursts of activity with volatility clustering (e.g., between the 1.396 and 1.398 mark in the time axis) – i.e., many trades over short periods of time followed by relatively high volatility; and periods of very little activity and low volatility (e.g., around the 1.408 mark in the time axis) – which could be interpreted as no news arriving in the market.
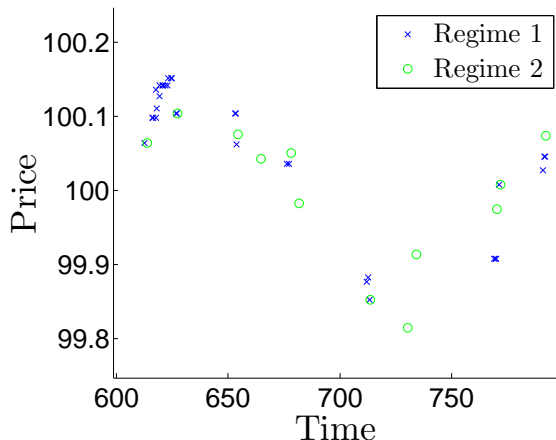
Figure 2: A sample price path generated by our model together with the state of the hidden Markov chain. The large and small circles indicate trades that occurred while the Markov chain was in regime 1 and 2 respectively. The model parameters used to generate these paths are recorded in Table 1 and were estimated using the PCP Feb 2008 data with 2 regimes.

3 Model Estimation & Data

In this section we describe our approach to estimating the parameters of our model and the data sets that we used.

3.1 The EM-algorithm

We employ the Baum-Welch EM algorithm for the HMM to estimate the transition probability matrix A, the within regime model parameters $\theta = \{\lambda, \, p, \, \sigma\}$, and the initial distribution of the regimes π, for details see Baum, Petrie, Soules, and Weiss (1970). The methodology amounts to maximizing the log-likelihood

$$\ln L = \sum_{t=1}^{n} \sum_{j=1}^{K} \ln f_{\theta_j}(\{(\tau_t, X_t)\}) \mathbb{I}(Z_t = j)$$
$$+ \sum_{t=1}^{n-1} \sum_{j=1}^{K} \sum_{k=1}^{K} \ln A_{jk} \mathbb{I}(Z_t = j, \, Z_{t+1} = k) + \sum_{j=1}^{K} \ln \pi_j \mathbb{I}(Z_1 = j)$$

of the sequence of observations $\{(\tau_t, X_t)_{t=1,\ldots,n}\}$. Here, $f_{\theta_j}(\{(\tau_t, X_t)\})$ denotes the joint probability density of the observation (τ_t, X_t) given that the chain is in state j with parameters θ_j. Since the durations between trades have been recorded to the nearest second, we adopt a censored version of the density and for our specific model write

$$f_{\theta_j}(\tau_t, X_t) = e^{-\lambda_j \tau_t}(1 - e^{-\lambda_j}) \times (p_j \mathbb{I}(X_t = 0) + (1 - p_j)\mathbb{I}(X_t \neq 0)\phi(X_t; \sigma_j)) \, , \qquad (3)$$

where $\mathbb{I}(\cdot)$ is the indicator function, X_t is the log-price innovation at time t and τ_t is the duration since the last trade. The initial starting parameters for the HMM learning were estimated assuming that the duration/price innovation pairs are independent and drawn from the related mixture model

$$f_{X,\tau}^{(0)} = \sum_{j=1}^{K} \alpha_j \, e^{-\lambda_j \tau_t}(1 - e^{-\lambda_j}) \times (p_j \mathbb{I}(X_t = 0) + (1 - p_j)\mathbb{I}(X_t \neq 0)\phi(X_t; \sigma_j)) \, .$$

The estimated mixture weights α_j were used to provide an initial estimate for the transition probability matrix A by assuming that only transitions between neighboring regimes can occur. The EM

7

algorithm was then run until a relative tolerance of 10^{-6} was achieved. A review of the Baum-Welch approach for fitting HMMs with the EM algorithm is provided in Appendix A together with the updating rule for our specific within regime model.

3.2 The Data

We used TAQ data for several mid-cap and large-cap stocks for the months of February 2001 and February 2008. Trade data during the normal trading hours between 9:30am and 4:00pm were analyzed. The data were cleaned by deleting entries with a non-zero Field Correction flag and entries with a Field Condition flag of Z. Furthermore, the data were filtered to remove any data points that were outside 15 standard deviations because we assume that these are errors in the tape. Unlike many previous works, we keep all other reported trades, and in particular do not throw away trades which reported a price revision of zero nor do we throw away trades which reported a duration of zero. Deleting such trades results in well over 30% reduction in the data and there are two important reasons why discarding these trades is undesirable. First, from an estimation point of view, deleting these trades destroys the auto-correlation structure of the data and consequently biases the estimation. From a financial point of view, trades with zero price revision or with zero duration convey key information that is valuable for certain types of strategies that AT and in particular HFTs employ regularly (we discuss such strategies in Section 5).

One of the reasons why in previous studies zero duration trades were deleted is because it was assumed that trades arrive at a rate where it is not (mathematically) possible to have two trades arrive at the same point in time. For instance, if trades arrive according to a Poisson process or any other counting process where the arrival rate is finite there can only be at most one trade over an infinitesimally small time-step. In or model we are able to keep these trades for two reasons: (i) the model for price revisions is a mixture model, in which zero price revisions are captured separately from non-zero price revisions (ii) we use censoring to account for the fact that data are reported only to the nearest smallest second which allows us to effortlessly include zero waits. In Table 2, we report some relevant statistics concerning data deletion for each data set.

Markets tend to be more active during the morning and afternoon than in the middle of the day. Thus, one expects that durations are shorter around the hours when the market opens and closes, and longer around midday. Depending on the goal of the model for stock dynamics one option is to diurnally adjust durations to account for this intra-day seasonal pattern, eg. Engle (2000), or to employ the duration data without adjustments, eg. Cartea and Meyer-Brandis (2010). The results we obtain are qualitatively the same whether we estimate the HMM using diurnally adjusted durations or do not make any adjustments for intraday seasonality. In what follows we show the results when no adjustments are made because in the two examples we discuss in reference to HF trading and AT, the HMM parameters must be estimated online and it seems more plausible to assume that the duration data are not adjusted as it is processed in real time.

3.3 Picking the Number of States

Since we are utilizing an HMM, one key step is to estimate the number of hidden regimes. One often used performance measure is the Bayesian Information criterion (BIC). That is,

$$BIC = \ln L^* - \frac{\nu_K}{2} \ln n \ ,$$

where $\nu_K = 4K + K * (K - 1)$ is the number of model parameters for a model with K regimes, n is the number of observations, and L^* is the maximum log-likelihood (in this context, since we are using the EM algorithm, it is our best estimate of the maximum log-likelihood, see Appendix A for more details). Another often used performance measure is the Integrated Completed Likelihood

Symbol	FEB 01				FEB 08			
	Raw Data	Correc	Std Dev	Data	Raw Data	Correc	Std Dev	Data
AA	35,137	2,623	0	32,514	979,211	16	165	979,030
AMZN	163,400	229	2	163,169	1,144,832	39	445	1,144,348
HNZ	14,786	29	0	14,757	232,983	1	33	232,949
IBM	98,311	343	26	97,942	805,380	609	344	805,380
KO	41,877	130	3	41,744	777,876	26	231	777,619
PCP	5,149	4	0	5,145	197,784	7	67	197,710

Table 2: This table summarizes how data were cleaned. Column 'Raw Data' shows all the trades reported on the TAQ database; column 'Correc' are trades that were deleted because the Field Correction was different from 0 and the Field Condition was equal to Z; column 'Std Dev' shows the total number of log-returns outside 15 standard deviations that were deleted; and column 'Data' shows the number of trades that we use in the empirical analysis.

year	criteria	AA	AMZN	HNZ	IBM	KO	PCP
2001	BIC	4	5	3	5	4	2
	ICL	4	3	2	3	3	1
2008	BIC	6	7	6	7	6	7
	ICL	3	2	2	2	3	2

Table 3: The preferred number of regimes using the BIC and ICL criteria based on estimation of all data sets.

(ICL). Biernacki, Celeux, and Govaert (2001) propose to use a BIC-like approximation of the ICL leading to the criterion

$$ICL = \sum_{t=1}^{n} \ln f_{\theta_{\widehat{Z}_t}} (\tau_t, X_t) - \frac{\nu_K}{2} \ln n \, ,$$

where the sequence of missing states are replaced by the most probable value $\widehat{Z}_t$ based on the estimated parameters (as computed for example from the Viterbi (1967) algorithm). The optimal number of states is the one which maximizes the criterion. However, as described in Celeux and Durand (2008), the BIC criterion tends to overestimate the number of hidden states while the ICL criterion tends to underestimate the number of hidden states.

In our implementation for assessing the number of states we use the following cross validation approach

1. The parameters for a fixed number of regimes were estimated using all but one single day's data – this provided 19 (for 2001) or 20 (for 2008) parameter estimates.

2. The performance criterion (both BIC and ICL) were computed for the missing day's data only – providing 19 (for 2001) or 20 (for 2008) measures of BIC and ICL.

3. These measures were then averaged and the process repeated from step 1 with an increased number of regimes.

Table 3 shows the results of this estimation procedure. For the 2001 data, the average number of regimes is 3 while in 2008 the average number of regimes is 4. In the remainder of the article we use 4 regimes in our HMM.

Below in Section 4 we present and interpret the parameter estimates of the HMM for each stock we study. But before proceeding we discuss how the HMM is able to capture the empirical distribution of the waiting times. When looking at data that involve the random arrival of trades
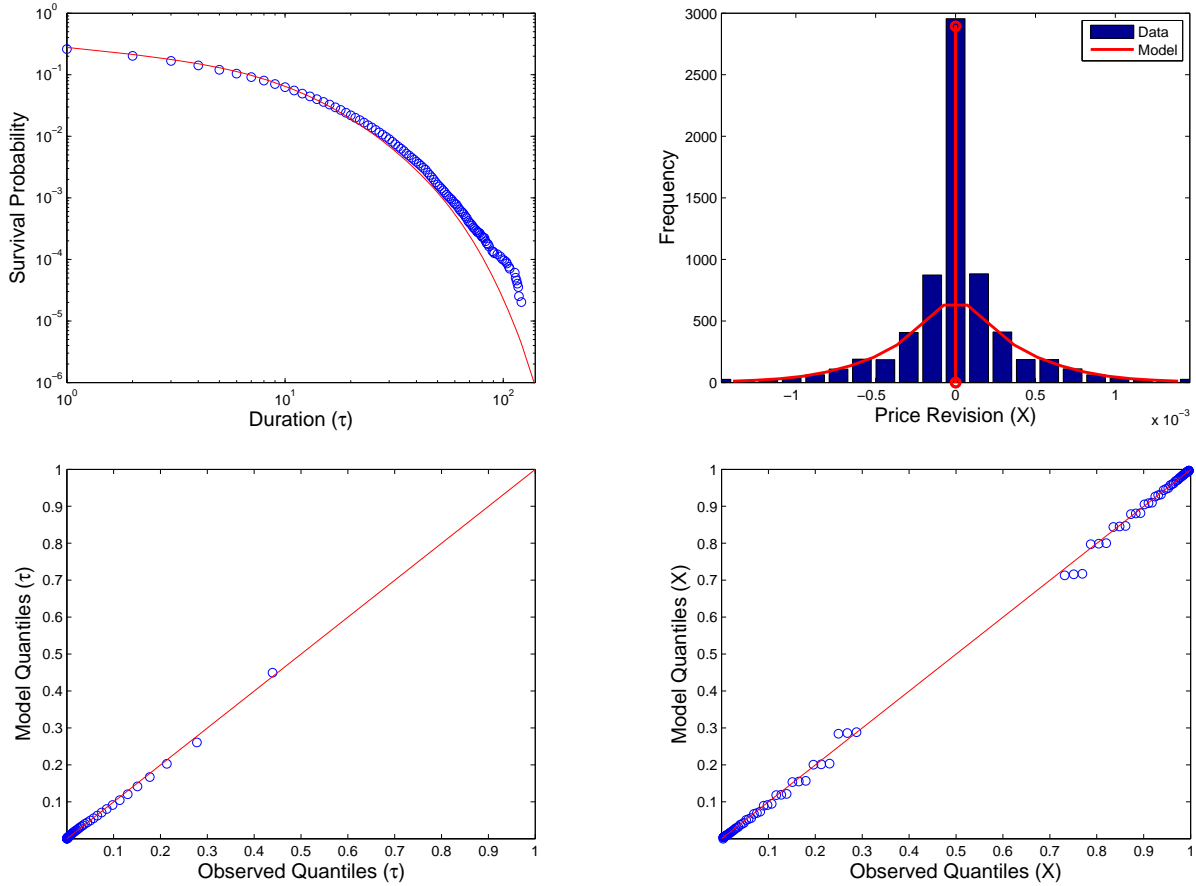
Figure 3: The model fit to the empirical distribution of duration and price revision based on four regimes for Feb 2008. The estimated model parameters are provided in Table 4.

it is customary to look at the survival function, which represents the probability that the waiting-time between two consecutive trades is greater than t. One of the empirical features of durations in tick-by-tick data is that the unconditional survival function is not exponential. The common assumption that durations are exponentially distributed fails because the tail of the exponential distribution decays too fast and in the market we frequently observe long durations, see Cartea and Meyer-Brandis (2010). In our HMM model we have assumed that within the intra-day state the waiting time distribution is exponential, but the transit from one state to another state (with state dependent parameters) allows us to capture the unconditional survival function extremely well. As an example, in Figure 3 we show the empirical fit to the PCP data for both the trade duration and the price revisions – which illustrate the model's goodness of fit.

4   Discussion of results

The estimated parameters for the HMM with 4 regimes for the PCP dataset are reported in Table 4 – the remaining results for 6 other stocks are reported in the same format in Appendix D. The standard errors, computed through a bootstrap procedure,[3] are reported in the braces below each parameter. In Table 4 we organized the intra-day regimes starting with the fastest by trade arrival

---

[3]The bootstrap was performed by simulating data from the estimated model. The simulated data had the same number of segments (days) as the original data, and the same number of trades on each day as the original data. Given the simulated data, the model was then re-estimated, and this procedure repeated ten times. The sample 95% confidence intervals (based on student-t with 9-degrees of freedom) are reported.

(or equivalently with the shortest durations) which is given by the highest estimate of the within regime hazard function λ. The last three columns of the table provide information about the distribution of price innovations. Column p denotes the probability that the trade arriving within that state occurs at the same price as the previous trade; column σ ($\times 10^{-4}$) contains the volatility of the price revision conditioned on the price innovation being different from zero; and column $\sigma\sqrt{1-p}$ ($\times 10^{-4}$) provides the within regime unconditional volatility of the price revision.

PCP – FEB 2001

Regime	Transition Probability Matrix A				Conditional Parameters			
	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	22.89%	33.61%	24.38%	19.13%	0.232	99.97%	5.917	0.099
	(5.72%)	(14.00%)	(11.94%)	(6.80%)	(0.048)	(4.57%)	(15.261)	
2	15.63%	69.40%	1.85%	13.13%	0.019	18.66%	21.817	19.677
	(4.72%)	(7.16%)	(5.77%)	(3.66%)	(0.003)	(6.20%)	(1.385)	
3	3.72%	3.40%	87.67%	5.22%	0.015	44.64%	8.529	6.346
	(1.92%)	(3.50%)	(5.49%)	(2.10%)	(0.001)	(2.51%)	(0.758)	
4	6.70%	9.11%	13.77%	70.43%	0.006	23.84%	12.534	10.938
	(4.89%)	(5.57%)	(8.56%)	(7.10%)	(0.001)	(3.56%)	(1.131)	

PCP – FEB 2008

Regime	Transition Probability Matrix A				Conditional Parameters			
	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	66.22%	16.17%	0.62%	16.99%	1.803	91.32%	1.355	0.399
	(0.81%)	(0.95%)	(0.22%)	(0.41%)	(0.045)	(0.41%)	(0.079)	
2	6.29%	75.27%	2.16%	16.28%	1.112	29.23%	3.004	2.527
	(0.86%)	(0.74%)	(0.24%)	(0.55%)	(0.009)	(1.06%)	(0.031)	
3	3.70%	8.11%	81.72%	6.48%	0.635	21.68%	11.559	10.230
	(1.18%)	(1.59%)	(1.14%)	(0.69%)	(0.011)	(1.11%)	(0.219)	
4	24.25%	18.51%	0.38%	56.86%	0.123	16.87%	5.121	4.669
	(0.54%)	(0.68%)	(0.19%)	(0.80%)	(0.002)	(0.34%)	(0.021)	

Table 4: Estimated 4-regime model parameters on PCP data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.

Tables 6 to 10 in the appendix show the parameter estimates for the other stocks we study. We find that across all stocks in February 2008: the regime where trading occurs at the highest (lowest) activity is regime 1 (regime 4); the lowest volatility of price revisions (last column of tables) is in regime 1; the highest probability of observing a zero price revision is in regime 1; and the least persistent is regime 4. In most cases we find that the lowest probability of seeing a zero price revision is in regime 4. The results for the same stocks in February 2001 are less clear cut in terms of visible patterns across different stocks. The intra-day states with lowest durations are not necessarily the ones with lowest volatility of price revisions; in half the cases the most persistent state is regime 3; in most cases the state with highest (lowest) probability of observing zero price revisions is regime 1 (regime 4); and there is no one state which is the least or most persistent.

In Table 5 we show the total number of trades for each stock and the proportion of trades[4] that took place in every intra-day state during February 2001 and February 2008. As expected, the number of trades for each stock increased considerably between the two dates, implying that the

---

[4]Calculated using the most likely path of the Markov chain, see Viterbi (1967).

overall trading pace has also increased and average durations decreased. This increase in pace is also observed at the intra-day regime level, where we see that for all stocks durations have become shorter – i.e., the hazard rate λ for every state increases from 2001 to 2008.

From the tables that report the HMM parameter estimates and from Table 5 we also observe that in Feb 2001 regime 1 is both the fastest and the least visited across all stocks. Contrastingly, in Feb 2008 it is the slowest regimes where intra-day trading spends the least amount of time, with the exception of IBM where approximately 63% of trades occurred in regimes 2 and 3. Furthermore, if we look at all stocks combined, in 2001 less than 10% of trades occurred in the fastest state and less than 25% in the second fastest state, whereas in 2008 more than 30% of trades occurred in the fastest state and more than 36% in the second fastest state.

| Regime | FEB 2001 | | | | | FEB 2008 | | | | |
	$Z=1$	$Z=2$	$Z=3$	$Z=4$	Total trades	$Z=1$	$Z=2$	$Z=3$	$Z=4$	Total trades
AA	0.007	0.038	0.564	0.391	32,514	0.315	0.418	0.166	0.101	979,010
AMZN	0.176	0.463	0.218	0.143	163,150	0.241	0.672	0.018	0.069	1,144,327
HNZ	0.011	0.033	0.697	0.260	14,738	0.374	0.135	0.274	0.218	232,930
IBM	0.032	0.068	0.552	0.347	97,923	0.222	0.056	0.627	0.094	804,427
KO	0.035	0.037	0.716	0.212	41,725	0.461	0.275	0.039	0.225	777,600
PCP	0.089	0.198	0.595	0.118	5,126	0.337	0.366	0.053	0.245	197,691

Table 5: Proportion of trades per intra-day state.

Undoubtedly, the recent increase in volume of trades in equity markets is mainly due to AT. In our sample of data we see that the number of trades between 9.30am and 4.00pm for all stocks has seen an explosion in the last years. For instance, Table 5 shows that trading volume for KO increased from 41,725 trades in February 2001 to 777,600 in the same month of 2008. Other qualitative changes that we observe in the data, which are most certainly a consequence of AT, are: i) From 2001 to 2008 we observe that for most stocks the intra-day states have become less persistent.[5] ii) In Table 5 we see that the fastest regime(that with the shortest average durations) in 2008 is also an intra-day state where a great deal of trades take place which contrasts with the 2001 results where the fastest regime was where the least amount of trades took place. One plausible explanation is that competition among different superfast computer-based algorithmic traders (which include HF trading) is very active in regime 1. This also confirms the theoretical predictions of Cvitanić and Kirilenko (2010) who show that the introduction of HFTs increases trading activity (by reducing the waiting time between trades) and modifies the distribution of price revisions by increasing mass around the center and thinning the tails.

We can also view our results in light of the microstructure literature. This literature has mixed results concerning the link between durations and volatility. One of the conclusions in the early work of Diamond and Verrechia (1987) is that long durations should be positively correlated with price volatility. Admati and Pfleiderer (1988) also conclude that slow trading means high volatility. This is confirmed by the empirical results of Dufour and Engle (2000) who find that short durations and thus fast trading follow large returns and large trades; and those of Manganelli (2005) who finds that for frequently traded stocks short durations increase the price variance of the next trade. On the other hand, Easley and O'Hara (1992) find that periods of low variance tend to occur in periods where there is little trading, i.e. low variance is linked to long durations. This is empirically verified by Engle (2000) who finds evidence that longer durations and longer expected durations are associated with lower volatilities.

---

[5]From the six stocks, we see that only 6 out of 24 (recall that there are four states per stock) regimes became less persistent in 2008.

Our empirical findings clearly indicate that for the 2008 data set the regime where trading is most active is always the one where the volatility of price revisions is lowest. In this sense our findings confirm the theoretical predictions of Diamond and Verrechia (1987) and Admati and Pfleiderer (1988) and the empirical findings in Dufour and Engle (2000) and (for frequently trade stocks) Manganelli (2005). The slowest regimes on the other hand, are not necessarily the ones with highest volatility of price revisions.

5 What the states say about potential Algorithmic and High Frequency Trades

One of the key aspects of AT is how the arrival of information is processed in order to make trading decisions. Information are marks associated to the trade and quote flow (prices, duration, volume, seller initiated trade, buyer initiated trade, etc.) as well as other pieces of news (announcement of firm specific information; and macroeconomic variables such as unemployment, growth, etc.) that are released to the market and trading activity reacts until this new information is impounded in stock prices. Therefore, if the objective is to design trading algorithms, one of the challenges is how can these algorithms incorporate this information as soon as it arrives. The HMM we propose here has the advantage that the model parameters and the states can be estimated simultaneously and "online" (see e.g. Mongillo and Deneve (2008)). Consequently, trading algorithms can use all of this information and in particular 'know' the intra-day state of the market as well as the parameters relating to price revisions, duration, and probability of migrating to another state. Below we discuss two trading strategies that can be implemented based on the HMM.[6]

5.1 High frequency trading for liquidity rebates

Within AT there are activities that are carried out by what is known in the market as HFTs. These traders are different from the rest because they: i) submit a vast number of orders over short time intervals and, more importantly, a large number of these orders are canceled immediately if they are not executed in a split second. For example, February 5 2008 is a typical day for AA in Nasdaq where 96% of all orders were cancelled. More interestingly, 12% of all orders were cancelled within 100 milliseconds of being sent, 25% were cancelled within 500 milliseconds, and 33% within 1 second. ii) they aim at being flat, that is to hold no inventories, ideally within the day or at most at the end of the day, see Cvitanić and Kirilenko (2010). HFTs' inventories quickly mean revert to zero throughout the day because of the time scale over which the HF strategies are designed to profit from buying and selling assets. HFTs use their superior speed to process information and act ahead of other slower traders. Admittedly, there are a great deal of HF strategies and all we know is that their success depends on being able to profit from roundtrip trades. Therefore, because HFTs' competitive edge is speed, their strategies seek opportunities to enter and exit the market very quickly (milliseconds, seconds, or minutes) and, as a result, holding periods are extremely short, see Cartea and Jaimungal (2012). Furthermore, HFTs aim at ending the day with no inventories to avoid having to post collateral overnight and to avoid the risk of adverse price movements when trading resumes the following day.

HFTs deploy different strategies depending on market conditions and depending on what the aim of the set of trades is. For instance, HFTs may trade with the sole purpose of making what is known as 'liquidity rebates'. Some exchanges incentivise liquidity provision by paying a rebate of up to 0.3 cents per share. Exchanges typically charge a somewhat higher access fee than the amount of their liquidity rebates but these access fees are paid by those who hit a bid or lift the

---

[6]See Bouchard, Dang, and Lehalle (2011) for a framework that discusses when and how to use different trading algorithms.

offer posted by the liquidity provider because they are aggressive order types, i.e. they are liquidity takers. Sometimes, however, exchanges have offered "inverted" pricing and pay a liquidity rebate that exceeds the access fee, see SEC (2010).

To illustrate exactly how an HFT may take advantage of rebates, consider the following example of a rebate trade: Assume that the exchange offers 0.25 cents per share to dealers who post orders. If this particular order is filled, the liquidity provider takes the 0.25 cents rebate and the trader that lifted the offer or hit the bid pays the access fee. One of the many ways in which the HFT spots a rebate opportunity is to 'observe' that a big buy order that has been broken up in small batches is being put through the market by an algorithmic trader. The current price is $10.00 per share and the HFT uses her speed advantage and sends out a buy order for $10.01 per share. This posting is considered as providing liquidity because it ups the price by one cent and sits there until it is hit by another party (presumably those that were initially selling at $10.00 to the AT). After the HFT's buy order is filled, she immediately turns around and posts an order to sell them for $10.01 per share (again the HFT is providing liquidity) which is lifted by the algorithmic trader who is still liquidating his position. This round trip trade generates 0.5 cents profit per share as a result of the rebates despite the fact that the HFT makes zero profit on the shares themselves.[7]

In the set of rebate trades discussed above the HFT had to up the buy price by one cent to be treated as a liquidity provider by the exchange. Had the HFT got ahead of the AT and bought shares at $10.00 she would have been seen as a liquidity taker (aggressive order) and would have incurred an access fee. Even if she made the rebate on the second leg of the trade by selling at $10.00 per share the one way rebate trip would have delivered a loss of 0.05 cents per share (assuming an access fee of 0.3 cents per share). However, if exchanges offer an inverted pricing scheme to 'attract' liquidity, then even in trades where only one leg of the round trip earns the rebate, the HFT post positive net profits.

Collecting rebates is not risk-free, since there are scenarios where the risk is adverse move in prices. However there are regimes in which the risk of these adverse moves are lower. The information provided by our HMM can be used to assess how likely a rebate trade, or set of rebate trades, is able to produce a positive profit.[8] Take for example AA and the information in Table 6. There we can see that in February 2008 there are regimes that look 'safer' than others to execute rebate trades. There are three aspects we must consider: first, how persistent the regime is; second, what is the probability that trades within that regime have a zero price revision; and third, if the price revision is not zero, what is the volatility of the change in prices. For example regime 1 appears to be an ideal regime for HFTs to profit from rebates alone on all three accounts. The persistence of regime 1 is the highest across all regimes (80.67%), the probability of observing zero price revisions is also the highest across all regimes (99.97%), and if there is a price change in regime 1, the volatility of the price innovation is the lowest across all regimes (3.010×10^{-4}), and volatility of a price revision (without distinguishing between zero and non-zero price revision) is 0.050×10^{-4}. Moreover, 308,840 trades took place within this state which is around 31.5% of the total trades during that month, showing that rebate opportunities are not a rare occurrence. Therefore, an HFT that finds herself in regime 1 for AA shares can engage in rebate trading with a very high probability of making profits while bearing very little risk.

5.2 Limit order algorithmic trading

Another form of AT involves submitting buy and sell limit-orders around the mid-price in hopes of posting profits from the bid-ask spread. We pose this problem in a similar manner to Avellaneda

---

[7]For example, BATS pays a $0.0029 rebate per share for adding displayed liquidity, see https://batstrading.com/FeeSchedule/. And the NYSE Arca also pays a rebate of 0.0030 per share and this figure could vary depending on different characteristics, for more information see www.nyse.com.

[8]This example is for illustrative purposes where for simplicity we assume that the model parameters and intra-day states were estimated online, but use the expost results in Table 6 as reference.

and Stoikov (2008); however, here we use a continuous-time mid-price model based on our HMM to accurately reflect the autocorrelation of durations as well as the co-dependence of duration and price revisions.[9] Although the discrete HMM performs extremely well for empirical analysis, it poses mathematical difficulties when solving the optimal control problem arising in this AT setup, hence we utilize a continuous-time model counterpart (in Appendix B we show how to map between the two models). To this end, we assume that the mid-price S_t is a regime switching Brownian motion:

$$dS_t = \sigma_t \, dW_t \, . \tag{4}$$

Here, the volatility parameter $\sigma_t = \sigma^{(H_t)}$ is indexed by the continuous-time finite-state Markov chain H_t (taking on values $\{1, \ldots, K\}$) with generator matrix B. The process H_t determines the volatility of the mid-price, resulting in a regime-dependent volatility, and is the continuous-time counterpart of the discrete time Markov chain Z_t introduced earlier.

In this framework, the goal of the HF investor is to submit bid and ask limit orders (which are canceled shortly if not filled) at $(S_t - \delta_t^-)$ and $(S_t + \delta_t^+)$ respectively, so as to maximize her expected utility of terminal wealth at the end of the day (or, e.g., mid-day or hour which is a normal investment horizon for HFTs in one set of trades). We assume that the HFT is sufficiently small not to affect other market makers' strategies when sending limit orders to the book.[10] The investor has control over $\delta_\mp$, which represent the distance from the mid-price of the bid/ask orders. To achieve this goal, it is important for us to model the rate at which the orders are executed; consequently, we assume that if orders are placed at the mid-price, then the order is executed at a rate $\lambda_t = \lambda^{(H_t)}$. This rate of execution depends on the regime of the market and is the direct analog of the rate of arrival of trades in our discrete time HMM. However, as is well known, when orders are placed deeper into the limit-order queue (i.e. further away from the mid-price), the order is filled at a decreased rate. To account for this effect, we assume that the buy/sell limit-orders get filled at the rate $\Lambda_t^\mp = \lambda_t \, e^{-\kappa_{\mp,t} \delta_t^\mp}$ where, $\kappa_{\mp,t} = \kappa_\mp^{(H_t)}$ is a within-regime constant and is related to the shape of the limit-order book in the observed state H_t. In regimes when trades occur quickly, our earlier results imply that the volatility of trades is low and we expect that the limit-order book is concentrated near the mid-price; moreover, we expect this regime to have a small bid-ask spread. Therefore, in such regimes we expect that κ is large – because orders placed far from the mid-price are less likely to be filled. On the other hand, in regimes when trades occur slowly, our earlier results imply that the volatility of trades is high and we expect that the limit-order book is flatter – i.e. that as quotes move away from the mid-price the volume bid or offered does not change much; further, we expect this regime to have a larger bid-ask spread. Consequently, in such regimes we expect that κ is small – because orders placed deeper into the limit-order book are more likely to be executed in this regime.

The only parameter which does not have a counterpart in our discrete time HMM are the decay rates $\kappa_\mp^{(H_t)}$ which can in principle be estimated from level-II data[11] and is left for future work. An example of the form of this execution rate is show in Figure 4.

Having the same underlining Markov chain H_t drive both the volatility of the mid-price and the rate at which trades are executed allows us to capture the co-dependence between durations and price innovations just as in the discrete model. Furthermore, as can be seen from any of the

---

[9]For recent work on optimal market making see Guilbaud and Pham (2011); and for optimal liquidation and trade execution in limit order books see Alfonsi and Schied (2010), Kharroubi and Pham (2010) and Bayraktar and Ludkovski (2011).

[10]It is argued that HF market makers are constantly updating their quotes every time there is a change in the LOB. Thus, order flow information is key to the behavior of HF market participants and this explains in part the vast number of orders and cancellations that we see during trading hours.

[11]Level-II data contains the status of the entire limit-order book showing all current bid/sell offers and the number of shares being offered at that these price levels. This is in contrast to level-I data which contains only best-bid and best-ask. The shape of the limit-order book is directly related to the probability that a specific limit-order is executed and can be used to infer the decay factors $\kappa^{(k)}$ in our model.
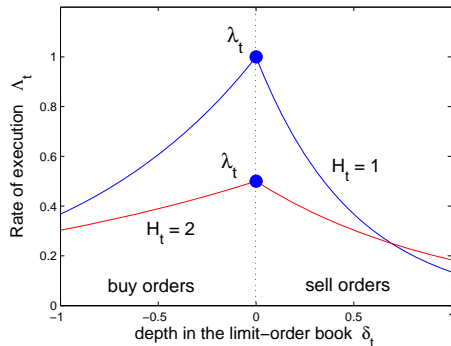
Figure 4: A sample plot of the rate at which limit buy/sell orders are executed as function of the distance to the mid-price. The dependence on the regime is also shown for a two-regime model. The second regime has slower rate of execution and a flatter limit-order book than the first regime.

calibrated parameters in the discrete model, the rate at which trades arrive is much larger than the rate at which the chain leaves a given state. This is an important point because one of the crucial elements in AT and HF trading in particular is to avoid having stale quotes in the book. In our model a quote becomes stale if the market migrates to another intra-day state or if a trade takes place. In states where the probability of migration is low relative to the arrival rate of the trade, coupled with the ability of submitting immediate-execution-or-cancel orders, makes it very unlikely for the AT to be filled right after the market has changed to another state or a trade takes place. The key dangers are both a change in the arrival rate of trades and the volatility of price revisions which are determinant variables for picking the optimal spread when submitting buy and sell orders to the book. Below we show how the optimal spread is chosen by the AT and how it depends on the volatility of prices and durations between trades.

To formalize the investor's problem we need to introduce some more notation. Let N_t^- and N_t^+ denote the counting processes for the executed buy and sell limit-orders (recall that buy/sell orders are executed at the rate $\Lambda_t^\mp$). Further, let $q_t = N_t^- - N_t^+$ denote the total inventory of the investor. Upon a buy/sell order being filled, the investor pays $(S_t - \delta_t^-)$ and gains $(S_t + \delta_t^+)$ respectively. Consequently, the investor's wealth X_t upon executing this strategy satisfies the SDE

$$dX_t = (S_t + \delta_t^+)\, dN_t^+ - (S_t - \delta_t^-)\, dN_t^- \;, \tag{5}$$

and the investor seeks the strategy $(\delta_s^\pm)_{t \leq s \leq T}$ which maximizes the expected utility of terminal wealth (e.g. for a HFT this would be at end of day, or end of hour). The investor's regime-dependent value function $V^{(k)}(t, x, S, q)$ is finally defined as

$$V^{(k)}(t, x, S, q) = \sup_{(\delta_u^+, \delta_u^-)_{t \leq u \leq T}} \mathbb{E}\,[\, u(X_T + q_T\, S_T)\,] \tag{6}$$

with exponential utility $u(x) = \frac{1}{\gamma}\, (1 - e^{-\gamma x})$.

Here γ is the risk-aversion parameter and we assume that algorithmic and HFTs executing these limit-order strategies are large enough to be considered as near risk-neutral investors with $\gamma \ll 1$. In this case, utility $u(x) \sim x - \frac{1}{2}\gamma x^2$, so that an HF investor who seeks to maximize (6) is essentially maximizing expected return while penalizing risk. As we discuss below, the optimal strategy induces a mean reversion toward zero in the inventories q_t, which is precisely one of the most revealing features of HFTs.

Proposition 1 *The optimal strategy for the HFTs with state dependent value function* (6) *is given*

16

by

$$\delta_t^{\pm} = \frac{1}{\kappa_{\pm,t}^{(H_t)}} + \gamma \left[-\frac{1}{2\,(\kappa_{\pm,t}^{(H_t)})^2} \mp (q_t \mp \tfrac{1}{2})\,b^{(H_t)}(t) \right] + o(\gamma)\,, \tag{7}$$

where the regime dependent function $b^{(k)}(t)$ *is*

$$\begin{pmatrix} b^{(1)}(t,T) \\ \vdots \\ b^{(K)}(t,T) \end{pmatrix} = V^{-1} diag \left((T-t),\ \frac{e^{d_2(T-t)} - 1}{d_2},\ \ldots,\ \frac{e^{d_K(T-t)} - 1}{d_K} \right) V \begin{pmatrix} (\sigma^{(1)})^2 \\ \vdots \\ (\sigma^{(K)})^2 \end{pmatrix}. \tag{8}$$

Here, $d_2, \ldots, d_K$ *are the non-zero eigenvalues*[12] *of the transition rate matrix* B *and* V *is the matrix of the eigenvectors.*

For a proof see Appendix C.

By inspecting (8) we see that $b^{(k)}(t,T) \geq 0$. This function plays a key role in setting the distance between the two limit-orders. An important point is that it is an increasing function of the volatility of the price revisions – therefore the higher the volatility of the price revision the wider the spread the investor posts. Further, as the transition rates between regimes increases, the non-zero eigenvalues become more negative implying that the function b approaches zero faster and the posted spreads are tighter.[13] Moreover, it is interesting to see that as the terminal time T approaches, the function $b^{(k)}(t,T)$ approaches zero implying that the optimal policy requires posting limit-orders with tighter spreads. This is once again a consequence of the investor's risk aversion which induces her to have a zero terminal inventory. Placing postings with tighter spreads increases the probability of being filled and increases the speed at which inventories revert to zero.

There are other interesting features of the bidding strategy in (7). Firstly, if the HF investor has no inventory and $\kappa_{+,t}^{(k)} = \kappa_{-,t}^{(k)}$, then the limit-orders are placed symmetrically around the mid-price. As the investor accumulates a long position, the investor's bid-price moves away from the mid-price and their ask-price moves in towards it – inducing the investor to sell assets. Contrastingly, as the investor accumulates a short position, the investor's ask price moves away from the mid-price and their bid-price moves in towards it – inducing the investor to buy assets. Therefore we see that the optimal strategy induces the HF investor's inventory q_t to mean revert towards zero.

Secondly, if the intra-day state of the market changes, the volatility of the price revisions will also change. If in the new state the volatility is higher (lower) the investor's bid-ask is adjusted via two channels: a larger (smaller) $b^{(k)}(t,T)$ and a smaller (larger) $\kappa_{\pm,t}^{(H_t)}$ both of which increase (decrease) the spread posted by the investor. As discussed above, the function $b^{(k)}(t,T)$ is responsible for adjusting the spread of the postings (from the mid-price) taking into account how much longer the investor has left before winding up her strategy, and the parameter $\kappa_{\pm,t}^{(H_t)}$ captures how likely a posting deep in the book is to be filled. On this last point, the intuition is that when volatility is high (low) it is more (less) likely to see trades occurring further (closer) away from the mid-price S_t, hence the optimal strategy is to post wider (tighter) spreads as a result of a smaller (larger) $\kappa_{\pm,t}^{(H_t)}$.

Finally, all else equal, as time approaches the investment horizon T, the investor submits buy and sell limit-orders which are tighter around the mid-price; a strategy that stresses the fact that the HF investor aims at holding zero inventories at time T.

---

[12]Recall that the generator matrix of an irreducible Markov chain must have a single zero eigenvalue, while the remaining eigenvalues have strictly negative real part. See e.g. Corollary 4.9, pg. 55 in Asmussen (2003).

[13]This point results from realizing that higher transition rates induces the Markov chain to reach its invariant distribution more quickly. Consequently, the system behaves more like a single regime model with a volatility equal to the (invariant weighted) average of regime specific volatilities.

5.2.1 Performance of strategies: informed and uninformed market making

We demonstrate some features of the market making strategy developed here by performing a simulation experiment in which sample paths of the mid-price for PCP are generated and HFTs make markets to profit from roundtrip trades. To simulate the mid-price of PCP we use a two-regime model where regime 1 is the fast regime (short waiting times between trades) with low volatility of price revisions and regime 2 is the slow regime (long waiting times between trades) with high volatility of price revisions. The model is calibrated to the discrete HMM in Table 1 which contains the parameters for the PCP Feb 2008 data set.[14]

To test the performance of the strategies we assume that there are two HFTs who use the same strategy to make markets at high frequencies, equation (7), but one HFT is better informed than the other. The better informed HFT knows that PCP trades in two regimes and she is able to correctly estimate the model parameters, whereas the other HFT is less informed because he assumes that there is only one regime in the market. We simulate 5,000 mid-price paths, both HFTs submit limit buy/sell orders which are cancelled an instant later if not filled, the trading horizon is one hour, and for every simulation we record the PnL of the strategy.

One sample path of this experiment is shown in the top panel of Figure 5. The picture shows the postings of both traders and the mid-price. We depict the mid-price with circles when PCP is in regime 1 (the fast regime with low volatility) and with rhomboids when PCP is in regime 2 (slow regime with high volatility). The dashed lines above and below the mid-price show the postings of the informed trader, and the solid lines above and below the mid-price show the postings of the less informed trader. By looking at the postings of the informed trader we notice that in regime 1 orders are placed closer to the mid-price because the HFT knows that PCP is in the fast regime with low volatility, while in regime 2 the spread is larger because the HFT knows that PCP is in the slow regime with high volatility of price revisions.

On the other hand, by looking at the postings of the less informed trader it is clear that he cannot differentiate which regime PCP is in so he is unable to adjust his posts in the same way that the informed trader does. Thus, this will affect the overall profitability of his market making activities, and, additionally, there will be many instances in which his limit orders will be taken advantage of by better informed market participants who adversely pick off his 'uninformed' limit orders – i.e. the less informed trader will be adversely selected.

Moreover, it is interesting to see how the optimal postings are adjusted every time the inventory changes. Let us focus on the postings of the informed HFT between seconds 50 and 55. During that five-second interval we see that two market buy orders were filled by the informed HFT's resting sell orders (the two stars on the sell side in that interval). Note that as soon as the HFT sells one share she immediately increases her sell half-spread and decreases her buy half-spread. This reflects the inventory management component of the strategy which is always exerting pressure on inventories so that they mean revert to zero. Finally, note that in business time, the chain spends most of its time in regime 1; however, in calendar time it spends most of its time in regime 2 – this is because the mean time to a trade in the slow regime is longer than in the fast regime.

In Figure 5 we also show the HFTs' PnLs resulting from the 5,000 simulations. We assume that the HFTs obtain zero rebates for providing liquidity and that their level of risk aversion is $\gamma = 1$. The left-hand picture of the bottom panel shows the distribution of the PnLs of both HFTs. The histogram in black shows the PnL distribution of the informed HF market maker (mean 13.30 and

---

[14]The calibrated transition rate matrix $B = \begin{pmatrix} -0.3193 & 0.3193 \\ 0.0980 & -0.980 \end{pmatrix}$ using the results in Appendix B, $\sigma_1 = 0.016\%$ and $\sigma_2 = 0.155\%$, and $\lambda_1 = 1.37$ and $\lambda_2 = 0.14$. We further assume $\kappa_+^{(1)} = \kappa_-^{(1)} = 100$ and $\kappa_+^{(2)} = \kappa_-^{(2)} = 50$ to reflect a flattening of the order book in regime 2 and that orders in the book of more than 5¢ from the mid-price occur with a probability of less than 0.1%. Further, an investment horizon of $T = 1$ hour is used and the investor is assumed to have a risk-aversion parameter of $\gamma = 1$. For the single regime case we use the invariant distribution of the Markov chain to compute the average vol $\sigma = 0.135\%$, intensity $\lambda = 0.44$ and fill probability parameter $\kappa = 62.33$.

standard deviation 0.61) and the histogram in grey shows the PnL distribution of the less informed HF market maker (mean 12.30 and standard deviation 0.59). The right-hand picture of the bottom panel shows the difference between the informed and less informed PnLs (mean 1.00, standard deviation 0.30 and the 5th percentile is 0.52). As expected, the less informed trader is less profitable because he trades on lesser quality information which precludes him from sending optimal orders to the LOB to profit from knowledge of PCP's market state and also exposes him to adverse selection costs.

To appreciate how the profitability of market making depends on the quality of information employed by the HFTs, Figure 6 shows the Sharpe ratio (left panel) and Risk/Return frontier (right panel) for $\gamma \in [0, 10]$.[15] As expected, for any level of risk aversion the Sharpe ratio of an informed strategy is always higher than that of a less informed strategy. In addition, it is interesting to note that for low values of the risk aversion coefficient γ the Sharpe ratio is increasing in γ, peaks at around $\gamma = 1$, and then is decreasing in γ. The right panel of the figure shows that for low levels of risk aversion there are clear gains from being a better informed market maker. And although it is always more profitable to be better informed, the risk/return frontier of the two HFTs become closer because risk aversion plays a key role in the optimal halfspreads.

6   Conclusions

We develop an HMM to understand the key behavior of stock dynamics at a tick-by-tick level. The HMM modulates different intra-day states of the high-frequency market dynamics and within every state we model price revisions and durations. As a whole, the model is able to capture the unconditional distribution of waiting times as well as the conditional (within regime) duration between trades and distribution (within regimes) of price revisions. An important feature of our model is that we are able to differentiate between trades with zero-price revision and trades that change prices relative to the previous observation. This distinction is important not only to correctly model the tick-by-tick dynamics of stock prices, but it is also crucial in the design of trading algorithms which these days are responsible for approximately 70% of the volume in US stocks.

Our approach allows us to discuss how the market has changed in recent years where the majority of trades are designed and executed by computer algorithms. Over the last decade the increasing presence of AT has changed not only the speed at which trades take place, but also other fundamental intra-day characteristics of stock price behavior have changed. We start by describing the characteristics that have changed only incrementally in the two periods, February 2001 and February 2008. i) For all but one asset, the states with shortest average durations is where we observe the highest probability of observing zero price innovations; and ii) The states with longest average durations are generally the ones where the probability of observing a zero price innovation is lowest. Some of the changes between the two periods are: i) Across all stocks we study in 2008, the intra-day states with shortest average durations are also the states with lowest volatility of price revisions. The same is not true for 2001 where there is no general connection between states of high activity and volatility. ii) For all stocks in 2001, the intra-day state with the shortest durations is also the state where the least amount of trades took place. On the other hand, in 2008 we find the opposite result where, generally, the intra-day states with the longest durations have the least number of trades.

Finally, we provide two concrete examples of how HF trading and AT strategies can be implemented based on the specific information derived from our model. The first example looks at rebate trading during February 2008 in AA stock. We discuss how given the large proportion of zero-price revisions (99.97%), and the low volatility of the non-zero-price revision of the remaining trades in

---

[15]The Sharpe ratio is calculated as $(\mu - r)/\sigma$ where μ is the mean PnL, σ is the standard deviation of the PnL, and the risk-free rate $r = 0$.
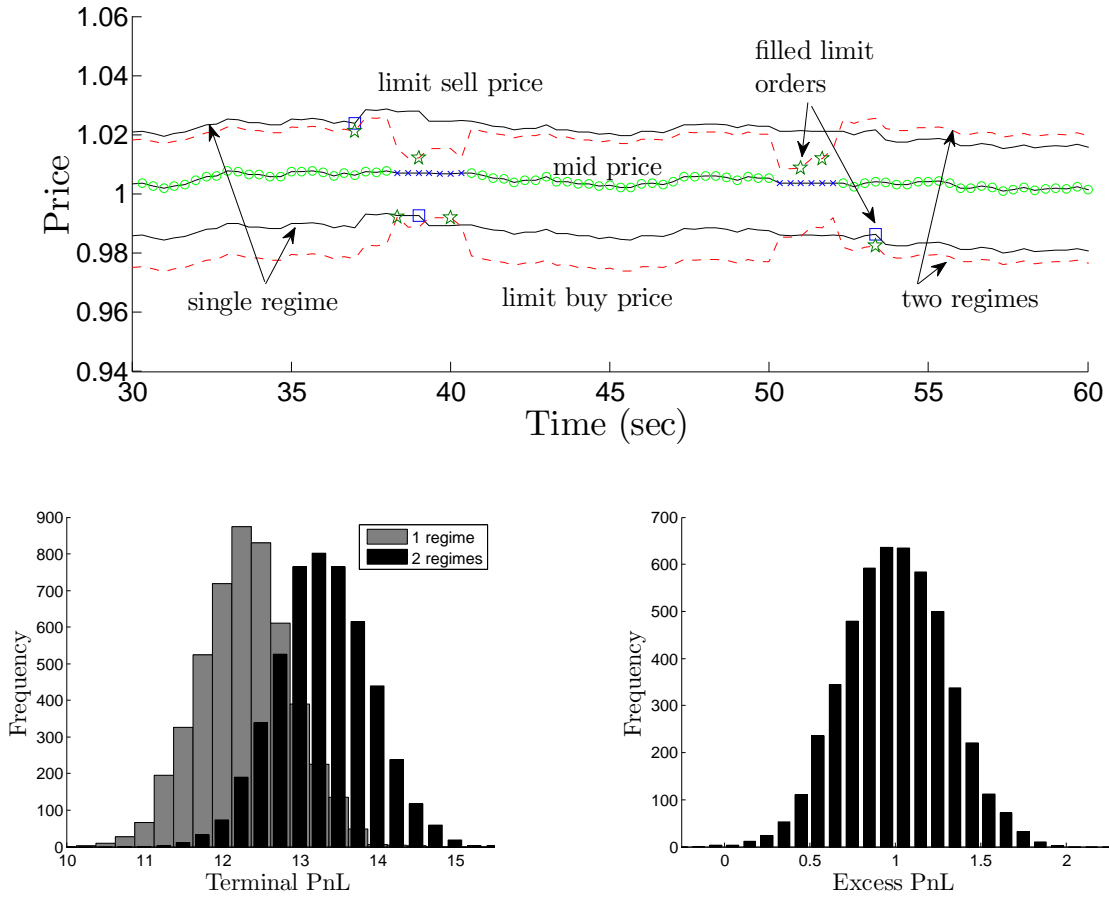
Figure 5: The top panel shows a sample path of the mid-price together with the optimal bid/ask strategy and the executed trades for a trader who uses two regimes (dashed-lines) and a trader who uses one regime (solid lines). The stars and boxes show filled limit orders events. The bottom left panel shows the distribution of the investors terminal PnL by investing in the two regime strategy, while the bottom right panel shows the excess PnL the two regime trader receives over the one-regime trader where both investors have a coefficient of risk aversion $\gamma = 1$.

that regime, coupled with the high persistence of the regime (80.67%), and the fact that over 30% of all AA trades during that month occurred in that state, trades with the sole purpose of collecting liquidity rebates are an important source of low-risk profits for HFTs.

In the second example of HF trading strategies, we first derive the optimal tick-by-tick strategy that an HF investor who uses limit-orders to profit from the bid-ask-spread should follow. In general, our analytical results provide the (immediate-or-cancel) buy and sell optimal strategy that the investor should post and how to update them every time a trade has occurred. These quantities depend on: the rate of arrival of trades, the intra-day-state of the market, the within state volatility of price revisions, the inventories which track the investor's accumulated stock, the shape of the limit-order book, and finally, the proximity to the investment horizon T. We show that the spread posted by the HF investor is wider (tighter) when the volatility of the price innovation is high (low). Moreover, as the investor accumulates a long (short) position, the investor's bid-price (ask-price) moves away from the mid-price and the ask-price (bid-price) moves in towards it – inducing the investor to sell (buy) assets and at the same time causing mean reversion toward zero in the inventories. The strategy also considers how likely a posting deep in the book is to be filled and thus

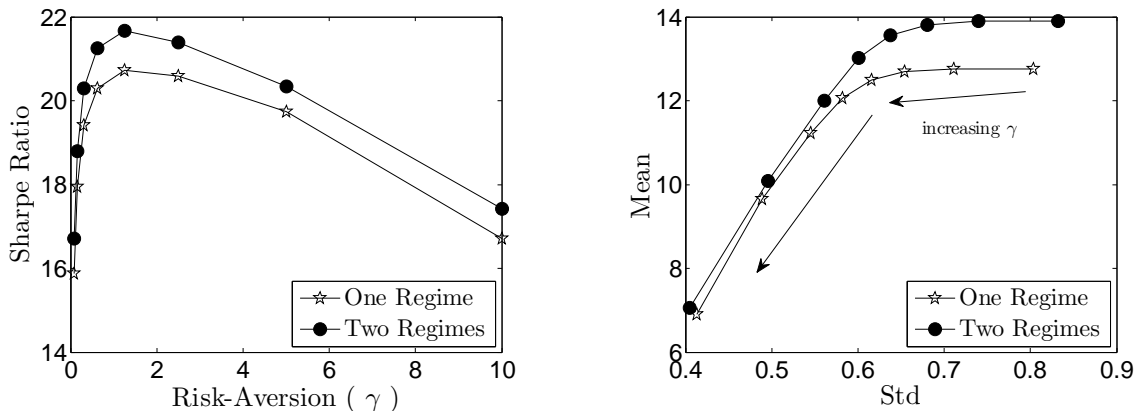Figure 6: The left-hand panel shows the Sharpe ratio as a function of the risk aversion parameter.

adjusts the buy and sell orders accordingly– which depend on the within-state arrival rate, volatility of trades and shape of the book. Finally, all else equal, as the investment horizon approaches T, the investor submits buy and sell limit-orders which are tighter around the mid-price – a strategy that stresses the fact that the HF investor aims at holding zero inventories at the end of investment horizon.

Moreover, we illustrate how the HF market making strategy performs under different assumptions about information and risk aversion. As expected, we show that better informed HFTs are more profitable and that those who make markets with lesser quality information see a reduction in their profits. This reduction in profits is a consequence of not being able to submit optimal limit orders to profit from periods of trade clustering or periods of heightened volatility and because some of the less informed limit orders can be picked off by better informed traders. Finally, we show that as the level of risk aversion increases, the gains from better quality information diminish because, everything else equal, the trader posts more conservative quotes in the book – i.e. limit orders are sent deeper into the book.

References

Admati, A. R. and P. Pfleiderer (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies 1*(1), 3–40.

Alfonsi, A. and A. Schied (2010). Optimal trade execution and absence of price manipulations in limit order book models. *SIAM Journal on Financial Mathematics 1*, 490–522.

Almgren, R. (2003, January). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance 10*(1), 1–18.

Almgren, R. (2009, June). Optimal trading in a dynamic market. Working Paper, New York University.

Asmussen, S. (2003). *Applied Probability and Queues* (2nd ed.). Springer.

Avellaneda, M. and S. Stoikov (2008, November). High-frequency trading in a limit order book. *Quantitative Finance 8*, 217–224.

Baum, L., T. Petrie, G. Soules, and N. Weiss (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics 41(1)*, 164–171.

Bauwens, L. and P. Giot (2000). The logarithmic ACD model: An application to the bid-ask quote process of three NYSE stocks. *Annales D'Economie et de Statistique* (60), 117–149.

Bauwens, L. and N. Hautsch (2009). *Modelling Financial High Frequency Data Using Point Processes*, pp. 953–979. Springer Berlin Heidelberg.

Bayraktar, E. and M. Ludkovski (2011). Liquidation in limit order books with controlled intensity. Working Paper.

Biernacki, C., G. Celeux, and G. Govaert (2001). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22(7)*, 719–725.

Bouchard, B., N.-M. Dang, and C.-A. Lehalle (2011). Optimal control of trading algorithms: A general impulse control approach. *SIAM Journal on Financial Mathematics 2*, 404–438.

Cappé, O., E. Moulines, and t. Rydén (2005). *Inference in Hidden Morkov Models*. Springer.

Cartea, Á. and S. Jaimungal (2012). Risk metrics and fine tuning of high frequency trading strategies. *Forthcoming Mathematical Finance*.

Cartea, Á., S. Jaimungal, and J. Ricci (2011). Buy low sell high: A high frequency trading perspective. *SSRN eLibrary, http://ssrn.com/abstract=1964781*.

Cartea, Á. and T. Meyer-Brandis (2010). How duration between trades of underlying securities affects option prices. *Review of Finance 14*(4), 749–785.

Cartea, Á. and J. Penalva (2012). Where is the value in high frequency trading? *Quarterly Journal of Finance 2*(3), 1–46.

Celeux, G. and J.-B. Durand (2008). Selecting hidden Markov model state number with cross-validated likelihood. *Computational Statistics 23(4)*, 541–564.

Cvitanić, J. and A. A. Kirilenko (2010). High Frequency Traders and Asset Prices. *SSRN eLibrary*.

de Jong, F. and B. Rindi (2009). *The Microstructure of Financial Markets* (1st ed.). Cambridge University Press.

Diamond, D. W. and R. E. Verrechia (1987). Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics 18*, 277–311.

Dufour, A. and R. F. Engle (2000, December). Time and the price impact of a trade. *The Journal of Finance LV*(6), 2467–2498.

Easley, D. and M. O'Hara (1992, June). Time and the process of security price adjustment. *The Journal of Finance XLVII*(2), 577–605.

Engle, R. F. (2000, January). The econometrics of ultra-high-frequency data. *Econometrica 68*(1), 1–22.

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

Fernandes, M. and J. Grammig (2005, July). Nonparametric specification tests for conditional duration models. *Journal of Econometrics 127*(1), 35–68.

Guilbaud, F. and H. Pham (2011). Optimal high frequency trading with limit and market orders. *SSRN eLibrary*.

Hujer, R., S. Vuletic, and S. Kokot (2002). The Markov switching ACD model. *SSRN eLibrary*.

Kharroubi, I. and H. Pham (2010). Optimal portfolio liquidation with execution cost and risk. *SIAM Journal on Financial Mathematics 1*, 897–931.

Latza, T., I. Marsh, and R. Payne (2012). Computer-based trading in the cross-section. Working Paper.

Lorenz, J. and R. Almgren (2011, April). Meanvariance optimal adaptive execution. Forthcoming Applied Mathematical Finance.

Maheu, J. M. and T. H. McCurdy (2000). Volatility dynamics under duration-dependent mixing. *Journal of Empirical Finance 7*(3-4), 345 – 372.

Manganelli, S. (2005). Duration, volume and volatility impact of trades. *Journal of Financial Markets 8*(4), 377 – 399.

Meitz, M. and T. Terasvirta (2006, January). Evaluating models of autoregressive conditional duration. *Journal of Business & Economic Statistics 24*, 104–124.

Mongillo, G. and S. Deneve (2008). Online learning with hidden Markov models. *Neural Computation 20*(7), 1706–1716.

Renault, E., T. van der Heijden, and B. J. Werker (2010, December). The dynamic mixed hitting-time model for multiple transaction prices and times. Working Paper.

SEC (2010, January). Concept release on equity market structure. Concept Release Release No. 34-61358; File No. S7-02-10, SEC. 17 CFR PART 242.

Viterbi, A. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory 13(2)*, 260–269.

Zhang, M. Y., J. R. Russell, and R. S. Tsay (2001, August). A nonlinear autoregressive conditional duration model with applications to financial transaction data. *Journal of Econometrics 104*(1), 179–207.

A   The EM Algorithm for HMMs

In this appendix, we provide a quick review of the EM algorithm for HMMs. More details on the Baum-Welch approach and HMMs in general can be found in for example Cappé, Moulines, and Rydén (2005).

- The E-step amounts to computing the conditional expectation of the *complete-data log-likelihood* given the current estimate of the full model parameters $\Theta^{(k-1)} = \{A^{(k-1)}, \pi^{(k-1)}, \theta^{(k-1)} = \{\lambda^{(k-1)}, \mathbf{p}^{(k-1)}, \mathbf{bf}\sigma^{(k-1)}\}\}$. That is compute

$$
\begin{aligned}
\mathcal{Q}(\Theta, \Theta^{(k-1)}) &= \mathbb{E}\left[\ln p(\{(\tau_t, X_t); Z_t\}_{t=1,\ldots,n} \mid \Theta) \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right] \\
&= \sum_{t=1}^{n}\sum_{j=1}^{K} \ln f_{\theta_j^{(k-1)}}(\{(\tau_t, X_t)\}) \, \mathbb{P}\left(Z_t = j \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right) \\
&\quad + \sum_{t=1}^{n-1}\sum_{j=1}^{K}\sum_{k=1}^{K} \ln A_{jk} \, \mathbb{P}\left(Z_t = j, Z_{t+1} = k \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right) \\
&\quad + \sum_{j=1}^{K} \ln \pi_j \, \mathbb{P}\left(Z_1 = j \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right) .
\end{aligned}
$$

The Baum-Welch forward-backward (or α-β) algorithm is used to compute the two types of conditional probabilities arising in the above expression: (i) the Markov chain *responsibilities* $r_{t,j} = \mathbb{P}\left(Z_t = j \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right)$ and (ii) the conditional transition probabilities $\xi_{t,jk} = \mathbb{P}\left(Z_t = j, Z_{t+1} = k \,\Big|\, \{(\tau_t, X_t)\}_{t=1,\ldots,n}; \Theta^{(k-1)} \right)$.

- In the M-step, $\mathcal{Q}(\Theta, \Theta^{(k-1)})$ is maximized (subject $\sum_{k=1}^{K} A_{jk} = 1$ and $\sum_j \pi_j = 1$). For our with-in continuous-mixture model of price-revisions and (censored) exponential durations, the resulting parameter update rules are

$$
\lambda_j^* = -\ln \frac{\sum_{t=1}^{n} r_{t,j}\, \tau_t}{\sum_{t=1}^{n} r_{t,j}\, (\tau_t + 1)} , \tag{9a}
$$

$$
p_j^* = \frac{\sum_{t:X_t=0} r_{t,j}}{\sum_t r_{t,j}} , \tag{9b}
$$

$$
\sigma_j^* = \sqrt{\frac{\sum_{t:X_t \neq 0} r_{t,j}\, X_t^2}{\sum_{t:X_t \neq 0} r_{t,j}}} , \tag{9c}
$$

$$
A_{jk}^* = \frac{\sum_{t=1}^{n-1} \xi_{t,jk}}{\sum_{t=1}^{n-1} r_{t,j}} , \tag{9d}
$$

$$
\pi_j^* = r_{1,j} . \tag{9e}
$$

The EM steps are then repeated until the relative increase in the complete-data-log-likelihood is less than 10^{-6}.

B   Matching Discrete & Continuous Models

In this appendix, we describe how to match the continuous HMM to any given estimated discrete HMM. The regime dependent rate of arrival of trades $\lambda^{(k)}$ are identical in both models. For the volatility matching we set the within regime volatility $\sigma_c^{(k)}$ of the continuous model such that its

24

variance (at the expected time of execution) is equal to the unconditional variance of the discrete model. Consequently,

$$\sigma_c^{(k)} = \sqrt{\frac{1 - p^{(k)}}{\lambda^{(k)}}} \, \sigma_d^{(k)} \,. \tag{10}$$

The only remaining parameters which require calibration are the transition rates B_{kl} of the continuous Markov chain H_t. For this purpose, we propose to match the probability that the chain begins in regime k, a single trade occurs, and the chain ends in regime l at time t. For the discrete time HMM this probability is

$$\begin{aligned}
P_{kl}^d(t) &\triangleq \mathbb{P}(N(t) = 1, \ Z_1 = l | Z_0 = k) \\
&= \int_0^t \left(\lambda_k e^{-\lambda_k u} \right) A_{kl} \left(e^{-\lambda_l(t-u)} \right) du \\
&= \begin{cases} \frac{\lambda^{(k)}}{\lambda^{(k)} - \lambda^{(l)}} A_{kl} \left(e^{-\lambda^{(l)}t} - e^{-\lambda^{(k)}t} \right), & k \neq l \,, \\[2mm] A_{kl} \, \lambda^{(k)} \, e^{-\lambda^{(k)}t} \, t, & k = l \,. \end{cases}
\end{aligned} \tag{11}$$

For the continuous time HMM this probability is

$$P_{kl}^c(t) \triangleq \mathbb{P}(N(t) = 1, \ H_t = l | H_0 = k) = \left(\Omega \, e^{(B-\Omega)t} \right)_{kl} t \,, \tag{12}$$

where $\Omega = \mathrm{diag}(\lambda^{(1)}, \dots, \lambda^{(K)})$ and the exponentiation is the matrix version.

It is not possible to match these two probabilities for every t; however, given that the trades arrive more quickly than transitions in the continuous time chain, we propose to match these probabilities at the expected time of a trade. Consequently, we choose the transition rates B_{jk} in the continuous time chain such that

$$P_{kl}^c \left(\tfrac{1}{\lambda^{(k)}} \right) = P_{kl}^d \left(\tfrac{1}{\lambda^{(k)}} \right) \qquad \forall k, l = 1, \dots, K \,.$$

This is a highly non-linear system of equations, but they pose no numerical difficulties. For our implementations, we used Matlab's fminsearch function.

C   Limit-Order Algorithmic Trading Strategy

In this Appendix we show that the feedback solution to the optimal control problem (6) is indeed given by (7). The dynamic programming principle implies that the value function $V^{(k)}(t, x, S, q)$ satisfies the HJB equation

$$\begin{cases}
V_t^{(k)}(t, x, s, q) + \tfrac{1}{2}(\sigma^{(k)})^2 \, V_{ss}^{(k)}(t, x, s, q) \\[2mm]
\quad + \max_{\delta^-} \left\{ \lambda^{(k)} e^{-\kappa_-^{(k)} \delta^-} \left(V^{(k)}(t, x - (s - \delta^-), s, q+1) - V^{(k)}(t, x, s, q) \right) \right\} \\[2mm]
\quad + \max_{\delta^+} \left\{ \lambda^{(k)} e^{-\kappa_+^{(k)} \delta^+} \left(V^{(k)}(t, x + (s + \delta^+), s, q-1) - V^{(k)}(t, x, s, q) \right) \right\} \\[2mm]
\qquad\qquad + \sum_{l=1}^{M} B_{kl} \left(V^{(l)}(t, x, s, q) - V^{(k)}(t, x, s, q) \right) \ = 0 \,, \\[4mm]
\qquad\qquad\qquad\qquad\qquad\qquad V^{(k)}(T, x, s, q) \ = u(x + sq) \,.
\end{cases} \tag{13}$$

Substituting the ansatz

$$V^{(k)}(T, x, s, q) = \frac{1}{\gamma} \left(1 - \exp \left\{ -\gamma \left(x + q \, S + g^{(k)}(t, q) \right) \right\} \right) ,$$

reduces the HJB equation to

$$
\begin{cases}
g_t^{(k)} - \frac{1}{2}(\sigma^{(k)})^2\,\gamma\,q^2 + \max_{\delta^-}\left\{ \lambda^{(k)}e^{-\kappa_-^{(k)}\delta^-}\dfrac{1-e^{-\gamma(\delta^-+\Delta^*g^{(k)})}}{\gamma} \right\} \\
\qquad\qquad + \max_{\delta^+}\left\{ \lambda^{(k)}e^{-\kappa_+^{(k)}\delta^+}\dfrac{1-e^{-\gamma(\delta^++\Delta_*g^{(k)})}}{\gamma} \right\} \\
\qquad\qquad\qquad + \sum_{l=1}^{M} B_{kl}\dfrac{1-e^{-\gamma(g^{(l)}-g^{(k)})}}{\gamma} \quad = 0\,, \\
\\
\qquad\qquad\qquad\qquad\qquad\qquad g^{(k)}(T,q) \;\; = 0\,.
\end{cases}
\tag{14}
$$

Here, the shift operators Δ^* and Δ_* act on functions $h(t,q)$ as follows

$$
\Delta^*h(t,q) = h(t,q+1) - h(t,q)\,, \qquad \text{and} \qquad \Delta_*h(t,q) = h(t,q-1) - h(t,q)\,.
$$

Applying the first order conditions provides us with the feedback control solutions

$$
\delta_t^+ = -\Delta_*g^{(H_t)}(t,q_t) + \frac{1}{\gamma}\ln\left(1 + \frac{\gamma}{\kappa_{+,t}}\right)\,, \quad \text{and}
\tag{15a}
$$

$$
\delta_t^- = -\Delta^*g^{(H_t)}(t,q_t) + \frac{1}{\gamma}\ln\left(1 + \frac{\gamma}{\kappa_{-,t}}\right)\,.
\tag{15b}
$$

Substituting the feedback controls into the HJB equation (14) then results in the non-linear integro-differential equation

$$
\begin{cases}
g_t^{(k)} - \frac{1}{2}(\sigma^{(k)})^2\,\gamma\,q^2 + \alpha_-^{(k)}e^{\kappa_-^{(k)}\Delta^*g^{(k)}} + \alpha_+^{(k)}e^{\kappa_+^{(k)}\Delta_*g^{(k)}} + \sum_{l=1}^{K} B_{kl}\dfrac{1-e^{-\gamma(g^{(l)}-g^{(k)})}}{\gamma} \quad = 0\,, \\
\\
\qquad\qquad\qquad\qquad\qquad\qquad\qquad g^{(k)}(T,q) \;\; = 0\,.
\end{cases}
\tag{16}
$$

Here the constant

$$
\alpha_\pm^{(k)} = \frac{\lambda^{(k)}}{\kappa_\pm^{(k)} + \gamma}\left(1 + \frac{\gamma}{\kappa_\pm^{(k)}}\right)^{-\frac{\kappa_\pm^{(k)}}{\gamma}}\,,
$$

is introduced to reduce notation. The remaining task is to solve (16). Once armed with its solution, the feedback controls given in (15) provide the investor with the optimal limit-order strategy.

We have not found an exact solution to this equation, however, it is possible to obtain a perturbation expansion. In contrast to Avellaneda and Stoikov (2008), who have a single regime model, have a different ansatz and perform a perturbation expansion[16] in the inventory level q, we perform an expansion in the risk-aversion parameter γ. For this purpose, first write $g^{(k)}(t,q) = g_0^{(k)}(t) + \gamma\,g_1^{(k)}(t,q) + o(\gamma)$. Notice that the first order term is assumed independent of q. Inserting

---

[16]It is not strictly correct to expand in the inventory level q, since q is an integer and can take on values significantly larger than 1.

this expression into (16) and collecting terms in powers of γ we find that

$$g_{0,t}^{(k)} + (\beta_+^{(k)} + \beta_-^{(k)}) + \sum_{l=1}^{K} B_{kl}(g_0^{(l)} - g_0^{(k)}) = 0 \,, \tag{17}$$

$$
\begin{aligned}
g_{1,t}^{(k)} &- \frac{1}{2}(\sigma^{(k)})^2 q^2 - \frac{1}{2}\left(\frac{\beta_-^{(k)}}{\kappa_-^{(k)}} + \frac{\beta_+^{(k)}}{\kappa_+^{(k)}}\right) + \sum_{l=1}^{K} B_{kl}(g_0^{(l)} - g_0^{(k)})^2 \\
&+ e^{-1}\lambda^{(k)}\left(\Delta^* g_1^{(k)} + \Delta_* g_1^{(k)}\right) + \sum_{l=1}^{K} B_{kl}(g_1^{(l)} - g_1^{(k)}) = 0\,,
\end{aligned}
\tag{18}
$$

where $\beta_\pm^{(k)} = \lambda^{(k)}/(e\kappa_\pm^{(k)})$. The solution for $g_1^{(k)}$ can further be decomposed as $g_1^{(k)}(t,q) = a^{(k)}(t) + b^{(k)}(t)\,q^2$ – this is not another approximation, rather it is the form of the exact solution. Moreover, since the optimal investment strategy, through the feedback controls (15), depend on $g^{(k)}$ only through $\Delta^* g^{(k)}$ and $\Delta_* g^{(k)}$, it is only necessary to solve for $b_t^{(k)}$ and not $a^{(k)}$ or $g_0^{(k)}$. To this end, we find that $b^{(k)}(t)$ solves the system of ODEs

$$b_t^{(k)} - \frac{1}{2}(\sigma^{(k)})^2 + \sum_l B_{kl} b^{(l)} = 0\,. \tag{19}$$

Standard techniques can be used to solve this system of ODEs. Let $D = \mathrm{diag}(d_1, \dots, d_K)$ denote the matrix of eigenvalues of the transition rate matrix B, and V be the matrix of eigenvectors so that $B = V^{-1}DV$. Since the transition matrix sums to zero along rows, there is one zero eigenvalue which we label as $d_1 = 0$. Assuming distinct eigenvalues,[17] then the solution is given by (8). On substituting the solution into the feedback control (15) one finds the result quoted in (7). This completes the proof.

---

[17]This assumption is easily removed if necessary, but it is likely that the eigenvalues will be distinct.

D The Estimated Model Parameters

Comparison of all data sets with four regimes.

	AA – FEB 2001							
	Transition Probability Matrix A				Conditional Parameters			
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	91.05%	8.95%	0.00%	0.00%	1.468	100.00%	3.238	0.000
	(3.93%)	(8.75%)	(8.89%)	(13.82%)	(1.200)	(34.95%)	(9.794)	
2	1.71%	85.37%	0.00%	12.92%	0.105	47.34%	24.338	17.661
	(2.33%)	(9.81%)	(27.05%)	(18.20%)	(0.023)	(9.05%)	(18.573)	
3	0.00%	0.30%	95.55%	4.15%	0.078	54.67%	4.363	2.937
	(1.31%)	(52.32%)	(62.51%)	(11.36%)	(0.006)	(14.83%)	(3.503)	
4	0.00%	1.68%	5.62%	92.70%	0.068	32.69%	9.712	7.968
	(2.00%)	(33.62%)	(8.02%)	(40.54%)	(0.001)	(10.24%)	(3.389)	

	AA – FEB 2008							
	Transition Probability Matrix A				Conditional Parameters			
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	80.67%	17.21%	2.12%	0.00%	2.871	99.97%	3.010	0.050
	(0.32%)	(15.06%)	(15.09%)	(0.00%)	(0.034)	(0.11%)	(0.995)	
2	14.88%	79.97%	1.71%	3.44%	1.782	47.32%	3.578	2.597
	(12.63%)	(10.79%)	(0.45%)	(23.10%)	(0.015)	(46.76%)	(16.079)	
3	2.14%	1.18%	69.39%	27.28%	1.779	98.85%	21.361	2.293
	(12.69%)	(0.36%)	(11.14%)	(23.82%)	(0.013)	(46.62%)	(16.377)	
4	8.08%	0.55%	32.78%	58.59%	0.389	36.86%	3.012	2.393
	(0.94%)	(34.54%)	(34.17%)	(0.34%)	(0.002)	(0.32%)	(0.020)	

Table 6: Estimated 4-regime model parameters on AA data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.

AMZN – FEB 2001							
	Transition Probability Matrix A				Conditional Parameters		
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	81.36%	17.71%	0.93%	0.00%	1.073	99.25%	264.549	22.897
	(0.83%)	(0.91%)	(0.32%)	(0.15%)	(0.017)	(0.15%)	(24.804)	
2	10.58%	85.09%	2.81%	1.51%	0.659	37.42%	47.748	37.771
	(0.76%)	(0.72%)	(0.30%)	(0.21%)	(0.004)	(0.78%)	(0.428)	
3	0.11%	4.11%	77.10%	18.68%	0.403	99.97%	46.685	0.805
	(0.07%)	(0.26%)	(0.71%)	(0.77%)	(0.009)	(0.32%)	(19.850)	
4	1.21%	1.97%	14.01%	82.81%	0.127	45.88%	47.656	35.059
	(0.33%)	(0.26%)	(1.00%)	(1.03%)	(0.002)	(1.25%)	(0.643)	

AMZN – FEB 2008							
	Transition Probability Matrix A				Conditional Parameters		
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	79.89%	3.05%	0.04%	17.03%	2.614	85.57%	1.810	0.688
	(0.36%)	(0.13%)	(0.02%)	(0.35%)	(0.018)	(0.19%)	(0.026)	
2	1.14%	94.48%	1.18%	3.20%	2.101	46.57%	2.931	2.143
	(0.17%)	(0.15%)	(0.07%)	(0.20%)	(0.011)	(0.15%)	(0.011)	
3	3.05%	18.64%	75.96%	2.36%	1.203	26.34%	11.480	9.853
	(0.79%)	(1.21%)	(1.19%)	(0.55%)	(0.018)	(0.95%)	(0.118)	
4	43.08%	0.60%	0.08%	56.24%	0.487	29.06%	2.496	2.102
	(0.50%)	(0.26%)	(0.03%)	(0.44%)	(0.004)	(0.41%)	(0.012)	

Table 7: Estimated 4-regime model parameters on AMZN data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.

	Transition Probability Matrix A				Conditional Parameters			

HNZ – FEB 2001

Regime	1	2	3	4	λ	p	$\sigma\ (\times 10^{-4})$	$\sigma\sqrt{1-p}\ (\times 10^{-4})$
1	86.61%	10.80%	0.84%	1.76%	0.786	100.00%	4.498	0.000
	(3.62%)	(1.40%)	(1.64%)	(4.75%)	(0.010)	(8.82%)	(2.625)	
2	0.69%	84.90%	0.01%	14.39%	0.052	45.99%	23.266	17.099
	(2.21%)	(32.16%)	(35.35%)	(10.94%)	(0.002)	(7.56%)	(0.411)	
3	0.06%	0.93%	92.86%	6.14%	0.038	57.17%	3.506	2.295
	(7.92%)	(17.99%)	(10.48%)	(18.70%)	(0.005)	(8.57%)	(1.038)	
4	0.04%	2.03%	10.13%	87.80%	0.025	31.54%	9.193	7.606
	(1.34%)	(12.38%)	(14.32%)	(4.16%)	(0.001)	(3.68%)	(0.396)	

HNZ – FEB 2008

Regime	1	2	3	4	λ	p	$\sigma\ (\times 10^{-4})$	$\sigma\sqrt{1-p}\ (\times 10^{-4})$
1	71.95%	3.02%	25.03%	0.00%	2.257	100.00%	2.566	0.012
	(0.71%)	(0.98%)	(0.92%)	(0.92%)	(0.041)	(0.07%)	(11.508)	
2	0.35%	54.26%	1.59%	43.80%	1.286	98.43%	17.567	2.198
	(0.23%)	(1.08%)	(0.97%)	(1.71%)	(0.019)	(0.16%)	(1.206)	
3	17.73%	2.92%	69.28%	10.08%	0.906	42.42%	2.828	2.146
	(1.26%)	(0.73%)	(1.34%)	(0.86%)	(0.018)	(1.20%)	(0.032)	
4	9.82%	31.09%	3.13%	55.96%	0.149	39.86%	2.975	2.307
	(1.07%)	(1.30%)	(0.57%)	(0.55%)	(0.002)	(0.43%)	(0.019)	

Table 8: Estimated 4-regime model parameters on HNZ data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.

	Transition Probability Matrix A				Conditional Parameters			
	\multicolumn{9}{c}{IBM – FEB 2001}							
Regime	1	2	3	4	λ	p	$\sigma\,(\times 10^{-4})$	$\sigma\sqrt{1-p}\,(\times 10^{-4})$
1	91.98%	5.41%	2.46%	0.15%	1.699	99.04%	57.840	5.676
	(1.11%)	(1.03%)	(0.71%)	(0.38%)	(0.080)	(0.47%)	(16.057)	
2	1.72%	88.45%	2.11%	7.72%	0.244	24.69%	17.123	14.859
	(0.60%)	(1.29%)	(0.55%)	(0.93%)	(0.007)	(0.87%)	(0.378)	
3	0.28%	0.45%	90.45%	8.82%	0.195	40.79%	1.991	1.532
	(0.10%)	(0.21%)	(0.63%)	(0.79%)	(0.004)	(0.70%)	(0.025)	
4	0.03%	2.06%	6.60%	91.31%	0.180	21.77%	6.019	5.324
	(0.07%)	(0.19%)	(0.53%)	(0.61%)	(0.002)	(0.43%)	(0.053)	

	Transition Probability Matrix A				Conditional Parameters			
	\multicolumn{9}{c}{IBM – FEB 2008}							
Regime	1	2	3	4	λ	p	$\sigma\,(\times 10^{-4})$	$\sigma\sqrt{1-p}\,(\times 10^{-4})$
1	77.55%	0.00%	6.88%	15.58%	2.039	89.92%	1.154	0.366
	(0.41%)	(0.00%)	(0.24%)	(0.29%)	(0.015)	(0.37%)	(0.018)	
2	0.80%	90.82%	6.93%	1.45%	1.851	31.70%	6.823	5.639
	(0.63%)	(0.41%)	(0.77%)	(0.30%)	(0.024)	(0.66%)	(0.069)	
3	4.62%	0.97%	91.65%	2.76%	1.769	42.24%	1.708	1.298
	(0.39%)	(0.05%)	(0.23%)	(0.21%)	(0.011)	(0.18%)	(0.006)	
4	28.73%	0.14%	0.98%	70.15%	0.413	33.15%	1.618	1.323
	(0.36%)	(0.02%)	(0.01%)	(0.36%)	(0.003)	(0.43%)	(0.010)	

Table 9: Estimated 4-regime model parameters on IBM data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.

	KO – FEB 2001							
	Transition Probability Matrix A				Conditional Parameters			
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	94.33%	4.43%	1.24%	0.00%	1.967	100.00%	1.898	0.000
	(7.18%)	(4.30%)	(1.51%)	(11.65%)	(1.994)	(40.15%)	(34.216)	
2	1.90%	86.26%	1.96%	9.89%	0.113	45.15%	20.635	15.282
	(2.14%)	(50.97%)	(39.79%)	(17.97%)	(0.025)	(8.73%)	(20.006)	
3	0.08%	0.65%	94.72%	4.55%	0.088	49.91%	2.602	1.842
	(0.64%)	(38.87%)	(55.46%)	(19.53%)	(0.004)	(9.75%)	(0.678)	
4	0.00%	2.20%	7.06%	90.74%	0.080	27.75%	6.702	5.697
	(3.42%)	(8.89%)	(8.55%)	(4.49%)	(0.002)	(1.47%)	(0.576)	

	KO – FEB 2008							
	Transition Probability Matrix A				Conditional Parameters			
Regime	1	2	3	4	λ	p	σ ($\times 10^{-4}$)	$\sigma\sqrt{1-p}$ ($\times 10^{-4}$)
1	79.25%	8.37%	0.01%	12.37%	2.417	99.99%	31.499	0.310
	(1.22%)	(8.53%)	(9.20%)	(0.57%)	(0.024)	(0.94%)	(30.033)	
2	11.75%	85.12%	0.40%	2.72%	1.729	54.19%	2.001	1.354
	(11.89%)	(48.51%)	(57.96%)	(2.45%)	(0.076)	(8.08%)	(0.286)	
3	1.87%	4.87%	91.67%	1.59%	1.422	39.21%	8.062	6.286
	(21.81%)	(39.58%)	(59.73%)	(1.70%)	(0.213)	(19.34%)	(5.024)	
4	28.79%	0.02%	0.08%	71.11%	0.387	43.55%	1.781	1.338
	(0.36%)	(0.34%)	(0.09%)	(0.44%)	(0.002)	(0.46%)	(0.017)	

Table 10: Estimated 4-regime model parameters on KO data for the months of February 2001 and 2008. The reported numbers in the braces are the 95% standard errors based on a bootstrap of the estimated model.