Econometric Models of Limit Order Executions

Andrew W. Lo
A. Craig MacKinlay
June Zhang

012-99

Econometric Models of Limit-Order Executions*

Andrew W. Lo†, A. Craig MacKinlay‡,

and June Zhang††

First Draft: September 25, 1997
Latest Revision: May 28, 1999

Abstract

Limit orders incur no price impact, however, their execution time is uncertain. We develop several econometric models of limit-order execution times using survival analysis, and estimate them with actual limit-order data. We estimate models for time-to-first-fill and time-to-completion, and for limit-sells and limit-buys, and incorporate the effects of explanatory variables such as the limit price, the limit size, the bid/offer spread, and market volatility. We find that execution times are very sensitive to limit price and several other explanatory variables, but not sensitive to limit size. We also show that hypothetical limit-order executions, constructed either theoretically from first-passage times or empirically from transactions data, are very poor proxies for actual limit-order executions.

Keywords: Market Microstructure; Transactions Costs; Portfolio Management.

JEL Classification: G23

*This research was partially supported by the MIT Laboratory for Financial Engineering, the Rodney White Center for Financial Research, and the National Science Foundation (Grant No. SBR–9709976). We are grateful to ITG for providing us with their limit-order dataset, and to the New York Stock Exchange for their TORQ dataset. We thank Jushan Bai, David Cushing, Robert Ferstenberg, Harvey Fram, Larry Harris, Joel Hasbrouck, Mauricio Karchmer, David Modest, Bill Schwert, George Sofianos, a referee, and participants at the 1998 Berkeley Program in Finance, the Western Finance Association's 1998 Annual Meeting, and the 1998 Conference on the International Competition for Order Flow for helpful comments and discussion.

†Harris & Harris Group Professor, MIT Sloan School of Management, 50 Memorial Drive, Cambridge, MA 02142–1347 (corresponding author).

‡Joseph P. Wargrove Professor of Finance, The Wharton School, Department of Finance, University of Pennsylvania, Philadelphia, PA 19104–6367.

††Research Associate, ITG, 44 Farnsworth Street, Boston, MA 02210.

Contents

1 Introduction

One of the most important tools for trading equity securities is the limit order, an order to transact a prespecified number of shares at a prespecified price. Indeed, limit orders comprise a significant fraction of stock-market trading activity, accounting for approximately 45% of the total orders on the New York Stock Exchange (NYSE).[1] The primary advantage of a limit order is the absence of price risk—a transaction occurs only if the *limit price* is attained. However, this advantage does not come without a cost: execution is not guaranteed, and the time-to-execution is a random function of many factors, e.g., limit price, number of shares, market conditions, private information. For some trades, the uncertainty in execution time is unimportant, but for others, the opportunity cost of waiting can be significant.

If immediacy is critical, the market order is the appropriate instrument to use. However, market orders can be subject to significant price risk, particularly for large orders and in volatile markets. In practice, traders submit both types of orders, with an eye towards balancing the opportunity cost of delaying execution against the risks associated with immediate execution.[2] A prerequisite for any quantitative approach to making such trade-offs is an econometric model of limit-order execution times.

Limit orders play another important role in determining trading costs: they influence bid/offer quotes and, therefore, spreads.[3] Therefore, limit-order execution times affect the frequency with which quotes are updated and comprise a significant component of the dynamics of the bid/offer spread. Moreover, limit-order execution times have been used to measure the overall quality of equity markets, e.g., Battalio et al. (1997) and SEC (1997), hence their determinants may have important implications for the economic consequences of market fragmentation, the practice of "preferencing", and the relative merits of specialist vs. multiple-dealer market structures.

In this paper, we propose and estimate several econometric models of limit-order ex-

[1]This figure is reported by Harris and Hasbrouck [1996] for orders submitted using the NYSE's SuperDOT system.

[2]See, for example, Cohen et al. (1981), O'Hara and Oldfield (1986), Glosten (1989, 1994), Easley and O'Hara (1991), Parlour (1995), Chakravarty and Holden (1995), Keim and Madhavan (1995), Belonsky (1996), Harris and Hasbrouck (1996), Kavajecz (1996), Rock (1996), and Seppi (1996).

[3]Using the NYSE's Transaction, Orders, and Quotes (TORQ) database, Chung et al. (1997) estimate that 21% of the quotes in their sample originate from limit orders on both the bid and offer sides without any direct participation from the specialist.

ecution times using actual historical limit-order data from Investment Technology Group, a brokerage firm specializing in electronic trading. Using survival analysis, a well-known statistical technique for modeling failure times and other non-negative random variables, we are able to estimate the conditional distribution of limit-order execution times as a function of economic variables such as limit price, order size, and current market conditions. Because limit-order execution times can be interpreted quite naturally as failure times—they are non-negative, random, and temporally ordered—survival analysis is the most appropriate method for modeling their evolution.

Moreover, survival analysis can accommodate an important feature of limit-order execution times that existing models have ignored: *censored* observations, i.e., limit orders that are cancelled before they are executed. There is great temptation to ignore censored observations since they seem to provide little information about execution times. However, the fact that a limit order is cancelled after, say, 30 minutes does yield one piece of useful information: the limit order "survived" for at least 30 minutes. Therefore, censored observations do affect the conditional distribution of execution times despite the fact that they are not executions. Alternatively, ignoring censored observations can bias the estimator of the conditional distribution of execution times dramatically.

Using a sample of limit orders for the 100 largest stocks in the S&P 500 Index from August 1994 to August 1995, we construct several econometric models of limit-order execution times based on survival analysis, and show that they fit the data remarkably well. In particular, we estimate separate models for limit-buy and limit-sell orders, and separate models for time-to-first-fill and time-to-completion, hence four models in all. Each of these four models yields a conditional distribution that closely matches the data's and passes several diagnostic measures of goodness-of-fit. The parameter estimates show that execution times can be quite sensitive to certain explanatory variables, e.g., market depth, spread between limit price and quote midpoint, market volatility, implying that the kind of strategic order-placement strategies described by Angel (1994), Foucalt (1996), Harris (1994), Hollifield, Miller, and Sandas (1996), Kumar and Seppi (1993), and Parlour (1994) may well be feasible in practice. Limit-order execution times can be accurately modeled, hence controlled.

In Section 2 we review the literature on limit orders, and in Section 3 we discuss some of the institutional features of limit orders and describe our limit-order dataset. We present a simple but powerful application of this dataset in Section 4 in which we compare actual limit-order execution times to their *hypothetical* counterpart, constructed theoretically (from the first-passage times of Brownian motion) and empirically (from transactions data). We present a brief review of survival analysis in Section 5 and turn to our empirical analysis in Section. 6. We conclude in Section 7.

2 Literature Review

There is a large and growing theoretical literature that considers the economic role of limit orders in the price discovery process. Foucault (1993), Glosten (1989, 1994), Easley and O'Hara (1991), Parlour (1995), Chakravarty and Holden (1995), Kavajecz (1996), Rock (1996), Sandas (1996), and Seppi (1996) are just a few recent examples. The general focus of these papers is the effect of limit orders on the market, the interaction between limit and market orders, and the role of the market maker. None of these studies are set in a continuous-trading environment, hence they provide little direct guidance for modeling limit-order execution times.

However, several studies do explore the probability of limit-order execution. For example, under a number of rather strong assumptions, Angel (1994) derives an analytical expression for the conditional probability of limit-order execution, conditional upon an investor's information set. His result applies to batch trading of one round-lot of the stock for informed traders, assuming that traders know the entire limit-order book. Within his analytical framework, Angel also conducts some simulations for continuous-trading environments.

Hollifield, Miller, and Sandas (1996) build a structural model of a pure limit-order market. The model captures the trade-off between order price and probability of execution. They estimate their model nonparametrically and derive implications for traders' order-submission strategies.

A number of studies have compared the use of market orders to limit orders empirically. In particular, using the NYSE TORQ data, Petersen and Fialkowski (1994) find that limit orders placed at the quote outperform market orders in markets that quote in \$1/8 increments, but underperform market orders in wider markets. Harris and Hasbrouck (1996) use the TORQ data to compare the profitability of order-submission strategies using limit orders versus market orders. They find that in some cases the use of limit orders can reduce execution costs. Handa and Schwartz (1996) also provide a comparison, assessing the profitability of limit-order trading by comparing unconditional expected returns of market orders versus limit orders. Their analysis is based on hypothetical limit-order executions, fictitious executions constructed from transactions data (see Section 4 for further discussion and an empirical critique). Biais, Hillion, and Spatt (1995) present an empirical analysis of the order flow of the Paris Bourse which is a pure limit-order market. They find that traders strategies vary with market conditions, using more limit orders at times when spreads are wide and using more market orders at times when spreads are narrow.

Other empirical studies have focused on the role of limit orders in determining execution costs and overall market quality. Keim and Madhavan (1995) examine a unique dataset containing information for 62,000 equity orders (each order generating one or more trades) of 21 institutional investors from January 1991 to March 1993, and one of the many issues they consider is the selection of order type (limit order, market order, working order, or crossing network). McInish and Wood (1995) argue that there are "hidden" limit orders on the NYSE, orders that would improve the posted quotes but which are not always displayed by specialists. Using the NYSE TORQ data, Chung et al. (1997) find that bid/offer spreads are heavily influenced by limit orders—posted spreads are widest when there is no competition from the limit-order book, and narrowest when the quotes originate exclusively from the limit-order book. Battalio et al. (1997) compare limit-order fill rates and execution times of primary and regional exchanges to gauge execution quality across markets. The SEC (1997) performs a similar analysis. And Belonsky (1996) provides an extensive cross-sectional analysis of limit orders submitted to the NYSE during the month of February 1994 for all common stocks with SuperDOT activity and prices between $1.00 and $150.00.

Although none of these papers attempt to model the determinants of limit-order execution times, it is apparent that such a model might provide important insights into each of the issues they address.

3 Limit-Order Data

Although limit orders differ slightly in their institutional features from one exchange to another, we shall focus on those characteristics that are common across the largest exchanges, e.g., the New York and American Stock Exchanges. Upon submission to a designated exchange, a limit order enters the specialist's display book, known as the *order book* or the *queue*. The queue gives the first priority to the highest limit-buy price and to the lowest limit-sell price. Limit orders with the same limit price are prioritized by time of submission, with the oldest order given the highest priority.[4] A order's execution often involves several *partial fills* before it is completed, but partial fills do not change the time priority. A limit order is not binding—it can be *cancelled* or *corrected* at any time.[5]

When a limit order is submitted a number of parameters must be specified, including: the limit price, whether the order is to buy or to sell, the order size (in shares), the designated exchange, and the *time-in-force*. The time-in-force is the period during which a limit order can be filled. For example, a *day-order* is a limit order that can be filled anytime until the market closes; a *good-till-cancelled* order is a limit that can be filled anytime prior to cancellation. The majority of limit orders are day-orders—82% in the NYSE's Transaction, Orders, and Quotes (TORQ) database considered by Harris and Hasbrouck (1996)—although the number of good-till-cancelled limit orders is also substantial, about 17% in the same sample.

In addition to the parameters of the order, we would expect the time-to-execution to depend on current market conditions for the stock itself as well as the market as a whole. Thus, in modeling the time-to-execution it is necessary to specify relevant measures to capture the interaction between time-to-execution and market conditions. We shall return to these issues in Section 6.

[4]On the NYSE, time priority applies only to orders in the limit-order book; a floor trader can trade ahead of the limit-order book at the same price even if he arrives at the post after the limit order was submitted. Also, time priority is given only to the first order at a given price—after that, all members in the crowd have equal time priority, and the limit-order book may be considered a single floor trader. We are grateful to the referee for pointing out this interesting aspect of the NYSE order-handling rules.

[5]The term "correction" does not imply that a mistake has been made, but merely that the original limit order has been revised—either in price or size or both—and resubmitted (and, as a result, has lost its time priority).

3.1 ITG Limit-Order Data

The limit-order data used in this study was provided by Investment Technology Group (ITG), an institutional brokerage firm that provides technology-based equity trading services such as POSIT (an electronic crossing system), QuantEX (a decision-support and routing system), and a full-service trading desk. The ITG limit-order dataset is comprised of all limit orders submitted through the ITG trading desk from August 1, 1994 to August 31, 1995 for the 100 largest stocks (in market capitalization as of the end of September 1995) in the S&P 500 Index.[6] This dataset is unique in several respects. Each limit order is time-stamped and tracked from submission to termination. After submission, a limit order may be partially or completely filled, may get cancelled or corrected by its initiator, or may expire if its time-in-force is reached. Every action relating to the order during its life is time stamped, reported, and recorded in the data set. The submission time is the time when the order departs electronically from its submitter, usually an Exchange-member firm, to the designated exchange. For example, the order may be submitted to the NYSE via the NYSE SuperDot System. The order is transmitted from the submitter to the exchange almost instantaneously, with a typical delay of less than a second. Once the order is received by the specialist, it is placed in the queue, ready for execution. When the specialist fills the order, a time-stamped report is sent to the submitter. This time stamp is the *report time* and considered the time of execution. When the investor requests cancellation or correction of an order, the submitter informs the exchange and the exchange sends back a cancellation report and the time of cancellation is recorded.

When a limit order is submitted to the NYSE, a number of things can happen. The order can be completely filled or partially filled. If partially filled, more fills may follow. Alternatively, the order may not be executed at all and expire or be cancelled. To illustrate the dynamics of typical limit orders, Figure 1 provides several examples of paths a limit order may follow from submission to termination using data for AT&T on December 29th, 1994. In the first panel the path of a buy order is followed. The order is first submitted at a limit price of $51.250 and then cancelled. It is resubmitted at $51.375, corrected, and resubmitted again at $51.500. It is executed at this price. In our analysis this sequence is

[6] ITG has graciously agreed to make their limit-order data available to all academics for research purposes. Please consult ITG's website for information about accessing the data: `http://www.itginc.com`.

treated as three observations: limit orders at \$51.250 and \$51.375 that are not executed, and a limit order at \$51.500 that is executed. In the second panel the path of a sell limit order is presented. The sell order is submitted at a limit price \$54.750, corrected and resubmitted at \$54.625, and then executed. In our analysis this sequence is treated as two observations, one at \$54.750 that is not executed and one at \$54.625 that is filled.

These examples illustrate three possible execution times that we shall distinguish in our subsequent analysis: (1) time-to-cancellation/correction; (2) time-to-first-fill; (3) time-to-completion. We shall develop separate models for (2) and (3)—they have markedly different properties—and incorporate (1) into our estimation procedures for both models.

In addition to presenting aggregate results for the entire sample of 100 stocks, we also provide detailed results for the following 16 individual stocks:

Ticker	Company Name
ABT	Abbott Labs
AXP	American Express Co
BUD	Anheuser Busch Cos Inc
C	Chrysler Corp
CL	Colgate Palmolive Co
DWD	Dean Witter Discover
GE	General Elec Co
GM	General Mtrs Corp
IBM	International Business Machine
JPM	Morgan J P & Co Inc
MOB	Mobil Corp
PAC	Pacific Telesis Group
PG	Procter & Gamble Co
SLE	Sara Lee Corp
VO	Seagram Ltd
XRX	Xerox Corp

The 16 stocks will be identified by their ticker symbols and we shall refer to the pooled sample of 100 stocks as 'POOL' in Tables 1–9 below.

3.2 TORQ vs. ITG

Because our limit-order dataset comes from a single source, ITG, it may contain certain biases that are not present in the TORQ dataset.[7] For example, ITG's clients are almost exclusively institutional investors and other broker/dealers, hence its trading desk sees few retail orders,

[7] We are grateful to the referee for some of these observations.

e.g., small or odd lots. Therefore, the ITG limit-order data reflects this institutional focus which, in turn, affects the econometric models estimated from the data.

Also, ITG's trading platform, QuantEX, is a real-time event-driven interface to multiple sources of liquidity, e.g., SuperDOT, POSIT, Instinet. A common function of the ITG trading desk is to handle residual (unfilled) trades from some of these liquidity sources—for example, orders that were unfilled in a POSIT match—hence a portion of the orders in the ITG limit-order dataset arise from an absence of liquidity.

QuantEX also provides data and analytics to support order-routing and order-handling decisions; it is an "expert system" that allows portfolio managers to program trading rules and optimization algorithms to fully automate their trading decisions. Therefore, the conditioning information behind an ITG limit order is likely to be richer than the typical TORQ limit order.

Despite these factors, ITG's limit-order dataset is likely to be an important one, particularly for institutional investors since virtually every major institution uses POSIT to some degree. Nevertheless, this dataset is only one of many possible datasets, each reflecting the style and customs of a particular set of traders. We hope to show that our application of survival analysis to limit-order execution times is promising enough to motivate others to apply the same techniques to their own datasets, and to compare their findings to ours.[8]

3.3 Summary Statistics

Summary statistics for the limit-order dataset are reported in Table 1. The number of limit orders per stock ranges from 1,160 (DWD) to 11,298 (GE), and are almost evenly split between buy orders (52.42%) and sell orders (48.58%) for the pooled sample of 375,998 limit orders. Among the sell orders, shortsales account for the majority (32.83%). Because shortsale orders are subject to the up-tick rule,[9] we expect their dynamics to differ from pure sell orders. For this reason, we omit them from our empirical analysis. Hereafter, by "sell order" we shall mean pure sell orders only.

[8] A fascinating study would be to compare the ITG limit-order dataset to a broader one like the TORQ dataset. We did conduct a preliminary comparison of the two datasets but quickly decided that it was unlikely to be very informative because of the difference in time periods (the TORQ dataset spans November 1, 1990 to January 31, 1991).

[9] Under the up-tick rule, a shortsale can be executed only if it occurs at a price higher than the preceding transaction at a different price.

Once an order is submitted, it can be partially filled, completely filled, or not filled at all due to cancellation or correction (we do not distinguish between these last two conditions). The last three columns of Table 1 report the percentage of orders which are partially filled, completely filled in the first fill, and completely filled, respectively. The orders not included in the "Partially Filled" category either expired or are cancelled. Approximately half the orders are at least partially filled and 37% are completely filled. About 30% are completely filled on the first fill.

Although most of the completed limit orders are completed with the first fill, a number do require multiple fills. Table 2 reports the percentage of completed limit orders that are completed with a given number of fills. Over 80% of completed orders are completed with the first fill, and only 1% are completed with 7 or more fills.

Summary statistics for time-to-execution and time-to-censoring are reported in Tables 3a–b. The buy orders are separated from the sell orders. Table 3a reports the mean and standard deviation for time-to-first-fill and time-to-completion and Table 3b reports the mean and standard deviation for time-to-censoring. The mean time-to-execution varies considerably across stocks. The average time-to-first-fill and time-to-completion of PG buy orders is 36.54 minutes and 37.88 minutes, respectively.[10] For PG sell orders the average times are 10.51 and 10.75 minutes. The PG numbers are representative, although there is some variability. The mean time-to-first-fill for the entire sample of buy orders is 29.22 minutes and for sell orders is 11.37 minutes. The corresponding completion averages are 30.40 minutes and 12.37 minutes.

The means and standard deviations for time until expiration or cancellation in Table 3b are presented for the orders not included in the time-to-execution statistics. The "No Fills" columns consists of the orders not executed at all and the "Partial Fills" columns consists of orders that are partially but not completely filled. As would be expected, the time-to-expiration or time-to-cancellation of the non-executed orders is longer than the fill times. For example, for PG, the average time-to-expiration or time-to-cancellation is 61.2 minutes for the buy orders versus an average time of 36.54 minutes to first fill. One consistent trend is that the average time for a buy limit order to be executed (first fill or completion) is longer

[10]Note that more observations are used to calculate the first fill numbers since it includes partially filled orders. Thus one should be cautious comparing the two times.

than that for a sell order. This result is consistent with sellers being more concerned about immediacy than buyers.

4 Hypothetical Limit Orders

Before turning to our econometric analysis of limit-order data in Sections 5 and 6, we explore the prospect of studying limit-order execution times indirectly via theoretical and empirical methods of constructing *hypothetical* limit-order executions. In our theoretical approach, described in Section 4.1, we model stock prices as a geometric Brownian motion and capture limit-order execution times as the first-passage time to the limit-price "boundary". The corresponding empirical approach, first proposed by Handa and Schwartz (1996) and which we describe in Section 4.2, is based on the same principle but uses transactions data to determine when the limit-price boundary is hit. Although both methods have the virtue of simplicity, a comparison with actual limit-order data reveals some severe biases that make hypothetical limit-order executions unreliable indicators of actual execution times.

4.1 A Theoretical Approach: First-Passage Times

From a purely statistical perspective, the execution time of a limit order can be modeled as a "first-passage time" of the stock price process $P(t)$, i.e., the first time the transaction price reaches or crosses the limit price.[11] By placing structure on the stochastic process for transaction prices, the statistical properties of execution times can be derived explicitly.

In particular, let the dynamics of $P(t)$ be given by the leading continuous-time specification for stock prices: geometric Brownian motion with drift:

$$dP(t) \;=\; \alpha P(t)dt \;+\; \sigma P(t)dW(t) \tag{4.1}$$

where α and σ are constants. Let t_0 denote the current time and P_0 denote the current stock price. Let $P_{\min}$ denote the lowest price observed in the time interval $[t_0, t_0 + t]$ (so that t is the length of the interval).

[11] Related financial applications of first-passage times include Gottlieb and Kalay (1985), Marsh and Rosenfeld (1986), Ball (1988), Cho and Frees (1988), and Harris (1990).

We assume that a limit-buy order with limit price P_l will be executed in the interval $[t_0, t_0 + t]$ if and only if $P_{\min}$ is less than or equal to P_l. Thus the probability of a limit-order execution is simply the probability that $P_{\min}$ is less than or equal to P_l in $[t_0, t_0 + t]$, i.e., the probability that the first passage of $P(t)$ to P_l occurs within $[t_0, t_0 + t]$. By modifying a formula given in Harrison (1990, p. 14) this probability can be derived exactly under (4.1) and is given by:

$$\Pr\Big(P_{min} \leq P_l | P(t_0) = P_0\Big) = 1 - \Phi\left(\frac{\log(P_0/P_l) + \mu t}{\sigma\sqrt{t}}\right) + \left(\frac{P_0}{P_l}\right)^{2\mu/\sigma^2} \Phi\left(\frac{\log(P_l/P_0) + \mu t}{\sigma\sqrt{t}}\right) \tag{4.2}$$

where $\mu \equiv \alpha - \frac{1}{2}\sigma^2$ and $\Phi(\cdot)$ is the standard normal cumulative distribution function (CDF). A similar expression can be obtained for limit-sell orders in the same manner.

Now if we denote by T the limit-order execution time—a non-negative real-valued random variable—then (4.2) yields the CDF, F(t), for T, i.e.,

$$\begin{aligned} F(t) &= \Pr(T \leq t | P(t_0) = P_0, P_l, \mu, \sigma) = \Pr(P_{\min} \leq P_l) \qquad \text{(Limit Buys)} \\ F(t) &= \Pr(T \leq t | P(t_0) = P_0, P_l, \mu, \sigma) = \Pr(P_{\max} \geq P_l) \qquad \text{(Limit Sells)}\ . \end{aligned}$$

The performance of the first-passage-time (FPT) model of limit-order executions can then be evaluated by comparing the theoretical CDF, $F(t)$, with the empirical distribution of actual limit-order execution times from our limit-order dataset.

In particular, if actual limit-order execution times T_i are distributed according to (4.2) with CDF $F(\cdot)$, then the random variables $F(T_i)$ must be uniformly distributed on the unit interval $[0, 1]$. Therefore, by tabulating the frequency counts of $F(T_i)$ within, say, each of the deciles of the uniform distribution on $[0, 1]$, i.e., $[0, .10)$, $[.10, .20), \ldots,$ $[.90, 1]$, we can see how closely the empirical behavior of limit-order execution times matches the theoretical predictions of the FPT model.

To do this, we require estimates of the parameters of $F(t)$, i.e., μ and σ. These parameters can be easily estimated from historical data via maximum likelihood:

$$\hat{\mu} = \frac{1}{N\tau}\sum_{j=1}^{N} r_j \quad , \quad \hat{\sigma}^2 = \frac{1}{N}\sum_{i=1}^{N}\frac{(r_j - \hat{\mu}\tau)^2}{\tau} \tag{4.3}$$

where N is the number of observations in the sample, $r_j \equiv \log P_j - \log P_{j-1}$ is the continuously compounded stock return over a time interval of τ units, and τ is a fixed sampling interval. Over the estimation period from August 1, 1994 to August 31, 1995, and for each of the 16 stocks in our individual-stock sample (see Table 1), we divide each trading day into 13 half-hour trading intervals, and calculate the continuously compounded return $r_j = \log P_j - \log P_{j-1}$ over each interval j, $j = 1, \ldots, 13$, where P_j is the average of bid and ask prices at the end of jth interval. Then for each stock we calculate

$$\hat{\mu} = \frac{1}{30N}\sum_{j=1}^{N} r_j \quad , \quad \hat{\sigma}^2 = \frac{1}{30N}\sum_{j=1}^{N}(r_j - 30\hat{\mu})^2$$

which are the maximum likelihood estimators of μ and σ^2, scaled by 30 to yield per-minute parameter estimates.

By inserting $\hat{\mu}$ and $\hat{\sigma}^2$ into (4.2), we obtain an estimate of the first-passage time CDF $F(t)$ as a function of t, P_0, and P_l. Therefore, for each limit order in our dataset that is executed,[12] we insert its parameters T_i, P_0, and P_l into F to obtain a random variable u_i which contributes towards the frequency count of one of the ten deciles of the Uniform$[0, 1]$.

But what about the many limit orders that are *not* executed, i.e., those that are cancelled or corrected (see Table 3b)? Eliminating them from our frequency count would clearly bias the empirical distribution towards shorter execution times (because we are discarding limit orders that have "survived"), but since they are unexecuted, we cannot evaluate the CDF for these *censored* observations. Fortunately, a well known technique for handling censored observations has been developed by Kaplan and Meier (1958), now known as the *Kaplan-Meier* estimator, and we use this procedure to incorporate limit-order cancellation/correction times into our decile counts.[13]

[12]We use time-to-first-fill only for the FPT model since this most closely matches the notion of a first-passage time. This underscores an important shortcoming of the FPT model: the inability to distinguish among time-to-first-fill, time-to-completion, and time-to-censoring.

[13]The Kaplan-Meier estimator is a nonparametric method of redistributing the probability mass of censored observations. Specifically, a censored observation indicates that the corresponding uncensored observation

Table 4 reports the percentage frequency counts of each of the ten uniform deciles for the limit orders of the 16 individual stocks in our sample. It is apparent from the entries in the last column—the 10th decile—that the limit-order data fits the FPT model very poorly. For example, 41.0% of the limit-order execution times of ABT fall into the 10th decile of the FPT model; if the FPT model were correct, this value should be close to 10%. For PG, the empirical value of the 10th decile is even higher: 74.6%! Even the smallest entry in this column—26.4% for GE—is still over twice the theoretical value of 10%, and all of the entries are statistically significantly different from 10%.[14]

The fact that there is a far higher proportion of execution times in the 10-th decile than predicted by the FPT distribution (and a correspondingly lower proportion of execution times in the lower deciles) implies that the FPT model vastly *underestimates* limit-order execution times. In fact, observe that what the FPT model predicts as the 90-th percentile of execution times is less than the empirical median execution time for DWD, PAC, PG, and SLE limit orders!

Of course, the FPT model (4.2) is predicated on the geometric Brownian motion specification (4.1) for stock prices, and if this specification is not appropriate, this can lead to the kind of inconsistencies documented in Table 4. If, for example, stock prices exhibited short-term mean reversion, e.g., an Ornstein-Uhlenbeck process (see Lo and Wang [1995]), this can lead to longer execution times than geometric Brownian motion. Unfortunately, explicit expressions for the distribution of first-passage times are unavailable for these more interesting stochastic processes.

The FPT model suffers from several other important limitations. It allows no role for limit-order size, makes no distinction between time-to-first-fill, time-to-completion, and time-

must lie to its right, but how far to the right is unknown due to the censoring. The Kaplan-Meier estimator redistributes the probability mass of the censored observation evenly over the portion of the empirical distribution function to the right of the censored observation. In the case of no censoring, the Kaplan-Meier estimator coincides with the conventional empirical distribution function, which assigns a mass of $1/n$ to each observation. To compute the frequency count for a censored sample, we first calculate the Kaplan-Meier estimator, $\hat{F}_n(x)$, of the true CDF using the transformed data $u_1, u_2, ..., u_n$ and the censoring indicators $\delta_1, \ldots, \delta_n$. The percentage frequency count for the i-th decile is given by $\hat{\pi}_i = \hat{F}_n(i/10) - \hat{F}_n((i-1)/10)$. See Kaplan and Meier (1958) and Miller (1981) for further discussion.

[14]The asymptotic z-statistics in Table 4 are calculated under the null hypothesis that the FPT model correctly describes the data. In that case, each of the percentage frequency counts $\hat{\pi}_j$ is a consistent estimator of the value 10%, and $N \times \hat{\pi}_j$ is a binomial random variable with mean $N \times 10\%$ and variance $N \times 10\% \times 90\%$ where N is the sample size. Therefore, the z-statistic $\sqrt{N}(\hat{\pi}_j - 10\%)/\sqrt{10\% \times 90\%}$ is asymptotically standard normal.

to-censoring, and cannot easily incorporate the effects of explanatory variables such as price volatility, spreads, and market conditions. Therefore, although the FPT model may be a natural theoretical framework in which to model limit-order executions, it leaves much to be desired from a practical point of view.

4.2 An Empirical Approach: Transactions Data

The empirical counterpart to the FPT model of Section 4.1 is based on first-passage times determined by the historical time series of transactions data. For example, consider a stock XYZ that trades at $50.875 at 10:37am on April 19th, 1995, and suppose that a limit-buy order for XYZ is submitted at that time at a limit price of $50.500. The first time after 10:37am that a transaction is observed at a price of $50.500 or lower, the limit order is considered executed, and the time between this transaction and 10:37am is considered the limit-order execution time. This approach has been used by Angel (1994), Handa and Schwartz (1996), Battalio et al. (1997) and others.

The primary advantage of such hypothetical limit-order executions over the FPT model is the fact that executions are determined by the historical time series of transactions data, not by geometric Brownian motion. Therefore, if the stochastic process for stock prices exhibits mean reversion or more complex forms of temporal dependence and heterogeneity, this will be incorporated into the empirical FPT model.

To compare actual limit orders with hypothetical ones generated by the empirical FPT model, we apply the following procedure to the limit orders of the 16 individual stocks from August 1994 to August 1995. For every limit-buy order in our limit-order database that had at least one fill, we create a matching hypothetical limit, i.e., the submission time and limit price are set to equal to those of the actual limit order. The time-to-execution of the hypothetical order is determined by the Transaction and Quotation (TAQ) database distributed by the NYSE, and involves searching for the first time after submission when the transaction price is less than or equal to the limit-buy price. The difference between this time and the submission time is recorded as the time-to-execution for the hypothetical limit order. This time-to-execution will obviously be a lower bound for the actual time-to-execution, hence we shall refer to it as the *lower-bound execution time*. It will equal the actual execution time only if the actual limit order is at the top of the queue or close enough to the top so that it is filled with the first incoming sell order. However, Handa and Schwartz (1996) treat this lower bound as the execution time itself.

If we continue to track the stock price after its first-passage time, we can obtain an upper bound to the execution time. The *upper-bound execution time* is either the first time during the day when the transaction price falls *below* the limit-buy price or the last time of the day the market price is equal to the limit buy price. If neither of these two conditions is met, we treat the observation as missing.

Table 5 reports the means and standard deviations of the lower-bound and upper-bound execution times, as well as those of the actual limit-order execution times (time-to-first-fill). In Figure 2 histograms of the times are presented. Together, they provide conclusive evidence that lower-bound and upper-bound execution times are poor proxies for actual limit-order execution times. In particular, the mean lower-bound execution time understates the mean actual execution time and the upper-bound execution time overstates it. For example, ABT's lower-bound mean is 15.58 minutes, its upper-bound mean is 60.12 minutes, yet its actual mean is 25.39 minutes. The standard deviations also disagree: ABT's lower-bound standard deviation is 50.61 minutes, its upper-bound standard deviation is 83.37 minutes, and the actual standard deviation is 55.84 minutes. Even for a very liquid stock such as IBM, the differences between the moments of hypothetical and actual execution times are substantial: its lower-bound mean is 16.80 minutes, its upper-bound mean is 43.26 minutes, and its actual mean is 23.41 minutes.

Table 5 also reports more formal statistical inferences in the last three columns in which the significance of the difference between the actual-time and lower-bound means are evaluated. The differences are strongly significant for all 16 stocks as the asymptotically standard normal z-statistics show—they range from 5.93 (VO) to 18.18 (GE). A similar test using differences between the upper-bound and actual-time means also yields strong rejections hence we omit them to conserve space.

Figure 2 plots the entire distributions of the lower-bound, upper-bound, and actual execution times, and a comparison of these three distributions for the 16 stocks reveals that they differ not only in one or two moments, but over their entire support. In fact, we have attempted to "shift" the distributions of the hypothetical execution times by using "n-th-passage" times in place of first-passage times[15] (as n increases, the mean of the hypothetical execution time increases also). But even selecting an n that minimizes the difference between the mean hypothetical execution time and the mean actual time does not yield similar distributions.

These results underscore several important weaknesses of the empirical FPT model, the most obvious being the assumption that the hypothetical limit order is executed when the limit price is first attained. Such an assumption implicitly presumes that there are no other limit orders with the same limit price and higher time priority, i.e., the hypothetical limit order is assumed to be at the "top of the queue". However, even intermediate hypothetical execution times such as the n-th-passage-time and lower-bound models cannot match the empirical distribution of actual limit-order execution times. Moreover, as in the theoretical FPT model, the empirical FPT model cannot easily handle varying limit-order sizes, explanatory variables, and the distinction between time-to-first-fill, time-to-completion, and time-to-censoring.

In summary, hypothetical limit-order execution times are very poor substitutes for actual limit-order data.

5 Survival Analysis

In developing an econometric model of limit-order execution times, it is important to distinguish between the various execution possibilities, to incorporate all the characteristics of the order, and to capture the influence of market conditions. We are able to incorporate all of these aspects through the application of a well-known statistical technique called *survival analysis*, which we shall review in Sections 5.1–5.4.

Since a limit order may require multiple fills to complete it (see Section 3), we must distinguish between time-to-first-fill and time-to-completion. Recognizing this distinction, we estimate two separate models, one for first fills and one for completions. Moreover, since

[15]That is, instead of determining the execution time as the first time the transaction price reaches the limit price, let it be the n-th time that the transaction price reaches the limit price. This is tantamount to assuming a lower position in the queue, and yields intermediate executions to the lower-bound (top of the queue) and upper-bound (bottom of the queue) cases.

market conditions may affect execution times differently for limit-buy orders and limit-sell orders, we also estimate separate models for buy orders and sell orders. Thus, we estimate four separate models in all.

For each model, we seek to estimate the following conditional probability, essentially the CDF of the execution time T_k of the k^{th} limit order:

$$\Pr(T_k \leq t \mid \mathbf{X}_k,\ P_{lk},\ S_k,\ I_k) \tag{5.1}$$

where $\mathbf{X}_k$ is a vector of "explanatory" variables that captures market conditions and other conditioning information at the time of submission for the k^{th} limit order, and P_{lk}, S_k, and I_k are the limit-order price, size (in shares), and side indicator (buy or sell), respectively, of the k^{th} limit order.

5.1 A Brief Review of Survival Analysis

Survival analysis is a statistical technique for analyzing positive-valued random variables such as lifetimes, failure times, or, in our case, time-to-execution. It is particularly useful for modeling the time-to-execution of limit orders because censored observations (orders terminated prior to execution) can be easily and correctly accommodated. In this section we present a brief review of survival analysis; readers interested in a more detailed exposition should consult Cox and Oakes (1984), Kalbfleisch and Prentice (1980), and Miller (1981).

Let T denote a non-negative random variable which represents the lifetime of an item, also known as the *failure time*—in our application, it is a limit-order execution time. Let $f(t)$ and $F(t)$ denote the probability density function (PDF) and CDF, respectively, of T. The *instantaneous failure rate* or *hazard rate* of T at time t, denoted by $h(t)$, is defined as

$$h(t) \;=\; \frac{f(t)}{1-F(t)}$$

since $h(t)dt$ is the probability that an item that has survived through time t will fail in the interval $[t, t+dt)$. Alternatively, we can define the *survivor* function, $S(t) \equiv 1 - F(t)$, which is the probability that an item's lifetime will be at least t. Any one of these four quantities—the PDF, the CDF, the hazard rate, and the survival function—uniquely determines the other three, and all are the focus of survival analysis.

In particular, there are two general approaches to estimating these functions: parametric and nonparametric. Parametric survival analysis, described in Section 5.2, assumes a specific parametric family for the distribution of failure times, e.g., the generalized gamma distribution, for which maximum likelihood estimation can be performed. Nonparametric survival analysis, described in Section 5.3, involves estimating the survival function nonparametrically, i.e., without resorting to any parametric assumptions.

5.2 Parametric Methods

The parametric approach to survival analysis begins with the specification of the distribution of the random variable T, from which the likelihood function is obtained. Let $(t_1, ..., t_n)$ denote a sequence of n realizations of T, possibly with censoring. We assume that we know which observations have been censored (limit-order cancellations and corrections are reported) and let $(\delta_1, ..., \delta_n)$ denote censoring indicators:

$$\delta_i \equiv \begin{cases} 1 & \text{if observation } i \text{ is censored} \\ 0 & \text{if observation } i \text{ is not censored} \end{cases} . \tag{5.2}$$

If the pairs (t_i, δ_i) are statistically independent, then the likelihood function for the data is given by:

$$\prod_{i=1}^{n} f(t_i)^{\delta_i} S(t_i)^{1-\delta_i} = \prod_{U} f(t_i) \prod_{C} S(t_i) \tag{5.3}$$

where U and C denote the indexes of the uncensored and censored observations, respectively. Given the likelihood function (5.3), the parameters of the distribution of T can be estimated via maximum likelihood.

The assumption of independence of (t_i, δ_i) is a restrictive one for limit-order execution times. Therefore, it is useful to note that the likelihood function (5.3) is appropriate under more general assumptions for the dependence structure of the data. In particular, as long as the censoring mechanism for each observation (t_i, δ_i) is conditionally independent of the probability that the limit order is executed (conditional on a vector explanatory variables

$\mathbf{X}_i$), the likelihood function of the sample is given by:

$$\prod_{i=1}^{n} f(t_i;\mathbf{X}_i)^{\delta_i} S(t_i;\mathbf{X}_i)^{1-\delta_i} = \prod_U f(t_i\mathbf{X}_i) \prod_C S(t_i;\mathbf{X}_i) . \tag{5.4}$$

Therefore, execution times t_i may be dependent and related to $\mathbf{X}_i$, but at any time t and for a given $\mathbf{X}$, the censoring mechanism must be independent of the likelihood that the limit order is executed.[16] Kalbfleisch and Prentice (1980, Chapter 3.2) provide detailed discussion of the role that explanatory variables play in survival analysis including the specific assumptions underlying (5.4).

There are several widely used distributions for failure times such as the exponential, gamma, Weibull, lognormal, and inverse Gaussian (see, for example, Cox and Oakes [1984, Table 2.1]). We choose the *generalized gamma* distribution, which nests a number of other distributions as special cases. Given this nesting, we can test the restrictions imposed on the generalized gamma specification by the simpler cases to see if the other specifications are adequate.

The generalized gamma distribution has three parameters: two shape parameters κ and p, and one scale parameter λ. Its PDF is

$$f(t) = \frac{\lambda|p|\kappa^{\kappa}(\lambda t)^{p\kappa-1}\exp(-(\lambda t)^p\kappa)}{\Gamma(\kappa)} \tag{5.5}$$

and the corresponding survival function is

$$S(t) = \begin{cases} \Gamma(\kappa,(\lambda t)^p\kappa)/\Gamma(\kappa) & \text{if } p < 0 \\ 1-\Gamma(\kappa,(\lambda t)^p\kappa)/\Gamma(\kappa) & \text{if } p > 0 \end{cases} \tag{5.6}$$

where $\Gamma(a,b)$ denotes the incomplete gamma function and $\Gamma(a)$ denotes the complete gamma function.

When $\kappa = 1$, the generalized gamma distribution reduces to a Weibull distribution, which

[16]This type of censoring, known as *independent censoring*, includes cases where the censoring mechanism depends on previous execution times, previous censoring times, or on the explanatory variables $\mathbf{X}_i$. However, censoring as a result of prices moving away from the limit price would be a violation of the underlying assumption since prices at the time of censoring are not included in $\mathbf{X}_i$.

has PDF

$$f(t) = \lambda|p|(\lambda t)^{p-1} \exp(-(\lambda t)^p) . \tag{5.7}$$

When $\kappa = 1$ and $p = 1$, the generalized gamma reduces to an exponential distribution, and when $\kappa = 0$, it reduces to a lognormal distribution.

5.3 Nonparametric Methods

In the nonparametric approach to survival analysis, the functional form of the survival distribution is not specified—the survival distribution is approximated by its empirical distribution function which is estimated nonparametrically. When there are no censored observations, the empirical survival function is simply a step function:

$$\hat{S}(t) = \prod_{t_{(i)} \leq t} \left(\frac{n-i}{n-i+1} \right)$$

where the failure times $t_{(i)}$ are ordered (hence the parenthetical subscripts) such that $t_{(1)} < t_{(2)} < \cdots < t_{(n)}$. The function $\hat{S}(t)$ is simply equal to the proportion of observations greater than t.

When there are censored observations, Kaplan and Meier's (1958) method is used to calculate the empirical distribution. Suppose, for simplicity, there are no "ties". Then the Kaplan-Meier estimator is given by

$$\hat{S}(t) = \prod_{t_{(i)} \leq t} \left(\frac{n-i}{n-i+1} \right)^{\delta_{(i)}}$$

where $\delta_{(i)}$ is defined as in (5.2). In the presence of ties, further adjustments are necessary (see, for example, Miller [1981, p. 46–51]).

5.4 Incorporating Explanatory Variables

As we discussed in Section 5.2, the incorporation of explanatory variables into the likelihood function (5.4) poses no difficulties, and allows execution times to be dependent as long as restrictions are placed on the censoring mechanism.

There are two approaches for allowing failure times to depend on explanatory variables. One approach assumes that the effect of explanatory variables on failure times can be captured by rescaling time. The other assumes that the effect of explanatory variables can be captured by rescaling the hazard rate. The former is commonly called the *accelerated failure time* specification and the latter the *proportional hazard rate* specification. An exponential factor is often used to rescale time or the hazard rate.

Specifically, an accelerated failure time model has the form

$$T = e^{\mathbf{X}'\boldsymbol{\beta}} T_0$$

where T is the time-to-execution, $\mathbf{X}$ is a vector of explanatory variables, $\boldsymbol{\beta}$ is a parameter vector, and T_0 is called the *baseline failure time* and its distribution the *baseline distribution.* The time-to-execution T is then a scaled transformation of the baseline time T_0, where the explanatory variables and coefficients determine the scaling. Because the baseline distribution is typically specified parametrically, the accelerated failure time approach falls within the parametric framework.

In the second approach, the hazard rate $h(t)$ is assumed to satisfy

$$h(t; \mathbf{X}) = h_0(t) e^{-\mathbf{X}'\boldsymbol{\beta}} \tag{5.8}$$

where $h_0(t)$ is called the *baseline hazard rate.* For obvious reasons, this is known as the *proportional hazard rate* specification. In most applications, the functional form of $h_0(t)$ can be estimated nonparametrically, hence the proportional hazard rate specification falls within the nonparametric framework.

We shall investigate both of these specifications empirically in Section 6.

6 Empirical Analysis

In this section we turn to the empirical analysis of the limit-order data described in Section 3 using the econometric models of Section 5. We focus on two specifications in particular: the generalized gamma model for the accelerated failure time specification, and the Cox proportional hazard model for the proportional hazard rate specification. In Section 6.1 we define the explanatory variables, and present the parameter estimates of the models in Section 6.2. To compare the two specifications, we consider several measures of goodness-of-fit in Section 6.3, and we discuss the economic significance of our estimates in Section 6.4.

6.1 Explanatory Variables

The dependence of time-to-execution on the limit order's characteristics and on current market conditions is captured through the inclusion of explanatory variables. The included variables are measures of the limit order's price relative to the market price and quotes, the size of the limit order, measures of the market depth, and other stock specific measures relating to volatility and liquidity. In particular, let:

$$
\begin{aligned}
P &\equiv \text{Market Price} \\
P_l &\equiv \text{Limit Price} \\
P_b &\equiv \text{Bid Price} \\
P_o &\equiv \text{Offer Price} \\
P_q &\equiv \text{Mid-Quote Price} \\
S_o &\equiv \text{Offer Size} \\
S_b &\equiv \text{Bid Size} \\
S_l &\equiv \text{Limit-Order Size}
\end{aligned}
$$

Then the following are the explanatory variables included in the limit-buy models (all variables are measured at the time of submission):

Explanatory Variables for Limit-Buy Models

$$\text{MQLP} = P_q - P_l$$

$$\text{BSID} = \begin{cases} 1 & \text{if prior trade occurred above } P_q \\ 0 & \text{if prior trade occurred at } P_q \\ -1 & \text{if prior trade occurred below } P_q \end{cases}$$

$$\text{MKD1} = \begin{cases} (1 + P_b - P_l) \times \log S_b & \text{if } P_l \leq P_b \\ 0 & \text{if } P_l > P_b \end{cases}$$

$$\text{MKD1X} = \begin{cases} (P - P_l) \times \text{MKD1} & \text{if } P \geq P_l \\ 0 & \text{if } P < P_l \end{cases}$$

$$\text{MKD2} = \begin{cases} \log S_o/(1 + P_o - P_l) & \text{if } P_o \geq P_l \\ \log S_o & \text{if } P_o < P_l \end{cases}$$

$$\text{SZSD} = \begin{cases} \log(S_l) \times (1 + P_o - P_l) & \text{if } P_o > P_l \\ \log(S_l - S_o) & \text{if } P_o = P_l \text{ and } S_l > S_o \\ 0 & \text{otherwise} \end{cases}$$

STKV = # trades last half hour / # trades last one hour

TURN = log(# trades last one hour)

LSO = log(previous month-end shares outstanding, in thousands)

LPR = log(previous month's average daily closing price)

LVO = log(previous month's average daily share volume)

The first eight variables are designed to accommodate the dynamic nature of the market place by capturing current market conditions. These are updated on a real-time basis. In contrast the last three variables are designed to facilitate differences across stocks and are updated monthly.

The variable MQLP measures the distance the limit-buy price is away from the current quote midpoint. BSID is an indicator to measure whether the prior transaction was buyer-initiated or seller-initiated (see, for example, Hausman, Lo, and MacKinlay [1992]). MKD1 is a measure of the minimum number of shares that have higher priority for execution scaled by the distance the limit-buy price is below the bid price. The variable MKD1X is an interactive term to capture nonlinearities between the market depth and market price relative to the limit-buy price. MKD2 is a measure of the liquidity available from the selling side of the market. The measure is constructed to decline as the limit-buy price decreases below the offer price. SZSD is a measure of liquidity demanded by the limit order scaled by the distance the limit-buy price is from the offer price. STKV is a short term measure capturing shifts in trading activity. It is designed to proxy for high frequency changes in volatility. TURN is a trading activity measure providing an absolute measure of volatility. LSO is the logarithm of the number of shares outstanding, LPR is the logarithm of share price, and LVO is the logarithm of average daily volume. These are primitive variables included to capture differences across stocks. They can be combined to form a number of measures one might consider including. For example, the log of price plus the log of shares outstanding is the log of market value, the log of volume minus the log of shares outstanding is the log of turnover, and the log of price plus the log of volume is approximately the log of dollar volume.

Five of the explanatory variables are redefined for the limit-sell order models. The definitions are altered so that the underlying economic interpretation of these variables is retained (although the direction of the effect may be reversed). The redefined variables are listed below.

Redefined Explanatory Variables for Limit-Sell Models

$$\text{MQLP} = P_q - P_l$$

$$\text{MKD1} = \begin{cases} (1 + P_l - P_o) \times \log S_o & \text{if } P_l \geq P_o \\ 0 & \text{if } P_l < P_o \end{cases}$$

$$\text{MKD1X} = \begin{cases} (P - P_l) \times \text{MKD1} & \text{if } P \leq P_l \\ 0 & \text{if } P > P_l \end{cases}$$

$$\text{MKD2} = \begin{cases} \log S_b / (1 + P_l - P_o) & \text{if } P_o \leq P_l \\ \log S_b & \text{if } P_o > P_l \end{cases}$$

$$\text{SZSD} = \begin{cases} \log(S_l) \times (1 + P_l - P_b) & \text{if } P_l > P_b \\ \log(S_l - S_b) & \text{if } P_l = P_b \text{ and } S_l > S_b \\ 0 & \text{otherwise} \end{cases}$$

Summary statistics and correlation matrices for the explanatory variables are available in Lo, MacKinlay, and Zhang (1997), but for completeness we provide a brief overview here. For the 100 stocks in aggregate, there is considerable variation in the explanatory variables as well as differences across buy and sell orders. For example, consider the variable MQLP: on average the limit-buy price is almost one quarter below the quote midpoint. However, there is substantial variation: the standard deviation for this variable is over one quarter. In contrast, the mean of MQLP is -0.0373 indicating that the limit-sell price is only slightly above the quote midpoint on average. Also, there is much less variation for limit-sell orders: this variable has a standard deviation of only 0.0847. Similar observations hold for the other explanatory variables.

The cross-correlations of the explanatory variables are generally relatively small, with most being less than 30% in magnitude. For example, the highest correlation between the variable STKV which captures changing volatility and the variables related to the limit order is 8.6%. A similar observation holds for TURN, the other volatility-related variable. Exceptions to the low correlation are market-depth variables, which are more highly correlated with each other. For example, the correlation between MKD1 and MKD2 is 59.1%

for limit-sell orders. Most of the results are similar across limit-buy and limit-sell orders, with one exception: the correlation of BSID with the other market-depth variables, which are much higher in magnitude for limit-sells than for limit-buys.

In the next section, we consider the empirical results for the time-to-execution models using these proposed explanatory variables. Initially we consider the accelerated failure time specification and then the proportional hazard rate specification.

6.2 Parameter Estimates

In this section, we present estimation results for an accelerated failure time specification and for a proportional hazard rate specification, using the limit-order data described in Section 3. The generalized gamma distribution is used for the baseline distribution in the accelerated failure time approach and Cox's nonparametric method is used for estimation of the proportional hazard specification.

Recall that we estimate four different models for each specification: time-to-first-fill for limit-buy orders, time-to-first-fill for limit-sell orders, time-to-completion for limit-buy orders, and time-to-completion for limit-sell orders. We estimate each specification using the pooled sample of 100 stocks, and we perform specification checks for both the pooled sample as well as for 16 individual stocks (see Section 6.3 for further discussion). The specification checks using the individual stocks allow us to assess the models' abilities to capture cross-sectional differences in execution times.

Accelerated Failure Time Specification

As discussed in Section 5.4, the accelerated failure time specification assumes that the effect of explanatory variables on the time-to-execution is to rescale the failure time itself. The sign of the coefficient of an individual explanatory variable indicates the direction of the (partial) effect of that variable on the conditional probability of executing the limit order and on the expected time-to-execution. With this specification, the time-to-execution has a generalized gamma distribution and the maximum likelihood approach is used for estimation.

To complete the accelerated failure time specification, we must choose a baseline distribution. Using the generalized gamma distribution, we obtain $f(t;\mathbf{X})$, $S(t;\mathbf{X})$, and the likelihood function by replacing λ by $\exp(-\mathbf{X}'\boldsymbol{\beta})$ in (5.5), (5.6), and (5.3). The density function is given by:

$$f(t) = \frac{\exp(-\mathbf{X}'\boldsymbol{\beta})|p|\kappa^{\kappa}(\exp(-\mathbf{X}'\boldsymbol{\beta})t)^{p\kappa-1}\exp(-(\exp(-\mathbf{X}'\boldsymbol{\beta})t)^{p}\kappa)}{\Gamma(\kappa)}. \quad (6.9)$$

Under this specification the model has two parameters in addition to the parameter vector $\boldsymbol{\beta}$: κ and p. In our estimation procedure, we reparametrize the model with $\kappa = 1/\nu^2$ and $p = \nu/\sigma$,[17] and estimate it by maximizing the likelihood in (5.3).

Given the parameter values, we can easily calculate implications of the model for the time-to-execution. For example, the conditional mean of time-to-execution is

$$\mathrm{E}[T|\mathbf{X}] = \exp(\mathbf{X}'\boldsymbol{\beta})(\nu^2)^{(\sigma/\nu)}\frac{\Gamma(\nu^{-2}+\sigma\nu^{-1})}{\Gamma(\nu^{-2})} \quad (6.10)$$

and the τ-th conditional quantile q_τ is given by

$$q_\tau = \exp(X'\beta)(\nu^2)^{(\sigma/\nu)}\Big(G^{-1}(\tau,\nu^{-2})\Big)^{(\sigma/\nu)} \quad (6.11)$$

for $\nu > 0$ (for $\nu < 0$, replace τ by $1-\tau$ on the right side of 6.11), where $G^{-1}(\tau,\nu^{-2})$ is the τ-th quantile of a gamma-distributed random variable with parameter ν^{-2}. We shall make use of these formulas below.

The estimated parameters, along with their corresponding standard errors, are reported in Table 6. The estimates of the parameters associated with the conditioning variables, with only one exception, have the expected signs and are generally are statistically significant for all four of the models.

The coefficient on the variable of MQLP is positive with z-statistics of 197 and 190 for the buy models. This indicates that the larger of the gap between the quote-mid-price and the limit-buy price the longer the expected time-to-execution. The positive sign on the variable of BSID for buy orders indicates that if the prior transaction has been seller initiated a shorter time-to-execution is expected. The positive sign of the estimated coefficient of MKD1 is consistent with the expected time-to-execution increasing with the order size and decreasing with the limit order price. The negative sign on the variable of MKD2, on the other hand,

[17]This reparametrization entails no loss of generality and is purely an artifact of the SAS procedure LIFEREG.

indicates that the greater the depth of the opposite side of the market and the closer the limit-buy price is to the offer the shorter of the expected time. The variable of MKD1X captures a nonlinear between time-to-execution and the market price and its depth. The coefficient of SZSD is positive and statistically significance in three of the four models. In the first-fill buy model the coefficient is negative, but not large in magnitude. This is not surprising, since in the case of first fills, we would expect the order size to be less important. The negative signs for the variables STKV and TURN imply that shorter time-to-execution is expected when market conditions are more active and volatile.

The importance of the three variables included to capture cross-sectional differences is not consistent. This is not of concern, however, since these primitive variables are included to capture a number of composite cross-sectional effects including market value, turnover, and dollar volume. As far as the primitive variables are concerned, the log of share price is the most important. Its coefficient is consistently strongly negative. This is to be expected since higher price stocks tend to be more liquid. We go beyond the statistical significance of the estimates in Section 6.4 where we consider the economic significance.

Simplifications of the generalized gamma to the Weibull or exponential distribution are strongly rejected. In Table 6 the estimated shape parameter for all models is more than two standard errors from one, the value consistent with the simpler distributions. For example, with the first fill buy model, the estimate for the shape parameter is -0.404 with a standard error of 0.012. Thus the estimate is more than 117 standard errors from one. Given the strength of this result, we proceed using the generalized gamma.

The survival function can be easily estimated given the parameters of the model. For the generalized gamma, the estimate of the survival function is:

$$\hat{S}(t;X) = 1 - \Gamma(\hat{\kappa}, (\hat{\lambda}t)^{\hat{p}}\hat{\kappa})/\Gamma(\hat{\kappa}) \quad \text{if the estimated } p \text{ is positive} \tag{6.12}$$

$$\hat{S}(t;X) = \Gamma(\hat{\kappa}, (\hat{\lambda}t)^{\hat{p}}\hat{\kappa})/\Gamma(\hat{\kappa}) \quad \text{if the estimated } p \text{ is negative} \tag{6.13}$$

where $\hat{\lambda} = \exp(-\mathbf{X}'\hat{\boldsymbol{\beta}})$ for a given $\mathbf{X}$. We shall present diagnostics for each of these specifications in Section 6.3.

Proportional Hazard Rate Specification

We also consider Cox's proportional hazard model for the proportional hazard rate specification (see Cox and Oakes [1984]). Under this nonparametric approach, the distribution of the time-to-execution need not be specified, but instead the hazard function is assumed to be the product of a function of the explanatory variables and the baseline hazard function $h_0(t)$, as given in (5.8). The underlying conditional density is

$$f(t;\mathbf{X}) \;=\; h_0(t)e^{-\mathbf{X}'\boldsymbol{\beta}}\exp\Big(-e^{-\mathbf{X}'\boldsymbol{\beta}}\int_0^t h_0(s)ds\Big)\,. \tag{6.14}$$

Since $h_0(t)$ is not assumed to be known, the usual maximum likelihood method is not applicable. However, a partial likelihood method can be used.

Suppose $t_{(1)} < t_{(2)} < \cdots < t_{(k)}$ are the non-censored ordered times-to-execution, and let $\mathbf{X}_{(i)}$ denote the explanatory variables associated with $t_{(i)}$. Let $\mathcal{R}_{(i)}$ denote the set of observations still "alive" at time $t_{(i)}^-$, called the *risk set* at time $t_{(i)}$—it includes all observations with times-to-execution equal to or larger than $t_{(i)}$. The partial likelihood function is given by

$$L(\boldsymbol{\beta}) \;=\; \prod_{i=1}^{k}\frac{\exp(-\mathbf{X}_{(i)}'\boldsymbol{\beta})}{\sum_{j\in\mathcal{R}(i)}\exp(-\mathbf{X}_j'\boldsymbol{\beta})}\,. \tag{6.15}$$

Observe that the baseline hazard function $h_0(t)$ does not appear in the partial likelihood (6.15). This is why it is possible to estimate the parameters $\boldsymbol{\beta}$ without specifying the baseline hazard function.

There are a number of methods which make use of the estimate of $\boldsymbol{\beta}$ to construct a nonparametric estimate of the baseline survival function $S_0(t)$, e.g., Kalbfleisch and Prentice (1981, pp. 85–86). Once the baseline survival function is obtained, the survival function incorporating the explanatory variables $\mathbf{X}$ can be estimated using $S(t;\mathbf{X}) = S_0(t)^{\exp(-\mathbf{X}'\boldsymbol{\beta})}$.

We see from (6.15) that the intercept cannot be identified, since it appears in both the numerator and the denominator and cancels out. This expression assumes no ties among the observed times. When there are ties, a modification is necessary (see Cox and Oakes [1984, pp. 102–103]).

The results for the Cox proportional hazard model are reported in Table 7. Under this specification the signs of the coefficients for the explanatory variables can be interpreted as the direction of the effect of the given variable on the expected time-to-execution of a limit order. It is worth noting that the magnitude of the estimated parameters for the accelerated failure time and proportional hazard specifications are not directly comparable. However, the signs of the parameters for the two models do have the same interpretations for the expected execution time. For the gamma model, the conditioning variables affect the execution time through the distribution of the execution time itself whereas for the Cox model, the effect is through the hazard rate of the execution time.

In Table 7, the signs of the coefficients associated with the explanatory variables are generally consistent with economic intuition but there are exceptions. For example, the coefficients for variable MQLP are consistent. The positive sign for buys and the negative sign for sells indicates the lower the limit-buy price and the higher the limit sell price the longer the expected time-to-execution. An example of an inconsistency is the sign of the estimate of the coefficient for STKV in the first fill models. The estimates are .107 and .117 for the buy and sell models respectively, indicating that expected time-to-execution is longer when the stock has had increasing trading activity the past hour.

Overall, the estimates for the accelerated failure time model are more in line with our expectations than those for the proportional hazard specification. In the next section, we provide a more detailed comparison of the specifications using model diagnostics.

6.3 Assessing Goodness of Fit

To check the goodness of fit of the two specifications estimated in Section 6.2, we use two diagnostic measures: a graphical diagnostic (Q-Q plot), and a numerical diagnostic (deciles statistics). Both suggest that the generalized gamma model fits the limit-order data best.

Q-Q Plots For Pooled Data

If $S(t;\mathbf{X})$ is the true survival function of the random variable T, then $S(T;\mathbf{X})$ (S as a function of the random variable T) must be uniformly distributed on $[0,1]$. This implies that $-\log S(T;\mathbf{X})$ has an exponential distribution with density e^{-t} for $t>0$, hence we can regard $\{\eta_i \equiv -\log S(t_i;\mathbf{X}_i)\}$ as a sequence of realizations (with censoring) of an exponential random variable. Therefore, testing whether $\{\eta_i\}$ is drawn from an exponential distribution is equivalent to testing whether $S(t;\mathbf{X})$ is the true survival function.

In practice, the true survival function is unknown. For parametric models like the gamma, the survival function depends on unknown parameters, and for the Cox proportional hazard model, even the functional form is unspecified. However, we can use the estimated survival function in its place. If the model is correctly specified, the estimated survival function $\hat{S}$ will be close to the true survival function, and the sequence $\{\hat{\eta}_i \equiv -\log \hat{S}(t_i, \mathbf{X}_i)\}$ should have properties similar to $\{\eta_i\}$. That is, we can consider $\{\hat{\eta}_i\}$ as a (censored) sample from an exponential distribution, provided that the model is correctly specified. The sequence $\{\hat{\eta}_i\}$ is called *generalized residuals* (see Cox and Oakes [1985]).

To check this hypothesis, we use Q-Q plots in which the negative logarithm of the empirical survival function of the sample $\{\hat{\eta}_i\}$ is plotted against the negative logarithm of the theoretical survival function $(-\log e^{-t} = t)$. If the model is correctly specified, the plot should be a straight line with a unit slope. Because the generalized residuals can be censored (in particular, whenever the original survival time is censored), we use the Kaplan-Meier estimator. For the gamma model, $\hat{\eta}_i = -\log \hat{S}(t_i, \mathbf{X}_i)$ for $\hat{S}$ given in (6.12). For the proportional hazard model, $\hat{\eta}_i \;\; = \;\; -\exp(-\mathbf{X}_i'\hat{\boldsymbol{\beta}}) \log \hat{S}_0(t_i)$.

Since the empirical survival function is subject to sampling variation, we do not expect to see an exact straight line, however, if the model is correctly specified, the plot should show points closely clustered about the 45-degree line. Q-Q plots that deviate from the 45-degree line are an indication of model misspecification.

Figure 3 contains the Q-Q plots for the generalized gamma model and for the Cox proportional hazard model using the pooled sample of 100 stocks. It is apparent that the proportional hazard model does not fit as well as the gamma. In contrast to the relatively straight Q-Q plots for all four of the gamma models, the plots for the proportional hazards models all curve away from the 45-degree line. The fact that they all fall below the 45-degree line at the upper range of the plots indicates that it is the right tail of the distribution that is the main source of misspecification in the proportional hazard model.

It may seem surprising that the parametric approach yields a better fit than the nonparametric approach, after all, the nonparametric approach should be consistent with virtually *any* underlying distribution. The answer lies in the fact that the Cox specification assumes a proportional hazard rate, and our analysis suggests that limit-order execution times do not satisfy this restriction.

Q-Q Plots For Individual Stock Data

The Q-Q plots of Figure 3 suggest that the generalized gamma model fits the pooled data quite well, but this says little about the performance of the model from stock to stock. To address this issue, Figures 4a and 4b contain Q-Q plots of the time-to-first-fill generalized gamma models for limit-buy and limit-sell orders, respectively, using limit-order data from the 16 individual stocks listed in Table 1.[18] These Q-Q plots show that although there is some variation in the goodness-of-fit of the generalized gamma model across stocks, the pooled model fits individual limit-order data quite well. The only stock to exhibit a poor fit for both models is GE; for practical purposes it may be worthwhile to estimate a separate model for this one stock. Nevertheless, when compared to the Q-Q plots in Figure 3 for the proportional hazard model, the generalized gamma model performs admirably stock by stock.

The First Passage Time Model Revisited

For comparison, Figure 5 contains the estimated survival functions of the theoretical first-passage model (see Section 4.1) for the first four of the 16 individual stocks listed in Table 1—ABT, AXP, BUD, and C.[19] These functions are evaluated at two randomly selected limit orders for each of the four stocks, yielding the eight panels in Figure 5.comparison, the estimated survival function of the generalized gamma model (evaluated for the same two randomly selected limit orders) for the time-to-first-fill of limit-buy orders is also plotted.

Figure 5 shows that when the limit-buy price is close to or at the market price, the theoretical model underpredicts the time-to-execution. The FPT model predicts that such an order is executed almost immediately and this manifests itself in Figure 5 as a horizontal line along the horizontal axis. In practice, such an order is typically not executed immediately.

[18] Note that the generalized gamma models are estimated with the pooled data, *not* with individual stock data. The Q-Q plots are constructed stock by stock by calculating generalized residuals for each stock using the pooled model and stock-specific limit-order data.

[19] The results for the other 12 stocks are similar hence we omit them to conserve space. In contrast to the generalized gamma model which was estimated on the entire pooled sample, the FPT survival function was estimated individually based on each of the four stocks' estimated drift and diffusion coefficients.

For example, transactions occurring at the market price may have been trades on the other side of the market, i.e., sells.

Moreover, because the generalized gamma model incorporates market information into its model of survival probabilities, it yields more realistic execution times than the FPT model. Interestingly, the two methods have similar predictions when the limit-buy price is one tick below the market. But the predictions from the two models diverge again as the limit price moves away from the market price. The assumption of a geometric Brownian motion tends to imply smaller price changes over short intervals than are observed in the data.

Assessing Statistical Significance

Although the Q-Q plots in Figures 3 and 4a–d suggest that the generalized gamma model fits best, graphical diagnostics are, of course, meant to be indicative, not conclusive. To gauge the performance of the generalized gamma model quantitatively, we follow the same procedure outlined in Section 4.1 in constructing decile statistics. In particular, we tabulate the frequency counts of the estimated CDF (evaluated at each of the failure times in our sample) for each of the deciles of the uniform distribution on $[0,1]$, i.e., $[0,.1),\ldots,[.9,1]$. If the specification is correct, these frequency counts should be close to their theoretical value of 10%.

We report decile statistics for the 16 individual stocks for the limit-buy and limit-sell time-to-first-fill models in Tables 8a and 8b, respectively.[20] Despite the large sample sizes, there are few decile statistics significantly different from 10% (asymptotic z-statistics are reported in parentheses). For example, in Table 8a the decile statistics range from 9.1% (decile 10) to 12.0% (decile 2) for ABT and although the decile 2 statistic is statistically significant (with a z-statistic of 2.6), the difference between 10% and 12% is not very meaningful from an economic standpoint. Moreover, when compared to the decile statistics of Table 4 for the FPT model, the statistics in Table 8 show that the generalized gamma model fits very well indeed.

[20] Results for time-to-completion models have been omitted to conserve space. See Lo, MacKinlay, and Zhang (1997) for the complete set of results.

Summary

In summary, using the pooled limit-order data for the 100 largest stocks in the S&P 500 universe, we find that the generalized gamma model is a much better specification for limit-order execution times than the Cox proportional hazard model. The estimated coefficients $\hat{\boldsymbol{\beta}}$ are quite similar across the two models, indicating that these coefficients are stable and robust, and not easily influenced by model specification. Nevertheless, the Q-Q plots in Figures 3 and 4 show that the generalized gamma model does provide a better fit, particularly in the right tail of the distribution.

6.4 Implications of the Generalized Gamma Model

Since the empirical analysis of Section 6.2 points to the generalized gamma model as the best specification for limit-order execution times, we focus exclusively on this specification in presenting some of the empirical properties of the parameter estimates.

To see if there is much variation in the estimated survival function from one limit order to another and as the **X**'s change, we plot in Figure 6 the estimated survival function $\hat{S}(t)$ of the limit-buy/time-to-completion model for three randomly selected limit-buy orders for each of four stocks: ABT, AXP, BUD, and C. Each plot also contains the survival function evaluated at the average **X** (averaged across the **X**'s for the three randomly chosen limit orders). From these plots, it is apparent that the estimated survival functions vary considerably from one observation to the next, implying that the conditional distribution of execution times are quite sensitive to conditioning information represented by the explanatory variables.

Figures 7 and 8 illustrate the sensitivity of the estimated survival function to the limit price and limit shares, respectively, and Table 9 documents the sensitivity of the forecast median execution time to the limit price. In Figure 7, the estimated survival function is plotted for a single randomly selected limit order for each stock, and the limit price is varied from -2 ticks to $+2$ ticks, holding all other explanatory variables fixed. Figure 7 shows that, as expected, the higher limit-buy price the higher the probability of execution over any given time interval. Moreover, the plots show that the survival time is quite sensitive to the limit price, with survival-time probabilities doubling or tripling with just a one- or two-tick change in the limit price. For example, the probability of an ABT limit-buy order surviving

20 minutes drops from about 95% to just over 20% when the limit price changes from one tick below to one tick above the original limit price. This limit-price sensitivity is common to most of the limit orders we have examined.

Table 9 contains related results, reporting the sensitivity of the forecast median execution time to the limit price. This table is based on an actual limit order for each stock. The median time is reported for the actual limit order price and for prices within two ticks in each direction. There can be substantial price sensitivity. For example, the median time for a limit-buy order for ABT submitted at the offer price of $31 + 3/8$ is 0.128 minutes. In contrast, if the buy order is submitted with a limit price of $30 + 1/8$ the median time is 100.557 minutes, dramatically longer. Similar sensitivities exist across the other orders.

A similar experiment is conducted with limit shares in Figure 8: the estimated survival function is plotted for a single randomly selected limit order for each stock, and the limit shares is varied from its original value to 10 times the original value, holding all other explanatory variables fixed. In contrast to the limit-price graphs of Figure 7, Figure 8 shows that the estimated survival functions are much less sensitive to the limit-shares variable. This somewhat surprising finding is even more striking in view of the fact that Figure 8 is based on the time-to-completion model—common intuition suggests (and the empirical evidence confirms) that the time-to-first-fill model is even less sensitive to the magnitude of limit shares. This may have important practical implications, for it implies that the size of a limit order has relatively little impact on its time-to-completion (holding other explanatory variables constant). Therefore, adjusting the size of a limit order is a relatively inefficient means for controlling execution times.

Alternatively, the insensitivity of execution times to limit size may be a symptom of a selection bias in our sample: traders may avoid submitting very large limit orders that they judge to be difficult to complete in a timely manner, choosing instead to break up large blocks into smaller orders to be submitted sequentially. Since we are conditioning on limit shares as a regressor, we have no simple way of accounting for this type of censoring in our dataset. We hope to obtain more refined data in the near future to be able to distinguish this possible explanation from others.

7 Conclusion

The behavior of limit-order execution times is critical to the price-discovery process of most market microstructure models, and we have shown that it can be quantified to a large extent by econometric models based on survival analysis and estimated with actual limit-order data using the ITG limit-order dataset. Survival analysis is designed to model lifetime data and incorporates many of the subtleties that characterize such data, e.g., skewness and censoring. We find that the generalized gamma model with an accelerated failure time specification fits the data remarkably well, and that execution times are quite sensitive to some explanatory variables, e.g., limit price, but not to others, e.g., limit shares. Despite the fact that we pool the limit orders of 100 stocks to estimate an aggregate model of execution times, our diagnostics show that such aggregate models fit reasonably well stock by stock.

We also explore the properties of hypothetical limit-order executions, constructed theoretically from the first-passage times of geometric Brownian motion and empirically from transactions data. Although such models have a certain elegance due to their parsimony, they perform very poorly when confronted with actual limit-order data.

Our findings support the practical feasibility of sophisticated dynamic order-submission strategies, strategies that trade off the price impact of market orders against the opportunity costs inherent in limit orders. We hope to explore such strategies in future research.

References

Abrahamowicz, M., MacKenzie, T. and J. Esdaile, 1996, "Time-Dependent Hazard Ratio: Modeling and Hypothesis Testing with Application in Lupus Nephritis", *Journal of the American Statistical Assocation* 91, 1432–1439.

Angel, J., 1994, "Limit Versus Market Orders", Working Paper No. FINC-1377-01-293, School of Business Administration, Georgetown University.

Ball, C., 1988, "Estimation Bias Induced by Discrete Security Prices", *Journal of Finance* 43, 841–865.

Battalio, R., Greene, J., Hatch, B. and R. Jennings, 1997, "A Comparison of Equity Limit Order Execution Quality Across Trading Venues", unpublished working paper, School of Business, Indiana University.

Biais, B., P. Hillion, and C. Spatt, 1995, An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance* 50, 1655–1689.

Belonsky, G., 1996, "Contributions, Characteristics and Rewards of Limit Orders at the NYSE", in *An Analysis of Market Organization, Order Strategy and Market Quality*, Ph.D. dissertation, UC Berkeley.

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

Cho, D. and E. Frees, 1988, "Estimating the Volatility of Discrete Stock Prices", *Journal of Finance* 43, 451–466.

Cohen, K., Maier, S., Schwartz, R. and D. Whitcomb, 1981, "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread", *Journal of Political Economy* 89, 287–305.

Chung, K., Van Ness, B. and R. Van Ness, 1997, "Limit Orders and the Bid-Ask Spread", unpublished working paper, Fogelman Colelge of Business and Economics, University of Memphis.

Cox, D.R. and D. Oakes, 1984, *Analysis of Survival Data*, New York: Chapman and Hall.

Easley, D., and M. O'Hara, 1991, "Order Form and Information in Securities Markets," *Journal of Finance* 46, 905–927.

Engle, R. and J. Russell, 1995, "Forecasting Transaction Rates: The Autoregressive Conditional Duration Model," UCSD Discussion Paper.

Fahrmeir, L. and S. Wagenpfeil, 1996, "Smoothing Hazard Functions and Time-Varying Effects in Discrete Duration and Competing Risks Models", *Journal of the American Statistical Association* 91, 1584–1594.

Foster, D. and S. Viswanathan, 1990, "A Theory Of The Interday Variations In Volume, Variance, And Trading Costs In Securities Markets", *Review of Financial Studies* 3, 593–624.

Foucault, T., 1993, "Price Formation in a Dynamic Limit Order Market", working paper, HEC.

Gamerman, D., 1991, "Dynamic Bayesian Models for Survival Data", *Applied Statistics* 40, 63–79.

Glosten, L., 1994, "Is the Electronic Open Limit Order Book Inevitable?" *Journal of Finance* 49, 1127–1161.

Gottlieb, G. and A. Kalay, 1985, "Implications of the Discreteness of Observed Stock Prices", *Journal of Finance* 40, 135–154.

Habib, M. and D. Thomas, 1986, "Chi-square Goodness-of-Fit Tests for Randomly Censored Data", *Annals of Statistics* 14, 759–765.

Handa, P., and R. Schwartz, 1996, "Limit Order Execution," *Journal of Finance* 51, 1835–1861.

Harris, L., 1986, "'A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns," *Journal of Financial Economics* 16, 99–118.

Harris, L., 1990, "Estimation of Stsock Variances and Serial Covariances from Discrete Observations", *Journal of Financial and Quantitative Analysis* 25, 291–306.

Harris, L., 1994, "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems", Working Paper No. 94–8, School of Business Administration, University of Southern California.

Harris, L. and J. Hasbrouck, 1996, "Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy", *Journal of Financial and Quantitative Analysis* 31, 213–231.

Harrison, M., 1990, *Brownian Motion and Stochastic Flow Systems.* New York: John Wiley and Sons.

Hausman, J., Lo, A. and A. C. MacKinlay, 1992, "An Ordered Probit Analysis of Transaction Stock Prices", *Journal of Financial Economics* 31, 319–379.

Johnson, Kotz, and Balakrishnan, 1995, *Continuous Univariate Distributions.* New York: Wiley.

Kalbfleisch, J.D. and R.L. Prentice, 1980, *The Statistical Analysis of Failure Time Data.* New York: Wiley.

Kaplan, E. and P. Meier, 1958, "Nonparametric Estimation from Incomplete Observations", *Journal of the American Statistical Association* 53, 457-481.

Kavajecz, K., 1996"A Specialist's Quoted Depth as a Strategic Choice Variable," Working Paper 12-96, Rodney White Center for Financial Research, Wharton School, University of Pennsylvania.

Keim, D. and A. Madhavan, 1995, "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders", *Journal of Financial Economics* 37, 371–398.

Kim, J., 1993, "Chi-Square Goodness-of-Fit Tests for Randomly Censored Data", *Annals of Statistics* 21, 1621–1639.

Kumar, P. and D. Seppi, 1993, "Limit and Market Orders with Optimizing Traders", working paper, Graduate School of Industrial Administration, Carnegie-Mellon University.

Lo, A., MacKinlay, C. and J. Zhang, 1997, "Econometric Models of Limit-Order Executions", MIT Laboratory for Financial Engineering Working Paper No. LFE–1031–97.

Lo, A. and J. Wang, 1995, "Implementing Option Pricing Models When Asset Returns Are Predictable, *Journal of Finance* 50, 87–129.

Madhavan, A., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance* 47, 607–641.

McInish, T. and R. Wood, 1995, "Hidden Limit Orders on the NYSE", *Journal of Portfolio Management* 21, 19–26.

Miller, R.G., 1981, *Survival Analysis.* New York: Wiley.

Mudholkar, G., Srivastava, D. and G. Kollia, 1996, "A Generalization of the Weibull Distribution with Application to the Analysis of Survival Data", *Journal of the American Statistical Association* 91, 1575–1583.

O'Hara, M. and G. Oldfield, 1986, "The Microeconomics of Market Making", *Journal of Financial and Quantitative Analysis* 21, 361–376.

Parlour, C., 1995, "Price Dynamics in Limit Order Markets," working paper, Queens University.

Petersen, M. and D. Fialkowski, 1994, "Posted Versus Effective Spreads", *Journal of Financial Economics* 35, 269–292.

Sandas, P., 1996, "Adverse Selection and Competitive Market Making: Empirical Evidence from a Pure Limit Order Market", working paper, Graduate School of Industrial Administration, Carnegie-Mellon University.

Seppi, D., 1996, "Liquidity Provision with Limit Orders and a Strategic Specialist," *Review of Financial Studies*, forthcoming.

U.S. Securities and Exchange Commission, 1997, *Report on the Practice of Preferencing.*

Zucker, D. and A. Karr, 1990, "Nonparametric Survival Analysis with Time-Dependent Covariate Effects: A Penalized Partial Likelihood Approach", *Annals of Statistics* 18, 329–353.

Table 1

Summary statistics for limit-order data from August 1994 to August 1995 for a pooled sample of 100 stocks (POOL), and for 16 individual stocks.

Stock	Number of Observations	% Buy Orders	% Sell Orders	% Short Sales	% Partially Filled	% Completed One Fill	% Completed Multiple Fills
POOL	375,998	52.42	14.75	32.83	53.85	30.51	37.47
ABT	4,208	52.23	17.21	30.56	55.44	34.74	41.24
AXP	3,600	49.44	16.31	34.25	51.08	31.81	37.98
BUD	2,640	50.98	16.93	32.08	49.97	31.62	38.32
C	5,606	48.59	12.67	38.74	46.83	23.01	29.32
CL	4,544	51.74	8.19	40.07	43.74	29.38	34.67
DWD	1,160	56.98	32.16	10.86	73.60	35.20	45.84
GE	11,298	50.24	10.28	39.48	48.25	22.44	27.98
GM	6,284	51.00	13.88	35.12	55.70	27.15	34.04
IBM	8,331	55.80	10.89	33.31	52.00	28.29	35.78
JPM	5,485	43.92	20.42	35.66	62.26	35.59	44.15
MOB	6,524	54.52	10.06	35.42	48.94	34.18	39.09
PAC	1,457	56.62	33.70	9.68	70.06	38.60	47.80
PG	6,619	52.97	9.13	37.91	47.74	26.76	33.09
SLE	2,207	60.13	17.90	21.98	58.30	28.16	34.79
VO	1,618	49.94	11.25	38.81	48.69	30.20	36.46
XRX	8,646	54.73	3.19	42.08	33.83	23.92	27.66

Table 2

Percentage breakdown of the total number of completed orders by the number of fills required for completion, for a pooled sample of 100 stocks (POOL) and for 16 individual stocks. The sample period of the data is August 1994 to August 1995.

Stock	Number of Fills to Completion						
	1	2	3	4	5	6	≥ 7
POOL	81.42	10.79	3.64	1.66	0.92	0.51	1.05
ABT	84.23	8.96	3.15	1.66	0.66	0.33	1.00
AXP	83.76	8.68	3.34	1.89	1.00	0.89	0.44
BUD	82.53	11.79	3.20	1.02	0.58	0.58	0.29
C	78.45	10.92	4.27	2.38	1.39	0.89	1.69
CL	84.75	10.28	2.75	1.38	0.21	0.21	0.42
DWD	76.79	13.08	5.06	1.90	0.63	1.05	1.48
GE	80.19	9.51	3.66	1.78	1.05	1.20	2.61
GM	79.76	10.37	3.53	1.51	1.51	0.94	2.38
IBM	79.07	11.72	4.48	2.21	1.11	0.40	1.01
JPM	80.62	11.17	3.98	2.12	0.83	0.64	0.64
MOB	87.43	8.14	2.49	1.03	0.36	0.24	0.30
PAC	80.76	10.33	3.18	2.38	1.75	0.95	0.64
PG	80.88	10.00	3.38	2.50	1.18	0.81	1.25
SLE	80.97	8.51	3.67	1.00	2.00	0.50	3.34
VO	82.83	10.80	3.60	1.66	0.55	0.28	0.28
XRX	86.50	10.25	2.02	0.79	0.14	0.14	0.14

Table 3a

Summary statistics for limit-order execution times for a pooled sample of 100 stocks (POOL) and for 16 individual stocks, for the sample period from August 1994 to August 1995. Columns labeled 'First Fill' report statistics for the time-to-first-fill (in minutes) for limit orders with at least one fill. Columns labeled 'Completions' report statistics for the time-to-completion (in minutes) for completed limit orders.

Stock	First Fills				Completions			
	Buy Orders		Sell Orders		Buy Orders		Sell Orders	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
POOL	27.92	54.91	11.30	26.84	29.07	55.59	12.29	28.10
ABT	25.39	55.84	8.66	16.11	25.79	56.18	9.98	21.00
AXP	29.61	61.58	15.47	37.20	31.42	62.09	16.57	38.58
BUD	28.48	49.02	11.21	23.39	29.38	49.20	12.76	24.78
C	27.99	54.01	13.66	33.76	29.52	55.06	14.50	35.88
CL	31.84	56.42	9.25	17.81	32.84	56.97	10.81	21.17
DWD	9.71	16.29	13.79	27.24	11.27	18.78	15.71	30.64
GE	39.41	65.75	9.47	20.87	40.31	65.72	10.93	22.80
GM	32.98	56.17	12.01	27.56	34.58	57.29	12.69	28.26
IBM	23.41	50.48	6.25	17.09	24.27	51.07	6.51	17.35
JPM	27.11	54.86	7.51	18.56	28.52	55.67	8.99	22.39
MOB	34.50	64.32	7.25	17.23	34.79	63.97	7.92	17.90
PAC	12.25	30.62	11.69	24.65	14.33	31.28	12.22	25.58
PG	36.42	59.94	10.10	28.51	37.64	60.00	10.33	24.74
SLE	21.14	37.32	16.83	34.51	23.28	38.39	20.42	42.34
VO	40.88	73.97	14.28	32.21	42.84	75.00	16.85	36.24
XRX	51.90	67.23	6.43	13.94	53.10	67.37	7.23	15.72

Table 3b

Summary statistics for limit-order time-to-censoring for a pooled sample of 100 stocks (POOL) and for 16 individual stocks, for the sample period from August 1994 to August 1995. Columns labeled 'No Fills' report statistics for the time-to-censoring (in minutes) for limit orders without any fills. Columns labeled 'Partial Fills' report statistics for the time-to-censoring (in minutes) for limit orders partially but not completely filled.

Stock	No Fills				Partial Fills			
	Buy Orders		Sell Orders		Buy Orders		Sell Orders	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
POOL	46.92	72.31	34.15	53.94	41.14	55.65	49.16	65.71
ABT	44.50	71.33	35.43	48.06	63.57	75.54	17.92	24.93
AXP	32.34	52.71	47.92	66.82	28.19	22.25	36.63	37.00
BUD	62.04	83.26	39.93	63.61	57.65	86.04	29.38	26.21
C	60.29	80.34	41.13	60.46	33.27	60.83	24.33	21.23
CL	45.57	70.20	21.45	38.43	19.31	13.43	54.10	60.58
DWD	30.78	38.71	40.07	59.78	33.85	24.82	72.06	63.79
GE	31.55	62.27	22.13	36.27	39.46	54.88	49.64	79.32
GM	63.99	88.62	39.93	59.49	49.36	63.33	99.54	65.14
IBM	48.88	77.51	16.82	27.72	32.67	50.93	6.23	5.56
JPM	36.40	62.89	29.61	47.86	31.74	44.38	51.01	55.56
MOB	52.89	77.26	23.10	37.11	80.55	103.36	38.55	39.72
PAC	36.13	46.44	49.70	65.21	23.82	20.55	64.18	73.19
PG	71.41	93.91	37.76	65.70	42.52	70.44	48.90	100.54
SLE	41.92	70.81	36.66	50.77	47.55	38.69	43.04	10.43
VO	48.45	65.97	45.00	57.74	29.33	22.09	65.14	77.39
XRX	51.48	78.14	32.31	55.71	43.82	71.24	35.95	40.12

Table 4

Goodness-of-fit diagnostics for the first-passage-time (FPT) model for a sample of 16 individual stocks, for the sample period from August 1994 to August 1995. For each stock, the percentage of execution times that fall within each of the 10 theoretical deciles of the FPT model are tabulated. If the FPT model is correct, the expected percentage falling in each decile is 10%. Test statistics that are asymptotically standard normal under the FPT model are given in parentheses.

Stock	Decile:									
	1	2	3	4	5	6	7	8	9	10
ABT	10.0 (0.1)	5.6 (−8.2)	5.9 (−7.3)	6.5 (−6.1)	6.2 (−6.7)	6.5 (−6.1)	6.1 (−6.9)	6.1 (−6.9)	5.7 (−7.9)	41.0 (26.9)
AXP	11.1 (1.3)	6.0 (−6.6)	7.0 (−4.7)	6.0 (−6.6)	6.2 (−6.1)	6.0 (−6.5)	6.0 (−6.6)	5.8 (−7.0)	5.6 (−7.4)	40.3 (24.0)
BUD	10.6 (0.7)	7.7 (−2.8)	6.0 (−5.4)	6.3 (−4.9)	5.6 (−6.1)	6.4 (−4.8)	6.5 (−4.6)	5.2 (−6.9)	5.4 (−6.5)	39.4 (19.2)
C	5.6 (−6.9)	4.6 (−9.2)	4.7 (−9.1)	5.3 (−7.6)	5.4 (−7.3)	7.0 (−4.3)	8.3 (−2.3)	8.7 (−1.6)	7.5 (−3.4)	42.9 (24.0)
CL	5.0 (−10.2)	5.2 (−9.6)	5.7 (−8.2)	7.1 (−5.0)	7.4 (−4.5)	8.5 (−2.3)	9.1 (−1.4)	9.3 (−1.1)	9.4 (−0.9)	33.3 (21.9)
DWD	5.9 (−3.8)	1.8 (−13.0)	3.0 (−8.9)	3.3 (−8.1)	4.1 (−6.4)	4.8 (−5.2)	3.5 (−7.5)	3.3 (−8.0)	2.9 (−9.0)	66.9 (25.9)
GE	12.1 (4.0)	7.9 (−4.7)	8.4 (−3.6)	7.7 (−5.3)	8.0 (−4.6)	8.2 (−4.0)	7.2 (−6.7)	7.1 (−7.0)	7.0 (−7.4)	26.4 (23.1)
GM	11.2 (1.6)	6.5 (−6.1)	6.1 (−7.0)	6.5 (−6.2)	5.7 (−8.1)	6.6 (−6.0)	6.9 (−5.2)	6.5 (−6.0)	6.0 (−7.2)	38.2 (24.8)
IBM	8.6 (−2.6)	5.8 (−9.6)	5.9 (−9.3)	6.5 (−7.5)	6.8 (−6.8)	6.9 (−6.4)	6.8 (−6.7)	6.2 (−8.3)	6.4 (−7.9)	40.2 (32.5)
JPM	9.0 (−1.5)	6.0 (−7.3)	6.6 (−5.9)	6.3 (−6.6)	5.4 (−8.8)	5.4 (−8.8)	5.6 (−8.2)	5.5 (−8.4)	5.1 (−9.6)	45.1 (30.4)
MOB	10.3 (0.6)	6.7 (−6.6)	7.1 (−5.7)	6.3 (−7.5)	6.9 (−6.2)	6.3 (−7.5)	7.1 (−5.6)	7.2 (−5.3)	7.3 (−5.2)	34.8 (25.9)
PAC	7.0 (−2.9)	2.3 (−12.7)	2.2 (−12.8)	2.1 (−13.4)	3.2 (−9.3)	3.2 (−9.3)	2.6 (−11.5)	2.9 (−10.3)	2.5 (−11.9)	71.9 (33.5)
PG	0.0 (—)	0.0 (—)	0.3 (−82.1)	0.5 (−66.6)	1.1 (−42.2)	1.6 (−33.3)	2.4 (−25.0)	4.6 (−12.8)	13.8 (5.4)	74.6 (73.4)
SLE	7.6 (−2.9)	6.4 (−4.7)	5.5 (−6.4)	5.4 (−6.6)	5.6 (−6.1)	5.5 (−6.3)	4.9 (−7.6)	4.7 (−8.0)	5.0 (−7.4)	49.1 (25.0)
VO	9.2 (−0.8)	7.7 (−2.2)	6.0 (−4.4)	8.5 (−1.4)	5.4 (−5.3)	6.7 (−3.5)	6.7 (−3.4)	6.0 (−4.4)	6.3 (−4.0)	36.9 (14.6)
XRX	7.2 (−5.7)	6.1 (−8.7)	6.7 (−7.0)	7.6 (−4.8)	9.1 (−1.6)	9.3 (−1.2)	10.3 (0.6)	10.3 (0.4)	11.1 (1.8)	22.3 (15.6)

Table 5

Comparison of hypothetical time-to-first-fill (lower and upper bounds, in minutes) for limit orders simulated using TAQ data with actual time-to-first-fill for limit orders for 16 stocks, for the sample period from August 1994 to August 1995. The 'Actual Minus TAQ' column reports the difference between the actual time-to-first-fill and the TAQ lower bound. The z-statistics are asymptotically standard normal under the null hypothesis that the expected difference is zero.

Stock	TAQ Hypothetical				Actual		Actual Minus TAQ		
	Lower Bound		Upper Bound						
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	z
ABT	15.58	50.61	60.12	83.37	25.39	55.84	9.81	22.88	12.18
AXP	18.21	57.34	66.12	89.05	29.61	61.58	11.41	23.42	13.21
BUD	18.04	41.34	67.69	85.99	28.48	49.02	10.44	23.68	9.34
C	18.44	48.59	56.48	82.54	27.99	54.01	9.55	24.90	9.94
CL	25.88	51.51	66.71	77.51	31.84	56.42	5.96	19.31	8.44
DWD	5.05	10.49	44.33	61.65	9.71	16.29	4.66	12.35	6.52
GE	24.27	57.46	65.08	82.44	39.41	65.75	15.14	31.93	18.18
GM	19.97	49.47	55.56	74.76	32.98	56.17	13.02	26.95	14.95
IBM	16.80	44.57	43.26	72.17	23.41	50.48	6.61	20.37	12.52
JPM	20.27	49.31	55.41	77.27	27.11	54.86	6.84	22.35	9.67
MOB	28.49	61.53	61.38	83.18	34.50	64.32	6.01	21.94	9.40
PAC	1.82	4.82	57.09	81.41	12.25	30.62	10.43	28.64	7.83
PG	27.65	55.10	66.32	81.25	36.42	59.94	8.77	24.80	11.35
SLE	6.21	24.02	67.25	86.86	21.14	37.32	14.94	27.22	12.25
VO	32.64	70.98	89.32	100.99	40.88	73.97	8.24	23.29	5.93
XRX	44.47	66.47	70.81	80.69	51.90	67.23	7.43	13.46	8.84

Table 6

Parameter estimates of the accelerated-failure-time specification of limit-order executions under the generalized gamma distribution for limit orders of a pooled sample of 100 stocks from August 1994 to August 1995. The variable 'INTCP' denotes the intercept and the definitions of the remaining explanatory variables are given in the text. z-statistics are asymptotically standard normal under the null hypothesis that the corresponding coefficient is 0.

Variable	Buy-Limit-Order Model			Sell-Limit-Order-Model		
	Estimate	S.E.	z	Estimate	S.E.	z
			Time-to-First-Fill			
INTCP	6.507	0.207	31.365	4.979	0.308	16.181
MQLP	8.989	0.046	197.180	−13.674	0.161	−85.034
BSID	−5.613	0.076	−74.168	6.852	0.154	44.543
MKD1	0.641	0.005	127.608	0.476	0.008	59.106
MKD1X	−0.920	0.012	−79.882	0.903	0.058	15.464
MKD2	−0.353	0.005	−66.409	−0.171	0.008	−22.617
SZSD	−0.015	0.005	−3.250	0.091	0.007	13.308
STKV	−0.414	0.052	−7.984	−0.563	0.080	−7.048
TURN	−0.252	0.012	−21.217	−0.331	0.018	−18.757
LSO	0.278	0.014	19.969	0.187	0.021	8.939
LPR	−0.529	0.019	−28.101	−0.272	0.028	−9.872
LVO	−0.082	0.015	−5.563	−0.000	0.021	−0.022
SCALE	1.927	0.006	344.736	1.804	0.008	224.781
SHAPE	−0.404	0.012	−33.781	−0.526	0.018	−29.574
			Time-to-Completion			
INTCP	6.468	0.212	30.560	5.052	0.317	15.959
MQLP	8.744	0.046	189.979	−13.307	0.163	−81.713
BSID	−5.517	0.077	−71.582	6.766	0.158	42.892
MKD1	0.620	0.005	121.052	0.457	0.008	55.189
MKD1X	−0.895	0.012	−76.409	0.943	0.060	15.708
MKD2	−0.334	0.005	−61.798	−0.148	0.008	−19.122
SZSD	0.069	0.005	14.581	0.186	0.007	25.974
STKV	−0.394	0.053	−7.451	−0.568	0.082	−6.911
TURN	−0.259	0.012	−21.409	−0.327	0.018	−18.045
LSO	0.281	0.014	19.787	0.181	0.022	8.427
LPR	−0.498	0.019	−25.950	−0.229	0.028	−8.069
LVO	−0.092	0.015	−6.139	−0.021	0.022	−0.940
SCALE	1.960	0.006	338.053	1.854	0.008	221.060
SHAPE	−0.410	0.012	−32.965	−0.566	0.018	−30.901

Table 7

Parameter estimates of the Cox proportional-hazard model of limit-order executions for limit orders of a pooled sample of 100 stocks from August 1994 to August 1995. The definitions of the explanatory variables are given in the text. z-statistics are asymptotically standard normal under the null hypothesis that the corresponding coefficient is 0.

Variable	Buy-Limit-Order Model			Sell-Limit-Order-Model		
	Estimate	S.E.	z	Estimate	S.E.	z
			Time-to-First-Fill			
MQLP	5.312	0.032	164.153	−5.349	0.058	−92.184
BSID	−2.980	0.041	−72.266	2.213	0.078	28.434
MKD1	0.324	0.003	95.811	0.283	0.005	51.872
MKD1X	−0.451	0.009	−50.591	0.082	0.045	1.824
MKD2	−0.216	0.003	−64.034	−0.172	0.005	−36.805
SZSD	0.017	0.003	5.712	0.083	0.004	18.966
STKV	0.104	0.032	3.253	−0.009	0.050	−0.185
TURN	−0.024	0.007	−3.316	−0.117	0.011	−10.916
LSO	0.119	0.009	13.418	0.078	0.013	5.851
LPR	−0.355	0.012	−29.942	−0.168	0.018	−9.410
LVO	−0.143	0.009	−15.806	−0.019	0.014	−1.380
			Time-to-Completion			
MQLP	5.178	0.033	155.638	−5.208	0.063	−82.936
BSID	−2.925	0.042	−69.522	2.146	0.079	27.101
MKD1	0.304	0.003	89.433	0.259	0.006	46.956
MKD1X	−0.426	0.009	−46.763	0.085	0.046	1.875
MKD2	−0.198	0.003	−58.344	−0.149	0.005	−31.536
SZSD	0.078	0.003	25.847	0.150	0.004	33.331
STKV	0.117	0.032	3.611	−0.048	0.051	−0.942
TURN	−0.037	0.007	−5.087	−0.106	0.011	−9.720
LSO	0.128	0.009	14.253	0.056	0.014	4.127
LPR	−0.334	0.012	−27.776	−0.146	0.018	−8.088
LVO	−0.155	0.009	−16.782	−0.033	0.014	−2.343

Table 8a

Goodness-of-fit diagnostics for the accelerated-failure-time specification of the limit-buy time-to-first-fill model under the generalized gamma distribution for a sample of 16 individual stocks, for the sample period from August 1994 to August 1995. For each stock, the percentage of execution times that fall within each of the 10 theoretical deciles of the accelerated-failur-time specification are tabulated. If this specification is correct, the expected percentage falling in each decile is 10%. Test statistics which are asymptotically standard normal under this specification are given in parentheses.

Stock	Decile:									
	1	2	3	4	5	6	7	8	9	10
ABT	11.5	12.0	10.4	10.0	10.0	9.2	8.9	9.5	9.6	9.1
	(2.0)	(2.6)	(0.5)	(−0.0)	(0.0)	(−1.2)	(−1.7)	(−0.7)	(−0.7)	(−1.3)
AXP	8.0	10.0	10.2	8.5	10.0	10.8	9.5	11.5	11.1	10.6
	(−2.9)	(0.1)	(0.2)	(−2.2)	(−0.1)	(1.0)	(−0.7)	(1.8)	(1.3)	(0.8)
BUD	8.9	9.4	9.9	10.0	10.2	9.3	10.8	10.3	10.4	10.9
	(−1.3)	(−0.6)	(−0.2)	(−0.0)	(0.2)	(−0.8)	(0.8)	(0.3)	(0.4)	(0.9)
C	7.1	9.1	11.8	9.1	11.4	10.8	10.2	11.1	9.5	10.0
	(−4.1)	(−1.2)	(2.0)	(−1.2)	(1.6)	(0.9)	(0.3)	(1.2)	(−0.6)	(0.0)
CL	9.0	9.4	8.8	9.5	11.3	10.0	10.0	10.9	10.0	11.0
	(−1.5)	(−0.9)	(−1.9)	(−0.7)	(1.9)	(−0.0)	(−0.0)	(1.3)	(0.0)	(1.4)
DWD	11.0	9.9	9.8	11.6	9.2	9.5	9.3	9.1	10.2	10.4
	(0.7)	(−0.1)	(−0.1)	(1.1)	(−0.6)	(−0.4)	(−0.5)	(−0.7)	(0.2)	(0.3)
GE	10.6	11.2	10.2	11.3	9.9	10.3	9.4	9.5	9.0	8.8
	(1.1)	(2.3)	(0.4)	(2.5)	(−0.2)	(0.6)	(−1.2)	(−1.1)	(−2.2)	(−2.7)
GM	9.2	9.7	10.5	10.3	11.2	10.2	10.4	9.1	9.2	10.3
	(−1.3)	(−0.4)	(0.7)	(0.4)	(1.6)	(0.3)	(0.5)	(−1.4)	(−1.2)	(0.5)
IBM	7.3	10.7	11.8	10.7	10.3	10.8	9.7	9.3	9.3	10.0
	(−5.4)	(1.2)	(2.9)	(1.2)	(0.4)	(1.4)	(−0.6)	(−1.2)	(−1.2)	(0.1)
JPM	9.0	11.2	10.9	11.5	11.2	9.9	10.1	9.0	8.7	8.5
	(−1.5)	(1.7)	(1.3)	(2.0)	(1.6)	(−0.1)	(0.1)	(−1.5)	(−2.0)	(−2.3)
MOB	11.1	11.9	10.2	9.6	9.2	8.9	9.2	9.8	10.0	10.0
	(1.8)	(2.9)	(0.4)	(−0.6)	(−1.4)	(−2.0)	(−1.3)	(−0.3)	(0.0)	(−0.0)
PAC	10.8	12.1	10.2	10.9	8.5	10.9	9.6	9.2	8.8	9.0
	(0.7)	(1.6)	(0.2)	(0.7)	(−1.3)	(0.7)	(−0.4)	(−0.7)	(−1.1)	(−0.9)
PG	11.1	9.0	9.6	10.1	9.9	9.5	10.3	9.6	10.2	10.8
	(1.8)	(−1.7)	(−0.8)	(0.1)	(−0.2)	(−0.9)	(0.5)	(−0.8)	(0.4)	(1.3)
SLE	8.9	9.7	10.4	11.0	10.1	10.7	9.4	10.6	9.9	9.3
	(−1.2)	(−0.3)	(0.5)	(1.0)	(0.1)	(0.7)	(−0.6)	(0.6)	(−0.1)	(−0.8)
VO	11.3	9.8	10.4	10.5	10.5	9.7	9.3	9.7	9.6	9.1
	(1.1)	(−0.2)	(0.4)	(0.4)	(0.4)	(−0.2)	(−0.6)	(−0.2)	(−0.4)	(−0.8)
XRX	9.4	9.4	8.9	10.1	9.4	10.3	9.9	10.3	10.8	11.6
	(−1.0)	(−1.0)	(−2.1)	(0.2)	(−1.1)	(0.5)	(−0.3)	(0.5)	(1.3)	(2.6)

Table 8b

Goodness-of-fit diagnostics for the accelerated-failure-time specification of the limit-sell time-to-first-fill model under the generalized gamma distribution for a sample of 16 individual stocks, for the sample period from August 1994 to August 1995. For each stock, the percentage of execution times that fall within each of the 10 theoretical deciles of the accelerated-failur-time specification are tabulated. If this specification is correct, the expected percentage falling in each decile is 10%. Test statistics which are asymptotically standard normal under this specification are given in parentheses.

Stock	Decile:									
	1	2	3	4	5	6	7	8	9	10
ABT	9.9 (−0.1)	10.6 (0.4)	13.1 (2.2)	11.7 (1.2)	10.2 (0.2)	9.6 (−0.3)	10.3 (0.2)	9.3 (−0.6)	7.2 (−2.5)	8.0 (−1.7)
AXP	7.8 (−1.8)	9.8 (−0.2)	9.2 (−0.6)	9.3 (−0.5)	11.0 (0.7)	10.2 (0.1)	8.8 (−0.9)	11.8 (1.2)	9.9 (−0.1)	12.3 (1.5)
BUD	12.7 (1.5)	9.7 (−0.2)	6.4 (−2.7)	9.7 (−0.2)	11.1 (0.6)	10.2 (0.1)	10.9 (0.6)	9.0 (−0.6)	10.4 (0.2)	9.9 (−0.1)
C	8.8 (−0.9)	7.6 (−2.0)	12.4 (1.6)	9.4 (−0.4)	11.4 (1.0)	8.4 (−1.2)	11.6 (1.1)	8.4 (−1.3)	10.4 (0.3)	11.6 (1.1)
CL	7.2 (−1.8)	9.6 (−0.2)	10.7 (0.4)	8.8 (−0.7)	11.8 (0.9)	9.3 (−0.4)	8.9 (−0.6)	13.9 (1.8)	11.8 (0.9)	5.7 (−3.0)
DWD	11.0 (0.5)	12.7 (1.3)	9.6 (−0.2)	7.1 (−1.9)	8.7 (−0.7)	7.2 (−1.8)	11.0 (0.5)	11.4 (0.7)	9.7 (−0.2)	11.6 (0.8)
GE	9.0 (−0.8)	10.0 (−0.0)	10.1 (0.0)	10.7 (0.6)	10.7 (0.6)	11.7 (1.3)	9.8 (−0.1)	10.0 (0.0)	9.6 (−0.3)	8.4 (−1.5)
GM	7.7 (−2.1)	9.0 (−0.9)	9.5 (−0.4)	10.4 (0.4)	9.9 (−0.1)	10.1 (0.1)	9.0 (−0.9)	10.8 (0.6)	11.9 (1.5)	11.8 (1.4)
IBM	9.4 (−0.6)	10.0 (0.0)	11.1 (1.0)	11.6 (1.4)	7.9 (−2.1)	10.6 (0.6)	10.6 (0.5)	10.6 (0.5)	8.8 (−1.1)	9.4 (−0.6)
JPM	8.4 (−1.6)	9.0 (−1.0)	8.5 (−1.5)	12.9 (2.4)	9.7 (−0.2)	10.4 (0.4)	8.9 (−1.0)	10.2 (0.2)	11.1 (1.0)	10.8 (0.7)
MOB	12.0 (1.4)	12.2 (1.5)	10.6 (0.4)	9.6 (−0.3)	9.9 (−0.1)	9.5 (−0.4)	10.6 (0.5)	7.4 (−2.2)	9.5 (−0.4)	8.7 (−1.1)
PAC	8.0 (−1.5)	7.5 (−1.9)	8.4 (−1.1)	11.3 (0.8)	12.2 (1.3)	9.4 (−0.4)	11.0 (0.6)	9.5 (−0.3)	11.8 (1.1)	10.9 (0.6)
PG	12.0 (1.3)	13.5 (2.2)	12.3 (1.4)	10.2 (0.1)	11.1 (0.7)	7.6 (−1.8)	10.7 (0.5)	8.0 (−1.6)	7.5 (−2.0)	7.0 (−2.4)
SLE	12.0 (1.0)	9.7 (−0.2)	9.4 (−0.3)	9.7 (−0.2)	12.5 (1.2)	9.1 (−0.5)	10.0 (−0.0)	10.7 (0.3)	9.8 (−0.1)	5.5 (−3.1)
VO	9.5 (−0.2)	11.2 (0.4)	13.1 (1.1)	6.8 (−1.5)	13.4 (1.1)	7.2 (−1.2)	11.2 (0.4)	10.8 (0.3)	7.3 (−1.2)	9.6 (−0.2)
XRX	12.0 (0.8)	12.1 (0.9)	9.2 (−0.4)	6.9 (−1.6)	11.7 (0.7)	12.3 (1.0)	7.3 (−1.4)	10.5 (0.2)	7.4 (−1.4)	10.5 (0.2)

Table 9

Sensitivity of forecasts of median time-to-completion to limit-order price, using one randomly selected limit-buy order from the sample of completed limit orders for each stock. Forecasts are obtained from the accelerated-failure-time specification of the limit-buy time-to-completion model using the generalized gamma distribution for a pooled sample of 100 stocks for the sample period from August 1994 to August 1995. Forecasts are reported for the limit-order price actually submitted (0 Ticks), for the limit-order price minus one and two ticks, and for the limit-order price plus one and two ticks. The completion time is the actual time-to-completion for the limit order. The order size is in units of round lots (100 shares) and the market conditions used are those at the time the order was submitted.

Stock	Limit-Buy Order		Market Conditions			Conditional Median Execution Time					Completion Time
	Price	Shares	Bid	Offer	Market	−2 Ticks	−1 Tick	0 Tick	+1 Tick	+2 Ticks	
ABT	31+3/8	18	31+1/4	31+3/8	31+1/4	100.557	23.144	0.128	0.066	0.013	0.167
AXP	30+1/2	2	30+1/2	30+5/8	30+1/2	347.512	111.846	30.434	0.351	0.071	18.150
BUD	50+1/4	5	50+1/4	50+3/4	50+1/2	3031.245	582.587	111.970	21.520	4.136	30.250
C	47+7/8	1	47+5/8	48+1/2	48	221.718	43.855	8.074	1.716	0.339	8.133
CL	55+7/8	8	55+3/4	56+1/8	55+7/8	141.439	26.897	5.115	0.973	0.176	11.133
DWD	47+1/8	9	47	47+3/8	47+1/8	119.146	22.595	4.285	0.813	0.147	2.983
GE	49+3/8	4	49+3/8	49+1/2	49+1/2	224.268	62.424	15.586	0.602	0.121	13.967
GM	46+3/8	2	46+5/8	46+7/8	46+3/4	369.071	101.893	69.428	47.406	15.670	30.650
IBM	67+3/4	50	67+5/8	68+1/8	67+7/8	548.279	99.641	18.108	3.291	0.598	10.350
JPM	60+5/8	2	60+3/4	60+7/8	60+7/8	903.572	242.474	60.485	14.002	0.913	7.867
MOB	80+1/8	19	79+3/4	80+1/4	80	21.162	3.861	0.701	0.118	0.024	0.183
PAC	28+3/8	2	28+3/8	28+1/2	28+1/2	256.649	77.549	20.356	0.624	0.126	22.217
PG	58+1/8	40	58+1/8	58+1/4	58+1/4	626.419	195.124	51.723	0.661	0.133	26.883
SLE	25+1/4	119	25+1/8	25+1/4	25+1/4	470.946	139.102	1.472	0.296	0.060	3.483
VO	29+7/8	14	29+5/8	29+7/8	29+7/8	100.931	12.559	2.053	0.413	0.083	1.317
XRX	107+1/4	1	107+1/4	107+5/8	107+1/2	750.138	181.122	41.588	5.480	1.057	60.467

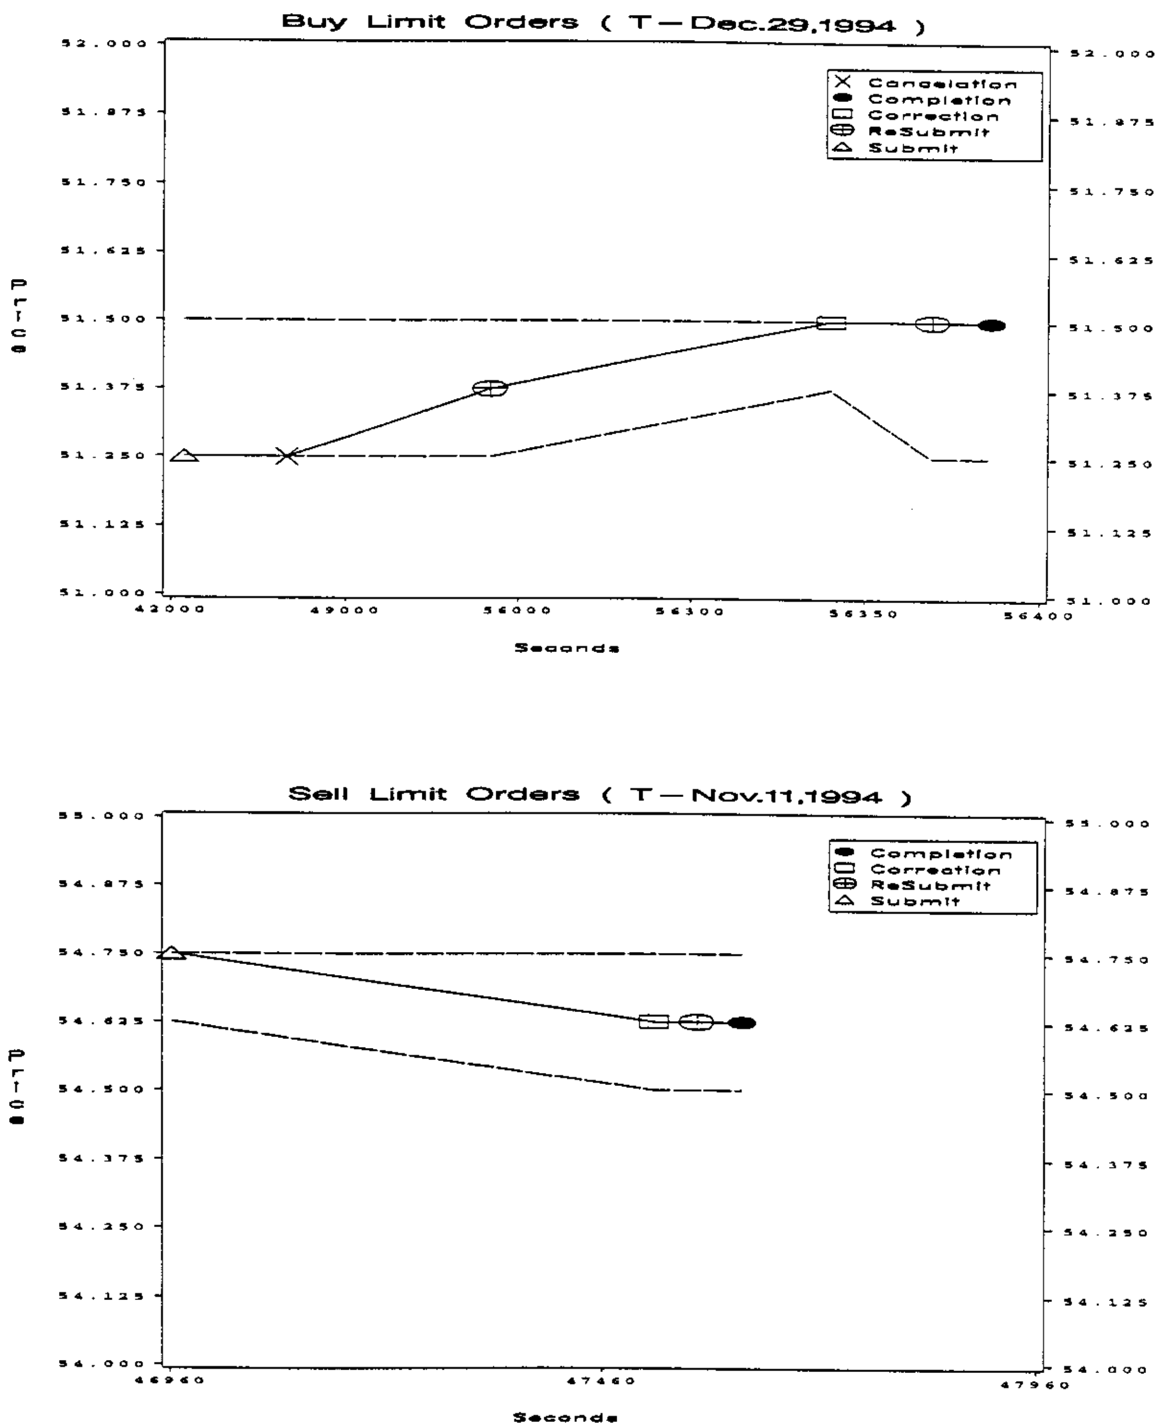

Buy Limit Orders (T—Dec.29,1994)
Cancelation
Completion
Correction
ReSubmit
Submit
52.000
51.875
51.750
51.625
51.500
51.375
51.250
51.125
51.000
42000
49000
56000
56300
56350
56400
Seconds
Price
Sell Limit Orders (T—Nov.11,1994)
Completion
Correction
ReSubmit
Submit
55.000
54.875
54.750
54.625
54.500
54.375
54.250
54.125
54.000
46960
47460
47960
Seconds
Price

Figure 1

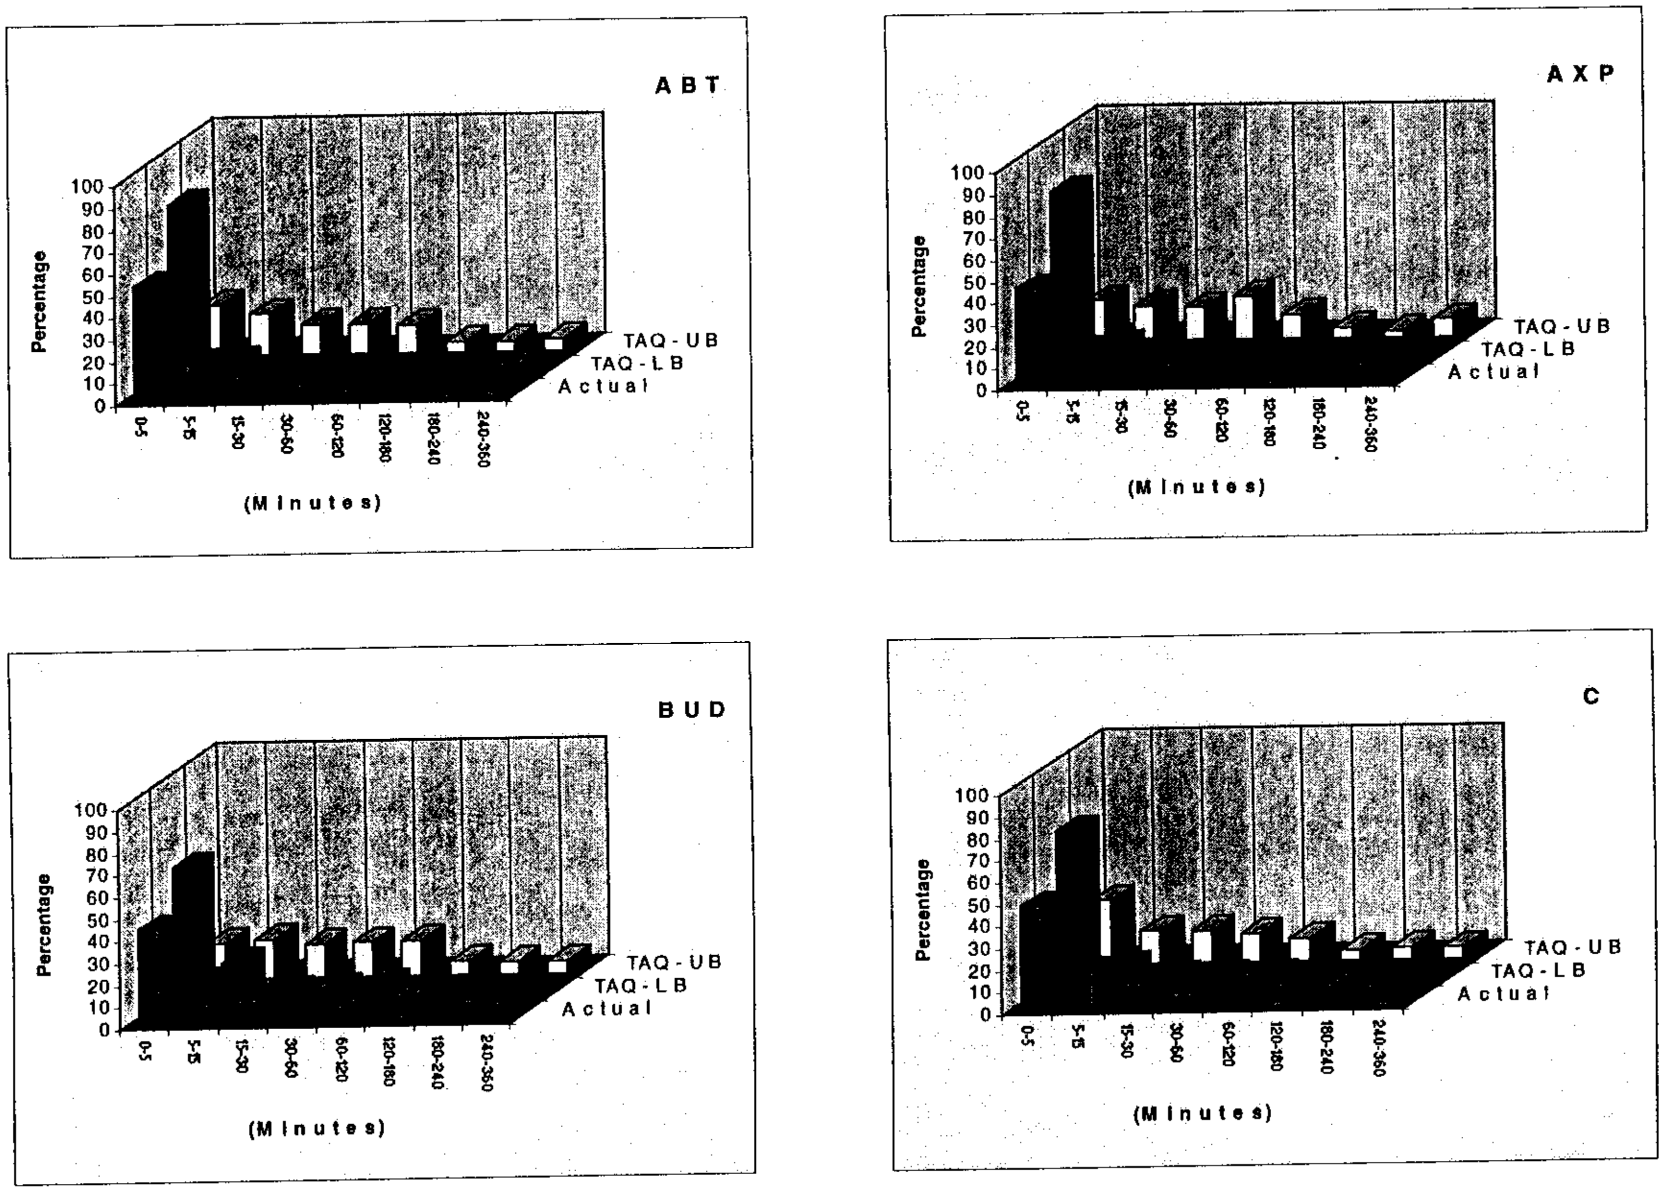
ABT
Percentage
100
90
80
70
60
50
40
30
20
10
0
0-5
5-15
15-30
30-60
60-120
120-180
180-240
240-360
(Minutes)
TAQ - UB
TAQ - LB
Actual
AXP
Percentage
(Minutes)
TAQ - UB
TAQ - LB
Actual
BUD
Percentage
(Minutes)
TAQ - UB
TAQ - LB
Actual
C
Percentage
(Minutes)
TAQ - UB
TAQ - LB
Actual

Figure 2

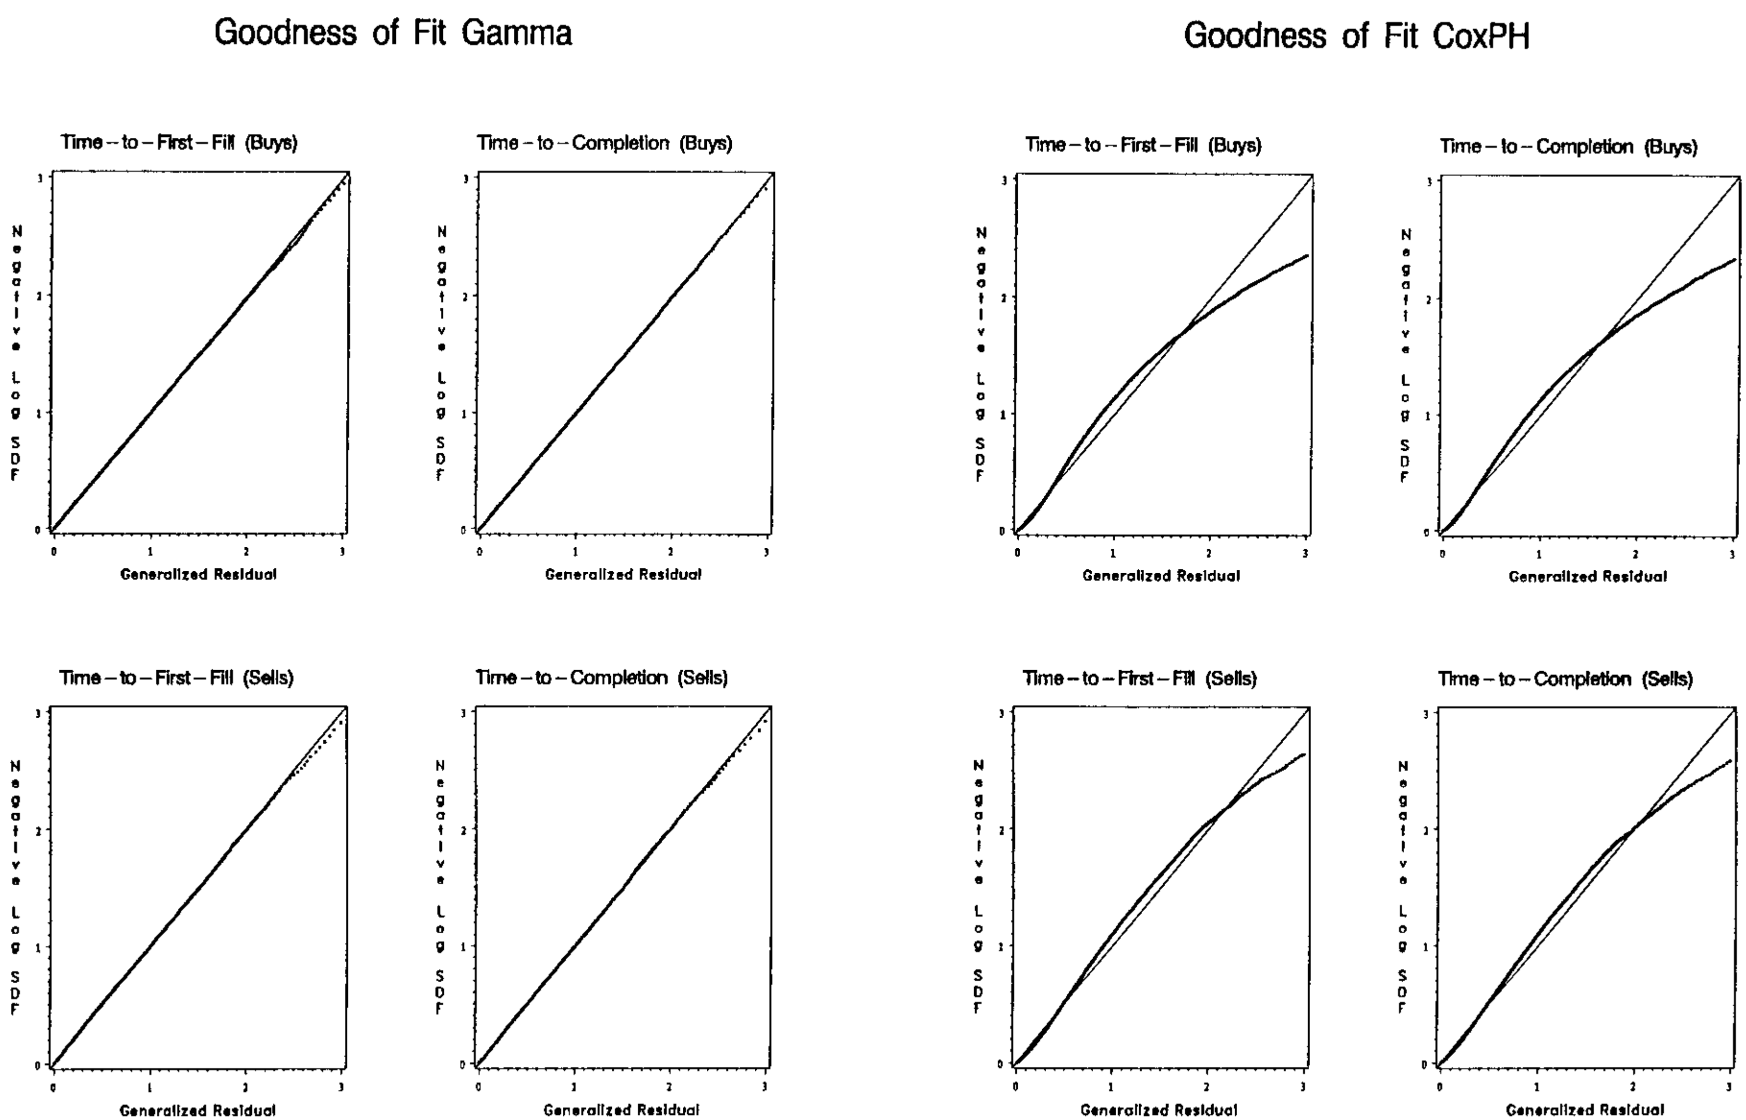
Goodness of Fit Gamma
Goodness of Fit CoxPH
Time-to-First-Fill (Buys)
Time-to-Completion (Buys)
Time-to-First-Fill (Sells)
Time-to-Completion (Sells)
Negative Log SDF
Generalized Residual
0
1
2
3

Figure 3

Diagnostics For Individual Stocks

Time – to – First – Fill (Buys)

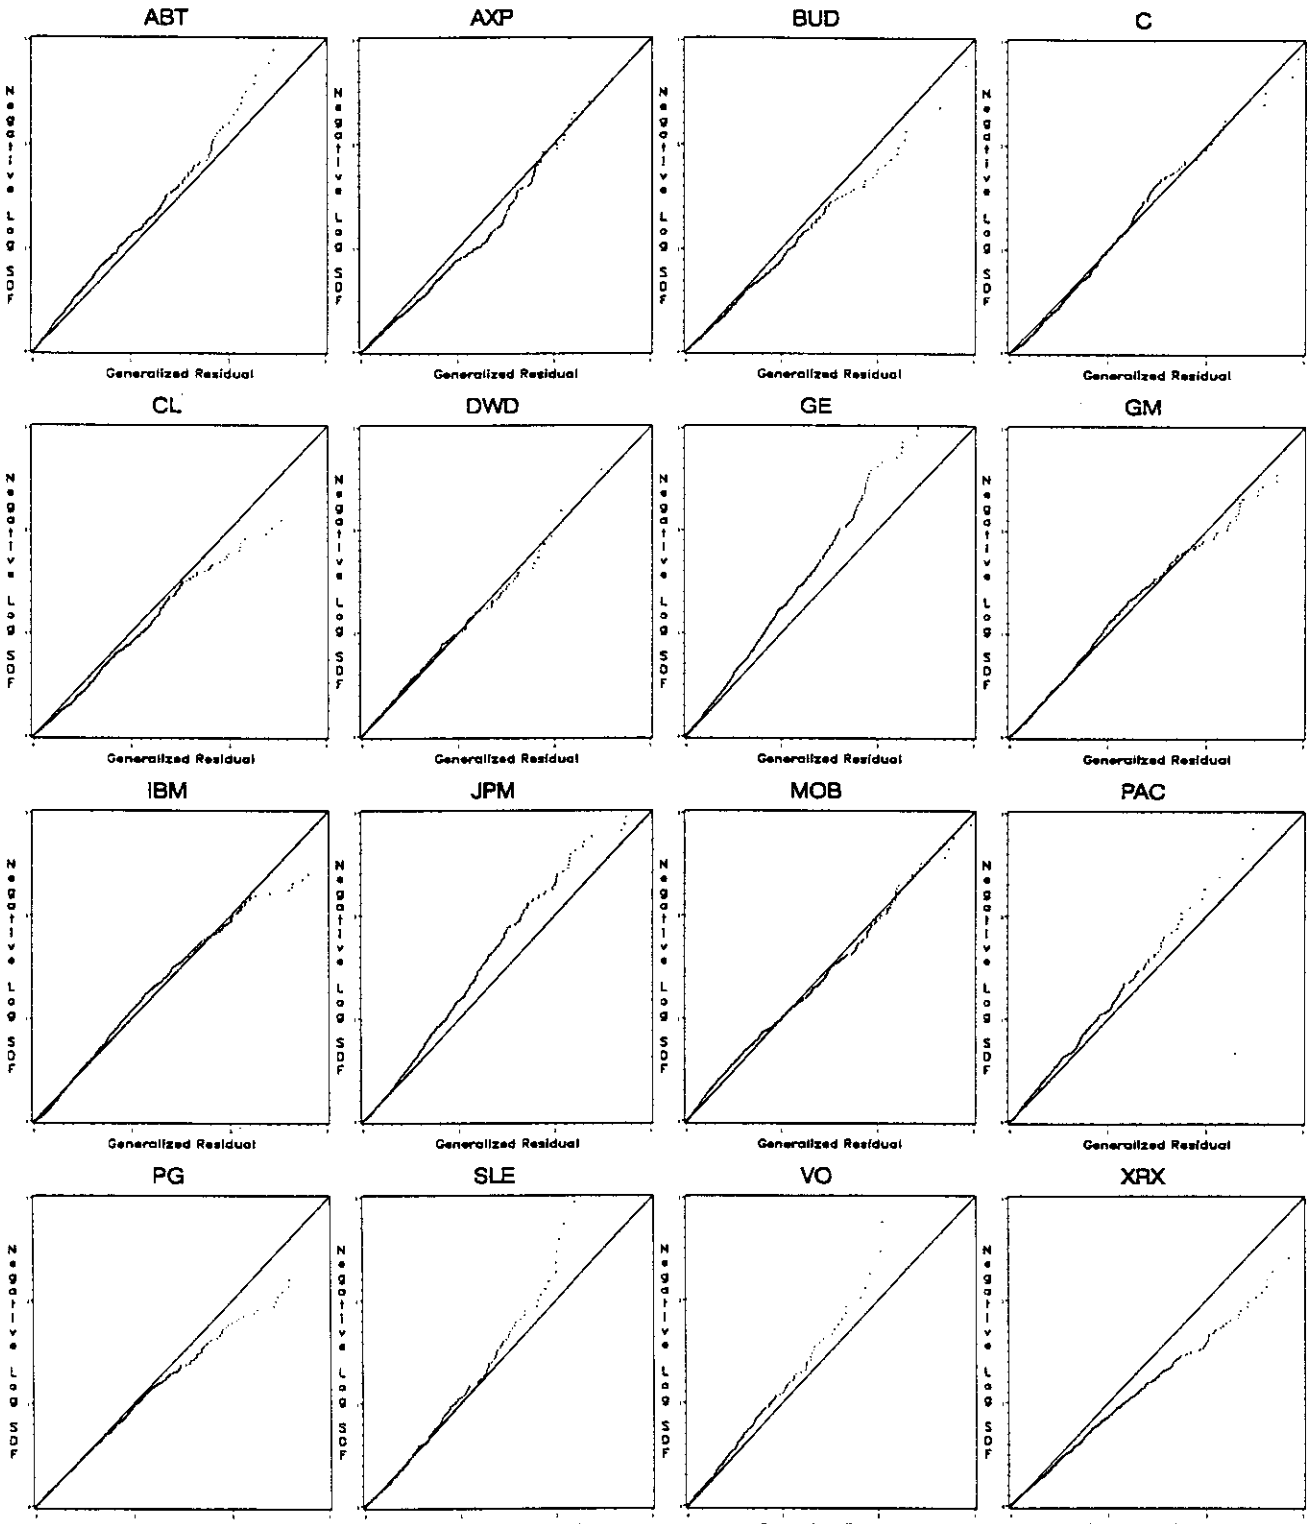

Figure 4a

Diagnostics For Individual Stocks

Time-to-First-Fill (Sells)

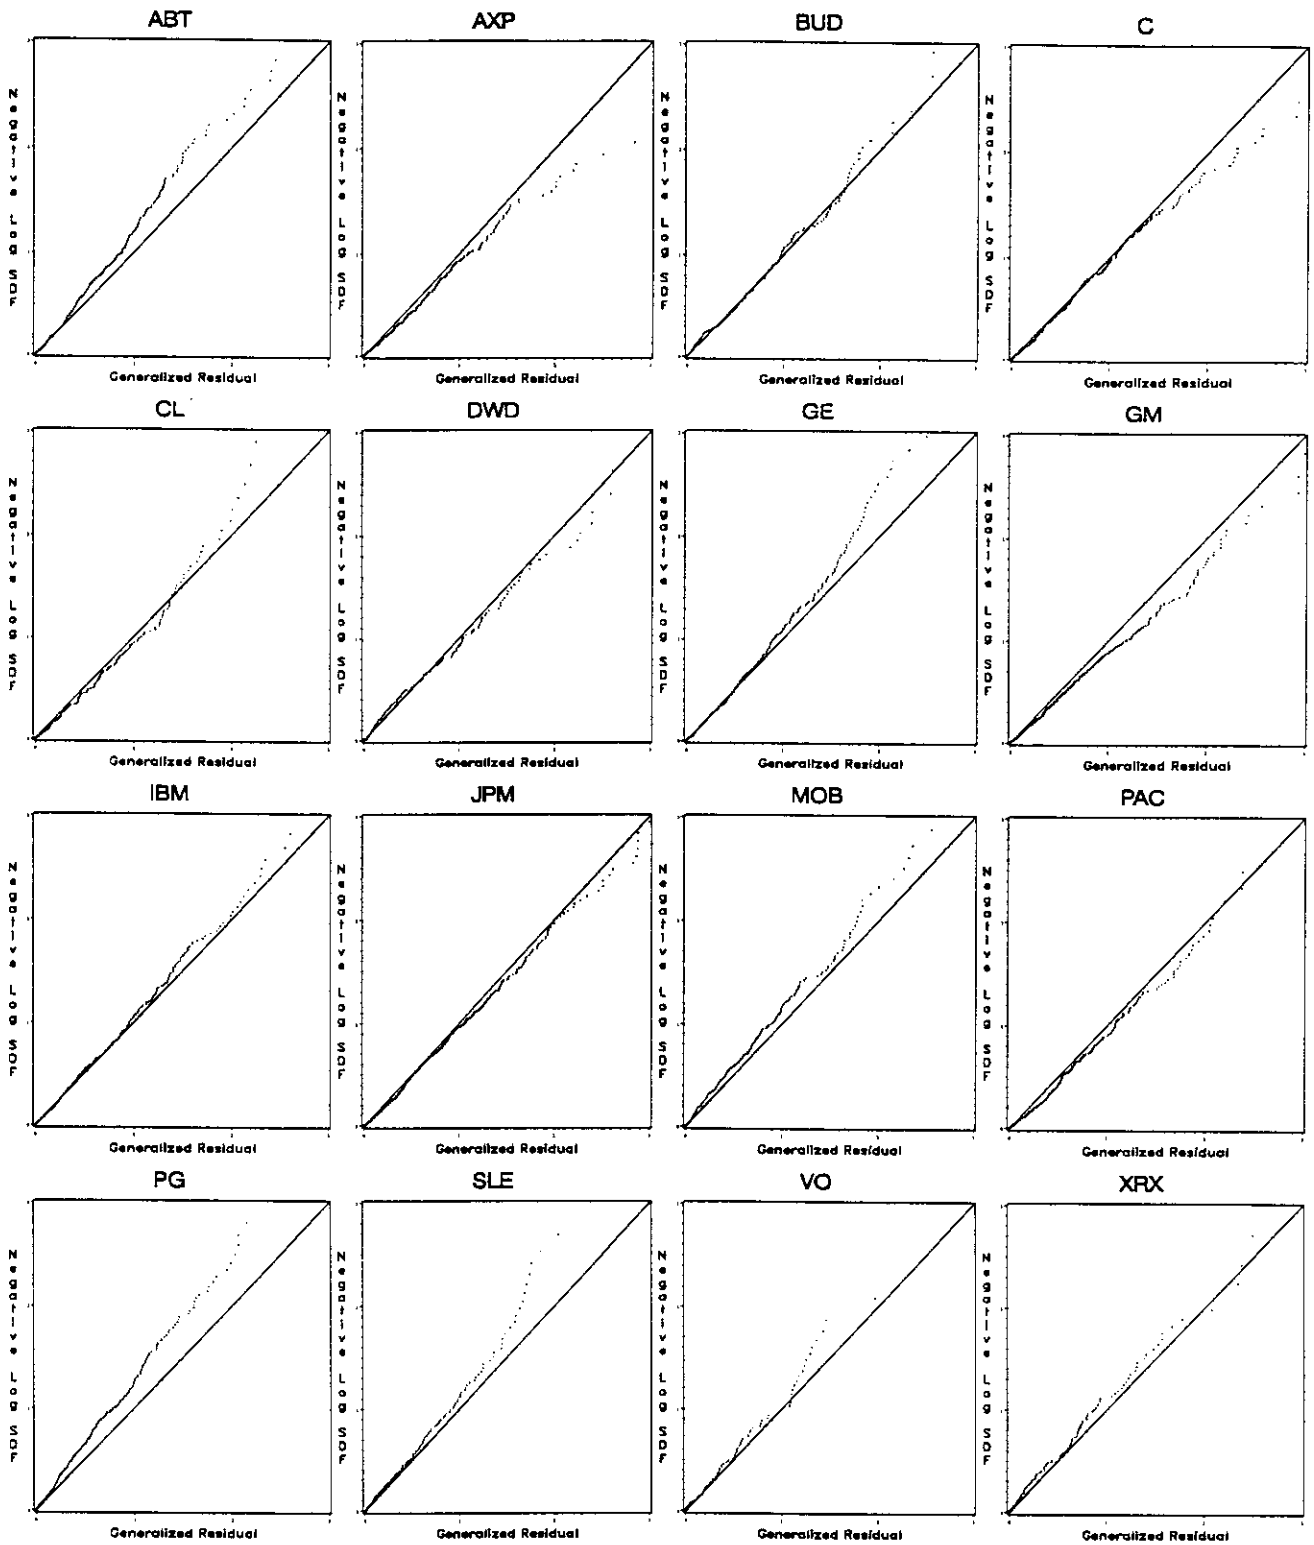

Figure 4b

Estimated Survival Functions, Gamma vs FPT*

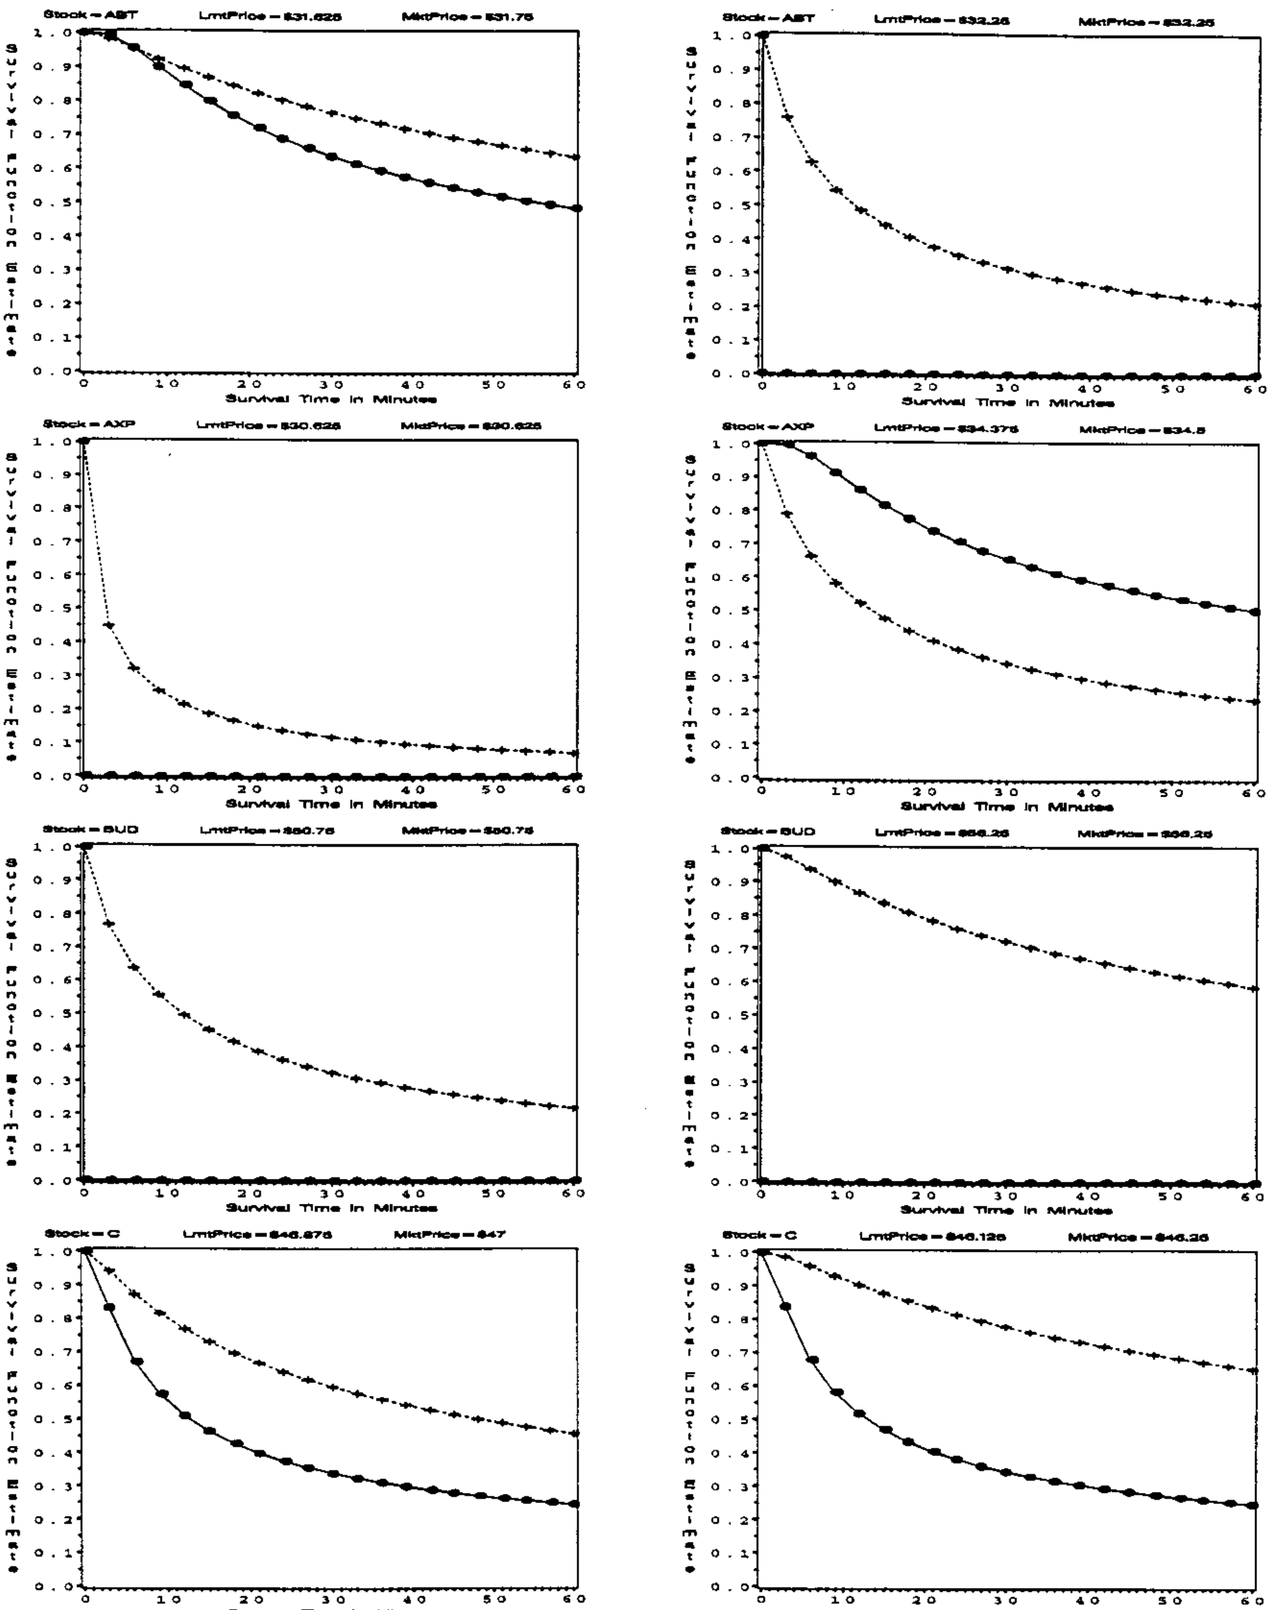

+ Gamma   Figure 5   ● FPT*: FirstPassageTime

Sensitivity to Market Condition, BCP

ABT

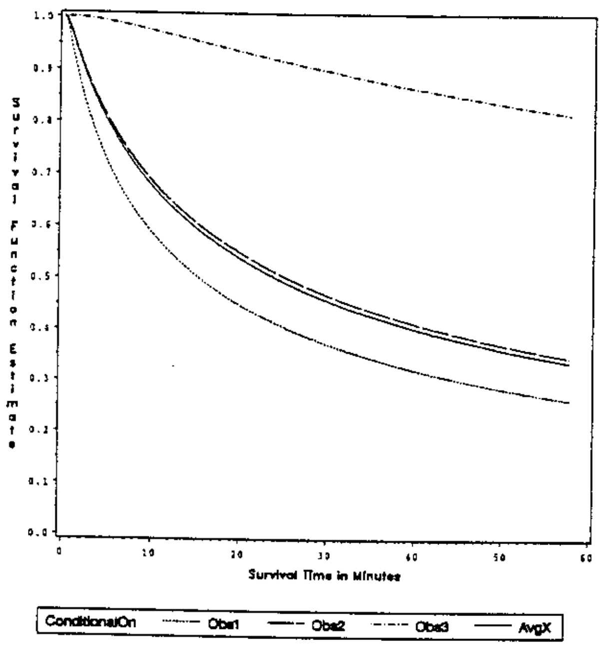

Obs1: LmtP=27.750 LmtS=19 OfrP=27.875 OfrS=19 BidP=27.625 BidS=42 MktP=27.750
Obs2: LmtP=31.625 LmtS=9 OfrP=31.875 OfrS=5 BidP=31.500 BidS=5 MktP=31.625
Obs3: LmtP=36.875 LmtS=1 OfrP=37.000 OfrS=115 BidP=36.875 BidS=350 MktP=36.875
AvgX: LmtP=35.875 LmtS=5 OfrP=36.000 OfrS=12 BidP=35.750 BidS=10 MktP=35.875

AXP

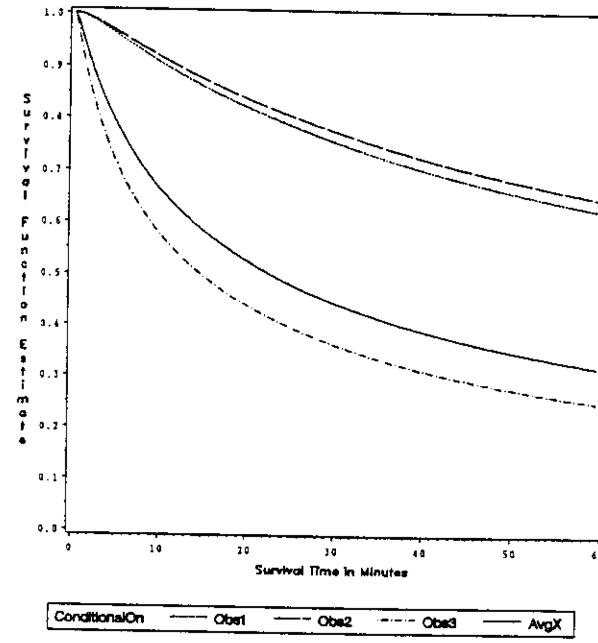

Obs1: LmtP=27.250 LmtS=1 OfrP=27.750 OfrS=1 BidP=27.250 BidS=1 MktP=27.375
Obs2: LmtP=30.125 LmtS=10 OfrP=30.375 OfrS=5 BidP=30.125 BidS=5 MktP=30.250
Obs3: LmtP=34.500 LmtS=4 OfrP=34.625 OfrS=5 BidP=34.375 BidS=5 MktP=34.500
AvgX: LmtP=34.875 LmtS=4 OfrP=35.000 OfrS=9 BidP=34.750 BidS=10 MktP=34.875

BUD

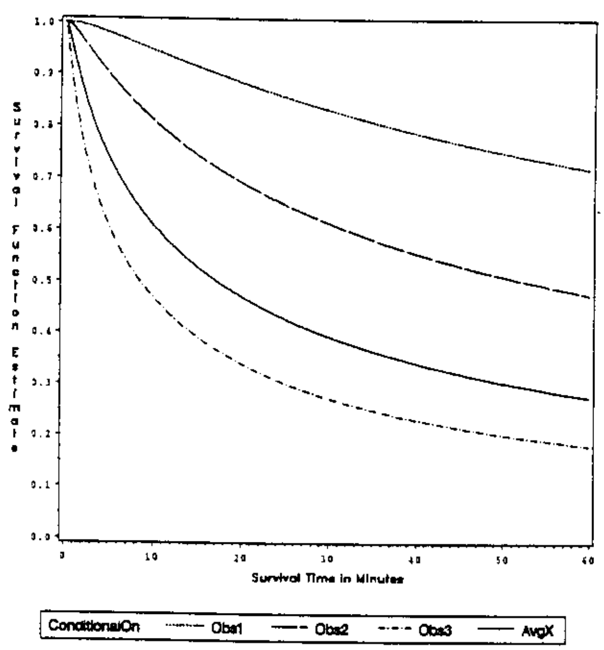

Obs1: LmtP=52.250 LmtS=1 OfrP=52.750 OfrS=3 BidP=52.250 BidS=3 MktP=52.500
Obs2: LmtP=56.625 LmtS=8 OfrP=56.750 OfrS=90 BidP=56.625 BidS=30 MktP=56.625
Obs3: LmtP=56.000 LmtS=2 OfrP=56.125 OfrS=5 BidP=55.750 BidS=5 MktP=56.000
AvgX: LmtP=55.625 LmtS=2 OfrP=55.750 OfrS=7 BidP=55.500 BidS=6 MktP=55.625

C

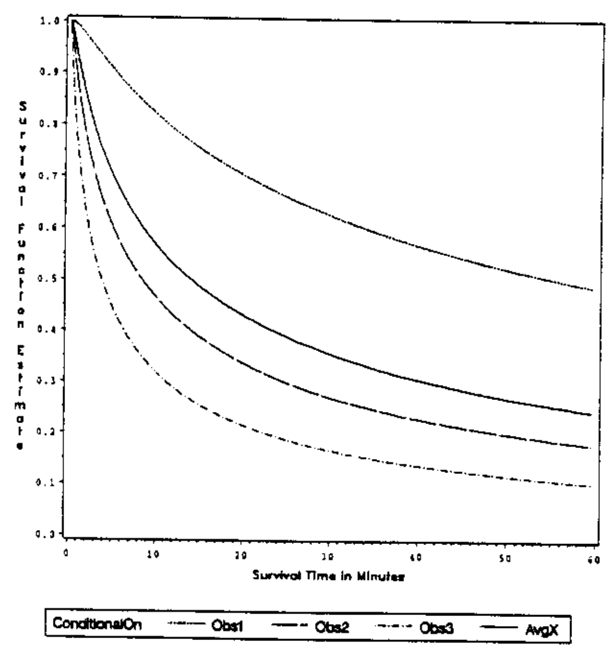

Obs1: LmtP=45.625 LmtS=3 OfrP=45.875 OfrS=5 BidP=45.500 BidS=3 MktP=45.875
Obs2: LmtP=48.000 LmtS=4 OfrP=48.125 OfrS=15 BidP=47.875 BidS=17 MktP=48.000
Obs3: LmtP=45.000 LmtS=4 OfrP=45.000 OfrS=50 BidP=44.875 BidS=250 MktP=45.000
AvgX: LmtP=47.500 LmtS=7 OfrP=47.750 OfrS=17 BidP=47.500 BidS=15 MktP=47.750

Figure 6

Sensitivity to Limit Price, BCP

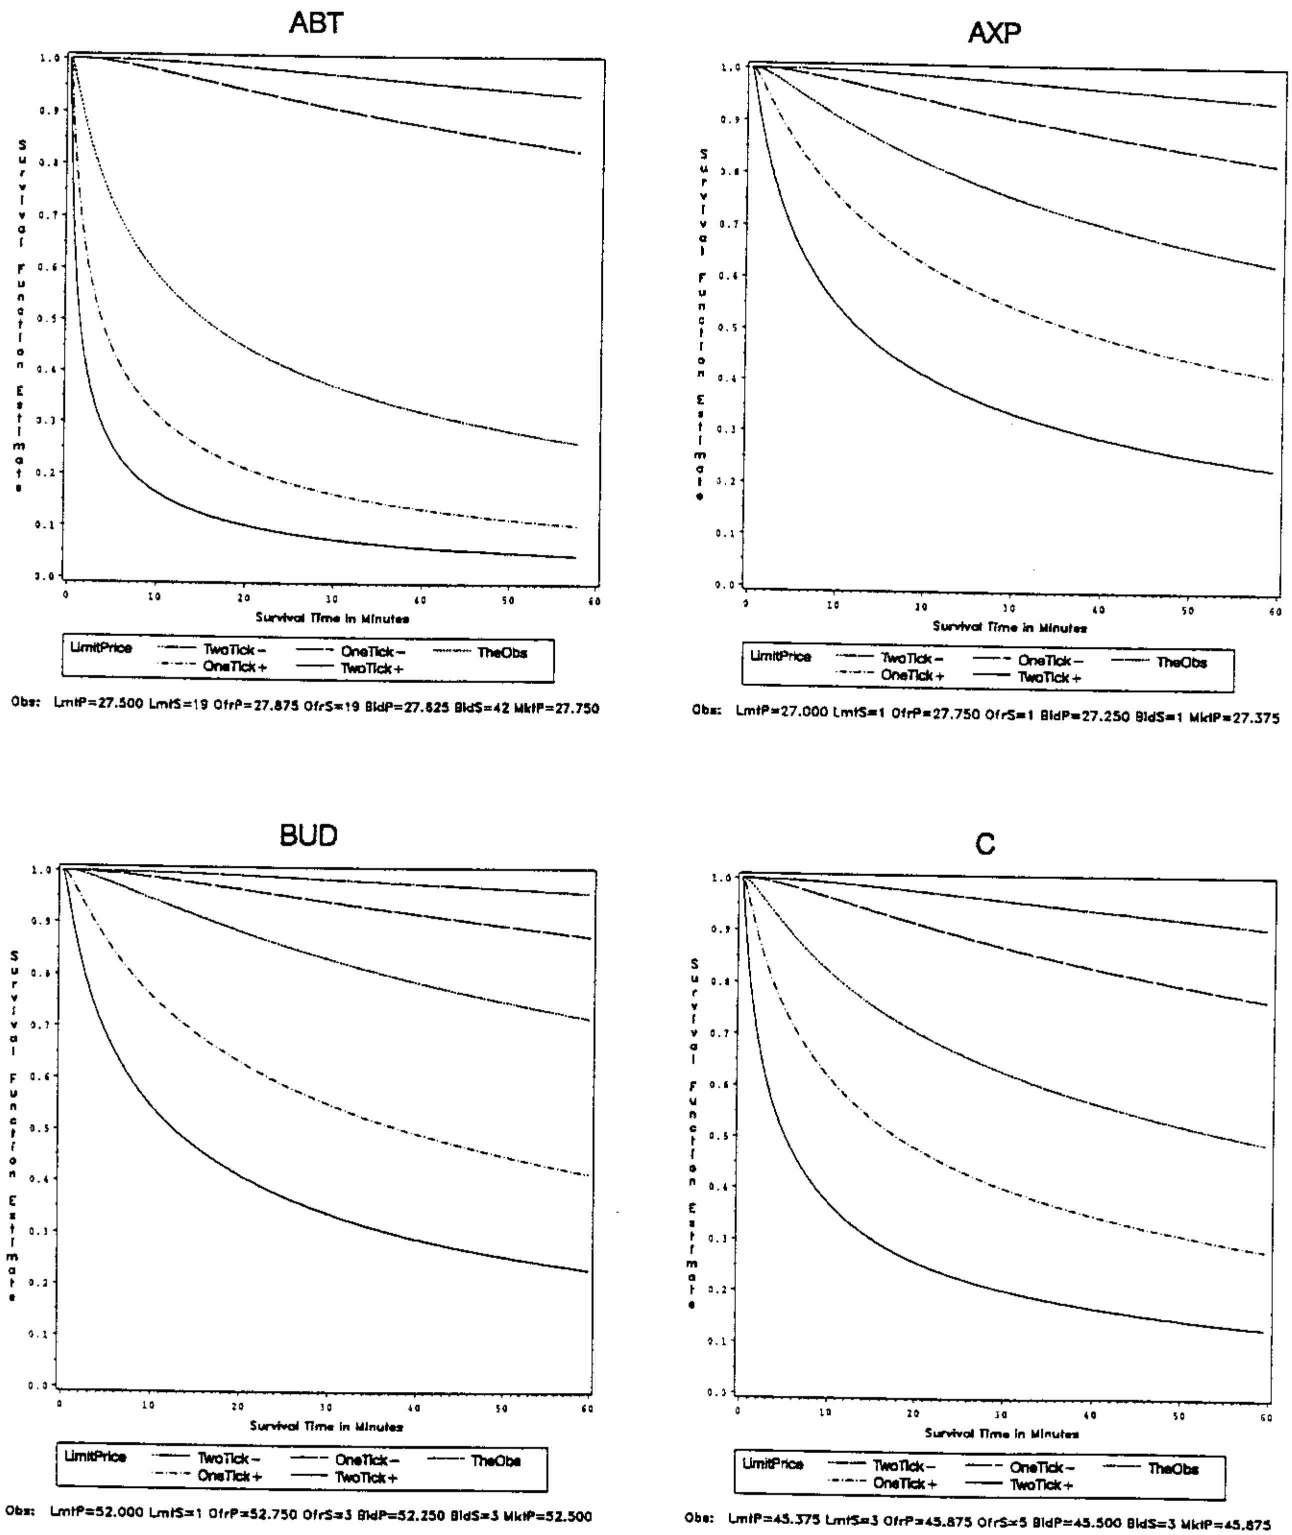

Figure 7

Sensitivity to Limit Shares, BCP

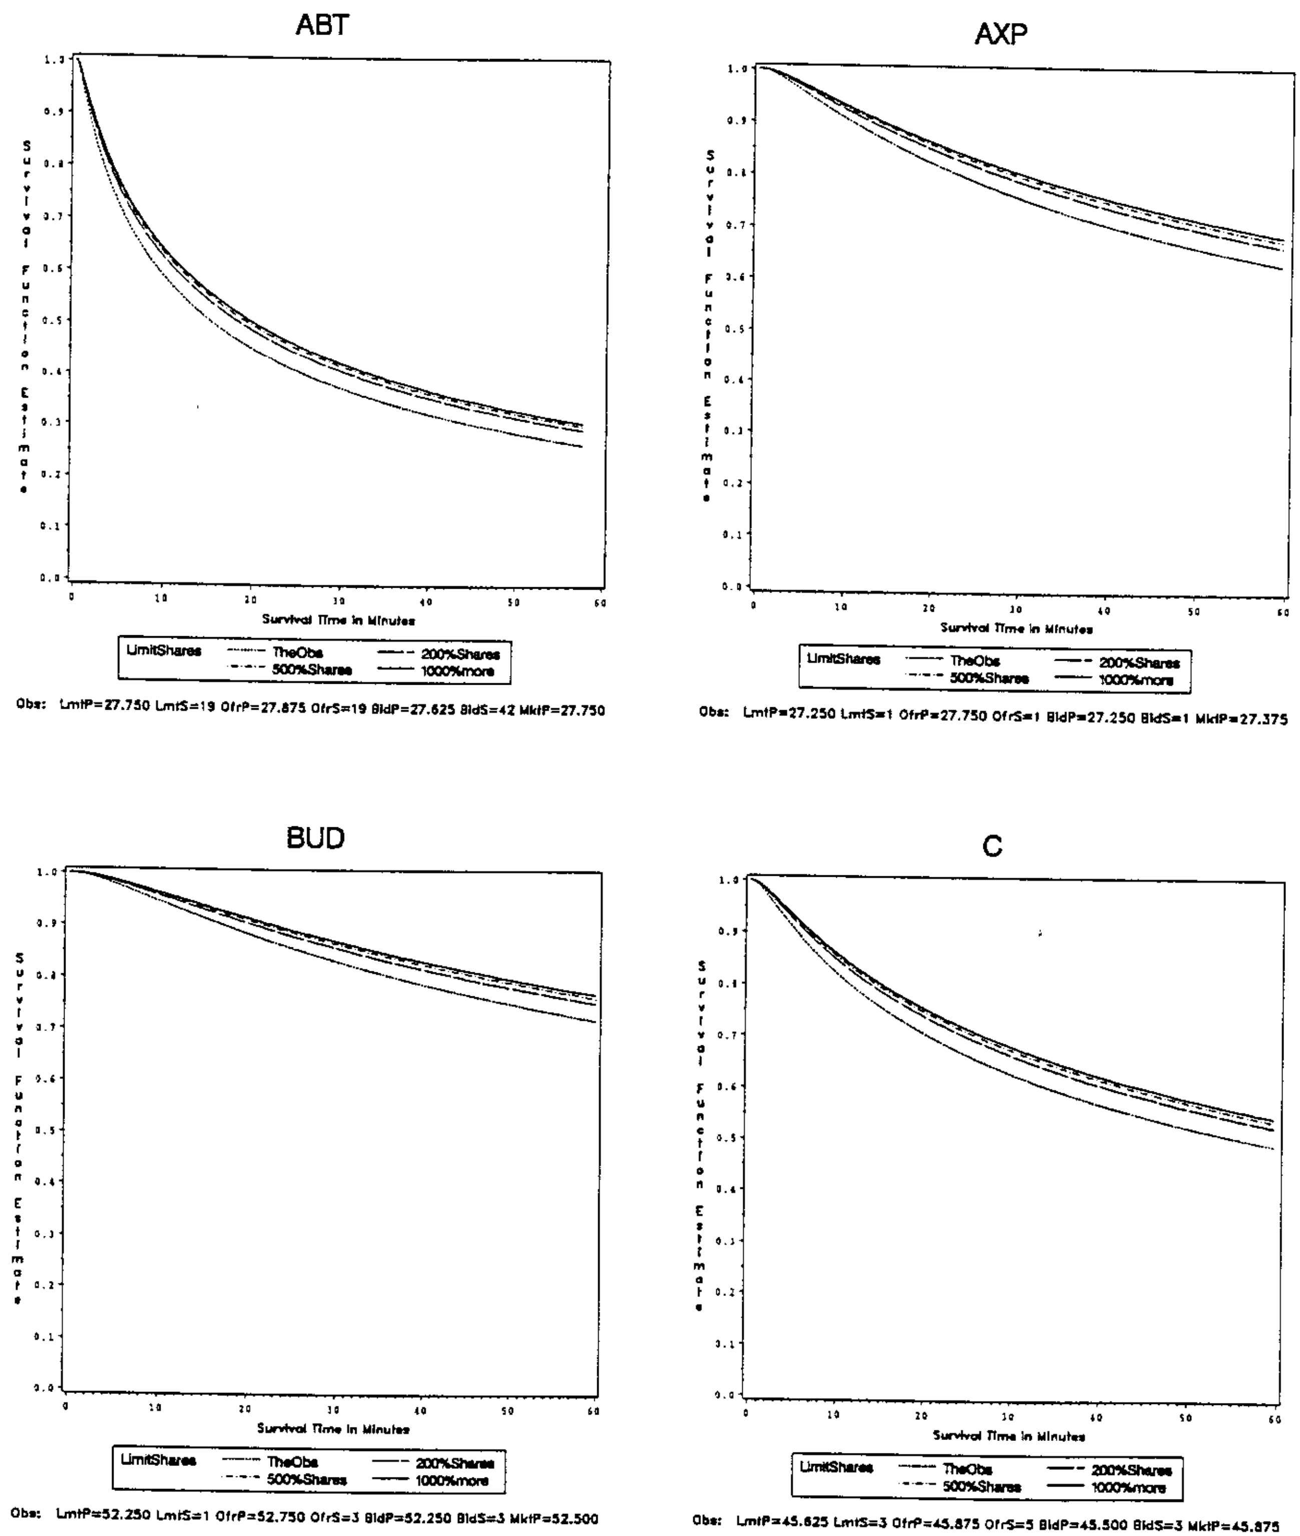

Figure 8