

The Cost of Latency*

Ciamac C. Moallemi
Graduate School of Business
Columbia University
email: ciamac@gsb.columbia.edu

Mehmet Sağlam
Graduate School of Business
Columbia University
email: msaglam13@gsb.columbia.edu

Initial Version: November 18, 2009

Current Revision: July 19, 2011

Abstract

Modern electronic markets have been characterized by a relentless drive towards faster decision making. Significant technological investments have led to dramatic improvements in latency, the delay between a trading decision and the resulting trade execution. We describe a theoretical model for the quantitative valuation of latency. Our model measures the trading frictions created by the presence of latency, by considering the optimal execution problem of a representative investor. Via a dynamic programming analysis, our model provides a closed-form expression for the cost of latency in terms of well-known parameters of the underlying asset. We implement our model by estimating the latency cost incurred by trading on a human time scale. Examining NYSE common stocks from 1995 to 2005 and show that median latency cost across our sample more than tripled during this time period. Furthermore, using the same data set, we compute a measure of implied latency, and conclude that the median implied latency decreased by approximately two orders of magnitude. Empirically calibrated, our model suggests that the reduction in cost achieved by going from trading on a human time scale to a low latency time scale is comparable with other execution costs faced by the most cost efficient institutional investors, and is consistent with the rents that are extracted by ultra low latency agents, such as providers of automated execution services or high frequency traders.

1. Introduction

In the past decade, electronic markets have become pervasive. Technological advances in these markets have led to dramatic improvements in latency, or, the delay between a trading decision and the resulting trade execution. In the past 30 years, the time scale over which a trade is processed has gone from minutes¹ to milliseconds² — “low latency” in a contemporary electronic market would be qualified as under 10 milliseconds, “ultra low latency” as under 1 millisecond. This change represents a dramatic reduction by *five orders of magnitude*. To put this in perspective, human reaction time is thought to be in the hundreds of milliseconds.

*The first author would like to thank Jim Gatheral for a helpful discussion that motivated this work. The authors would like to thank Albert Menkveld, Larry Glosten, and Jialin Yu for their helpful comments.

¹NYSE, pre-1980 upgrade (Easley et al., 2008).

²“The value of a millisecond: Finding the optimal speed of a trading infrastructure,” TABB Group, April 2008.

One factor behind this trend has been competition between exchanges, as one mechanism for differentiation between exchanges is latency. This competition is driven by a significant demand amongst a class of investors, sometimes called “high frequency” traders, for low latency trade execution. High frequency traders are thought to account for more than half of all US equity trades.³ They expend significant resources in order to develop algorithms and systems that are able to trade quickly. For example, on the time scale of milliseconds, the speed of light can become a binding constraint on the delay in communications. Hence, traders seeking low latency will “co-locate”, or house their computers in the same facility as the exchange, in order eliminate delays due to a lack of physical proximity. This co-location comes at a significant expense, however it has been stated that a 1 millisecond advantage can be worth \$100 million to a major brokerage firm.⁴

There has been much discussion of the importance of latency among various market participants, regulators, and academics. Despite the significant amount of recent interest, however, latency remains poorly understood from a theoretical perspective. For example, how does latency relate to transaction costs? Is latency only relevant to investors with short time horizons, such as high frequency traders, or does latency also affect long term investors such as pension funds and mutual funds? Many of these important questions have been considered in anecdotal or ad hoc discussions. Our goal here is to provide a framework for quantitative analysis of these issues.

In particular, we wish to understand the benefit to a single trader in the marketplace of lowering their latency, while holding everything else fixed. This is a different question than understanding whether the collective marketplace is better or worse off given lower latency. We believe that understanding the former, however, is an important first step which will inform our ultimate understanding of the latter.

The cost that a trader bears due to latency can take many different forms, depending on the precise trading strategy. However, we can identify a number of broad themes,⁵ sometimes overlapping, as to why the ability to trade with low latency might be valuable to an investor:

1. **Contemporaneous decision making.** A trader with significant latency will be making trading decisions based on information that is stale.

For example, consider an automated trader implementing a market-making strategy in an electronic limit order book. The trader will maintain active limit orders to buy and sell. The prices at which the trader is willing to buy or sell will naturally depend on, say, the limit orders submitted by other investors, the price of the asset on other exchanges, the price of related assets, overall market factors, etc. If the trader cannot update his orders in a timely fashion in response to new information, he may end up trading at disadvantageous prices.

2. **Comparative advantage/disadvantage.** The ability to trade with low latency in absolute terms may not be as important as the ability to trade with low *relative* latency, that is, as

³“Stock traders find speed pays, in milliseconds,” *New York Times*, July 23, 2009.

⁴“Wall Street’s quest to process data at the speed of light,” *Information Week*, April 21, 2007.

⁵See Cespa and Foucault (2008) for a related discussion.

compared to competitors.

For example, consider a program trader implementing an index arbitrage strategy, seeking to profit on the difference between an index and its underlying components. There may be many market participants pursuing such strategies and identifying the same discrepancies. The challenge for the trader is to be able to act in the marketplace to exploit a discrepancy *before* a price correction takes place, i.e., before competitors are able to act. The means having a low relative latency.

3. **Time priority rules.** Many modern markets treat orders differentially based on the time of arrival, and favor earlier orders.

For example, in an electronic limit order book, the limit orders on each side of the market are prioritized in a particular way. When a market order to buy arrives, it is matched against the limit orders to sell according to their priorities. Priority is first determined by price, i.e., limit orders with more lower prices receive higher priority. In many markets, however, limit orders at the same price are prioritized according to the time of their arrival. While a trader can always increase the priority of his orders (and, therefore, the likelihood of execution) by decreasing price, this comes at an obvious cost. If a trader can submit orders in a faster fashion, however, he can increase priority while maintaining the *same* price.

In this paper, we will quantify the cost of latency due to the first effect, a lack of contemporaneous decision making. It is an open question as to whether the other effects are more or less significant than the first, and their relative importance may depend on the particular investor and their trading strategy. Our analysis does not speak to this point. However, in what follows we will demonstrate that, by itself, the lack of contemporaneous decision making can induce trading costs that are of the same order of magnitude as other execution costs faced by large investors, and hence cannot be neglected.

Further, the importance of contemporaneous decision making will certainly vary from investor to investor. We will focus on an aspect of this that is universal, however, which is the importance of timely information for the execution of *contingent orders*. A contingent order, such as a limit order in an electronic limit order book or a resting order in a dark pool, presents the possibility of uncertain execution over an interval of time in exchange for price improvement relative to a market order, which executes immediately and with certainty. Specifically, when an investor employs a contingent order, the investor may be exposed to the realization of new information (for example, in the form of price movements, news, etc.) over the lifespan of the order. Latency, which prevents the investor from continuously and instantaneously accessing the market so as to update the order, can thus adversely impact the investor.

As a broad proxy for understanding the importance of latency in contingent order execution, we consider the effects of latency in an extremely simple yet fundamental trade execution problem: that of a risk-neutral investor who wishes to sell 1 share of stock (i.e., an atomic unit) over a fixed,

short time horizon (i.e., seconds) in a limit order book, and must decide between market orders and limit orders. Our problem formulation is reminiscent of barrier-diffusion models for limit order execution (e.g., Harris, 1998). It captures the fundamental *cost of immediacy* of trading (e.g., Grossman and Miller, 1988; Chacko et al., 2008), that is, the premium due to a patient liquidity supplier (who submits limit orders) relative to an impatient demander of liquidity (who submits market orders). While this problem is quite stylized, we will argue that it is broadly relevant since, at some level, *all* investors make such a choice of immediacy. For example, it may not seem at first glance that our execution problem is relevant for a pension fund that trades large blocks of stock over multiple days. However, the execution of a block trade via algorithmic trading involves the division of a large “parent” order into many atomic orders over the course of a day, each of these atomic “child” orders can be executed as limit orders or as market orders.

In our problem, in the absence of latency, the optimal strategy of the seller is a “pegging” strategy: the seller maintains a limit order at a constant spread above the bid price at any instant in time. We consider this case as a benchmark. In the presence of latency, the seller can no longer maintain continuous contact with the market so as to track the bid price in the market. The seller is forced to deviate from the benchmark policy in order to take into account the uncertainty introduced by the latency delay by incorporating a safety margin and lowering his limit order prices. The friction introduced by latency thus results in a loss of value to the seller. We will establish the difference in value to the seller between the case with latency and the benchmark case via dynamic programming arguments, and thus provide a quantification of the effects of latency.

The contributions of this paper are as follows:

- *We mathematically quantify the cost of latency.*

The trading problem we consider (deciding between limit and market orders) is faced by all large investors in modern equity markets, either directly (e.g., high frequency traders) or indirectly (e.g., pension funds who execute large trades via providers of automated execution services). Our analysis suggests that latency impacts all of these market participants, and that, all else being equal, the ability to trade with low latency results in quantifiably lower transaction costs. Further, when calibrated with market data, the latency cost we measure can be significant. It is of the same order of magnitude as other trading costs (e.g., commissions, exchange fees, etc.) faced by the most cost efficient large investors. Moreover, it is consistent with the rents that are extracted by agents who have made the requisite technological investments to trade with ultra low latency. For example, the latency cost of our model is comparable to the execution commissions charged by providers that offer algorithmic trade execution services on an agency basis. It is also comparable to the reported profits of high frequency traders.

To our knowledge, our model is the first to provide a quantification of the costs of latency in trade execution.

- *We provide a closed-form expression for the cost of latency as a function of well-known parameters of the asset price process.*

The cost of latency in our model can be computed numerically via dynamic programming. However, in the regime of greatest interest, where the latency is close to zero, we provide a closed-form asymptotic expression. In particular, define the *latency cost* associated with an asset as the costs incurred due to latency as a fraction of the overall cost of immediacy (the premium paid to a patient liquidity supplier by an impatient demander of liquidity). Given a latency of Δt , a price volatility of σ , and a bid-offer spread of δ , the latency cost takes the form

$$(1) \quad \frac{\sigma\sqrt{\Delta t}}{\delta} \sqrt{\log \frac{\delta^2}{2\pi\sigma^2\Delta t}}$$

as $\Delta t \rightarrow 0$.

- *Our method provides clear qualitative insight into the importance of latency.*

From (1), it is clear that the latency cost is an increasing function of the ratio of the standard deviation of prices over the latency interval (i.e., $\sigma\sqrt{\Delta t}$) to the bid-offer spread. Latency has a more important role when trading assets that are either more volatile (σ large) or, alternatively, more liquid (δ small). Further, as the latency approaches 0, the marginal benefit of latency reduction is increasing.

- *We empirically demonstrate that latency cost incurred by trading on a human time scale has dramatically increased for U.S. equities and the implied latency of a representative trader in this market decreased by approximately two orders of magnitude.*

We consider the cost due to the latency of trading on the time scale of human interaction. Using the data-set of Ait-Sahalia and Yu (2009), we estimate the latency cost of NYSE common stocks over the 1995–2005 period. We show that the median latency cost more than tripled in this time. This coincides with a period of decreasing tick sizes and increasing algorithmic and high frequency trading activity (Hendershott et al., 2010).

An alternative perspective is to consider a hypothetical investor who fixes a target level of cost due to latency, relative to the overall cost-of-immediacy. The representative trader maintains this target over time through continual technological upgrades to lower levels of latency. We determine the requisite level of *implied latency* for such a trader, over time and across the aggregate market. Using the same data-set, we observe that the median implied latency decreased by approximately two orders of magnitude over this time frame.

The rest of this paper is organized as follows: In Section 1.1, we review the related literature. In Section 2, as a starting point, we present a stylized, continuous-time trade execution problem in the absence of latency. We develop a variation of the model with latency in Section 3. In Section 4, we

provide a mathematical analysis of the optimal policy for our problem. By contrasting the results in the presence and absence of latency, we are able to quantitatively assess the cost of latency. In Section 5, we consider some empirical applications of the model. Finally, in Section 6 we conclude and discuss some future directions.

1.1. Related Literature

There has been a significant empirical literature studying, broadly speaking, the effects of improvements in trading technology. Closest to the aspect we consider is the work of Easley et al. (2008). They empirically test the hypothesis that latency affects asset prices and liquidity by examining the time period around an upgrade to the New York Stock Exchange technological infrastructure that reduced latency. Hendershott et al. (2010) explore the more general, overall effects of algorithmic and high frequency trading. Hasbrouck and Saar (2009) provide different evidence of changes in investor trading strategies that may be a result of improved technology. In subsequent work, they further consider the impact of measurements of low latency on market quality (Hasbrouck and Saar, 2010). Hendershott and Riordan (2009) analyze the impact of algorithmic trading on the price formation process using a data set from Deutsche Börse and conclude that algorithmic trading assists in the efficient price discovery without increasing the volatility. Kirilenko et al. (2010) consider the impact of high frequency trading on the ‘flash crash’ of 2010, while Brogaard (2010) more broadly examines the impact of high frequency traders on market quality.

On the theoretical front, Cespa and Foucault (2008) consider a rational expectations equilibrium between investors with different access to past transaction data. Some investors observe transactions in real-time, while others only observe transactions with a delay. This model of latency focuses on latency of the price ticker of past transactions, as opposed to latency in execution, which we consider here. Moreover, the goals of the two models differ significantly: Cespa and Foucault (2008) seek to build intuition regarding the equilibrium welfare implications of differential access to information via a structural model. We, on the other hand, seek a reduced form model that can be used to directly estimate the value of execution latency in a particular real world instance, given readily available data. Also related is the work of Ready (1999) and Stoll and Schenzler (2006), who consider the ability of intermediaries (e.g., specialists or dealers) to delay customer orders for their own benefit, thus creating a “free option” in the presence of execution latency. Finally, Cvitanic and Kirilenko (2010) and Jarrow and Protter (2011) consider the effect of high frequency traders on asset prices.

The trade execution problem we consider is that of an investor who wishes to sell a single share of and must decide between market and limit orders. This problem has been considered by many others (e.g., Angel, 1994; Harris, 1998; Lo et al., 2002). Our formulation is similar to the class of barrier-diffusion models considered by these authors; Hasbrouck (2007) provides a good account of this line of work. For a broad survey on limit order markets, see Parlour and Seppi (2008). In our model, the inability to trade continuously gives a limit order an option-like quality that

relates execution cost, order duration, and asset volatility. This idea goes as far back as the work of Copeland and Galai (1983). Closely related is the concept of the cost of immediacy, or, the premium paid by a liquidity demander via a market order to a liquidity supplier who posts a limit order. Grossman and Miller (1988) and Chacko et al. (2008) develop theoretical explanations of the cost of immediacy. For empirical evidence of the demand for immediacy in capital markets, see Bacidore et al. (2003) and Werner (2003).

Finally, also related is work on the discrete-time hedging of contingent claims with or without transaction costs (e.g., Boyle and Emanuel, 1980; Leland, 1985; Bertsimas et al., 2000). This literature addresses a different problem and draws different conclusions than our paper, however both relate to implications of a lack of continuous access to the market.

2. A Stylized Execution Model without Latency

Our goal is to understand the impact on the trade execution of latency. To this end, we will first describe a trade execution problem in the absence of latency. In Section 3, we will revisit this model in the presence of latency, so as to understand the resulting trade friction that is introduced. The spirit of our model is to consider an investor who wants to trade, but at a price that depends on an informational process that evolves stochastically and must be monitored continuously. We could directly consider such an abstract model of investor behavior. Instead, however, we will motivate the informational dependence of the trader through a specific optimal execution problem.

Consider the following stylized execution problem of an uninformed trader who must sell exactly one share⁶ of a stock over a time horizon $[0, T]$. At any time $t \in [0, T)$, the trader can take one of two actions:

1. The trader can submit a market order to sell. This order will execute at the best bid price at time t , denoted by S_t . We assume that the bid price evolves according to

$$(2) \quad S_t = S_0 + \sigma B_t,$$

where the process $(B_t)_{t \in [0, T]}$ is a standard Brownian motion and $\sigma > 0$ is an (additive) volatility parameter. Here, the choice of Brownian motion is made for simplicity; our model can be extended to the more general class of Markovian martingales, as is discussed in Section 4.4.

2. The trader can choose to submit a limit order to sell. In this case, the trader must also decide the limit price associated with the order, which we denote by L_t .

Once the trader sells one share, he exits the market. If the trader is not able to sell 1 share before time T , however, we assume that he is forced sell via a market order at time T , and therefore

⁶Note that the trade quantity of a single share is meant to represent an atomic unit of the asset, or the smallest commonly traded lot size. The underlying assumption is that the desired trade execution will ultimately be accomplished by a single transaction. In typical U.S. equity markets, for example, this atomic unit might be a block of 100 shares.

receives S_T . Here, we imagine the time horizon T to be small, on the order of the typical trade execution time (i.e., seconds).

2.1. Limit Order Execution

It remains to describe the execution of limit orders. In our setting, a limit order can execute in one of the following two ways:

1. We assume that there are impatient buyers who arrive to the market according to a Poisson process with rate μ . Denote by $(N_t)_{t \in [0, T]}$ the cumulative arrival process for impatient buyers. Each impatient buyer seeks to buy a single share. An arriving impatient buyer arriving at time t has a reservation price $S_t + z_t$, expressed as a premium $z_t \geq 0$ above the bid price S_t that the buyer is willing to forgo in order to achieve immediate execution. We assume that the premium z_t is independent and identically distributed with cumulative distribution function $F: \mathbb{R}_+ \rightarrow [0, 1]$. In this setting, the instantaneous arrival rate of impatient buyers at time t willing to pay a limit order price of L_t is given by

$$(3) \quad \lambda(u_t) \triangleq \mu(1 - F(u_t)),$$

where $u_t \triangleq L_t - S_t$ is the instantaneous price premium of the limit order. In what follows, we will be particularly interested in the special case where

$$(4) \quad \lambda(u_t) \triangleq \begin{cases} \mu & \text{if } u_t \leq \delta, \\ 0 & \text{otherwise.} \end{cases}$$

Here, we assume that every impatient buyer is willing to pay a price premium of at most $\delta > 0$. We assume that δ will be specific to the security and fixed for the trading horizon. We will discuss the extension to the general case (3) in Section 4.4.

Given (4), an impatient buyer is willing to buy 1 share at a fixed premium $\delta > 0$ to the bid price at the time of their arrival. Hence, if a buyer arrives at time $\tau \in [0, T]$, and the trader has placed a limit order with price L_τ , the limit order will execute if $L_\tau \leq S_\tau + \delta$.

2. Alternatively, a limit order will also execute at time τ if the bid price crosses the limit order price, i.e., $S_\tau \geq L_\tau$.

The execution of limit orders in the model is illustrated in Figure 1.

The limit order execution dynamics above can also be economically interpreted in the spirit of the non-informational trade model of Roll (1984). In particular, imagine that the asset has a fundamental value V_t at time t , and that V_t evolves exogenously according to the additive random walk

$$V_t = V_0 + \sigma B_t.$$




Figure 1: An illustration of the limit order execution in the stylized model over the time horizon $[0, T]$. Here, we assume the trader leaves a limit order with the (constant) price L_t and S_t is the bid price process. If market orders arrive at times τ_1 and τ_2 , the limit order would execute at time τ_2 but not time τ_1 , since the limit order price is in excess of δ to the best bid price. The limit order would also execute at time τ_3 in the absence of a market order arrival, since the bid price crosses the limit order price at this time.

If all investors observe this underlying value process and are symmetrically informed, competitive market makers will always be willing to sell shares at a price of $\delta/2$ above the fundamental value or buy shares at a spread of $\delta/2$ below the fundamental value. Here, the quantity δ captures the per share operating costs of trade to the market makers. The liquidating trader can thus sell at the bid price $S_t = V_t - \delta/2$ at any time t . We assume that all other traders in the market are impatient, and that these traders arrive according to the Poisson dynamics described above. An arriving impatient buyer will choose to purchase from the liquidating trader only at a price lower than that provided by the market makers, i.e., only below the price of $V_t + \delta/2 = S_t + \delta$. In this way, we can interpret the parameter δ as the *prevailing bid-offer spread*, that is, the bid-offer spread in the absence of the liquidating trader.

2.2. Optimal Solution

Let P denote the random variable associated with the sale price. We assume the trader is risk-neutral and seeks to maximize the expected sale price. Equivalently, we assume the trader seeks to solve the optimization problem

$$(5) \quad \bar{h}_0 \triangleq \text{maximize } \mathbb{E}[P] - S_0.$$

Here, the maximization is over policies of market orders and limit orders which are non-anticipating, i.e., which are adapted to the filtration generated by $(B_t, N_t)_{t \in [0, T]}$. This objective is equivalent to minimizing implementation shortfall (Perold, 1988).

Note that, while this stylized problem may seem quite simplified, it seeks to answer a fundamental question: at the level of an atomic unit of stock and over a short time horizon, how should

an risk-neutral investor choose between limit orders and market orders? This problem is a central ingredient in more sophisticated optimal execution problems involving risk averse investors selling large quantities over longer time horizons.⁷ This is because, in a typical algorithmic trading setting, a large “parent” order will be scheduled across time into many very small “child” orders. Each of these “child” orders need to be executed optimally. Since each child order is small and since there are many such child orders, it is reasonable to view the investor as risk-neutral with respect to each child order.

The following lemma characterizes a simple strategy that is optimal for the execution problem we have described:

Lemma 1. *An optimal strategy is to employ only limit orders at times $t \in [0, T)$, with limit price $L_t = S_t + \delta$. In other words, the limit order price is be “pegged” at a constant premium δ above the bid price. This pegging strategy achieves the optimal value*

$$(6) \quad \bar{h}_0 = \delta \left(1 - e^{-\mu T}\right).$$

Proof. Consider a trader using an arbitrary strategy, and denote by $\tau \in [0, T]$ the (random) time at which the trader sells the share, and by $\tau_1 \in [0, \infty)$ the time at which the first impatient buyer arriving to the market. Let \mathcal{E} be the event that the trader sells via a limit order to an impatient buyer at the price L_τ . Then, under the event \mathcal{E}^c , the trader sells at the bid price S_τ . Then, the sale price P can be written as⁸

$$(7) \quad P = S_\tau \mathbb{I}_{\mathcal{E}^c} + L_\tau \mathbb{I}_{\mathcal{E}} \leq S_\tau \mathbb{I}_{\mathcal{E}^c} + (S_\tau + \delta) \mathbb{I}_{\mathcal{E}} \leq S_\tau + \delta \mathbb{I}_{\{\tau_1 < T\}}.$$

Here, for the first inequality, we used the fact that an impatient buyer will only buy at time τ is $L_\tau \leq S_\tau + \delta$, and, for the second inequality, we used the fact that the event \mathcal{E} can only occur if an impatient buyer arrives in the time interval $[0, \tau)$. Denote by \bar{h}_0 the value under an optimal strategy. Using the fact that τ is a bounded stopping time and the fact that S_t is a martingale, by the optional sampling theorem,

$$\bar{h}_0 \leq \mathbb{E}[P] - S_0 \leq \mathbb{E}[S_\tau + \delta \mathbb{I}_{\{\tau_1 < T\}}] - S_0 = \delta \mathbb{P}(\tau_1 < T) = \delta \left(1 - e^{-\mu T}\right).$$

On the other hand, the hypothesized strategy results in equality in (7). Thus, the result follows. ■

The optimal pegging strategy suggested by Lemma 1 is illustrated in Figure 2. This policy can be interpreted intuitively as follows: since the trader is risk-neutral and the bid price process is a martingale, the trader is indifferent between trading at time 0 at the bid price or trading at any other time at the bid price. Via a limit order, however, the trader can receive a price which is in

⁷For example, see Bertsimas and Lo (1998) or Almgren and Chriss (2001). These questions have also recently been addressed by Back and Baruch (2007) and Pagnotta (2010) in equilibrium settings.

⁸We denote by $\mathbb{I}_{\mathcal{E}}$ the indicator function of the event \mathcal{E} .




Figure 2: An illustration of an optimal strategy with no latency, over the time horizon $[0, T]$. The trader uses only limit orders prior to end of the time T . The limit order price L_t is pegged to the bid price S_t , with an additional premium corresponding to the bid-offer spread δ .

excess of the bid price. The excess premium is limited to δ , since an impatient buyer will not pay more than this. Hence, the trader maintains a single limit order in the book, and continuously updates the price to track bid price, plus an additional premium of δ .

3. A Model for Latency

The optimal policy for the stylized execution problem of Section 2 relied on the ability of a trader to continuously track an informational process, namely, the bid price in the market, and to update his order as the process evolves. Here, we will consider a variation of that problem where the trader is unable to continuously participate in the market, but faces a fixed latency $\Delta t > 0$. We are interested in quantifying the cost of this latency by comparing the expected payoff in this model to that in the stylized model without latency. Note that the model at hand is quite basic with regards to some of primitives (e.g., the stochastic process describing the evolution of bid prices), we will discuss a number of tractable extensions in Section 4.4, including more complicated models of the bid price process and of limit order execution.

In general, latency that a trader experiences can take many forms. Minimally, for example, there is the delay of the data feeds that deliver market price information to the trader. There is the delay of the trader's own decision making. Finally, there is the delay of the trader's resulting order reaching the marketplace. We assume that the trader makes decisions instantaneously — we will see that this is reasonable since the optimal decision rule for the trader will take a very simple form. Further, from the trader's perspective, the roundtrip delay (the total delay for an order to be processed by an exchange and reflected in the data feeds observed by the trader) cannot be decomposed into a delay to the exchange and a delay from the exchange. Hence, without loss of generality, we will assume that the trader is able to observe market price information with no delay or latency,⁹ but that the trader's orders experience a latency Δt before they are processed by the

⁹Equivalently, we can assume that our definition of time corresponds to the trader's clock.




Figure 3: An illustration of the model of latency. Here, the time horizon $[0, T]$ is divided into n slots, each of duration equal to the latency Δt . The limit order price ℓ_i is decided at the start of the i th time slot, i.e., at time T_i . This price only takes effect Δt units of time later, and is active during the subsequent time interval $[T_{i+1}, T_{i+2})$.

exchange. This latency is meant to capture, for example, networking or routing delays that are specific to the trader, and that might be reduced through colocation or additional investment in networking technology.

In our latency model, we consider an investor who maintains a limit order to sell one share over the time horizon $[0, T]$ (the possibility of market orders will be discussed shortly), so that once the limit order is executed, the investor immediately exits the market. The time horizon $[0, T]$ is divided into n slots each of length Δt , i.e., $T = n\Delta t$. For each $i \in \{0, 1, \dots, n\}$, define $T_i \triangleq i\Delta t$.

At each time T_i , based on all information observed thus far, we assume that the trader can instantaneously decide to update the limit order with a new price ℓ_i . Due to a latency of Δt , the updated price does not reach the market and take effect until the beginning of the next time slot, i.e., T_{i+1} . This limit order price remains active until time T_{i+2} , at which point it is superseded¹⁰ by the next price ℓ_{i+1} . This sequence of events is illustrated in Figure 3. Between the time T_i , when the price ℓ_i is decided, and the time T_{i+1} , when the updated order reaches the market, the following events can occur:

- $\mathcal{E}_i^{(1)}$: An impatient buyer arrives in the time interval (T_i, T_{i+1}) and $\ell_{i-1} \leq S_{T_i} + \delta$, i.e., the *prior* limit price ℓ_{i-1} , which is active at that time, is within a margin δ of the bid price at the start of the interval. In this case, the limit order executes at the price ℓ_{i-1} , and the investor leaves the market. Note that the updated limit price ℓ_i never takes effect.

We assume that the probability that an impatient buyer arrives in any given time slot is $\mu\Delta t$, and that these arrivals occur independently of everything else.¹¹ We assume that $\Delta t < 1/\mu$ so that this probability is well-defined. The bid price process evolves according to the random walk (2).

¹⁰In practice, this ordering scheme might be achieved a sequence of cancel-and-replace limit orders, each of which cancels the prior limit order, and inserts a new limit order with the updated price. If the prior limit order has already been filled when a subsequent cancel-and-replace order arrives, the new order will fail. Hence, the investor is guaranteed to sell at most one share.

¹¹Note that this is simply a discrete-time Bernoulli arrival process that is analogous to the the Poisson arrival process of Section 2.

- $\mathcal{E}_i^{(2)}$: Otherwise, if $S_{T_{i+1}} \geq \ell_i$, i.e., the bid price has crossed the order price ℓ_i at the instant the order reaches the market, then the order immediately executes at price $S_{T_{i+1}}$.
- $\mathcal{E}_i^{(3)}$: Otherwise, the limit order price ℓ_i is active over the time interval $[T_{i+1}, T_{i+2})$.

In order to consider the possibility of market orders, we allow the limit price $\ell_i = -\infty$. By picking this price, the trader can guarantee that the bid price at time T_{i+1} will cross the order price, i.e., $S_{T_{i+1}} \geq \ell_i$ with probability 1. Thus, the choice of $\ell_i = -\infty$ corresponds to a certain execution at the bid price $S_{T_{i+1}}$, i.e., a market order. Similarly, the trader can make the decision at time T_i not to trade by setting $\ell_i = \infty$. As in the model of Section 2, if the investor has been unable to sell the share by the end of the time horizon T , the investor is forced to sell via a ‘clean-up’ trade, i.e., a market order at time T . This is accomplished by enforcing the constraint that $\ell_{n-1} = -\infty$, which we will assume implicitly in what follows.

As before, if P is the random variable associated with the sale price, the trader is risk-neutral and seeks to solve the optimization problem

$$(8) \quad h_0(\Delta t) \triangleq \underset{\ell_0, \dots, \ell_{n-1}}{\text{maximize}} \mathbf{E}[P] - S_0.$$

Here, the maximization is over the choice of limit order prices $(\ell_0, \ell_1, \dots, \ell_{n-1})$. We assume that the price decisions are non-anticipating, i.e., each ℓ_i is adapted to the filtration generated by the bid price process and the arrival of impatient buyers up to and including time T_i . Our goal is to analyze $h_0(\Delta t)$, which is the value under an optimal trading strategy when the latency is Δt .

Note that, as compared to the model of Section 2, our present model with latency differs in two ways: First, the trader makes decisions at the beginning of discrete-time intervals of length Δt , as opposed to continuously. Second, the orders of the trader incur a latency or delay of length Δt before they reach the marketplace. We are interested in studying the impact of the latter feature, latency, and we adopt the former feature, discrete-time decision making, so as to admit a tractable dynamic programming analysis. In Section 4.3, however, we will see that in the low latency regime in which we are most interested, the discrete-time nature of our model has a negligible impact.

4. Analysis

In this section, we solve for the optimal policy for the trader in the latency model of Section 3. This problem can be solved via a dynamic programming decomposition that is presented in Section 4.1. While the exact dynamic programming solution can be computed numerically, in Section 4.2 we will present an asymptotic analysis that provides a closed-form analytic expression for the cost of latency in the low latency regime, where $\Delta t \rightarrow 0$. In Section 4.3, we will consider the implications of the discrete-time nature of our latency model. Finally, in Section 4.4, we will discuss a number of extensions of our latency model.

4.1. Dynamic Programming Decomposition

The standard approach to solving the optimal control problem (8) is to employ dynamic programming arguments. In Appendix A, we formally derive the optimal control policy using these methods. In order to focus on the high level picture, however, for the moment we will be content with summarizing those results.

In particular, assume a fixed latency of Δt . For each decision time T_i with $0 \leq i < n$, define \mathcal{U}_i to be the event that the trader's limit order remains unfulfilled prior to time T_{i+1} , i.e., none of the orders submitted at prices $\ell_0, \dots, \ell_{i-1}$ are executed. Note that if the event \mathcal{U}_i *does not* hold, then the limit order price ℓ_i to be decided at time T_i is irrelevant. This is because, by the time that order arrives to the market, the trader would have already sold a share. Define the quantity

$$(9) \quad h_i \triangleq \underset{\ell_i, \dots, \ell_{n-1}}{\text{maximize}} \mathbb{E}[P \mid S_{T_i}, \mathcal{U}_i] - S_{T_i}.$$

Note that $h_0 = h_0(\Delta t)$, where $h_0(\Delta t)$ is defined in (8), and thus our notation is consistent. More generally, for $i > 0$, we can interpret h_i to be the trader's expected payoff at time T_i relative to the current bid price S_{T_i} under the optimal policy, the order does not get filled prior to time T_{i+1} . Thus, h_i can be interpreted as a *continuation value* in the dynamic programming context.

The continuation values $\{h_i\}$ quantify the remaining value for a trader at each time period if his order remains unfulfilled. Given the continuation values, at each time T_i , the investor can make an optimal decision as to the limit order price ℓ_i by balancing the benefits of execution in the time slot $[T_{i+1}, T_{i+2})$ with the value h_{i+1} that will be obtained if the order is not executed. Moreover, the optimal decisions and continuation values can be jointly computed via backward induction of a Bellman equation. This result is captured in the following theorem. The proof, which is provided in Appendix A, follows from formal dynamic programming arguments.

Theorem 1. *Suppose $\{h_i\}$ satisfy, for $0 \leq i < n - 1$,*

$$(10) \quad h_i = \max_{u_i} \left\{ \mu \Delta t \left[u_i \left(\Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) + \sigma \sqrt{\Delta t} \left(\phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) \right] \right. \\ \left. + h_{i+1} \left[(1 - \mu \Delta t) \Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) + \mu \Delta t \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right] \right\},$$

and

$$(11) \quad h_{n-1} = 0.$$

Here, ϕ and Φ are, respectively, the p.d.f. and c.d.f. of the standard normal distribution. Then, $\{h_i\}$ correspond to the continuation values under the optimal policy.

Suppose further that, for $0 \leq i < n - 1$, u_i^* is a maximizer of (10). Then, a policy which chooses

limit order prices which are pegged to the bid prices according to the premia defined by $\{u_i^*\}$, i.e.,

$$\ell_i^* = S_{T_i} + u_i^*, \quad \forall 0 \leq i < n - 1,$$

is optimal.

Theorem 1 suggests a computational strategy for determining continuation values and an optimal policy. Starting with the terminal condition $h_{n-1} = 0$, one proceeds via backward induction, solving the single variable optimization problem (10) over the decision variable u_i once per time slot. So long as optimal solutions exist, they will determine the continuation values and optimal policy. Moreover, the optimal policy is a pegging strategy. That is, the limit order price is pegged at a deterministic (but time varying) premium above the current bid price. These limit order premia are given by the maximizers $\{u_i^*\}$.

In the following theorem, whose proof is provided in Appendix B, we establish the existence and uniqueness of the optimal solutions to (10) and provide upper and lower bounds for the resulting limit price premia, for small values of latency Δt .

Theorem 2. *Fix $\alpha > 1$. If Δt is sufficiently small, then there exists a unique optimal solution $\{h_i\}$ to the dynamic programming equations (10)–(11). Moreover, the corresponding optimal policy $\{u_i^*\}$ is unique. For $0 \leq i < n - 1$, this strategy chooses limit prices in the range*

$$\ell_i^* \in \left(S_i + \delta - \sigma \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, S_i + \delta - \sigma \sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}} \right),$$

where

$$L \triangleq \frac{\delta^2}{2\pi\sigma^2}, \quad R(\Delta t) \triangleq \frac{\delta^2(1 - \mu\Delta t)^{2n}}{2\pi\sigma^2}.$$

Figure 4 illustrates the intuition behind Theorem 2, by considering the situation of a trader at time $t = 0$, when the bid price is S_0 . In the absence of latency, the trader would peg the limit order price at a fixed premium of δ , i.e., $\ell_0 = S_0 + \delta$. This would result in a trade with the next impatient buyer with probability 1. If there is latency present, however, this limit price is not optimal. To see this, suppose that an impatient trader will arrive at time $\tau_1 \in (\Delta t, 2\Delta t)$. If the limit order price is set at ℓ_0 , the probability that the trade *does not* get executed is

$$\mathbf{P}(\ell_0 \geq S_{\Delta t} + \delta) = \mathbf{P}(S_0 \geq S_{\Delta t}) = 1/2.$$

When Δt is small, the probability of missing an execution can be significantly lowered at a small cost by lowering ℓ_0 by an additional safety margin. If we set this safety margin to be C standard deviations of the one-period price change, i.e., $\ell_0 = S_0 + \delta - C\sigma\sqrt{\Delta t}$, then the probability of missing




Figure 4: An illustration of the optimal policy of Theorem 2. In the absence of latency, at time $t = 0$, the trader would set the limit price at a premium of δ , i.e., $\ell_0 = S_0 + \delta$. In an environment with latency, the trader might set the limit price to be ℓ'_0 , which lowers ℓ_0 by an additional safety margin of C standard deviations. This serves to increase the likelihood of trade execution in the interval $(\Delta t, 2\Delta t)$. The optimal limit price ℓ^*_0 utilizes a safety margin that is slightly larger.

execution becomes

$$\mathbb{P}(\ell'_0 \geq S_{\Delta t} + \delta) = \mathbb{P}(S_0 - C\sigma\sqrt{\Delta t} \geq S_{\Delta t}) = \Phi(-C).$$

This probability can be made close to 0 by the choice of C . However, given a fixed choice of C independent of Δt , the probability remains constant (i.e., independent of Δt) and non-zero. The additional safety margin corresponding to the log term in Theorem 2 is a second order adjustment. This is introduced so that, given the optimal limit price ℓ^*_0 , the probability of execution tends to 1 as $\Delta t \rightarrow 0$.

4.2. Asymptotic Analysis

The dynamic programming decomposition developed in Section 4.1 allows the exact numerical computation of the value $h_0(\Delta t)$, the value under an optimal policy of the latency model introduced in Section 3, when the latency is Δt . As discussed earlier, the latency observed in modern electronic markets is extremely small, often on the time scale of milliseconds. Thus, we are most interested in the qualitative behavior of $h_0(\Delta t)$ in the asymptotic regime where $\Delta t \rightarrow 0$. The main result of this section is the following theorem, whose proof is provided in Appendix C. It provides a closed-form expression for $h_0(\Delta t)$, which holds asymptotically¹² as $\Delta t \rightarrow 0$.

¹²In what follows, given arbitrary functions f and g , and a positive function q , we will say that $f(\Delta t) = g(\Delta t) + O(q(\Delta t))$ if $\limsup_{\Delta t \rightarrow 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) < \infty$, i.e., if the difference between f and g , as $\Delta t \rightarrow 0$, is asymptotically bounded above by *some* positive multiple of q . Similarly, we will say that $f(\Delta t) = g(\Delta t) + o(q(\Delta t))$

Theorem 3. As $\Delta t \rightarrow 0$,

$$h_0(\Delta t) = \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right) + o(\sqrt{\Delta t}),$$

where

$$\bar{h}_0 = \delta \left(1 - e^{-\mu T} \right)$$

is the optimal value for the stylized model without latency, i.e., the value defined by (5).

Theorem 3 is not surprising when considered in the context of Theorem 2. In the stylized model without latency, the optimal strategy is to peg the limit order price at a premium of δ , and this yields a value of \bar{h}_0 . On the other hand, Theorem 2 suggests a trader facing latency Δt will lower this limit price premium by a factor of, approximately,

$$\frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} + o(\sqrt{\Delta t}).$$

If this lowers the ultimate value proportionally, then the value of the optimal policy in the presence of latency Δt should approximately be

$$\bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right) + o(\sqrt{\Delta t}).$$

The proof of Theorem 3, provided in Appendix C, makes this intuition precise.

One implication of Theorem 3 is that $h_0(\Delta t) \rightarrow \bar{h}_0$ as $\Delta t \rightarrow 0$, i.e., the value of the latency model converges to that of the stylized model without latency of Section 2. This suggests the following definition:

Definition 1. Define the *latency cost* associated with latency Δt by

$$(12) \quad \text{LC}(\Delta t) \triangleq \frac{\bar{h}_0 - h_0^*(\Delta t)}{\bar{h}_0}.$$

Latency cost has an easy interpretation. Using \bar{h}_0 , the value obtained in the stylized model without latency as a benchmark, the numerator of (12) is the *lost revenue* incurred due the the presence of latency. On the other hand, we can regard the denominator as the *cost of immediacy* for an impatient investor in a time horizon of length T . This is because, in the stylized model without latency, it is the difference in revenue obtained by a risk-neutral investor willing to patiently provide

if $\lim_{\Delta t \rightarrow 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) = 0$, i.e., if the difference between f and g , as $\Delta t \rightarrow 0$, is asymptotically dominated by *every* positive multiple of q . Finally, we will say that $f(\Delta t) = g(\Delta t) + \Theta(q(\Delta t))$ if $0 < \liminf_{\Delta t \rightarrow 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) \leq \limsup_{\Delta t \rightarrow 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) < \infty$, i.e., if the difference between f and g is asymptotically bounded above and below by positive multiples of q .

liquidity by employing limit orders over the length of the time horizon, and an impatient investor who demands immediate liquidity and sells at the bid price at time $t = 0$, cf. (5). Therefore, we can describe the latency cost as the amount a trader forgoes due to latency, as a percentage of the cost of immediacy.

The following corollary restates the asymptotic approximation of Theorem 3 in terms of latency cost.

Corollary 1. *As $\Delta t \rightarrow 0$,*

$$\text{LC}(\Delta t) = \frac{\sigma\sqrt{\Delta t}}{\delta} \sqrt{\log \frac{\delta^2}{2\pi\sigma^2\Delta t}} + o(\sqrt{\Delta t}).$$

There are a number of interesting observations that can be made regarding the asymptotic approximation of Corollary 1. First of all, asymptotically, latency cost *does not* depend on the length of the time horizon T or the arrival rate of impatient traders μ . As a function of the remaining parameters, the asymptotic latency cost depends only on a composite parameter that is the ratio the one-period standard deviation of price changes $\sigma\sqrt{\Delta t}$ to the bid-offer spread δ . Both of these quantities are readily measurable empirically. Corollary 1 suggests that the latency cost increasing in this ratio. Thus, at the same level of latency, the latency cost is most significant for assets which are very volatile or very liquid. Further, Corollary 1 suggests that, when latency is low, there are increasing marginal benefits to further reductions in latency, i.e., $\text{LC}''(\Delta t) < 0$. In Section 5.1, we illustrate some of facts numerically, as well as considering the accuracy of our approximation, as compared to the exact latency cost.

4.3. Discreteness of Time vs. Latency

The latency model introduced in Section 3 differs from the the stylized model without latency of Section 2 in two principal ways: (i) the trader faces a delay or latency between the time that trading decisions are made and when they reach the marketplace, and (ii) the latency model is formulated in discrete-time rather than continuous time. The latter point refers to the facts that, in the model with latency, a trader is only able to update his limit order at discrete intervals of time rather than continuously, impatient buyers arrive according to a Bernoulli process rather than a Poisson process, etc. In order to disentangle these two effects, in this section we will briefly describe a trading model that is formulated in discrete time but *without* latency. By considering this model, we will demonstrate that the asymptotic latency cost derived in Section 4.2 is indeed due to latency effects and not due to the discreteness of time.

To this end, consider a model in the discrete-time setting of Section 3 but with no latency. Here, at each time $T_i \triangleq i\Delta t$, for $i = 0, 1, \dots, n$, the investor sets a limit order price ℓ_i . This limit order price takes effect immediately. Between time T_i and time T_{i+1} the following events can occur:

- If $S_{T_i} \leq \ell_i$, i.e., the bid price is less than the limit order price, the limit order immediately

executes at the price S_{T_i} .

- Otherwise, suppose that an impatient buyer arrives in the time interval (T_i, T_{i+1}) and $\ell_i \leq S_{T_i} + \delta$, i.e., the limit price ℓ_i is within a margin δ of the bid price at the start of the interval. In this case, the limit order executes at the price ℓ_i . We assume that an impatient buyer arrives with probability $\mu\Delta t$, independent of everything else.

As before, if the investor is unable to sell the share by the end of the time interval, he is forced to sell via a market order, i.e., $\ell_n = -\infty$. If P is the sale price, the optimal value for the trader in this discrete model is given by

$$h_0^D(\Delta t) \triangleq \underset{\ell_0, \dots, \ell_n}{\text{maximize}} \mathbf{E}[P] - S_0.$$

We have the following result, whose proof is identical to the martingale argument used to establish Lemma 1.

Lemma 2. *An optimal strategy for the discrete model is to place limit orders at the price $\ell_i = S_{T_i} + \delta$, for $i = 0, 1, \dots, n - 1$. This strategy achieves the value*

$$h_0^D(\Delta t) \triangleq \delta(1 - (1 - \mu\Delta t)^n).$$

Now, note that, for all $0 < \Delta t < 1/\mu$,

$$e^{-\mu T - \frac{1}{2}\mu^2 T \Delta t} \leq (1 - \mu\Delta t)^{T/\Delta t} \leq e^{-\mu T}.$$

Therefore, the difference in value between the continuous model of Section 2 and the discrete model considered here is at most

$$|h_0^D(\Delta t) - \bar{h}_0| \leq \delta e^{-\mu T} \left(1 - e^{-\frac{1}{2}\mu^2 T \Delta t}\right) \leq \frac{1}{2}\delta \mu^2 T e^{-\mu T} \Delta t.$$

In other words, this difference is asymptotically $O(\Delta t)$. By Theorem 3, however, the difference between the continuous model and the latency model is asymptotically

$$\Theta(\sqrt{\Delta t \log(1/\Delta t)}).$$

Hence, the asymptotic effect of latency dominates the asymptotic effect of the discreteness of time.

4.4. Extensions

The analysis of the latency model that we have presented proceeded according to two high level steps:

- (i) First, in Section 4.1, a simplified dynamic programming decomposition was developed. In this decomposition, at each time, the trader's value function is parameterized by a single scalar, rather than being an arbitrary function of state. This allows the Bellman equation to be solved through a system of n equations in n unknowns, given by (10)–(11).
- (ii) Second, in Section 4.2, an asymptotic analysis of the simplified dynamic programming equations (10)–(11) was performed. This gave rise to the asymptotic latency cost expression of Corollary 1.

The dynamic programming decomposition step (i) that is at the heart of our analysis can be extended to a much broader set of stochastic primitives than the present setting. In each of these cases, a different set of simplified dynamic programming equations, analogous to (10)–(11) would arise, and would require a customized variation of asymptotic analysis step (ii). In particular, consider the following tractable generalizations:

- **Bid price process.** In our model, the bid price process S_t is a Brownian motion. Our dynamic programming decomposition only requires that the S_t be a Markov process and a martingale. It would be straightforward to consider other Markovian martingales, for example, allowing for non-Gaussian processes, time-inhomogeneous volatility, or for jump processes.
- **Limit order execution.** In our model, the execution of a limit order in the time slot (T_i, T_{i+1}) required that the limit order price ℓ_{i-1} be within a spread δ of the bid price S_{T_i} , and that an impatient trader arrive. More generally, our dynamic programming decomposition only requires that the execution of this limit order, conditional on the price difference $\ell_{i-1} - S_{T_i}$, be independent of everything else. This can accommodate a number of generalizations, for example, the arrival rate of impatient buyers can be time-varying. Further, the maximum premium above the bid price S_t that an impatient buyer is willing to pay can be randomly distributed, as in (3). This would allow models where an limit order that is priced aggressively low has a much higher probability of execution. Such models could alternatively be interpreted, as discussed in Section 2, as cases where the prevailing bid-offer spread is not constant, but is independent and identically distributed, varying from period to period.

5. Empirical Estimation of Latency Cost

In this section, we will consider empirical applications of our model. First, we will illustrate the optimal trading policy and the corresponding value function when the model parameters are estimated from high frequency market data for a single stock. We will also compare the exact latency cost (numerically computed via dynamic programming) to the approximation provided by Corollary 1 in order to assess the quality of our approximation. Subsequently, we show the historical

evolution of latency cost and implied latency across a range of U.S. equities using cross-sectional data on volatilities and bid-offer spreads during the 1995–2005 period.

Our empirical analysis should be regarded as a first-order study to obtain a rough calibration of our model. It will allow us to analyze the model in relevant parameter regimes, as well as gaining a broad understanding the implications of our model for the trading of U.S. equities. Our empirical measurement of latency cost requires estimates of, in particular, the high frequency price volatility σ and the prevailing bid-offer spread δ . Here, we make a number of simplifications and rely on the recent empirical work of Ait-Sahalia and Yu (2009) to obtain these quantities:

- We estimate price volatility σ using the maximum likelihood estimates of the volatility of returns provided by Ait-Sahalia and Yu (2009). Note that this estimation of high frequency volatility aims to filter out the impact of microstructure noise and obtain an unbiased estimate of daily volatility. However, for an investor with a trading horizon of 1 second, microstructure noise needs to be incorporated as well. Therefore, the high frequency volatility estimate that is used in our empirical analysis underestimates the actual volatility faced by a high frequency trader with a very short trading horizon.
- Recall that the prevailing bid-offer spread, δ , equals the bid-offer spread in the absence of the liquidating trader. In the empirical data, it is impossible to disentangle the presence of liquidating traders. Moreover, the bid-offer spread will not be constant, but will vary over the course of the trading day. As a proxy for δ , we use the average bid-offer spread over the trading day.

Despite these shortcomings, we believe that our empirical analysis provides interesting, high-level conclusions, and sheds light on the importance of latency in the trading of U.S. equities.

5.1. The Optimal Policy and the Approximation Quality

In what follows, we will numerically evaluate the optimal policy in our model, the corresponding value function, and the latency cost approximation. These numerical experiments are meant to be illustrative of our model. We will use realistic model parameters estimated from recent market data for a single stock. Our methodology here is not meant to be authoritative — there are many subtleties in the analysis of high frequency data; these are beyond the scope of the work at hand. However, we do seek to demonstrate that our model parameters can be readily derived from commonly available data.

Specifically, the model parameters herein are estimated from trade-and-quote (TAQ) data for a stock that is a representative example of a liquid name, Goldman Sachs Group, Inc. (NYSE: GS), on the trading day of January 4, 2010. This data was obtained from the Wharton Research and Data Services (WRDS) consolidated TAQ database. Only trades and quotes originating from the primary exchange (NYSE) during regular trading hours were considered. The model parameters were estimated as follows:




Figure 5: An illustration of the optimal strategy for GS, expressed in terms of limit price premium over the course of the time, for different choices of latency. In each case, the dashed line illustrates the relative distance below the bid-offer spread δ of the price premium of the final limit order, as a multiple of the standard deviation of prices over the latency interval.

- Initial bid price: $S_0^{\text{GS}} = \$170.00$. This was chosen to be the first transaction price on the trading day.
- Bid-offer spread: $\delta^{\text{GS}} = \$0.058$, i.e., equivalently, 3.4 basis points relative to the initial price S_0^{GS} . This was estimated by computing the average spread between bid and offer quotes over the course of the trading day and rounding to the nearest cent.
- Arrival rate of market orders: $\mu^{\text{GS}} = 12.03$ (per minute). This was estimated by dividing the total number of NYSE trades by the length of the trading day.
- Price volatility: $\sigma^{\text{GS}} = \$1.92$ (daily), i.e., approximately equivalent to an annualized volatility of returns of 17.9%. These were estimated from the time series of transaction prices over the course of the trading day, using maximum likelihood estimation as described in Ait-Sahalia and Yu (2009).
- Trading horizon: $T = 10$ (seconds).

Figure 5 illustrates the optimal limit order policy for GS under different values of latency. If there is no latency, the limit orders are submitted at a constant premium of δ . When there is latency, the optimal order policy is obtained using the exact dynamic programming solution of (10)–(11). As the latency increases, the limit order premium is reduced below δ so as to account for the increasing uncertainty of price movements over the latency interval. Theorem 2 suggests




Figure 6: An illustration for the evolution of the continuation value of the optimal policy over time for GS, for different choices of latency. The expected value of the trader decreases as latency increases or as the end of the trading horizon approaches. As the latency increases from 0 ms to 500 ms, the trader loses more than 0.01 of the 0.05 cent spread, i.e., more than 20% of the spread.

that this reduction is approximately equal to

$$(13) \quad \sigma \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}}.$$

In Figure 5, we see that, with a latency of 500 ms, this adjustment is up to approximately $1.4\sigma\sqrt{\Delta t}$, i.e., 1.4 times the standard deviation of prices over the latency interval. When the latency is reduced to 250 ms and to 50 ms, the adjustment increases to 1.6 and 2.1 standard deviations, respectively. The fact that this adjustment, when measured as a multiple of the uncertainty over the latency period, increases as the latency decreases is consistent with (13).

In Figure 5, we also observe that as t increases and the trading deadline approaches, the limit order premium u_t^* becomes lower. This makes intuitive sense: the trader faced with a terminal value of 0 since he is required to sell using market order at the end of the period. As the deadline approaches, the trader is more willing to sacrifice the potential profits of a limit order in order to increase the probability of execution.

Figure 6 illustrates the corresponding continuation value under the optimal policy for GS, for different values of latency. Clearly, the trader's expected payoff decreases as latency increases or the end of the trading horizon approaches.

Finally, Figure 7 illustrates the latency cost as a function of latency. Both the exact value of the latency cost, computed numerically via the dynamic programming decomposition (10)–(11), and the asymptotic latency cost approximation provided by Corollary 1 are shown. The latency costs decrease from approximately 20% of the cost of immediacy to 5% of the cost of immediacy, as the latency decreases from 500 ms to 5 ms. Further, the marginal benefit of reducing latency




Figure 7: An illustration of the latency cost as a function of the latency. both the exact latency cost and the asymptotic approximation are shown. The approximate latency cost closely aligns with the exact latency cost across the entire range of latency values. This illustrates that our closed-form formula can accurately approximate the exact latency cost for low values of latency.

increases as the latency approaches zero. Finally, we note that the approximate and exact latency costs are quite close across the entire range of latency values. This suggests that the approximation is of very high quality in this case.

5.2. Historical Evolution of Latency Cost

In this section, we will examine the historical evolution of latency cost in U.S. equities. Here, we consider the situation of a hypothetical investor with a fixed latency of 500 milliseconds. This choice of latency is made to approximately corresponds to reaction time of a very fast human trader. We will use this as a proxy for the fastest possible trading on a “human time scale”. By analyzing the evolution of the associated latency cost, we will get a sense of the importance of latency over time.

Our empirical analysis relies on the data set of Ait-Sahalia and Yu (2009). Their data set contains estimates for various liquidity measures for all NYSE common stocks on a daily basis during the sample period of June 1, 1995 to December 31, 2005. The estimates are derived from intraday transaction prices and quotes from the NYSE TAQ database. We utilized only the volatility and bid-offer spread data as we have seen both analytically (Corollary 1) and numerically (Figure 7) that latency cost can be approximated accurately for low values of latency using only these two measures.

The data set contain volatility and bid-offer spread estimates for given stock on a particular day if the number of transactions on that day exceeds 200. The minimum, average, and maximum number of stocks in the sample on any day are 61, 653, and 1,278, respectively. In particular, earlier periods in the data set contain fewer stocks due to a smaller number of firms and a lower volume

of transactions. In this data set, the bid-offer spread is estimated using only NYSE quotes in the regular trading hours. The volatility estimate is obtained using maximum likelihood estimation in the presence of market microstructure noise. Maximum likelihood estimation is preferred over other nonparametric estimation methods (e.g., “Two Scales Realized Volatility”) as a simulation study shows that maximum likelihood estimation provides robust estimators under reasonable stochastic volatility and jump models in the underlying asset. The reader is urged to consult to Section 2.1 of Aït-Sahalia and Yu (2009) for full details of their estimation procedure.

For each stock in the data set, on a daily basis, we compute the latency cost facing an investor with a fixed latency of 500 ms using the asymptotic approximation of Corollary 1. These daily latency costs are then averaged over each month. Figure 8 displays percentiles of the monthly averages of latency cost over all of the stocks in the sample, as a function of time. As a representative example of a liquid name, we also report the monthly averages of latency cost of Goldman Sachs Group, Inc. (NYSE: GS). Note that the time series for GS begins from its initial public offering in 1999. For reference, we have added an additional point to this time series based on our estimation in Section 5.1 of the latency cost for GS on January 4, 2010.

Figure 8 illustrates that latency costs have had an increasing trend over the 1995–2005 period. In particular, we observe that the median latency cost incurred by trading on a human time scale roughly tripled, by increasing from approximately 3% to approximately 10%. One important factor in this increase has been the reduction of bid-offer spreads over this time period. Instances during the period when the NYSE reduced the tick size (from $\$1/8$ to $\$1/16$ in June 1997, and from $\$1/16$ to $\$0.01$ in January 2001) coincide with spikes in latency cost. This is consistent with bid-offer spreads decreasing significantly and volatility maintaining the same level at these times. This suggests that any future reduction in tick sizes will result in increased latency costs.

Using a data set in a similar time-frame, from February 2001 to December 2005, Hendershott et al. (2010) conclude that in the post-decimalization era, the increase in algorithmic trading activity had a positive impact on the level of liquidity. This result suggests that the increase in algorithmic trading in and of itself elevated the importance of low latency trading and increased the cost of latency.

5.3. Historical Evolution of Implied Latency

An alternative perspective on the historical importance of latency comes from considering a hypothetical investor with a target level for the cost of latency, relative to the overall cost-of-immediacy. The representative trader maintains this target over time through continual technological upgrades to lower levels of latency. We determine the requisite level of latency for such a trader, over time and across the aggregate market. In other words, fixing the latency cost percentage LC to the target level, we can solve the asymptotic approximation (12) for the level of latency required at each time to achieve latency cost LC . We call this the *implied latency*.

Figure 9 illustrates the implied latency values over the 1995–2005 period assuming that the

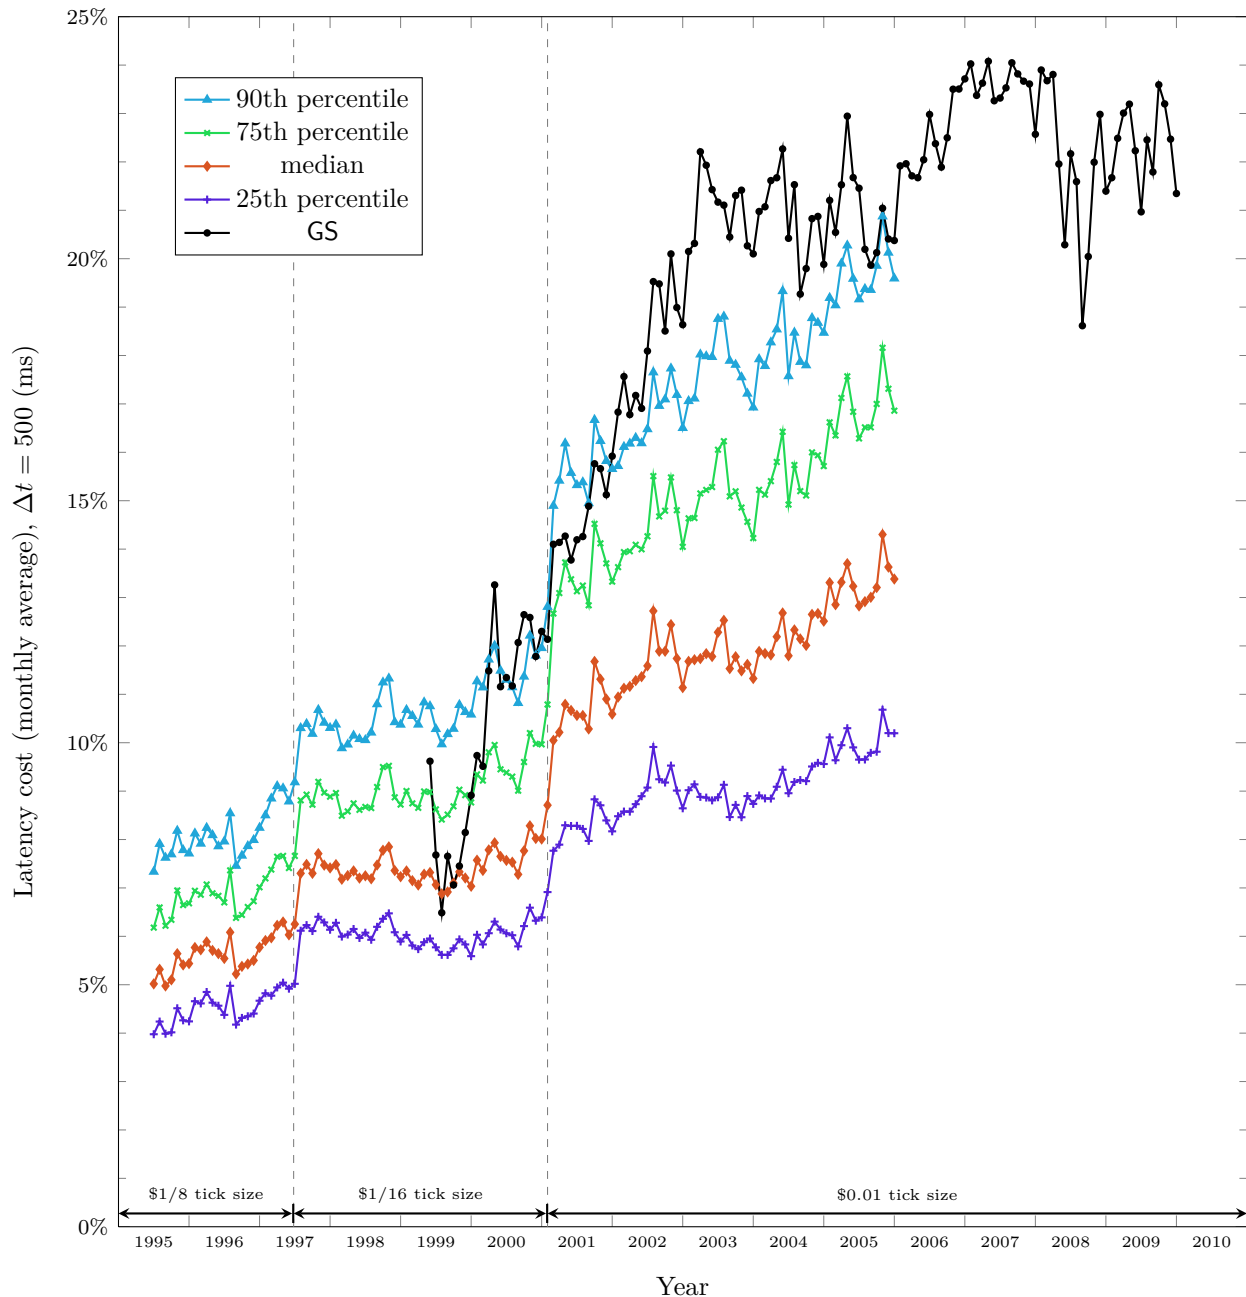


Figure 8: An illustration of the historical evolution of latency cost over the 1995–2005 time period. Here, we consider a hypothetical “human time scale” investor with a fixed latency of $\Delta t = 500$ (ms). Percentiles for the resulting latency cost are reported across NYSE common stocks. The latency costs are computed from data set of Ait-Sahalia and Yu (2009). The latency cost for GS is also reported, beginning from its IPO. The dashed lines correspond to dates where the NYSE tick size was reduced. We observe that latency cost had a consistent increasing trend over the 1995–2005 period. Specifically, the median latency cost approximately increased three-fold by reaching roughly to 10% from 3%.




Figure 9: An illustration of the historical evolution of implied latency over the 1995–2005 time period. Here, we consider a hypothetical investor who makes sufficient technological investments to ensure a constant latency cost of 10%. The implied latency is the level of latency required to achieve this latency cost. Percentiles for the implied latency are reported across NYSE common stocks. The implied latencies are computed from data set of Ait-Sahalia and Yu (2009). The implied latency for GS is also reported, beginning from its IPO. We observe that implied latency has had a decreasing trend over the 1995–2005 period. Specifically, the median implied latency decreased by approximately two orders of magnitude over this time frame.

target level $LC = 10\%$ of overall transaction costs result from latency. We observe that the median implied latency decreased by approximately two orders of magnitude over this time frame. The 90th percentile of U.S. equities, for example, went from an implied latency on the scale of seconds to an implied latency on the scale of tens of milliseconds.

5.4. Empirical Importance of Latency

Our model captures the cost of latency due to a lack of contemporaneous information. Figure 8 suggests that, when our model is calibrated to the topmost quartile of U.S. equities, a investor with latency on the human time scale faces a latency cost of at 15% to 25%. In order to assess the significance of this, we can compare it to other trading costs. Suppose we normalize the cost of immediacy to \$0.01, which is the typical bid-offer spread for a liquid U.S. equity. Then, our model suggests that the benefit of reducing latency from a human time scale of 500 ms to an ultra low latency time scale of less than 1 ms is approximately \$0.0015–\$0.0025 per share traded.

While this might seem very small as an absolute number, note that is of the same order of magnitude as other trading costs faced by the most cost efficient institutional investors. For example, a hedge fund would pay an average commission of \$0.0007 per share for market access.¹³ Furthermore, investors may pay an SEC fee of \$0.0005 per share traded,¹⁴ and exchange fees or rebates of \$0.0020–\$0.0030 per share traded. To the extent that a sophisticated institutional investor is cost sensitive and wishes to optimize these other execution costs, they should also be concerned with latency. This isn't to suggest that latency cost is important to all investors. A typical retail investor, for example, may pay a brokerage fee that is up to \$0.10 per share traded.¹⁵ For this latter type of investor, the cost of latency as described here is not a significant component of overall trading costs.

Alternatively, we can compare the \$0.0015–\$0.0025 per share traded latency cost to the rents extracted by agents that have made the required technological investments to trade on an ultra low latency time scale. For example, providers of automated algorithmic trade execution services charge an average commission of \$0.0033 per share traded for their execution services, which leverage sophisticated low latency technological infrastructure.¹⁶ Note that this cost is comparable to the latency cost. Another class of agents with ultra low latency trading capabilities are high frequency traders. Reported net profit numbers for high frequency traders are in the range of \$0.0010–\$0.0020

¹³“U.S. Equity Trading: Low Touch Trends,” TABB Group, July 2010.

¹⁴As of January 21, 2011, the SEC fee is a fraction \$0.0000192 of the proceeds of an equity sale. If we assume a typical stock price of \$50, this is approximately \$0.0010 per share sold. Amortizing this cost equally between buys and sells results in \$0.0005 per share traded.

¹⁵For example, at the time of writing, the brokerage firm E-TRADE charges \$10 per trade. Assuming a typical trade of 100 shares, this cost is \$0.10 per share traded.

¹⁶“U.S. Equity Trading: Low Touch Trends,” TABB Group, July 2010. Note that some institutional investors pay significantly larger commissions for trade execution in order to compensate their brokers for trading ideas or research services. The commission we quote here is for “non-idea driven” services that relate purely to trade execution using the algorithms and technological platform of the broker.

per share traded.¹⁷ This is of the same order of magnitude as the latency cost.

6. Conclusion and Future Directions

This paper provides a highly tractable model to quantify the cost of latency on transaction costs. We consider a stylized execution problem, where a trader must sell an atomic unit of stock over a fixed time horizon. We consider this model in the absence of latency as a benchmark, and we incorporate latency by not allowing the trader to continuously participate in the market. Orders submitted by the trader reach the market with a fixed latency, and the trader is forced to deviate from the benchmark policy in order to take into account the uncertainty introduced by this delay. We quantify the cost of latency as the normalized difference in expected payoffs between this model and the stylized model without latency.

Since the latency values observed in modern electronic markets are on the order of milliseconds, we provide an asymptotic analysis for the low latency regime, in which we obtain an explicit closed-form solution. In order to compute this asymptotic latency cost empirically, we only need to estimate the volatility and the average bid-offer spread of the stock. This is an elegant and practical result as data sets and estimation procedures for these quantities are readily abundant in the literature. Indeed, using an existing data set, we show that the cost of latency incurred by trading on a human time scale (500 ms) increased three-fold over the 1995–2005 time-frame. In addition, using the alternative approach of keeping a fixed level of latency cost through continuous technological improvements, we compute the various percentiles of the implied latency over this time frame. Using the same data set, we observe that the median implied latency decreased by approximately two orders of magnitude.

Our empirical analysis can also be utilized to compare the magnitude of latency cost to other trading costs incurred by institutional investors. Our results suggest that the difference in payoff between trading with a human time scale (500 ms) and an automated trading platform with ultra low latency (1 ms) is approximately of the same order of magnitude as other trading costs faced by institutional investors. This observation certainly underlines the significance of latency for such investors. In conclusion, our model is the first theoretical approach in the literature to concretely quantify the impact of latency on the optimal order submission policy and its resulting cost to the trader.

There are a number of interesting future directions for research. First, as discussed in Section 4.4, there are a number of tractable extensions to the present model that can be analyzed. One particularly interesting case would be where the bid price process is a jump process. Here, our suspicion is that the cost of latency would *decrease*. This is because, even in the absence of any latency, the trader cannot adjust his limit prices ahead of a jump.

More generally, in the introduction, we identified a number of broad themes to the costs that

¹⁷“Tradeworx, Inc. Public Commentary on SEC Market Structure Concept Release,” Tradeworx, Inc., April 2010.

arise from latency. The model we have presented captures mainly costs due to a lack of contemporaneous decision making. It does not capture the latency costs due to strategic effects (i.e., comparative advantage/disadvantage relative to other investors) or due to time priority rules. These remain important questions for future research.

References

- Y. Ait-Sahalia and J. Yu. High frequency market microstructure noise estimates and liquidity measures. *The Annals of Applied Statistics*, 3(1):422–457, 2009.
- R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- J. Angel. Limit versus market orders. Working Paper No. FINC-1377-01-293, School of Business Administration, Georgetown University, Washington, D.C., 1994.
- J. Bacidore, R. Battalio, and R. Jennings. Order submission strategies, liquidity supply, and trading in pennies on the New York Stock Exchange. *Journal of Financial Markets*, 6:337–362, 2003.
- K. Back and S. Baruch. Working orders in limit order markets and floor exchanges. *The Journal of Finance*, 62(4):1589–1621, 2007.
- D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.
- D. Bertsimas, L. Kogan, and A. W. Lo. When is time continuous? *Journal of Financial Economics*, 55:173–204, 2000.
- P. P. Boyle and D. Emanuel. Discretely adjusted option hedges. *Journal of Financial Economics*, 8:259–282, 1980.
- J. Brogaard. High frequency trading and its impact on market quality. Working paper, 2010.
- G. Cespa and T. Foucault. Insiders-outsiders, transparency, and the value of the ticker. Working Paper No. 628, Department of Economics, Queen Mary, University of London, April 2008.
- G. C. Chacko, J. W. Jurek, and E. Stafford. The price of immediacy. *Journal of Finance*, 63(3):1253–1290, 2008.
- T. Copeland and D. Galai. Information effects and the bid-ask spread. *Journal of Finance*, 38(5):1457–1469, 1983.
- J. Cvitanić and A. Kirilenko. High frequency traders and asset prices. Working paper, March 2010.
- R. Durrett. *Probability: Theory and Examples*. Duxbury Press, 3rd edition, 2004.
- D. Easley, T. Hendershott, and T. Ramadorai. The price of latency. Working paper, May 2008.
- S. Grossman and M. Miller. Liquidity and market structure. *Journal of Finance*, 43:617–633, 1988.
- L. Harris. Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions & Instruments*, 7(2):1–76, April 1998.
- J. Hasbrouck. *Empirical Market Microstructure*. Oxford University Press, 2007.
- J. Hasbrouck and G. Saar. Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12:143–172, 2009.

- J. Hasbrouck and G. Saar. Low-latency trading. Working paper, 2010.
- T. Hendershott and R. Riordan. Algorithmic trading and information. Working paper, September 2009.
- T. Hendershott, C. M. Jones, and A. Menkveld. Does algorithmic trading improve liquidity? To appear in *Journal of Finance*, April 2010.
- R. A. Jarrow and P. Protter. A dysfunctional role of high frequency trading in electronic markets. Working paper, 2011.
- A. Kirilenko, A. Kyle, M. Samadi, and T. Tuzun. The flash crash: The impact of high frequency trading on an electronic market. Working paper, 2010.
- H. Leland. Option pricing and replication with transactions costs. *Journal of Finance*, 40:1283–1301, 1985.
- A. W. Lo, A. C. MacKinlay, and J. Zhang. Econometric models of limit order execution. *Journal of Financial Economics*, 65:31–71, 2002.
- E. Pagnotta. Information and liquidity trading at optimal frequencies. Technical report, Working paper, NYU Stern, 2010.
- C.A. Parlour and D.J. Seppi. Limit order markets: A survey. *Handbook of Financial Intermediation and Banking*, 5, 2008.
- A. F. Perold. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14(3): 4–9, 1988.
- M. J. Ready. The specialist’s discretion: Stopped orders and price improvement. *Review of Financial Studies*, 12(5):1075–112, Winter 1999.
- R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39:1127–1139, 1984.
- H. R. Stoll and C. Schenzler. Trades outside the quotes: Reporting delay, trading option or trade size? *Journal of Financial Economics*, 79:615–653, 2006.
- I. Werner. NYSE order flow, spreads, and information. *Journal of Financial Markets*, 6:309–335, 2003.

Technical Appendices

A. Dynamic Programming Decomposition

In order to solve the optimal control problem (8) via dynamic programming, note that we can equivalently consider the objective of maximizing the sale price P . Consider a decision time T_i with $0 \leq i < n$, and assume that the trader's limit order remains unfilled at time T_i . The state of the system consists of the current price, S_{T_i} as well as the previously chosen limit price,¹⁸ ℓ_{i-1} , since this price will become active at time T_i . We can define an optimal value function $J_i(S_{T_i}, \ell_{i-1})$, as a function of this state, by optimizing the eventual sale price over all future decisions. In other words,

$$(14) \quad J_i(S_{T_i}, \ell_{i-1}) \triangleq \underset{\ell_i, \dots, \ell_{n-1}}{\text{maximize}} \mathbb{E}[P \mid S_{T_i}, \ell_{i-1}].$$

At time $T = T_n$, the trader must sell via a market order, hence

$$(15) \quad J_n(S_{T_n}, \ell_{n-1}) = S_{T_n}.$$

Now, for $0 \leq i < n$, there are three mutually exclusive events one of which must occur between time T_i and time T_{i+1} . These are the events $\mathcal{E}_i^{(1)}$, $\mathcal{E}_i^{(2)}$, and $\mathcal{E}_i^{(3)}$ described in Section 3. By considering cases corresponding to these events, we have the Bellman equation

$$(16) \quad J_i(S_{T_i}, \ell_{i-1}) \triangleq \max_{\ell_i} \mathbb{E} \left[\mathbb{I}_{\mathcal{E}_i^{(1)}} \ell_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(2)}} S_{T_{i+1}} + \mathbb{I}_{\mathcal{E}_i^{(3)}} J_{i+1}(S_{T_{i+1}}, \ell_i) \mid S_{T_i}, \ell_{i-1} \right].$$

Here, the first term corresponds to an execution at the prior price ℓ_{i-1} , the second term corresponds to the price ℓ_i being crossed by the bid price upon arrival to the market, and the third term corresponds to all other cases.

Define the function Q_i , for $0 \leq i \leq n$, by

$$Q_i(S_{T_i}, v_{i-1}) \triangleq J_i(S_{T_i}, S_{T_i} + v_{i-1}) - S_{T_i}.$$

The function Q_i is the premium of the value at time T_i , *relative* to the current bid price S_{T_i} . Similarly, $v_{i-1} \triangleq \ell_{i-1} - S_{T_i}$ is the premium of limit price decided at time T_{i-1} *relative* to the

¹⁸We will assume that $\ell_{-1} = \infty$, i.e., there is no limit order active at the beginning of the time horizon.

current bid price at time T_i . Then, applying (16), we have for $0 \leq i < n$,

$$\begin{aligned} & Q_i(S_{T_i}, v_{i-1}) \\ &= \max_{\ell_i} \mathbf{E} \left[\mathbb{I}_{\mathcal{E}_i^{(1)}}(S_{T_i} + v_{i-1}) + \mathbb{I}_{\mathcal{E}_i^{(2)}} S_{T_{i+1}} + \mathbb{I}_{\mathcal{E}_i^{(3)}} J_{i+1}(S_{T_{i+1}}, \ell_i) \mid S_{T_i}, v_{i-1} \right] - S_{T_i} \\ &= \max_{u_i} \mathbf{E} \left[\mathbb{I}_{\mathcal{E}_i^{(1)}} v_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(2)} \cup \mathcal{E}_i^{(3)}} X_{i+1} + \mathbb{I}_{\mathcal{E}_i^{(3)}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \mid S_{T_i}, v_{i-1} \right]. \end{aligned}$$

Here, $X_{i+1} \triangleq S_{T_{i+1}} - S_{T_i} \sim N(0, \sigma^2 \Delta t)$ is the change in bid price from time T_i to time T_{i+1} . We define $u_i \triangleq \ell_i - S_{T_i}$ as the premium of the limit price at time T_i (i.e., the decision variable) *relative* to the current bid price S_{T_i} . Note that the price change X_{i+1} is zero mean under the event

$$\mathcal{E}_i^{(2)} \cup \mathcal{E}_i^{(3)} = (\mathcal{E}_i^{(1)})^c,$$

by our assumption that the arrival of impatient buyers is independent of the bid price process, hence

$$(17) \quad Q_i(S_{T_i}, v_{i-1}) = \max_{u_i} \mathbf{E} \left[\mathbb{I}_{\mathcal{E}_i^{(1)}} v_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(3)}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \mid S_{T_i}, v_{i-1} \right].$$

Finally, by (15),

$$(18) \quad Q_n(S_{T_n}, v_{n-1}) = 0.$$

As should be clear from the above discussion, the Bellman equation (16) with terminal condition (15) and the backward recursion (17) with terminal condition (18) are completely equivalent, up to a change in variables. Expressing these equations in the latter form, however, brings significant simplifications, as the following lemma shows.

Lemma 3. *Suppose a collection of functions $\{Q_i\}$ satisfies the dynamic programming equations (17)–(18). Then, for each $0 \leq i < n$, Q_i does not depend on the price S_{T_i} , and takes the form*

$$(19) \quad Q_i(v_{i-1}) = \mathbb{I}_{\{v_{i-1} \leq \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i,$$

where the scalar h_i satisfies

$$(20) \quad h_i = \max_{u_i} \mathbf{P}(X_{i+1} < u_i) \mathbf{E}[Q_{i+1}(u_i - X_{i+1}) \mid X_{i+1} < u_i].$$

Proof. Observe that, for $0 \leq i < n$, (17) can be simplified according to

$$(21) \quad \begin{aligned} & Q_i(S_{T_i}, v_{i-1}) \\ &= \max_{u_i} \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} \\ & \quad + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}}\right) \mathbf{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \mid S_{T_i} \right], \end{aligned}$$

where we have used the definitions of the events $\mathcal{E}_i^{(1)}$ and $\mathcal{E}_i^{(3)}$.

Now, we proceed by backward induction. For the terminal case $i = n - 1$, from (21) and the fact that $Q_n = 0$ and $u_{n-1} = -\infty$ (i.e., the trader must use a market order at the last time slot), we have that

$$Q_{n-1}(S_{T_{n-1}}, v_{n-2}) = \mu \Delta t v_{n-2} \mathbb{I}_{\{v_{n-2} \leq \delta\}}.$$

In other words, Q_{n-1} satisfies the hypotheses of the lemma, with $h_{n-1} = 0$.

Now, suppose that the result holds for some $0 \leq i + 1 < n$. By (21), and since Q_{i+1} does not depend on $S_{T_{i+1}}$,

$$\begin{aligned} & Q_i(S_{T_i}, v_{i-1}) \\ &= \max_{u_i} \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}}\right) \mathbf{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1}) \right] \\ &= \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}}\right) h_i \\ &= \mathbb{I}_{\{v_{i-1} \leq \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t v_{i-1}) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i. \end{aligned}$$

Here, in the second equality, we define h_i through (20). The result then follows. \blacksquare

Notice that, at the beginning of the trading horizon, there is no active limit order, i.e., $u_{-1} = \infty$. From Lemma 3, we have that

$$h_0 = Q_0(\infty) = \underset{\ell_0, \dots, \ell_{n-1}}{\text{maximize}} \mathbf{E}[P \mid S_0] - S_0.$$

In other words, $h_0 = h_0(\Delta t)$, as defined in (8), and our notation is consistent. More generally, for $i > 0$, from (20), we can interpret h_i to be the trader's expected payoff at time T_i relative to the current bid price under the optimal policy, assuming that the limit order does not get executed in that time slot. Thus, h_i can be interpreted as a *continuation value* in the dynamic programming context, as in (9).

The continuation values $\{h_i\}$ allow for a compact representation of the value function, since they consist of only a single real number for each time slot, rather than a function of the entire state space. Theorem 1 directly expresses the dynamic programming equations (17)–(18) in terms of this representation. The proof follows by explicitly computing the expectations in Lemma 3.

Theorem 1. Suppose $\{h_i\}$ satisfy, for $0 \leq i < n - 1$,

$$(22) \quad h_i = \max_{u_i} \left\{ \mu \Delta t \left[u_i \left(\Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) + \sigma \sqrt{\Delta t} \left(\phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) \right] \right. \\ \left. + h_{i+1} \left[(1 - \mu \Delta t) \Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) + \mu \Delta t \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right] \right\},$$

and

$$(23) \quad h_{n-1} = 0.$$

Here, ϕ and Φ are, respectively, the p.d.f. and c.d.f. of the standard normal distribution. Then, $\{h_i\}$ correspond to the continuation values under the optimal policy. In other words, the value functions $\{Q_i\}$ defined by $\{h_i\}$ via (19) solve the dynamic programming equations (17)–(18).

Suppose further that, for $0 \leq i < n - 1$, u_i^* is a maximizer of (22). Then, a policy which chooses limit prices according to the premia defined by $\{u_i^*\}$, i.e.,

$$\ell_i^* = S_{T_i} + u_i^*, \quad \forall 0 \leq i < n - 1,$$

is optimal.

Proof. Suppose that we are given $\{h_i\}$ that satisfy the hypotheses of the theorem. Define $\{Q_i\}$ by setting, for $0 \leq i \leq n - 1$,

$$(24) \quad Q_i(v_{i-1}) \triangleq \mathbb{I}_{\{v_{i-1} \leq \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i,$$

and $Q_n \triangleq 0$. We wish to show that $\{Q_i\}$ solve the dynamic programming equations (17)–(18).

Note that (18) holds by definition. For $0 \leq i < n$, we have that (17) is equivalent to (21). Define \hat{Q}_i to be the right hand side of (21), i.e.,

$$\hat{Q}_i(v_{i-1}) \triangleq \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} \right) \max_{u_i} \mathbb{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1}) \right].$$

Comparing with (24), in order that the dynamic programming equation (21) hold (i.e., that $\hat{Q}_i = Q_i$), we must have that

$$(25) \quad h_i = \max_{u_i} \mathbb{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1}) \right]$$

Using the definition of Q_{i+1} from (24), this is equivalent to

$$\begin{aligned}
& h_i \\
&= \max_{u_i} \mathbb{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} \left(\mathbb{I}_{\{u_i - X_{i+1} \leq \delta\}} [\mu\Delta t(u_i - X_{i+1}) + (1 - \mu\Delta t)h_{i+1}] + \mathbb{I}_{\{u_i - X_{i+1} > \delta\}} h_{i+1} \right) \right] \\
&= \max_{u_i} \mathbb{E} \left[\mathbb{I}_{\{0 < u_i - X_{i+1} \leq \delta\}} \mu\Delta t(u_i - X_{i+1}) + \mathbb{I}_{\{X_{i+1} < u_i\}} (1 - \mu\Delta t)h_{i+1} + \mathbb{I}_{\{u_i - X_{i+1} > \delta\}} \mu\Delta t h_{i+1} \right].
\end{aligned}$$

For the first term in the expectation, we have

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{I}_{\{0 < u_i - X_{i+1} \leq \delta\}} \mu\Delta t(u_i - X_{i+1}) \right] \\
&= \mu\Delta t \int_{-\infty}^{u_i} (u_i - x) \mathbb{I}_{\{u_i - x \leq \delta\}} \frac{1}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx \\
&= \mu\Delta t \int_{u_i - \delta}^{u_i} (u_i - x) \frac{1}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx \\
&= \mu\Delta t \left[u_i \left(\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right) \right) + \int_{u_i - \delta}^{u_i} \frac{-x}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx \right] \\
&= \mu\Delta t \left[u_i \left(\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right) \right) + \sigma\sqrt{\Delta t} \left(\phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right) \right) \right].
\end{aligned}$$

For the second term in the expectation, we have

$$\mathbb{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} (1 - \mu\Delta t)h_{i+1} \right] = (1 - \mu\Delta t)h_{i+1} \Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right).$$

finally, for the last term in the expectation, we have

$$\mathbb{E} \left[\mathbb{I}_{\{u_i - X_{i+1} > \delta\}} \mu\Delta t h_{i+1} \right] = \mu\Delta t h_{i+1} \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)$$

Combining all the terms, we obtain the desired recursion for h_i .

The balance of the theorem (i.e., the optimality of the $\{u_i^*\}$ policy) follows from standard dynamic programming arguments. \blacksquare

B. Proof of Theorem 2

We begin with a preliminary lemma.

Lemma 4. *Suppose that $\{h_i : 0 \leq i < n\}$ solves the dynamic programming recursion (10)–(11). Then, for $0 \leq i < n$,*

$$(26) \quad 0 \leq h_i \leq \delta(1 - (1 - \mu\Delta t)^n) < \delta.$$

Proof. First, note that the result is trivially true for $i = n-1$, since $h_{n-1} = 0$. Now, if $0 \leq i < n-1$, we can always choose $u_i = -\infty$, i.e., a market order, and this results in a continuation value of 0.

Thus, $h_i \geq 0$.

For the upper bound, consider the discrete model without latency described in Section 4.3. Any strategy for the latency model is also feasible for the discrete model, since the trader can simply delay the implementation of trading decisions by one period. Therefore, at time T_i (with $0 \leq i < n - 1$), a policy with latency cannot achieve more value than the optimal policy for the discrete model without latency. At time T_i , there are $n - i - 1$ trading decisions remaining. This corresponds to the initial time of a discrete model with a total time horizon of $(n - i - 1)\Delta t$. Then, with reference to Lemma 2, we have that

$$h_i \leq \delta(1 - (1 - \mu\Delta t)^{n-i-1}).$$

The result immediately follows. ■

Theorem 2. *Fix $\alpha > 1$. If Δt is sufficiently small, then there exists a unique optimal solution $\{h_i\}$ to the dynamic programming equations (10)–(11). Moreover, the corresponding optimal policy $\{u_i^*\}$ is unique. For $0 \leq i < n - 1$, this strategy chooses limit prices in the range*

$$\ell_i^* \in \left(S_i + \delta - \sigma \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, S_i + \delta - \sigma \sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}} \right),$$

where

$$L \triangleq \frac{\delta^2}{2\pi\sigma^2}, \quad R(\Delta t) \triangleq \frac{\delta^2(1 - \mu\Delta t)^{2n}}{2\pi\sigma^2}.$$

Proof. Assume that, for some $0 \leq i < n - 1$, a solution $\{h_j : i + 1 \leq j < n\}$ exists to (10)–(11). We will establish that, for Δt sufficiently small (and not dependent on i), a solution h_i also exists and satisfies the conditions of the theorem. The result will follow by backward induction. Note that the base case of our induction (i.e., the existence of h_{n-1}) is trivial.

To this end, define the auxiliary function f by

$$(27) \quad \begin{aligned} f(u, h) \triangleq & \mu\Delta t \left[u(\Phi(A_u) - \Phi(B_u)) + \sigma\sqrt{\Delta t}(\phi(A_u) - \phi(B_u)) \right] \\ & + h \left[(1 - \mu\Delta t)\Phi(A_u) + \mu\Delta t\Phi(B_u) \right], \end{aligned}$$

where

$$(28) \quad A_u \triangleq \frac{u}{\sigma\sqrt{\Delta t}}, \quad B_u \triangleq \frac{u - \delta}{\sigma\sqrt{\Delta t}}.$$

Then, from Theorem 1, for $0 \leq i < n - 1$, the dynamic programming recursion is given by

$$(29) \quad h_i = \max_{u_i} f(u_i, h_{i+1}),$$

and we can establish the present theorem by proving that, for Δt sufficiently small, (29) has a unique maximizer $u_i^* \in (\hat{u}_L, \hat{u}_R)$, where

$$(30) \quad \hat{u}_L \triangleq \delta - \sigma \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, \quad \hat{u}_R \triangleq \delta - \sigma \sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}}.$$

Note that

$$(31) \quad R(0) \triangleq \lim_{\Delta t \rightarrow 0} R(\Delta t) = \lim_{\Delta t \rightarrow 0} L(1 - \mu\Delta t)^{2T/\Delta t} = Le^{-2\mu T} < \alpha L.$$

Hence, there exists some $\overline{\Delta t} > 0$ so that if $0 < \Delta t < \overline{\Delta t}$, then

$$\delta/2 < \hat{u}_L < \hat{u}_R < \delta, \quad \text{and} \quad 0 < 1 - \mu\Delta t < 1.$$

For the balance of the theorem, we will assume that $0 < \Delta t < \overline{\Delta t}$, in addition to whatever other assumptions are made regarding the magnitude of Δt .

The first and second derivatives of $f(\cdot, h)$ are given by

$$(32) \quad \begin{aligned} f_u(u, h) &= \mu\Delta t \left[\Phi(A_u) - \Phi(B_u) + A_u(\phi(A_u) - \phi(B_u)) - \frac{u\phi(A_u) + (\delta - u)\phi(B_u)}{\sigma\sqrt{\Delta t}} \right] \\ &\quad + \frac{h}{\sigma\sqrt{\Delta t}} \left[(1 - \mu\Delta t)\phi(A_u) + \mu\Delta t\phi(B_u) \right] \\ &= \frac{(1 - \mu\Delta t)h}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t \left[\Phi(A_u) - \Phi(B_u) - \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(B_u) \right] + \frac{h\mu\sqrt{\Delta t}}{\sigma} \phi(B_u) \\ &= \frac{(1 - \mu\Delta t)h}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t [\Phi(A_u) - \Phi(B_u)] + \frac{\mu\sqrt{\Delta t}}{\sigma} \phi(B_u)(h - \delta), \\ f_{uu}(u, h) &= \frac{-u(1 - \mu\Delta t)h}{\sigma^3\Delta t\sqrt{\Delta t}} \phi(A_u) + \frac{\mu\sqrt{\Delta t}}{\sigma} [\phi(A_u) - \phi(B_u)] + \frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}} \phi(B_u)(h - \delta) \\ &= \phi(A_u) \left[\frac{\mu\sqrt{\Delta t}}{\sigma} - \frac{u(1 - \mu\Delta t)h}{\sigma^3\Delta t^{3/2}} \right] + \phi(B_u) \left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right]. \end{aligned}$$

First, we will show that, for Δt sufficiently small, $f(\cdot, h_{i+1})$ has a *local* maximum u_i^* in the interval (\hat{u}_L, \hat{u}_R) , and that this is the unique maximizer over the larger interval $(\delta/2, \delta)$. That is, $u \in (\delta/2, \delta)$ and $u \neq u_i^*$, then

$$(33) \quad f(u, h_{i+1}) < f(u_i^*, h_{i+1}), \quad \text{for all } u \in (\delta/2, \delta), u \neq u_i^*.$$

This is implied by the following claims, which we will demonstrate hold for Δt sufficiently small:

(i) $f_u(\hat{u}_L, h_{i+1}) > 0$.

(ii) $f_u(\hat{u}_R, h_{i+1}) < 0$.

(iii) $f_{uu}(u, h_{i+1}) < 0$, for all $u \in (\delta/2, \delta)$.

Claim (i): Note that

$$(34) \quad \begin{aligned} f_u(\hat{u}_L, h_{i+1}) &= \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_L}) + \mu\Delta t[\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})] + \frac{\mu\Delta t}{\delta\sqrt{\alpha}}(h_{i+1} - \delta) \\ &\geq \mu\Delta t[\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})] - \frac{\mu\Delta t}{\sqrt{\alpha}}, \end{aligned}$$

where we use the fact that $h_{i+1} \geq 0$ (cf. Lemma 4). In order to calculate a lower bound for $\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})$, we need the following standard bound on the tail probabilities of the normal distribution (see, e.g., Durrett, 2004). Define Q to be the tail probability of a standard normal distribution, i.e.,

$$Q(x) \triangleq 1 - \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{1}{2}u^2} du.$$

Then, for all $x > 0$,

$$(35) \quad \frac{x^2 - 1}{x^3\sqrt{2\pi}}e^{-\frac{1}{2}x^2} \leq Q(x) \leq \frac{1}{x\sqrt{2\pi}}e^{-\frac{1}{2}x^2}.$$

Applying this to (34),

$$\begin{aligned} f_u(\hat{u}_L, h_{i+1}) &\geq \mu\Delta t[1 - Q(A_{\hat{u}_L}) - Q(-B_{\hat{u}_L})] + \frac{\mu\Delta t}{\sqrt{\alpha}} \\ &= \mu\Delta t\left(1 - \frac{1}{\sqrt{\alpha}}\right) - \mu\Delta t[Q(A_{\hat{u}_L}) + Q(-B_{\hat{u}_L})] \\ &> \mu\Delta t\left(1 - \frac{1}{\sqrt{\alpha}}\right) - 2\mu\Delta tQ(-B_{\hat{u}_L}) \\ &\geq \mu\Delta t\left(1 - \frac{1}{\sqrt{\alpha}}\right) - \frac{2\mu\Delta t^{3/2}}{\sqrt{2\pi\alpha L \log \frac{\alpha L}{\Delta t}}} \\ &> 0, \end{aligned}$$

for sufficiently small Δt . Here, we have used the fact that $Q(-B_{\hat{u}_L}) > Q(A_{\hat{u}_L})$.

Claim (ii): Similarly, for the other endpoint of the interval, we have

$$\begin{aligned}
f_u(\hat{u}_R, h_{i+1}) &= \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_R}) + \mu\Delta t[1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] \\
&\quad + \frac{\mu\Delta t}{\delta(1 - \mu\Delta t)^n}(h_{i+1} - \delta) \\
&\leq \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_R}) + \mu\Delta t[1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] \\
&\quad + \frac{\mu\Delta t}{\delta(1 - \mu\Delta t)^n}[\delta(1 - (1 - \mu\Delta t)^n) - \delta] \\
&= \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_R}) + \mu\Delta t[1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] - \mu\Delta t \\
&= \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_R}) - \mu\Delta t[Q(A_{\hat{u}_R}) + Q(-B_{\hat{u}_R})] \\
&\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_{\hat{u}_R}) - \mu\Delta tQ(-B_{\hat{u}_R}),
\end{aligned}$$

where we have used the upper bound on h_{i+1} from Lemma 4. Using (31), for sufficiently small Δt , we have

$$\sqrt{\log \frac{R(\Delta t)}{\Delta t}} < \frac{\delta}{2\sigma\sqrt{\Delta t}},$$

and thus

$$A_{\hat{u}_R} \geq \frac{\delta}{2\sigma\sqrt{\Delta t}}.$$

On the other hand, using (35),

$$Q(-B_{\hat{u}_R}) \geq \left[\left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-1/2} - \left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-3/2} \right] \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}}.$$

Thus,

$$\begin{aligned}
f_u(\hat{u}_R, h_{i+1}) &\leq \frac{\delta}{\sigma\sqrt{2\pi\Delta t}} \exp\left(\frac{-\delta^2}{8\sigma^2\Delta t}\right) \\
&\quad - \mu\Delta t \left[\left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-1/2} - \left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-3/2} \right] \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}} \\
&< 0,
\end{aligned}$$

for sufficiently small Δt .

Claim (iii): Note that, for $u \in (\delta/2, \delta)$,

$$\phi\left(\frac{\delta}{\sigma\sqrt{\Delta t}}\right) < \phi(A_u) < \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) < \phi(B_u) < \phi(0).$$

Then, from (32), and using the fact that $0 \leq h_{i+1} < \delta$ (cf. Lemma 4), we have for Δt sufficiently small,

$$\begin{aligned} f_{uu}(u, h_{i+1}) &\leq \phi(A_u) \frac{\mu\sqrt{\Delta t}}{\sigma} + \phi(B_u) \left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \\ &\leq \phi(B_u) \frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) < 0. \end{aligned}$$

In order to complete the proof, it suffices to demonstrate that the *local* maximum $u_i^* \in (\hat{u}_L, \hat{u}_R)$ is the unique *global* maximum. Since u_i^* achieves a higher value than any other $u \in (\delta/2, \delta)$, we will analyze cases where $u \notin (\delta/2, \delta)$ as follows:

- $u \in [0, \delta/2]$.

Here,

$$\phi\left(\frac{\delta}{\sigma\sqrt{\Delta t}}\right) \leq \phi(B_u) \leq \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) \leq \phi(A_u) \leq \phi(0).$$

Further, for Δt sufficiently small,

$$\Phi(A_u) - \Phi(B_u) \geq \Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right) \geq \frac{1}{4}.$$

Then, for Δt sufficiently small,

$$\begin{aligned} (36) \quad f_u(u, h_{i+1}) &\geq \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}} \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) + \mu\Delta t \left[\Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right) \right] \\ &\quad + \frac{\mu\sqrt{\Delta t}}{\sigma} \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) (h_{i+1} - \delta) \\ &\geq \frac{h_{i+1}}{\sigma\sqrt{\Delta t}} \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) + \mu\Delta t \left[\Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right) \right] \\ &\quad - \frac{\delta\mu\sqrt{\Delta t}}{\sigma} \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) \\ &\geq \frac{\mu\Delta t}{4} - \frac{\mu\delta\sqrt{\Delta t}}{\sqrt{2\pi}\sigma} \exp\left(\frac{-\delta^2}{8\sigma^2\Delta t}\right) > 0. \end{aligned}$$

Here, we have used the fact that $h_{i+1} \geq 0$. Using (33) and the fact that $f(\cdot, h_{i+1})$ is continuous, this implies that

$$(37) \quad \sup_{u \in [0, \delta/2]} f(u, h_{i+1}) \leq f(\delta/2, h_{i+1}) < f(u_i^*, h_{i+1}).$$

- $u \in (-\infty, 0)$.

In this case, since $h_{i+1} \geq 0$ and $B_u < A_u < 0$,

$$\begin{aligned} f_u(u, h_{i+1}) &\geq \mu\Delta t [\Phi(A_u) - \Phi(B_u)] - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) \\ &= \mu\Delta t \int_{B_u}^{A_u} \phi(z) dz - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) \\ &> \mu\Delta t(A_u - B_u)\phi(B_u) - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) = 0. \end{aligned}$$

In conjunction with (37), this implies that

$$(38) \quad \sup_{u \in (-\infty, 0)} f(u, h_{i+1}) \leq f(0, h_{i+1}) < f(u_i^*, h_{i+1}).$$

- $u \in [\delta, \infty)$.

In this case, using the upper bound on h_{i+1} from Lemma 4,

$$(39) \quad \begin{aligned} f_u(u, h_{i+1}) &\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_u) + \mu\Delta t [\Phi(A_u) - \Phi(B_u)] \\ &\quad - \frac{\mu\delta(1 - \mu\Delta t)^n\sqrt{\Delta t}}{\sigma}\phi(B_u). \end{aligned}$$

Consider two cases. First, assume that $u > \delta + \sqrt{\Delta t}$. Then, applying (35),

$$\begin{aligned} f_u(u, h_{i+1}) &\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_u) + \mu\Delta t Q(B_u) - \frac{\mu\delta(1 - \mu\Delta t)^n\sqrt{\Delta t}}{\sigma}\phi(B_u) \\ &\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_u) + \frac{\mu\sigma\Delta t^{3/2}}{u - \delta}\phi(B_u) - \frac{\mu\delta(1 - \mu\Delta t)^n\sqrt{\Delta t}}{\sigma}\phi(B_u) \\ &\leq \phi(B_u) \left[\frac{\delta}{\sigma\sqrt{\Delta t}} \exp\left(\frac{-\delta^2}{2\sigma^2\Delta t}\right) + \mu\sigma\Delta t - \frac{\mu\delta(1 - \mu\Delta t)^n\sqrt{\Delta t}}{\sigma} \right]. \end{aligned}$$

Note that $(1 - \mu\Delta t)^n \rightarrow e^{-\mu T}$ as $\Delta t \rightarrow 0$. Then, for Δt sufficiently small,

$$(40) \quad \frac{1}{2}e^{-\mu T} < (1 - \mu\Delta t)^n.$$

Hence, for Δt sufficiently small,

$$f_u(u, h_{i+1}) \leq \phi(B_u) \left[\frac{\delta}{\sigma\sqrt{\Delta t}} \exp\left(\frac{-\delta^2}{2\sigma^2\Delta t}\right) + \mu\sigma\Delta t - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma} \right] < 0.$$

On the other hand, suppose that $u \in [\delta, \delta + \sqrt{\Delta t}]$. Then, from (39), (40), and since $0 < B_u <$

A_u , we have for Δt sufficiently small,

$$\begin{aligned}
f_u(u, h_{i+1}) &\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_u) + \frac{\mu\Delta t}{2} - \frac{\mu\delta(1 - \mu\Delta t)^n\sqrt{\Delta t}}{\sigma}\phi(B_u) \\
&\leq \frac{\delta}{\sigma\sqrt{\Delta t}}\phi(A_u) + \frac{\mu\Delta t}{2} - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma}\phi(B_u) \\
&\leq \frac{\delta}{\sigma\sqrt{2\pi\Delta t}}\exp\left(-\frac{\delta^2}{2\sigma^2\Delta t}\right) + \frac{\mu\Delta t}{2} - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma\sqrt{2\pi}}\exp\left(-\frac{1}{2\sigma^2}\right) < 0.
\end{aligned}$$

The above discussion, combined with (33) and the fact that $f(\cdot, h_{i+1})$ is continuous, implies that

$$(41) \quad \sup_{u \in [\delta, \infty)} f(u, h_{i+1}) \leq f(\delta, h_{i+1}) < f(u_i^*, h_{i+1}).$$

■

C. Proof of Theorem 3

We will establish Theorem 3 via a sequence of lemmas. First, recall the function $f(u, h)$ defined in (27) and the quantities \hat{u}_L and \hat{u}_R defined in (30).

Lemma 5. (i) As $\Delta t \rightarrow 0$,

$$\max_{\substack{u \in [\hat{u}_L, \hat{u}_R] \\ 0 \leq i < n-1}} |f_{uu}(u, h_{i+1})| = O\left(\sqrt{\Delta t \log \frac{1}{\Delta t}}\right).$$

(ii) For all $h \in \mathbb{R}$ and Δt sufficiently small,

$$0 \leq f_h(\hat{u}_R, h) \leq 1.$$

Proof. We begin with (i). Recall A_u and B_u from (28). Let u be in the interval $[\hat{u}_L, \hat{u}_R]$. Then, for $0 \leq i < n - 1$, from (32),

$$\begin{aligned}
|f_{uu}(u, h_{i+1})| &\leq \left| \phi(A_u) \left[\frac{\mu\sqrt{\Delta t}}{\sigma} - \frac{u(1 - \mu\Delta t)h_{i+1}}{\sigma^3\Delta t^{3/2}} \right] + \phi(B_u) \left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \right| \\
&\leq \phi(A_u) \frac{\mu\sqrt{\Delta t}}{\sigma} + \phi(A_u) \frac{\delta^2}{\sigma^3\Delta t^{3/2}} + \phi\left(\frac{\delta - u}{\sigma\sqrt{\Delta t}}\right) \left[\frac{\delta\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma} \right].
\end{aligned}$$

Here, we have used the fact that $0 \leq u \leq \delta$ and $0 \leq h_{i+1} < \delta$ (cf. Lemma 4). Note that, for Δt

sufficiently small, $\hat{u}_L \geq \delta/2$. Then,

$$\max_{u \in [\hat{u}_L, \hat{u}_R]} \phi(A_u) \leq \phi(A_{\delta/2}) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\delta^2}{8\sigma^2\Delta t}\right) \leq c_0\Delta t^2,$$

for an appropriately chosen constant c_0 . Thus,

$$\begin{aligned} |f_{uu}(u, h_{i+1})| &\leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \phi\left(\frac{\delta-u}{\sigma\sqrt{\Delta t}}\right) \left[\frac{\delta\mu(\delta-u)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \\ &\leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \phi\left(\frac{\delta-\hat{u}_R}{\sigma\sqrt{\Delta t}}\right) \left[\frac{\delta\mu(\delta-\hat{u}_L)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \\ &= \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}} \left[\frac{\delta\mu}{\sigma^2} \sqrt{\log \frac{\alpha L}{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \\ &\leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \frac{\mu\Delta t}{\sigma\sqrt{2\pi R(\Delta t)}} + \frac{\delta\mu}{\sigma^2\sqrt{2\pi R(\Delta t)}} \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}. \end{aligned}$$

Since $R(\Delta t) \rightarrow Le^{-2\mu T}$ as $\Delta t \rightarrow 0$, the last term asymptotically dominates and (i) follows.

For (ii), note that $\Phi(A_{\hat{u}_R}), \Phi(B_{\hat{u}_R}) \in (0, 1)$, so if $\Delta t < 1/\mu$, then for all h ,

$$f_h(\hat{u}_R, h) = (1 - \mu\Delta t)\Phi(A_{\hat{u}_R}) + \mu\Delta t\Phi(B_{\hat{u}_R}) \in (0, 1).$$

■

Lemma 6. As $\Delta t \rightarrow 0$,

$$\hat{u}_R - \hat{u}_L = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Proof. Note that

$$\hat{u}_R - \hat{u}_L = \sigma\sqrt{\Delta t} \left(\sqrt{\log \frac{\alpha L}{\Delta t}} - \sqrt{\log \frac{R(\Delta t)}{\Delta t}} \right) = \sigma\sqrt{\Delta t} [g(\alpha L) - g(R(\Delta t))],$$

where $g(x) \triangleq \sqrt{\log \frac{x}{\Delta t}}$. Then, by mean value theorem, for some $z \in [R(\Delta t), \alpha L]$,

$$\hat{u}_R - \hat{u}_L = \sigma\sqrt{\Delta t}g'(z)[\alpha L - R(\Delta t)] = \frac{\sigma}{2z}[\alpha L - R(\Delta t)]\sqrt{\frac{\Delta t}{\log \frac{z}{\Delta t}}} \leq \frac{\sigma\alpha L}{2R(\Delta t)}\sqrt{\frac{\Delta t}{\log \frac{R(\Delta t)}{\Delta t}}}.$$

The result follows since $R(\Delta t) \rightarrow R(0) \triangleq Le^{-2\mu T}$ as $\Delta t \rightarrow 0$. ■

Let $\{h_i : 0 \leq i < n-1\}$ be the optimal solution to the dynamic programming recursion (10)–(11), and let $\{u_i^* : 0 \leq i < n-1\}$ define the corresponding optimal policy. Define $\{\hat{h}_i : 0 \leq$

$i \leq n - 1$ by the recursion

$$\hat{h}_i \triangleq \begin{cases} f(\hat{u}_R, \hat{h}_{i+1}) & \text{if } 0 \leq i < n - 1, \\ 0 & \text{if } i = n - 1. \end{cases}$$

Note that \hat{h}_i is the continuation value of the suboptimal policy that always chooses $u_i = \hat{u}_R$, for $0 \leq i < n - 1$. We are interested in quantifying its difference to the optimal continuation value.

Lemma 7. As $\Delta t \rightarrow 0$,

$$0 \leq h_0 - \hat{h}_0 = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Proof. For $0 \leq i < n - 1$, define $\Delta_i \triangleq h_i - \hat{h}_i$. Clearly, $\Delta_i \geq 0$.

Using the mean value theorem,

$$\begin{aligned} \Delta_i &= f(u_i^*, h_{i+1}) - f(\hat{u}_R, \hat{h}_{i+1}) \\ &= [f(u_i^*, h_{i+1}) - f(\hat{u}_R, h_{i+1})] + [f(\hat{u}_R, h_{i+1}) - f(\hat{u}_R, \hat{h}_{i+1})] \\ &= -\frac{1}{2}f_{uu}(\bar{u}, h_{i+1})(\hat{u}_R - u_i^*)^2 + f_h(\hat{u}_R, \bar{h})\Delta_{i+1}. \end{aligned}$$

where \bar{u} is some point on the interval (u_i^*, \hat{u}_R) and \bar{h} is some point on the interval (\hat{h}_{i+1}, h_{i+1}) . Here, we have used the fact that the optimal solution u_i^* satisfies the first order condition $f_u(u_i^*, h_{i+1}) = 0$.

Using Lemmas 5 and 6, for Δt sufficiently small, there exist constants c_1 and c_2 so that

$$\max_{\substack{u \in [\hat{u}_L, \hat{u}_R] \\ 0 \leq i < n-1}} |f_{uu}(u, h_{i+1})| \leq c_1 \sqrt{\Delta t \log \frac{1}{\Delta t}}, \quad \hat{u}_R - \hat{u}_L \leq c_2 \sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.$$

Also, from Lemma 5, note that $0 \leq f_h(\hat{u}_R, \bar{h}) \leq 1$. Then, we obtain that, for Δt sufficiently small,

$$\Delta_i \leq \frac{c_1(\hat{u}_R - u_i^*)^2}{2} \sqrt{\Delta t \log \frac{1}{\Delta t}} + \Delta_{i+1} \leq \frac{c_1 c_2}{2} \frac{\Delta t^{3/2}}{\sqrt{\log \frac{1}{\Delta t}}} + \Delta_{i+1}.$$

Then, since $\Delta_{n-1} = 0$, we have that

$$\Delta_0 \leq \left(\frac{T}{\Delta t}\right) \frac{c_1 c_2}{2} \frac{\Delta t^{3/2}}{\sqrt{\log \frac{1}{\Delta t}}} = \frac{c_1 c_2 T}{2} \sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.$$

■

Define the sequence $\{\hat{\beta}_i : 0 \leq i \leq n - 1\}$ by the linear recursion

$$(42) \quad \hat{\beta}_i \triangleq \begin{cases} \mu \Delta t (\hat{u}_R - \hat{\beta}_{i+1}) + \hat{\beta}_{i+1} & \text{if } 0 \leq i < n - 1, \\ 0 & \text{if } i = n - 1. \end{cases}$$

Here, $\hat{\beta}_i$ is an approximation to the value \hat{h}_i . The next lemma bounds the approximation error.

Lemma 8. As $\Delta t \rightarrow 0$,

$$|\hat{h}_0 - \hat{\beta}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Proof. For $0 \leq i < n - 1$, define $\epsilon_i \triangleq \hat{h}_i - \hat{\beta}_i$. Recall the following definition from the proof of Theorem 2,

$$A_{\hat{u}_R} \triangleq \frac{\hat{u}_R}{\sigma\sqrt{\Delta t}}, \quad B_{\hat{u}_R} \triangleq \frac{\hat{u}_R - \delta}{\sigma\sqrt{\Delta t}} = -\sqrt{\log \frac{R(\Delta t)}{\Delta t}}.$$

Then, by the recursive definitions of \hat{h}_i and $\hat{\beta}_i$, $0 \leq i < n - 1$,

$$\begin{aligned} \epsilon_i &= (1 - \mu\Delta t)\epsilon_{i+1} - \mu\Delta t\hat{u}_R[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}[\phi(A_{\hat{u}_R}) - \phi(B_{\hat{u}_R})] \\ &\quad - (1 - \mu\Delta t)\hat{h}_{i+1} \left[1 - \Phi(A_{\hat{u}_R}) - \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right]. \end{aligned}$$

Since \hat{u}_R is not the optimal policy, we have $\hat{h}_{i+1} \leq h_{i+1} < \delta$ (cf. Lemma 4). Further, for Δt sufficiently small, $0 < \phi(A_{\hat{u}_R}) \leq \phi(B_{\hat{u}_R})$. This implies that

$$\begin{aligned} |\epsilon_i| &\leq (1 - \mu\Delta t)|\epsilon_{i+1}| + \delta\mu\Delta t[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}\phi(B_{\hat{u}_R}) \\ &\quad + \delta \left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right]. \end{aligned}$$

Note that, except for the first term, there is no dependence on i in the right hand side of this equality. Then, we can define

$$\begin{aligned} C(\Delta t) &\triangleq \delta\mu\Delta t[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}\phi(B_{\hat{u}_R}) \\ &\quad + \delta \left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right], \end{aligned}$$

and we have that

$$|\epsilon_i| \leq (1 - \mu\Delta t)|\epsilon_{i+1}| + C(\Delta t).$$

Since $\epsilon_{n-1} = 0$, it is easy to verify by backward induction on i that

$$|\epsilon_i| \leq \frac{1 - (1 - \mu\Delta t)^{n-i-1}}{\mu\Delta t} C(\Delta t).$$

Therefore,

$$\begin{aligned}
|\epsilon_0| &\leq \frac{1 - (1 - \mu\Delta t)^{n-1}}{\mu\Delta t} C(\Delta t) \leq \frac{C(\Delta t)}{\mu\Delta t} \\
(43) \quad &= \delta[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\sqrt{\Delta t}\phi(B_{\hat{u}_R}) + \frac{\delta}{\mu\Delta t} \left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right] \\
&= \delta[Q(A_{\hat{u}_R}) + Q(-B_{\hat{u}_R})] + \mu\sigma\sqrt{\Delta t}\phi(B_{\hat{u}_R}) + \frac{\delta}{\mu\Delta t} \left[Q(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}Q(-B_{\hat{u}_R})\right].
\end{aligned}$$

From (35), however,

$$Q(A_{\hat{u}_R}) \leq \frac{\sigma}{\hat{u}_R} \sqrt{\frac{\Delta t}{2\pi}} \exp\left(-\frac{\hat{u}_R^2}{2\sigma\Delta t}\right).$$

Since $\hat{u}_R \rightarrow \delta$ as $\Delta t \rightarrow 0$, for Δt sufficiently small, there exists constants a_1 and a_2 , with $0 < a_2 < \delta^2/2\sigma$, so that

$$Q(A_{\hat{u}_R}) \leq a_1\sqrt{\Delta t} \exp\left(-\frac{a_2}{\Delta t}\right).$$

Also by (35),

$$Q(-B_{\hat{u}_R}) \leq \sqrt{\frac{\Delta t}{2\pi R(\Delta t) \log \frac{R(\Delta t)}{\Delta t}}}.$$

Since and $R(\Delta t) \rightarrow R(0) \triangleq Le^{-2\mu T}$ as $\Delta t \rightarrow 0$, for Δt sufficiently small, there exists a constant a_3 so that

$$Q(-B_{\hat{u}_R}) \leq a_3\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.$$

Finally,

$$\phi(B_{\hat{u}_R}) = \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}},$$

so for Δt sufficiently small, there exists a constant a_4 with

$$\phi(B_{\hat{u}_R}) \leq a_4\sqrt{\Delta t}.$$

Applying these bounds to (43), the result follows. ■

Lemma 9. As $\Delta t \rightarrow 0$,

$$\hat{\beta}_0 = \hat{u}_R \left(1 - e^{-\mu T}\right) + O(\Delta t).$$

Proof. Note that the recurrence (42) can be explicitly solved to obtain

$$\hat{\beta}_0 = \sum_{i=0}^{n-2} (1 - \mu\Delta t)^i \mu\Delta t \hat{u}_R = \hat{u}_R \left(1 - (1 - \mu\Delta t)^{n-1}\right) = \hat{u}_R \left(1 - (1 - \mu\Delta t)^{T/\Delta t - 1}\right).$$

The result follows since $(1 - \mu\Delta t)^{T/\Delta t} = e^{-\mu T} + O(\Delta t)$ as $\Delta t \rightarrow 0$. ■

We are now ready to prove Theorem 3.

Theorem 3. As $\Delta t \rightarrow 0$,

$$h_0(\Delta t) = \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right) + o(\sqrt{\Delta t}),$$

where

$$\bar{h}_0 = \delta \left(1 - e^{-\mu T} \right)$$

is the optimal value for the stylized model without latency, i.e., the value defined by (5).

Proof. First, define

$$\hat{\gamma}_0 \triangleq (1 - e^{-\mu T}) \left(\delta - \sigma \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right).$$

Then,

$$(44) \quad \left| h_0 - \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right) \right| = |h_0 - \hat{\gamma}_0| \\ \leq |h_0 - \hat{h}_0| + |\hat{h}_0 - \hat{\beta}_0| + |\hat{\beta}_0 - \hat{\gamma}_0|.$$

We will bound each of the terms in the right hand side of (44). First, by Lemma 7,

$$(45) \quad |h_0 - \hat{h}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Next, by Lemma 8,

$$(46) \quad |\hat{h}_0 - \hat{\beta}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Finally, by Lemma 9, for Δt sufficiently small, there exists a constant c_1 so that

$$|\hat{\beta}_0 - \hat{\gamma}_0| \leq \sigma \left(1 - e^{-\mu T} \right) \left(\hat{u}_R - \delta + \sqrt{\Delta t \log \frac{L}{\Delta t}} \right) + c_1 \Delta t \\ \leq \sigma \left(1 - e^{-\mu T} \right) \left(\hat{u}_R - \delta + \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}} \right) + c_1 \Delta t, \\ = \sigma \left(1 - e^{-\mu T} \right) (\hat{u}_R - \hat{u}_L) + c_1 \Delta t,$$

where $\alpha > 1$ and L are defined by Theorem 2. Applying Lemma 6, we have that

$$(47) \quad |\hat{\beta}_0 - \hat{\gamma}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

By applying (45)–(47) to (44), we have that

$$\left| h_0 - \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}} \right) \right| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right),$$

which implies the desired result. ■