A Random Order Placement Model of Price Formation in the Continuous Double Auction

J. Doyne Farmer
László Gillemot
Giulia Iori
Supriya Krishnamurthy
D. Eric Smith
Marcus G. Daniels

Most modern financial markets use a continuous double auction mechanism to store and match orders and facilitate trading. In this chapter we use a microscopic dynamical statistical model for the continuous double auction under the assumption of IID random order flow. The analysis is based on simulation, dimensional analysis, and theoretical tools based on mean-field approximations. The model makes testable predictions for all the basic properties of markets, including price volatility, the depth of stored supply and demand, the bid-ask spread, the price impact function, and the time and probability of filling orders. These predictions are based on properties of order flow such as share volume of market and limit orders, cancellations, typical order size, and tick size. Because these quantities can all be measured directly in real data sets there are no free parameters. We show that the order size, which can be cast as a nondimensional granularity parameter, is in most cases a more significant determinant of market behavior than tick size. We also provide an explanation for the observed highly concave nature of the price impact function. On a broader level, this work demonstrates how stochastic models based on zero-intelligence agents may be useful in probing the structure of market institutions. Like the model of perfect rationality, a stochastic zero-intelligence model can be used to make strong predictions based on a compact set of assumptions. Preliminary evidence suggests that this model explains many aspects of real markets.

1 INTRODUCTION

1.1 MOTIVATION

In this chapter we analyze the continuous double auction trading mechanism under the assumption of random order flow, giving an overview of the work described in Daniels et al. [12] and Smith et al. [26]. This analysis produces quantitative predictions about the most basic properties of markets, such as volatility, depth of stored supply and demand, the bid-ask spread, the price impact, and probability and time to fill. These predictions are based on the rate at which orders flow into the market, and other parameters of the market, such as order size and tick size. The predictions are falsifiable with no free parameters. This extends the original random walk model of Bachelier [1] by providing a basis for the diffusion rate of prices. The model also provides a possible explanation for the highly concave nature of the price impact function. Even though some of the assumptions of the model are too simple to be literally true, preliminary results suggest that it explains several aspects of price formation and transaction costs in the London Stock Exchange [11] and the New York Stock Exchange [20]. Furthermore, the model provides a framework onto which more realistic assumptions may easily be added.

The model demonstrates the importance of financial institutions in setting prices, and how solving a necessary economic function such as providing liquidity can have unanticipated side effects. In a world of imperfect rationality and imperfect information, the task of demand storage necessarily causes persistence. Under perfect rationality, all traders would instantly update their orders with the arrival of each piece of new information, but this is clearly not true for real markets. The limit order book, which is the queue used for storing unexecuted orders, has long memory when there are persistent orders. It can be regarded as a device for storing supply and demand, somewhat like a capacitor is a device for storing charge. We show that even under completely random IID order flow and cancellations, the price process displays anomalous diffusion and interesting temporal structure. The converse is also interesting: For prices to be effectively random, incoming order flow must be non-random, in just the right way to compensate for the persistence. (See the remarks in section 4.3.)

This work is also of interest from a fundamental point of view because it suggests an alternative approach to doing economics. The assumption of perfect rationality has been popular in economics because it provides a parsimonious model that makes strong predictions. In the spirit of Gode and Sunder [17], we show that the opposite extreme of zero-intelligence random behavior provides another reference model that also makes very strong predictions. Like perfect rationality, zero-intelligence is an extreme simplification that is obviously not literally true. But as we show here, it provides a useful tool for probing the behavior of financial institutions. The resulting model may easily be extended by introducing simple boundedly rational behaviors. We also differ from standard treatments in that we do not attempt to understand the properties of prices from fundamental assumptions about utility. Rather, we split the problem in two. We attempt to understand how prices depend on order flow rates, leaving the problem of what determines these order flow rates for the future.

One of our main results concerns the average price impact function. The liquidity for executing a market order can be characterized by a price impact function $\Delta p = \phi(\omega, \tau, t)$. Δp is the shift in the logarithm of the price at time $t + \tau$ caused by a market order of size ω placed at time t. Understanding price impact is important for practical reasons such as minimizing transaction costs, and also because it is closely related to an excess demand function,[1] providing a natural starting point for theories of statistical or dynamical properties of markets [4, 15]. A naive argument predicts that the price impact $\phi(\omega)$ should increase at least linearly. This argument goes as follows: Fractional price changes should not depend on the scale of price. Suppose buying a single share raises the price by a factor $k > 1$. If k is constant, buying ω shares in succession should raise it by k^{ω}. Thus, if buying ω shares all at once affects the price at least as much as buying them one at a time, the ratio of prices before and after impact should increase at least exponentially. Taking logarithms implies that the price impact as we have defined it above should increase at least linearly.[2]

In contrast, from empirical studies $\phi(\omega)$ for buy orders appears to be concave [16, 18, 19, 20, 24, 28]. Lillo et al. have shown for that for stocks in the NYSE the concave behavior of the price impact is quite consistent across different stocks [20]. Our model produces concave price impact functions that are in qualitative agreement with these results. Furthermore, members of our group have recently begun to analyze data from the London Stock Exchange, which provides every action taken by each trader, and allows us to measure order flows in the way they are defined here. Preliminary results suggest that the model has a remarkable ability to predict the bid-ask spread, and that when plotted in the nondimensional coordinates defined here, the price impact collapses onto a universal function that is consistent through time and across stocks [11].

[1] In financial models it is common to define an excess demand function as demand minus supply; when the context is clear the modifier "excess" is dropped, so that demand refers to both supply and demand.

[2] This has practical implications. It is common practice to break up orders in order to reduce losses due to market impact. With a sufficiently concave market impact function, in contrast, it is cheaper to execute an order all at once.

Our work also demonstrates the value of physics techniques for economic problems. Our analysis makes extensive use of dimensional analysis, the solution of a master equation through a generating functional, and a mean-field approach that is commonly used to analyze non-equilibrium reaction-diffusion systems and evaporation-deposition problems.

1.2 BACKGROUND: THE CONTINUOUS DOUBLE AUCTION

Most modern financial markets operate continuously. The mismatch between buyers and sellers that typically exists at any given instant is solved via an order-based market with two basic kinds of orders. Impatient traders submit *market orders*, which are requests to buy or sell a given number of shares immediately at the best available price. More patient traders submit *limit orders*, or *quotes* which also state a limit price, corresponding to the worst allowable price for the transaction. (Note that the word "quote" can be used either to refer to the limit price or to the limit order itself.) Limit orders often fail to result in an immediate transaction, and are stored in a queue called the *limit order book*. Buy limit orders are called *bids*, and sell limit orders are called *offers* or *asks*. We use the logarithmic price $a(t)$ to denote the position of the best (lowest) offer and $b(t)$ for the position of the best (highest) bid. These are also called the *inside quotes*. There is typically a non-zero price gap between them, called the *spread* $s(t) = a(t) - b(t)$. Prices are not continuous, but rather have discrete quanta called *ticks*. Throughout this chapter, all prices will be expressed as logarithms, and, to avoid endless repetition, the word *price* will mean the logarithm of the price. The minimum interval that prices change on is the *tick size* dp (also defined on a logarithmic scale; note this is not true for real markets). Note that dp is *not* necessarily infinitesimal.

As market orders arrive they are matched against limit orders of the opposite sign in order of first price and then arrival time, as shown in figure 1.

Because orders are placed for varying numbers of shares, matching is not necessarily one-to-one. For example, suppose the best offer is for 200 shares at \$60 and the next best is for 300 shares at \$60.25; a buy market order for 250 shares buys 200 shares at \$60 and 50 shares at \$60.25, moving the best offer $a(t)$ from \$60 to \$60.25. A high density of limit orders per price results in high *liquidity* for market orders, in other words, it decreases the price movement when a market order is placed. Let $n(p, t)$ be the stored density of limit order volume at price p, which we will call the *depth profile* of the limit order book at any given time t. The total stored limit order volume at price level p is $n(p, t)dp$. For unit order size the shift in the best ask $a(t)$ produced by a buy market order is given by solving the equation

$$\omega = \sum_{p=a(t)}^{p\prime} n(p, t)dp \tag{1}$$

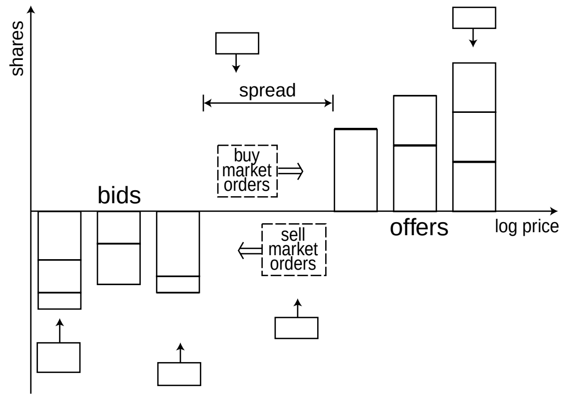

FIGURE 1 A schematic illustration of the continuous double auction mechanism and our model of it. Limit orders are stored in the limit order book. We adopt the arbitrary convention that buy orders are negative and sell orders are positive. As a market order arrives, it has transactions with limit orders of the opposite sign, in order of price (first) and time of arrival (second). The best quotes at prices $a(t)$ or $b(t)$ move whenever an incoming market order has sufficient size to fully deplete the stored volume at $a(t)$ or $b(t)$. Our model assumes that market order arrival, limit order arrival, and limit order cancellation follow a Poisson process. New offers (sell limit orders) can be placed at any price greater than the best bid, and are shown here as "raining down" on the price axis. Similarly, new bids (buy limit orders) can be placed at any price less than the best offer. Bids and offers that fall inside the spread become the new best bids and offers. All prices in this model are logarithmic.

for $p\prime$. The shift in the best ask is $p\prime - a(t)$, where eq. (1) is the instantaneous price impact for buy market orders. A similar statement applies for sell market orders, where the price impact can be defined in terms of the shift in the best bid. (Alternatively, it is also possible to define the price impact in terms of the change in the midpoint price.)

We will refer to a buy limit order whose limit price is greater than the best ask, or a sell limit order whose limit price is less than the best bid, as a *crossing limit order* or *marketable limit order*. Such limit orders result in immediate transactions, with at least part of the order immediately executed.

1.3 THE MODEL

This model, introduced in Daniels et al. [12], is designed to be as analytically tractable as possible while capturing key features of the continuous double auction. All the order flows are modeled as Poisson processes. We assume that market orders arrive in chunks of σ shares, at a rate of μ shares per unit time. The market order may be a "buy" order or a "sell" order with equal probability.

(Thus the rate at which buy orders or sell orders arrive individually is $\mu/2$.) Limit orders arrive in chunks of σ shares as well, at a rate of α shares per unit price and per unit time for buy orders and also for sell orders. Offers are placed with uniform probability at integer multiples of a tick size dp in the range of price $b(t) < p < \infty$, and similarly for bids on $-\infty < p < a(t)$. When a market order arrives it causes a transaction; under the assumption of constant order size, a buy market order removes an offer at price $a(t)$, and if it was the last offer at that price, moves the best ask up to the next occupied price tick. Similarly, a sell market order removes a bid at price $b(t)$, and if it is the last bid at that price, moves the best bid down to the next occupied price tick. In addition, limit orders may also be removed spontaneously by being canceled or by expiring, even without a transaction having taken place. We model this by letting them be removed randomly with constant probability δ per unit time.

While the assumption of limit order placement over an infinite interval is clearly unrealistic, it provides a tractable boundary condition for modeling the behavior of the limit order book near the midpoint price $m(t) = (a(t) + b(t))/2$, which is the region of interest since it is where transactions occur. Limit orders far from the midpoint are usually canceled before they are executed (we demonstrate this later in figure 5), and so far from the midpoint, limit order arrival and cancellation have a steady state behavior characterized by a simple Poisson distribution. Although under the limit order placement process the total number of orders placed per unit time is infinite, the order placement per unit price interval is bounded and thus the assumption of an infinite interval creates no problems. Indeed, it guarantees that there is always an infinite number of limit orders of both signs stored in the book, so that the bid and ask are always well-defined and the book never empties. (Under other assumptions about limit order placement this is not necessarily true, as demonstrated in Smith et al.) We are also considering versions of the model involving more realistic order placement functions (see the discussion in section 4.2).

In this model, to keep things simple, we are using the conceptual simplification of *effective market orders* and *effective limit orders*. When a crossing limit order is placed, part of it may be executed immediately. The effect of this part on the price is indistinguishable from that of a market order of the same size. Similarly, given that this market order has been placed, the remaining part is equivalent to a non-crossing limit order of the same size. Thus a crossing limit order can be modeled as an effective market order followed by an effective (non-crossing) limit order.[3] Working in terms of effective market and limit orders affects data analysis: The effective market order arrival rate μ combines both pure market orders and the immediately executed components of crossing limit orders, and similarly the limit order arrival rate α corresponds only to the components of limit orders that are not executed immediately. This is consistent

[3] In assigning independently random distributions for the two events, our model neglects the correlation between market and limit order arrival induced by crossing limit orders.

with the boundary conditions for the order placement process, since an offer with $p \leq b(t)$ or a bid with $p \geq a(t)$ would result in an immediate transaction, and thus would be effectively the same as a market order. Defining the order placement process with these boundary conditions realistically allows limit orders to be placed anywhere inside the spread.

Another simplification of this model is the use of logarithmic prices, both for the order placement process and for the tick size dp. This has the important advantage that it ensures that prices are always positive. In real markets price ticks are linear, and the use of logarithmic price ticks is an approximation that makes both the calculations and the simulation more convenient. We find that the limit $dp \to 0$, where tick size is irrelevant, is a good approximation for many purposes. We find that tick size is less important than other parameters of the problem, which provides some justification for the approximation of logarithmic price ticks.

Assuming a constant probability for cancellation is clearly *ad hoc*, but in simulations we find that other assumptions with well-defined timescales, such as constant duration time, give similar results. For our analytic model we use a constant order size σ. In simulations we also use variable order size, for example, half-normal distributions with standard deviation $\sqrt{\pi/2}\sigma$, which ensures that the mean value remains σ. As long as these distributions have thin tails, the differences do not qualitatively affect most of the results reported here, except in a trivial way. As discussed in section 4.2, decay processes without well-defined characteristic times and size distributions with power-law tails give qualitatively different results and will be treated elsewhere.

Even though this model is simply defined, the time evolution is not trivial. One can think of the dynamics as being composed of three parts: (1) the buy market order/sell limit order interaction, which determines the best ask; (2) the sell market order/buy limit order interaction, which determines the best bid; and (3) the random cancellation process. Processes (1) and (2) determine each other's boundary conditions. That is, process (1) determines the best ask, which sets the boundary condition for limit order placement in process (2), and process (2) determines the best bid, which determines the boundary conditions for limit order placement in process (1). Thus processes (1) and (2) are strongly coupled. It is this coupling that causes the bid and ask to remain close to each other, and guarantees that the spread $s(t) = a(t) - b(t)$ is a stationary random variable, even though the bid and ask are not. It is the coupling of these processes through their boundary conditions that provides the nonlinear feedback that makes the price process complex.

1.4 SUMMARY OF PRIOR WORK

There are two independent lines of prior work, one in the financial economics literature, and the other in the physics literature. The models in the economics literature are directed toward empirical analysis, and treat the order process as static. In contrast, the prior models in the physics literature are conceptual toy models, but they allow the order process to react to changes in prices, and are thus fully dynamic. Our model bridges this gap. This is explained in more detail below.

The first model of this type that we are aware of was due to Mendelson [23], who modeled random order placement with periodic clearing. This was developed along different directions by Cohen et al. [10], who used techniques from queuing theory, but assumed only one price level and addressed the issue of time priority with that price level (motivated by the hypothesized existence of a specialist who pins prices to make them stationary). Domowitz and Wang [13] and Bollerslev et al. [3] further developed this to allow more general order placement processes that depend on prices, but without solving the full dynamical problem. This allows them to get a stationary solution for prices. In contrast, in the physics models reviewed below, the prices that emerge make a random walk, and so are much more realistic. In our case, to get a solution for the depth of the order book we have to go into price coordinates that comove with the random walk. Dealing with the feedback between order placement and prices makes the problem much more difficult, but it is key for getting reasonable results.

The models in the physics literature incorporate price dynamics, but have tended to be conceptual toy models designed to understand the anomalous diffusion properties of prices. This line of work begins with a paper by Bak et al. [2] which was developed by Eliezer and Kogan [14] and by Tang [27]. They assume that limit orders are placed at a fixed distance from the midpoint, and that the limit prices of these orders are then randomly shuffled until they result in transactions. It is the random shuffling that causes price diffusion. This assumption, which we feel is unrealistic, was made to take advantage of the analogy to a standard reaction-diffusion model in the physics literature. Maslov [22] introduced an alternative model that was solved analytically in the mean-field limit by Slanina [25]. Each order is randomly chosen to be either a buy or a sell, and either a limit order or a market order. If a limit order, it is randomly placed within a fixed distance of the current price. This again gives rise to anomalous price diffusion. A model adding Poisson order cancellation was proposed by Challet and Stinchcombe [7]. Iori and Chiarella [9] have numerically studied a model including fundamentalists and technical traders.

The model studied in this chapter was introduced by Daniels et al. [12]. This adds to the literature by introducing a model that (like the other physics models above) treats the feedback between order placement and price movement, but unlike them is defined so that the parameters of the model can be measured and its predictions tested against real data. The prior models in the physics literature have tended to focus primarily on the anomalous diffusion of prices. While interesting and important for refining risk calculations, this is a second order effect. In contrast, we focus on the first order effects of primary interest to market participants, such as the bid-ask spread, volatility, depth profile, price impact, and the probability and time to fill an order. We demonstrate how dimensional analysis becomes a useful tool in an economic setting, and develop mean-field theories in a context that is more challenging than that of the toy models of previous work.

Subsequent to Daniels et al. [12], Bouchaud et al. [5] demonstrated that, under the assumption that prices execute a random walk, by introducing an additional free parameter they can derive a simple equation for the depth profile. In this chapter we show how to do this from first principles without introducing a free parameter.

2 OVERVIEW OF PREDICTIONS OF THE MODEL

In this section we give an overview of the phenomenology of the model. Because this model has five parameters, understanding all their effects would generally be a complicated problem in and of itself. This task is greatly simplified by the use of dimensional analysis, which reduces the number of independent parameters from five to two. Thus, before we can even review the results, we need to first explain how dimensional analysis applies in this setting. One of the surprising aspects of this model is that one can derive several powerful results using the simple technique of dimensional analysis alone.

Unless otherwise mentioned, the results presented in this section are based on simulations. A brief overview of the theoretical methods used and their agreement to the simulations is given in section 3.

2.1 DIMENSIONAL ANALYSIS

Because dimensional analysis is not commonly used in economics, we first present a brief review. For more details see Bridgman [6].

Dimensional analysis is a technique that is commonly used in physics and engineering to reduce the number of independent degrees of freedom by taking advantage of the constraints imposed by dimensionality. For sufficiently constrained problems it can be used to guess the answer to a problem without doing a full analysis. The idea is to write down all the factors that a given phenomenon can depend on, and then find the combination that has the correct dimensions. For example, consider the problem of deriving the correct formula for the period of a pendulum: The period T has dimensions of *time*. Obvious candidates that it might depend on are the mass of the bob m (which has units of *mass*), the length l (which has units of *distance*), and the acceleration of gravity g (which has units of $distance/time^2$). There is only one way to combine these to produce something with dimensions of *time*, i.e., $T \sim \sqrt{l/g}$. This determines the correct formula for the period of a pendulum up to a constant. Note that it makes it clear that the period does not depend on the mass, a result that is not obvious *a priori*. We were fortunate because in this problem there were three parameters and three dimensions, with a unique combination of the parameters having the

TABLE 1 The five parameters that characterize this model. α, μ, and δ are order flow rates, and dp and σ are discreteness parameters.

Parameter	**Description**	**Dimensions**
α	limit order rate	*shares/(price time)*
μ	market order rate	*shares/time*
δ	order cancellation rate	*1/time*
dp	tick size	*price*
σ	characteristic order size	*shares*

right dimensions. In general, dimensional analysis can only be used to reduce the number of free parameters through the constraints imposed by their dimensions.

For this problem the three fundamental dimensions in the model are *shares*, *price*, and *time*. Note that by *price*, we mean the logarithm of price; as long as we are consistent, this does not create problems with the dimensional analysis. There are five parameters: three rate constants and two discreteness parameters. The *order flow rates* are μ, the market order arrival rate, with dimensions of *shares per time*; α, the limit order arrival rate per unit price, with dimensions of *shares per price per time*; and δ, the rate of limit order decays, with dimensions of *1/time*. These play a role similar to rate constants in physical problems. The two *discreteness parameters* are the price tick size dp, with dimensions of *price*, and the order size σ, with dimensions of *shares*. This is summarized in table 1.

Dimensional analysis can be used to reduce the number of relevant parameters. Because there are five parameters and three dimensions (*price, shares, time*), and because in this case the dimensionality of the parameters is sufficiently rich, the dimensional relationships reduce the degrees of freedom, so that all the properties of the limit order book can be described by functions of two parameters. It is useful to construct these two parameters so that they are nondimensional.

We perform the dimensional reduction of the model by guessing that the effect of the order flow rates is primary to that of the discreteness parameters. This leads us to construct nondimensional units based on the order flow parameters alone, and take nondimensionalized versions of the discreteness parameters as the independent parameters whose effects remain to be understood. As we will see, this is justified by the fact that many of the properties of the model depend only weakly on the discreteness parameters. We can thus understand much of the richness of the phenomenology of the model through dimensional analysis alone.

There are three order flow rates and three fundamental dimensions. If we temporarily ignore the discreteness parameters, there are unique combinations of the order flow rates with units of shares, price, and time. These define a characteristic number of shares $N_c = \mu/2\delta$, a characteristic price interval $p_c = \mu/2\alpha$, and a characteristic timescale $t_c = 1/\delta$. This is summarized in table 2. The

TABLE 2 Important characteristic scales and nondimensional quantities. We summarize the characteristic share size, price, and times defined by the order flow rates, as well as the two nondimensional scale parameters dp/p_c and ϵ that characterize the effect of finite tick size and order size. Dimensional analysis makes it clear that all the properties of the limit order book can be characterized in terms of functions of these two parameters.

Parameter	**Description**	**Expression**
N_c	characteristic number of shares	$\mu/2\delta$
p_c	characteristic price interval	$\mu/2\alpha$
t_c	characteristic time	$1/\delta$
dp/p_c	nondimensional tick size	$2\alpha dp/\mu$
ϵ	nondimensional order size	$2\delta\sigma/\mu$

factors of two occur because we have defined the market order rate for either a buy or a sell order to be $\mu/2$. We can thus express everything in the model in nondimensional terms by dividing by N_c, p_c, or t_c as appropriate, for example, to measure shares in nondimensional units $\hat{N} = N/N_c$, or to measure price in nondimensional units $\hat{p} = p/p_c$.

The value of using nondimensional units is illustrated in figure 2.

Figure 2(a) shows the average depth profile for three different values of μ and δ with the other parameters held fixed. When we plot these results in dimensional units the results look quite different. However, when we plot them in terms of nondimensional units, as shown in figure 2(b), the results are indistinguishable. As explained below, because we have kept the nondimensional order size fixed, the collapse is perfect. Thus, the problem of understanding the behavior of this model is reduced to studying the effects of tick size and order size.

To understand the effects of tick size and order size it is useful to do so in nondimensional terms. The nondimensional scale parameter based on tick size is constructed by dividing by the characteristic price, that is, $dp/p_c = 2\alpha dp/\mu$. The theoretical analysis and the simulations show that there is a sensible continuum limit as the tick size $dp \to 0$, in the sense that there is non-zero price diffusion and a finite spread. Furthermore, the dependence on tick size is weak, and for many purposes the limit $dp \to 0$ approximates the case of finite tick size fairly well. As we will see, working in this limit is essential for getting tractable analytic results.

A nondimensional scale parameter based on order size is constructed by dividing the typical order size (which is measured in shares) by the characteristic number of shares N_c, i.e., $\epsilon \equiv \sigma/N_c = 2\delta\sigma/\mu$. ϵ characterizes the "chunkiness" of the orders stored in the limit order book. As we will see, ϵ is an important determinant of liquidity, and it is a particularly important determinant of volatility. In the continuum limit $\epsilon \to 0$ there is no price diffusion. This is because price

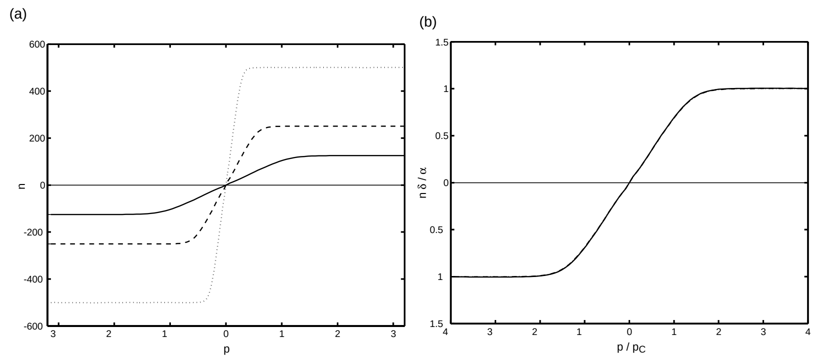

FIGURE 2 The usefulness of nondimensional units. (a) We show the average depth profile for three different parameter sets. The parameters $\alpha = 0.5$, $\sigma = 1$, and $dp = 0$ are held constant, while δ and μ are varied. The line types are: (dotted) $\delta = 0.001$, $\mu = 0.2$; (dashed) $\delta = 0.002$, $\mu = 0.4$ and (solid) $\delta = 0.004$, $\mu = 0.8$. (b) is the same, but plotted in nondimensional units. The horizontal axis has units of *price*, and so has nondimensional units $\hat{p} = p/pc = 2\alpha p/\mu$. The vertical axis has units of n *shares/price*, and so has nondimensional units $\hat{n} = np_c/N_c = n\delta/\alpha$. Because we have chosen the parameters to keep the nondimensional order size ϵ constant, the collapse is perfect. Varying the tick size has little effect on the results other than making them discrete.

diffusion can occur only if there is a finite probability for price levels outside the spread to be empty, thus allowing the best bid or ask to make a persistent shift. If we let $\epsilon \to 0$ while the average depth is held fixed, the number of individual orders becomes infinite, and the probability that spontaneous decays or market orders can create gaps outside the spread becomes zero. This is verified in simulations. Thus the limit $\epsilon \to 0$ is always a poor approximation of a real market. ϵ is a more important parameter than the tick size dp/p_c. In the mean-field analysis in section 3, we let $dp/p_c \to 0$, reducing the number of independent parameters from two to one, and in many cases find that this is a good approximation.

The order size σ can be thought of as the order *granularity*. Just as the properties of a beach with fine sand are quite different from those of one populated by fist-sized boulders, a market with many small orders behaves quite differently from one with a few large orders. N_c provides the scale against which the order size is measured, and ϵ characterizes the granularity in relative terms. Alternatively, $1/\epsilon$ can be thought of as the annihilation rate from market orders expressed in units of the size of spontaneous decays. Note that in nondimensional units the number of shares can also be written $\hat{N} = N/N_c = N\epsilon/\sigma$.

The construction of the nondimensional granularity parameter illustrates the importance of including a spontaneous decay process in this model. If $\delta = 0$ (which implies $\epsilon = 0$) there is no spontaneous decay of orders, and depending

TABLE 3 Estimates from dimensional analysis for the scaling of a few market properties based on order flow rates alone. α is the limit order density rate, μ is the market order rate, and δ is the spontaneous limit order removal rate. These estimates are constructed by taking the combinations of these three rates that have the proper units. They neglect the dependence on the order granularity ϵ and the nondimensional tick size dp/p_c. More accurate relations from simulation and theory are given in table 4.

Quantity	**Dimensions**	**Scaling relation**
Asymptotic depth	*shares/price*	$d \sim \alpha/\delta$
Spread	*price*	$s \sim \mu/\alpha$
Slope of depth profile	*shares/price*^{2}	$\lambda \sim \alpha^2/\mu\delta = d/s$
Price diffusion rate	*price*^{2}*/time*	$D_0 \sim \mu^2\delta/\alpha^2$

on the relative values of μ and α, generically either the depth of orders will accumulate without bound or the spread will become infinite. As long as $\delta > 0$, in contrast, this is not a problem.

For some purposes the effects of varying tick size and order size are fairly small, and we can derive approximate formulas using dimensional analysis based only on the order flow rates. For example, in table 3 we give dimensional scaling formulas for the average spread, the market order liquidity (as measured by the average slope of the depth profile near the midpoint), the volatility, and the asymptotic depth (defined below). Because these estimates neglect the effects of discreteness, they are only approximations of the true behavior of the model, which do a better job of explaining some properties than others. Our numerical and analytical results show that some quantities also depend on the granularity parameter ϵ and to a weaker extent on the tick size dp/p_c. Nonetheless, the dimensional estimates based on order flow alone provide a good starting point for understanding market behavior.

A comparison to more precise formulas derived from theory and simulations is given in table 4.

An approximate formula for the mean spread can be derived by noting that it has dimensions of *price*, and the unique combination of order flow rates with these dimensions is μ/α. While the dimensions indicate the scaling of the spread, they cannot determine multiplicative factors of order unity. A more intuitive argument can be made by noting that inside the spread, removal due to cancellation is dominated by removal due to market orders. Thus, the total limit order placement rate inside the spread, for either buy or sell limit orders αs, must equal the order removal rate $\mu/2$, which implies that spread is $s = \mu/2\alpha$. As we will see later, this argument can be generalized and made more precise within our mean-field analysis, which then also predicts the observed dependence on the granularity parameter ϵ. However, this dependence is rather weak and only causes a variation of roughly a factor of two for $\epsilon < 1$ (see fig. 10), and the factor

TABLE 4 The dependence of market properties on model parameters based on simulation and theory, with the relevant figure numbers. These formulas include corrections for order granularity ϵ and finite tick size dp/p_c. The formula for asymptotic depth from dimensional analysis in table 3 is exact with zero tick size. The expression for the mean spread is modified by a function of ϵ and dp/p_c, though the dependence on them is fairly weak. For the liquidity λ, corresponding to the slope of the depth profile near the origin, the dimensional estimate must be modified because the depth profile is no longer linear (mainly depending on ϵ) and so the slope depends on price. The formulas for the volatility are empirical estimates from simulations. The dimensional estimate for the volatility from Table 3 is modified by a factor of $\epsilon^{-0.5}$ for the early time price diffusion rate and a factor of $\epsilon^{0.5}$ for the late time price diffusion rate.

Quantity	**Scaling relation**	**Figure**
Asymptotic depth	$d = \alpha/\delta$	3
Spread	$s = (\mu/\alpha)f(\epsilon, dp/p_c)$	10
Slope of depth profile	$\lambda = (\alpha^2/\mu\delta)g(\epsilon, dp/p_c)$	3
Price diffusion ($\tau \to 0$)	$D_0 = (\mu^2\delta/\alpha^2)\epsilon^{-0.5}$	11, 14(c)
Price diffusion ($\tau \to \infty$)	$D_\infty = (\mu^2\delta/\alpha^2)\epsilon^{0.5}$	11, 14(c)

of 1/2 derived above is a good first approximation. Note that this prediction of the mean spread is just the characteristic price p_c.

It is also easy to derive the mean *asymptotic depth*, which is the density of shares far away from the midpoint. The asymptotic depth is an artificial construct of our assumption of order placement over an infinite interval; it should be regarded as providing a simple boundary condition so that we can study the behavior near the midpoint price. The mean-asymptotic depth has dimensions of *shares/price*, and is, therefore, given by α/δ. Furthermore, because removal by market orders is insignificant in this regime, it is determined by the balance between order placement and decay, and far from the midpoint the depth at any given price is Poisson distributed. This result is exact.

The average slope of the depth profile near the midpoint is an important determinant of liquidity, since it affects the expected price response when a market order arrives. The slope has dimensions of *shares/price*^{2}, which implies that in terms of the order flow rates it scales roughly as $\alpha^2/\mu\delta$. This is also the ratio of the asymptotic depth to the spread. As we will see later, this is a good approximation when $\epsilon \sim 0.01$, but for smaller values of ϵ the depth profile is not linear near the midpoint, and this approximation fails.

The last two entries in table 4 are empirical estimates for the price diffusion rate D, which is proportional to the square of the volatility. That is, for normal diffusion, starting from a point at $t = 0$, the variance v after time t is $v = Dt$. The volatility at any given timescale t is the square root of the variance at timescale t. The estimate for the diffusion rate based on dimensional analysis in terms of the order flow rates alone is $\mu^2\delta/\alpha^2$. However, simulations show that short time diffusion is much faster than long time diffusion, due to negative autocorrelations in the price process, as shown in figure 11. The initial and the asymptotic diffusion rates appear to obey the scaling relationships given in table 4. Though our mean-field theory is not able to predict this functional form, the fact that early and late time diffusion rates are different can be understood within the framework of our analysis, as discussed in section 3. Anomalous diffusion of this type implies negative autocorrelations in midpoint prices. Note that we use the term “anomalous diffusion” to imply that the diffusion rate is different on short and long timescales. We do not use this term in the sense that it is normally used in the physics literature, in other words, that the long-time diffusion is proportional to t^γ with $\gamma \neq 1$ (for long times $\gamma = 1$ in our case).

2.2 VARYING THE GRANULARITY PARAMETER ϵ

We first investigate the effect of varying the order granularity ϵ in the limit $dp \to 0$. As we will see, the granularity has an important effect on most of the properties of the model, and particularly on depth, price impact, and price diffusion. The behavior can be divided into three regimes, roughly as follows:

- **Large ϵ, i.e.,** $\epsilon \gtrsim 0.1$. This corresponds to a large accumulation of orders at the best bid and ask, nearly linear market impact, and roughly equal short and long time price diffusion rates. This is the regime in which the mean-field approximation used in the theoretical analysis works best.
- **Medium ϵ, i.e.,** $\epsilon \sim 0.01$. In this range the accumulation of orders at the best bid and ask is small, and near the midpoint price the depth profile increases nearly linearly with price. As a result, as a crude approximation the price impact increases as roughly the square root of order size.
- **Small ϵ, i.e.,** $\epsilon \lesssim 0.001$. The accumulation of orders at the best bid and ask is very small, and near the midpoint the depth profile is a convex function of price. The price impact is very concave. The short time price diffusion rate is much greater than the long time price diffusion rate.

Since the results for bids are symmetric with those for offers about $p = 0$, for convenience we only show the results for offers, that is, buy market orders and sell limit orders. In this subsection prices are measured relative to the midpoint, and simulations are in the continuum limit where the tick size $dp \to 0$. The results in this section are from numerical simulations. Also, bear in mind that far from the midpoint the predictions of this model are not valid due to the unrealistic assumption of an order placement process with an infinite domain. Thus the results are potentially relevant to real markets only when the price p is at most a few times as large as the characteristic price p_c.

2.2.1 Depth Profile. The *mean-depth profile*, in other words, the average number of shares per price interval, and the mean-cumulative depth profile are shown in

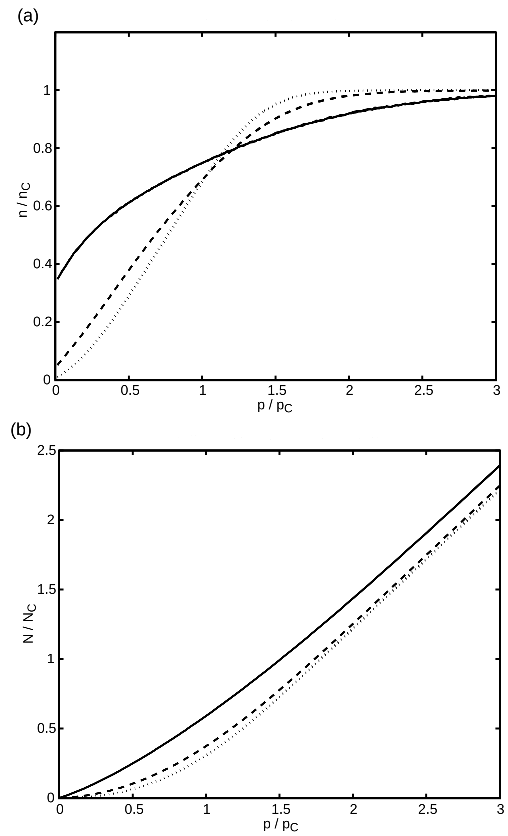

FIGURE 3 The mean-depth profile and cumulative depth versus $\hat{p} = p/p_c = 2\alpha p/\mu$. The origin $p/p_c = 0$ corresponds to the midpoint. (a) is the average depth profile n in nondimensional coordinates $\hat{n} = np_c/N_c = n\delta/\alpha$. (b) is nondimensional cumulative depth $N(p)/N_c$. We show three different values of the nondimensional granularity parameter: $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot), all with tick size $dp = 0$.

figure 3, and the standard deviation of the cumulative profile is shown in figure 4. Since the depth profile has units of *shares/price*, nondimensional units of depth profile are $\hat{n} = np_c/N_c = n\delta/\alpha$.

The cumulative depth profile at any given time t is defined as

$$N(p,t) = \sum_{\tilde{p}=0}^{p} n(\tilde{p},t)dp\,. \tag{2}$$

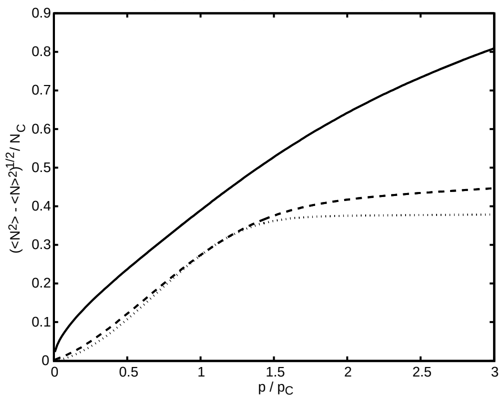

FIGURE 4 Standard deviation of the nondimensionalized cumulative depth versus nondimensional price, corresponding to figure 3.

This has units of shares and so in nondimensional terms is $\hat{N}(p) = N(p)/N_c = 2\delta N(p)/\mu = N(p)\epsilon/\sigma$.

In the high ϵ regime the annihilation rate due to market orders is low (relative to $\delta\sigma$), and there is a significant accumulation of orders at the best ask, so that the average depth is much greater than zero at the midpoint. The mean-depth profile is a concave function of price. In the medium ϵ regime the market order removal rate increases, depleting the average depth near the best ask, and the profile is nearly linear over the range $p/p_c \leq 1$. In the small ϵ regime the market order removal rate increases even further, making the average depth near the ask very close to zero, and the profile is a convex function over the range $p/p_c \leq 1$.

The standard deviation of the depth profile is shown in figure 4. We see that the standard deviation of the cumulative depth is comparable to the mean depth, and that as ϵ increases, near the midpoint there is a similar transition from convex to concave behavior.

The uniform order placement process seems at first glance one of the most unrealistic assumptions of our model, leading to depth profiles with a finite asymptotic depth (which also implies that there is an infinite number of orders in the book). However, orders far away from the spread in the asymptotic region almost never get executed and thus do not affect the market dynamics. To demonstrate this in figure 5 we show the comparison between the limit order depth profile and the depth n_e of only those orders which eventually get executed.[4]

[4]Note that the ratio n_e/n is not the same as the probability of filling orders (fig. 12) because in that case the price p/p_c refers to the distance of the order from the midpoint at the time when it was placed.

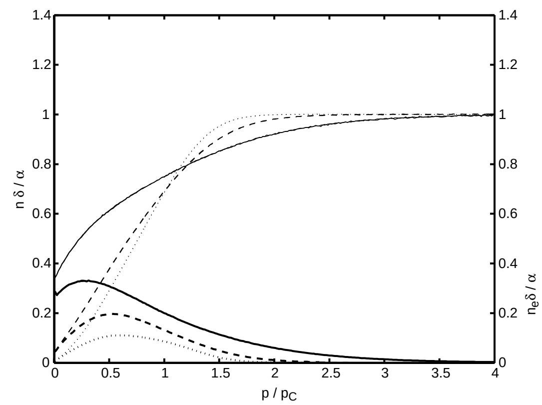

FIGURE 5 A comparison between the depth profiles and the effective depth profiles as defined in the text, for different values of ϵ. Heavy lines refer to the effective depth profiles n_e and the light lines correspond to the depth profiles.

The density n_e of executed orders decreases rapidly as a function of the distance from the mid-price. Therefore, we expect that near the midpoint our results should be similar to alternative order placement processes, as long as they also lead to an exponentially decaying profile of executed orders (which is what we observe above). However, to understand the behavior further away from the midpoint we are also working on enhancements that include more realistic order placement processes grounded on empirical measurements of market data, as summarized in section 4.2.

2.2.2 Liquidity for Market Orders: The Price Impact Function.

In this subsection we study the *instantaneous price impact* function $\phi(t, \omega, \tau \to 0)$. This is defined as the (logarithm of the) midpoint price shift immediately after the arrival of a market order in the absence of any other events. This should be distinguished from the asymptotic price impact $\phi(t, \omega, \tau \to \infty)$, which describes the permanent price shift. While the permanent price shift is clearly very important, we do not study it here. The reader should bear in mind that all prices—p, $a(t)$, etc.—are logarithmic.

The price impact function provides a measure of the liquidity for executing market orders. (The liquidity for limit orders, in contrast, is given by the probability of execution, studied in section 2.2.5.) At any given time t, the instantaneous ($\tau = 0$) price impact function is the inverse of the cumulative depth profile. This follows immediately from eqs. (1) and (2), which in the limit $dp \to 0$

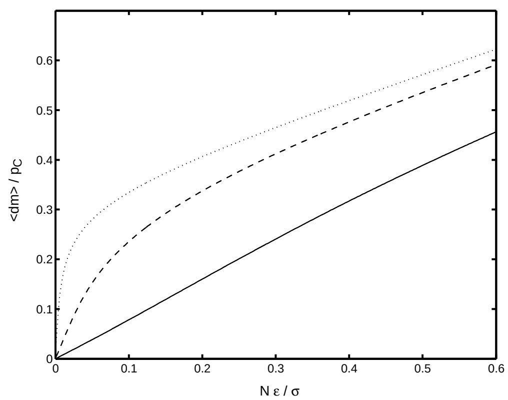

FIGURE 6 The average price impact corresponding to the results in figure 3. The average instantaneous movement of the nondimensional mid-price, $\langle dm \rangle / p_c$ caused by an order of size $N/N_c = N\epsilon/\sigma$. $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot).

can be replaced by the continuum transaction equation:

$$\omega = N(p,t) = \int_0^p n(\tilde{p},t) d\tilde{p}\,. \tag{3}$$

This equation makes it clear that at any fixed t the price impact can be regarded as the inverse of the cumulative depth profile $N(p,t)$. When the fluctuations are sufficiently small we can replace $n(p,t)$ by its mean value $n(p) = \langle n(p,t) \rangle$. In general, however, the fluctuations can be large, and the average of the inverse is not equal to the inverse of the average. There are corrections based on higher order moments of the depth profile, as given in the moment expansion derived in Smith et al. [26]. Nonetheless, the inverse of the mean-cumulative depth provides a qualitative approximation that gives insight into the behavior of the price impact function. Mean price impact functions are shown in figure 6 and the standard deviation of the price impact is shown in figure 7. The price impact exhibits very large fluctuations for all values of ϵ: The standard deviation has the same order of magnitude as the mean, or larger for small $N\epsilon/\sigma$ values. Note that these are actually *virtual price impact* functions. That is, to explore the behavior of the instantaneous price impact for a wide range of order sizes, we periodically compute the price impact that an order of a given size would have caused at that instant, if it had been submitted. We have checked that real price impact curves are the same, but they require a much longer time to accumulate reasonable statistics.

One of the interesting results in figure 6 is the scale of the price impact. The price impact is measured relative to the characteristic price scale p_c, which as we

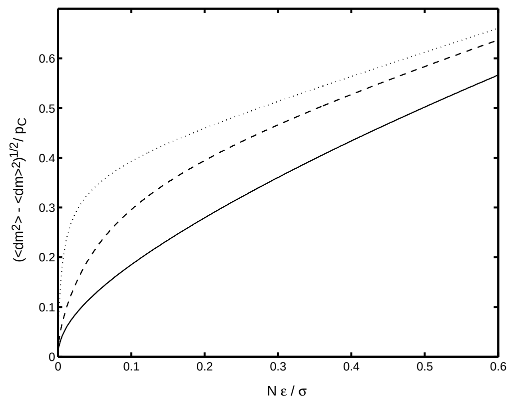

FIGURE 7 The standard deviation of the instantaneous price impact dm/p_c corresponding to the means in figure 6, as a function of normalized order size $\epsilon N/\sigma$. $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot).

have mentioned earlier is roughly equal to the mean spread. As we will argue in relation to figure 8, the range of nondimensional shares shown on the horizontal axis spans the range of reasonable order sizes. This figure demonstrates that throughout this range the price is the order of magnitude (and typically less than) the mean-spread size.

Due to the accumulation of orders at the ask in the large ϵ regime, for small p the mean-price impact is roughly linear. This follows from eq. (3) under the assumption that $n(p)$ is constant. In the medium ϵ regime, under the assumption that the variance in depth can be neglected, the mean-price impact should increase as roughly $\omega^{1/2}$. This follows from eq. (3) under the assumption that $n(p)$ is linearly increasing and $n(0) \approx 0$. (Note that we see this as a crude approximation, but there can be substantial corrections caused by the variance of the depth profile.) Finally, in the small ϵ regime the price impact is highly concave, increasing much slower than $\omega^{1/2}$. This follows because $n(0) \approx 0$ and the depth profile $n(p)$ is convex.

To get a better feel for the functional form of the price impact function, in figure 8 we numerically differentiate it versus log order size, and plot the result as a function of the appropriately scaled order size. (Note that because our prices are logarithmic, the vertical axis already incorporates the logarithm.) If we were to fit a local power law approximation to the function at each price, this corresponds to the exponent of that power law near that price. Notice that the exponent is almost always less than one, so that the price impact is almost always concave. Making the assumption that the effect of the variance of the depth is not too large, so that eq. (3) is a good assumption, the behavior of this figure can

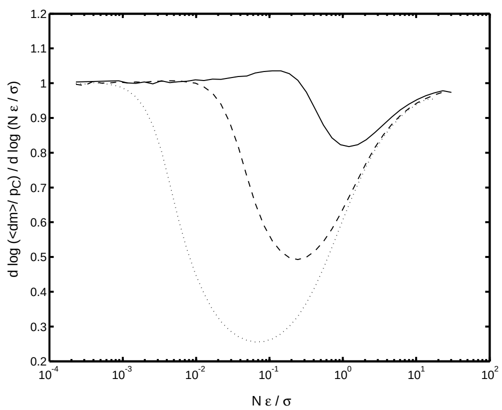

FIGURE 8 Derivative of the nondimensional mean-mid-price movement, with respect to logarithm of the nondimensional order size $N/N_c = N\epsilon/\sigma$, obtained from the price impact curves in figure 6.

be understood as follows: For $N/N_c \approx 0$ the price impact is dominated by $n(0)$ (the constant term in the average depth profile) and so the logarithmic slope of the price impact is always near to one. As N/N_c increases, the logarithmic slope is driven by the shape of the average depth profile, which is linear or convex for smaller ϵ, resulting in concave price impact. For large values of N/N_c, we reach the asymptotic region where the depth profile is flat (and where our model is invalid by design). Of course, there can be deviations to this behavior caused by the fact that the mean of the inverse depth profile is not in general the inverse of the mean, that is, $\langle N^{-1}(p)\rangle \neq \langle N(p)\rangle^{-1}$. This is discussed in more detail in Smith et al. [26].v

To compare to real data, note that $N/N_c = N\epsilon/\sigma$. N/σ is just the order size in shares in relation to the average order size, so by definition it has a typical value of one. For the London Stock Exchange, we have found that typical values of ϵ are in the range $0.001 - 0.1$. For a typical range of order sizes from $100 - 100,000$ shares, with an average size of $10,000$ shares, the meaningful range for N/N_c is therefore roughly 10^{-5} to 1. In this range, for small values of ϵ the exponent can reach values as low as 0.2. This offers a possible explanation for the previously mysterious concave nature of the price impact function, and contradicts the linear increase in price impact based on the naive argument presented in the introduction.

2.2.3 Spread. The probability density of the spread is shown in figure 9. This shows that the probability density is substantial at $s/p_c = 0$. (Remember that

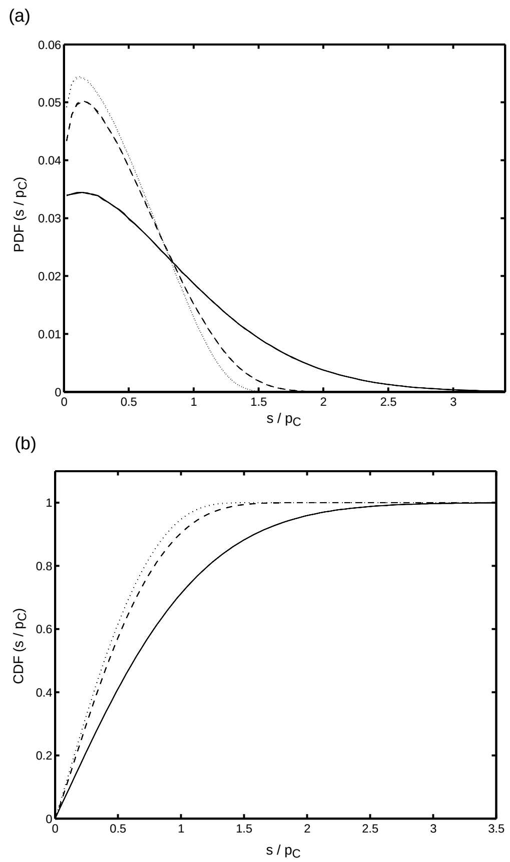

FIGURE 9 The probability density function (a), and cumulative distribution function (b) of the nondimensionalized bid-ask spread s/p_c, corresponding to the results in figure 3. $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot).

this is in the limit $dp \to 0$.) The probability density reaches a maximum at a value of the spread approximately $0.2p_c$, and then decays. It might seem surprising at first that it decays more slowly for large ϵ, where there is a large accumulation of orders at the ask. However, it should be borne in mind that the characteristic price $p_c = \mu/\alpha$ depends on ϵ. Since $\epsilon = 2\delta\sigma/\mu$, by eliminating μ this can be written $p_c = 2\sigma\delta/(\alpha\epsilon)$. Thus, holding the other parameters fixed, large ϵ corresponds to small p_c, and vice versa. So in fact, the spread is very small for large ϵ, and large for small ϵ, as expected. The figure just shows the small corrections to the large effects predicted by the dimensional scaling relations.

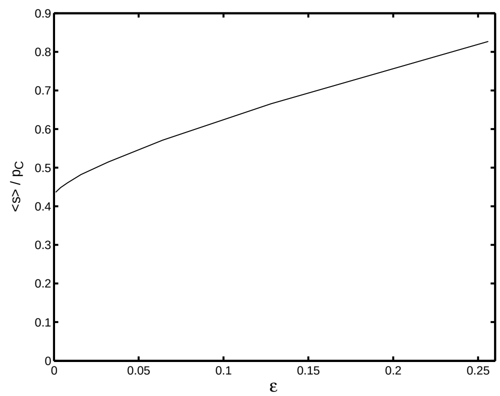

FIGURE 10 The mean value of the spread in nondimensional units $\hat{s} = s/p_c$ as a function of ϵ. This demonstrates that the spread only depends weakly on ϵ, indicating that the prediction from dimensional analysis given in table 3 is a reasonable approximation.

For large ϵ the probability density of the spread decays roughly exponentially moving away from the midpoint. This is because for large ϵ the fluctuations around the mean depth are roughly independent. Thus, the probability for a market order to penetrate to a given price level is roughly the probability that all the ticks smaller than this price level contain no orders, which gives rise to an exponential decay. This is no longer true for small ϵ. Note that for small ϵ the probability distribution of the spread becomes insensitive to ϵ, that is, the nondimensionalized distribution for $\epsilon = 0.02$ is nearly the same as that for $\epsilon = 0.002$.

It is apparent from figure 9 that in nondimensional units the mean spread increases with ϵ. This is confirmed in figure 10, which displays the mean value of the spread as a function of ϵ. The mean spread increases monotonically with ϵ. It depends on ϵ as roughly a constant (equal to approximately 0.45 in nondimensional coordinates) plus a linear term whose slope is rather small. We believe that for most financial instruments $\epsilon < 0.3$. Thus the variation in the spread caused by varying ϵ in the range $0 < \epsilon < 0.3$ is not large, and the dimensional analysis based only on rate parameters given in table 4 is a good approximation. We get an accurate prediction of the ϵ dependence across the full range of ϵ from the Independent Interval Approximation technique discussed in section 3.

2.2.4 Volatility and Price Diffusion. The price diffusion rate, which is proportional to the square of the volatility, is important for determining risk and is a property of central interest. From dimensional analysis in terms of the order flow rates the price diffusion rate has units of $price^2/time$, and so must scale as $\mu^2\delta/\alpha^2$. We

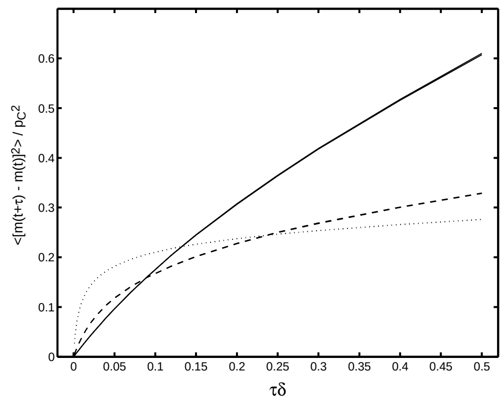

FIGURE 11 The variance of the change in the nondimensionalized midpoint price versus the nondimensional time delay interval $\tau\delta$. For a pure random walk this would be a straight line whose slope is the diffusion rate, which is proportional to the square of the volatility. The fact that the slope is steeper for short times comes from the nontrivial temporal persistence of the order book. The three cases correspond to figure 3: $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot).

can also make a crude argument for this as follows: The dimensional estimate of the spread (see table 4) is $\mu/2\alpha$. Let this be the characteristic step size of a random walk, and let the step frequency be the characteristic time $1/\delta$ (which is the average lifetime for a share to be canceled). This argument also gives the above estimate for the diffusion rate. However, this is not correct in the presence of negative autocorrelations in the step sizes. The numerical results make it clear that there are important ϵ-dependent corrections to this result, as demonstrated below.

In figure 11 we plot simulation results for the variance of the change in the midpoint price at time scale τ, $\text{Var}\,(m\,(t+\tau) - m\,(t))$. The slope is the diffusion rate, which at any fixed time scale is proportional to the square of the volatility. It appears that there are at least two time scales involved, with a faster diffusion rate for short time scales and a slower diffusion rate for long time scales. Such anomalous diffusion is not predicted by mean-field analysis. Simulation results show that the diffusion rate is correctly described by the product of the estimate from dimensional analysis based on order flow parameters alone, $\mu^2\delta/\alpha^2$, and a τ-dependent power of the nondimensional granularity parameter $\epsilon = 2\delta\sigma/\mu$, as summarized in table 4. We cannot currently explain why this power is $-1/2$ for short term diffusion and $1/2$ for long-term diffusion. However, a qualitative understanding can be gained based on the conservation law we discussed in section 3.

Note that the temporal structure in the diffusion process also implies non-zero autocorrelations of the midpoint price $m(t)$. This corresponds to weak negative autocorrelations in price differences $m(t) - m(t-1)$ that persist for time scales until the variance vs. τ becomes a straight line. The time scale depends on parameters, but is typically on the order of 50 market order arrival times. This temporal structure implies that there exists an arbitrage opportunity which, when exploited, would make prices more random and the structure of the order flow non-random.

2.2.5 Liquidity for Limit Orders: Probability and Time to Fill. The liquidity for limit orders depends on the probability that they will be filled, and the time to be filled. This obviously depends on price: Limit orders close to the current transaction prices are more likely to be filled quickly, while those far away have a lower likelihood to be filled. Figure 12 plots the probability Γ of a limit order being filled versus the nondimensionalized price at which it was placed (as with all the figures in this section, this is shown in the midpoint-price centered frame). Figure 12 shows that in nondimensional coordinates the probability of filling close to the bid for sell limit orders (or the ask for buy limit orders) decreases as ϵ increases. For large ϵ, this is less than 1 even for negative prices. This says that even for sell orders that are placed close to the best bid there is a significant chance that the offer is deleted before being executed. This is not true for smaller values of ϵ, where $\Gamma(0) \approx 1$. Far away from the spread the fill probabilities as a function of ϵ are reversed, that is, the probability for filling limit orders increases as ϵ increases. The crossover point where the fill probabilities are roughly the same occurs at $p \approx p_c$. This is consistent with the depth profile in figure 3 which also shows that depth profiles for different values of ϵ cross at about $p \sim p_c$.

Similarly Figure 13 shows the average time τ taken to fill an order placed at a distance p from the instantaneous mid-price. Again we see that though the average time is larger at larger values of ϵ for small p/p_c, this behavior reverses at $p \sim p_c$.

2.3 VARYING TICK SIZE DP/P_C

The dependence on discrete tick size dp/p_c, of the cumulative distribution function for the spread, instantaneous price impact, and mid-price diffusion, are shown in figure 14. We chose an unrealistically large value of the tick size, with $dp/p_c = 1$, to show that, even with very coarse ticks, the qualitative changes in behavior are typically relatively minor.

Figure 14(a) shows the cumulative density function of the spread, comparing $dp/p_c = 0$ and $dp/p_c = 1$. It is apparent from this figure that the spread distribution for coarse ticks "effectively integrates" the distribution in the limit $dp \to 0$. That is, at integer tick values the mean-cumulative depth profiles roughly match and, in between integer tick values, for coarse ticks the probability is smaller. This happens for the obvious reason that coarse ticks quantize the possible val-

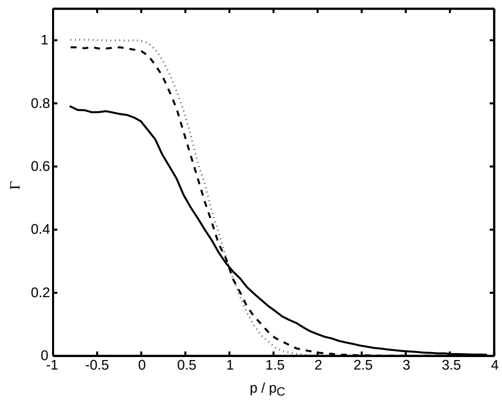

FIGURE 12 The probability Γ for filling a limit order placed at a price p/p_c where p is calculated from the instantaneous mid-price at the time of placement. The three cases correspond to figure 3: $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot).

ues of the spread, and place a lower limit of one tick on the value the spread can take. The shift in the mean spread from this effect is not shown, but it is consistent with this result; there is a constant offset of roughly 1/2 tick.

The alteration in the price impact is shown in figure 14(b). Unlike the spread distribution, the average price impact varies continuously. Even though the tick size is quantized, we are averaging over many events and the probability of a price impact of each tick size is a continuous function of the order size. Large tick size consistently lowers the price impact. The price impact rises more slowly for small p, but is then similar except for a downward translation.

The effect of coarse ticks is less trivial for mid-price diffusion, as shown in figure 14(c). At $\epsilon = 0.002$, coarse ticks remove most of the rapid short-term volatility of the midpoint, which in the continuous-price case arises from price fluctuations smaller than $dp/p_c = 1$. This lessens the negative autocorrelation of midpoint price returns, and reduces the anomalous diffusion. At $\epsilon = 0.2$, where both early volatility and late negative autocorrelation are smaller, coarse ticks have less effect. The net result is that the mid-price diffusion becomes less sensitive to the value of ϵ as tick size increases, and there is less anomalous price diffusion.

3 SUMMARY OF ANALYTIC RESULTS

This section summarizes our analytic results and discusses their agreement with simulations. For a more in-depth discussion with derivations see Smith et al. [26].

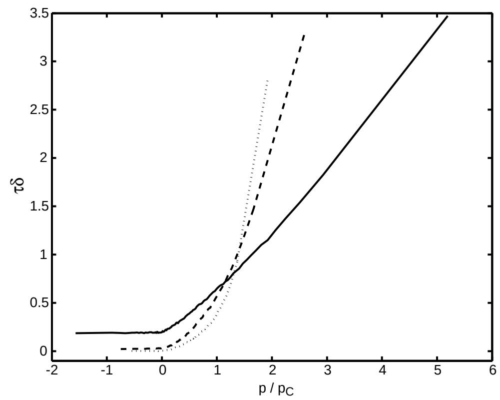

FIGURE 13 The average time τ nondimensionalized by the rate δ, to fill a limit order placed at a distance p/p_c from the instantaneous mid-price.

Below we describe three different theoretical approaches to understanding this model. A useful exact result can be derived from the requirement that all orders placed are eventually removed, which imposes global constraints on the mean-depth profile. Providing fluctuations at different prices are not too strongly correlated, an approximation to the mean-depth profile, the spread, and other properties can be obtained from an order-depth master equation. Alternatively, closed-form finite-difference expressions for the mean intervals separating orders (including the spread) may be obtained if the interval fluctuations have suitably regular distributions. These three levels of analysis are summarized in the following three subsections.

3.1 GLOBAL CONSERVATION RELATIONS

In this section we derive a useful global conservation relation. Because prices describe a random walk, in order to get stationary solutions we must use *comoving* coordinates. The resulting conservation law is slightly different depending on whether the coordinates are centered on the midpoint or on the best quote. (For convenience we derive the relation for sell orders, in which case the best quote is the best bid.)

Let $n(p,t)$ denote the number of shares in a half-closed logarithmic price interval $(p, p+dp)$ at time t, where dp is the logarithmic tick size, which may be infinitesimal. Then the share number in the bid-centered comoving frame is denoted n_b, and defined from the instantaneous bid price $b(t)$ as

$$n_b(p,t) \equiv n(p - b(t), t) . \tag{4}$$

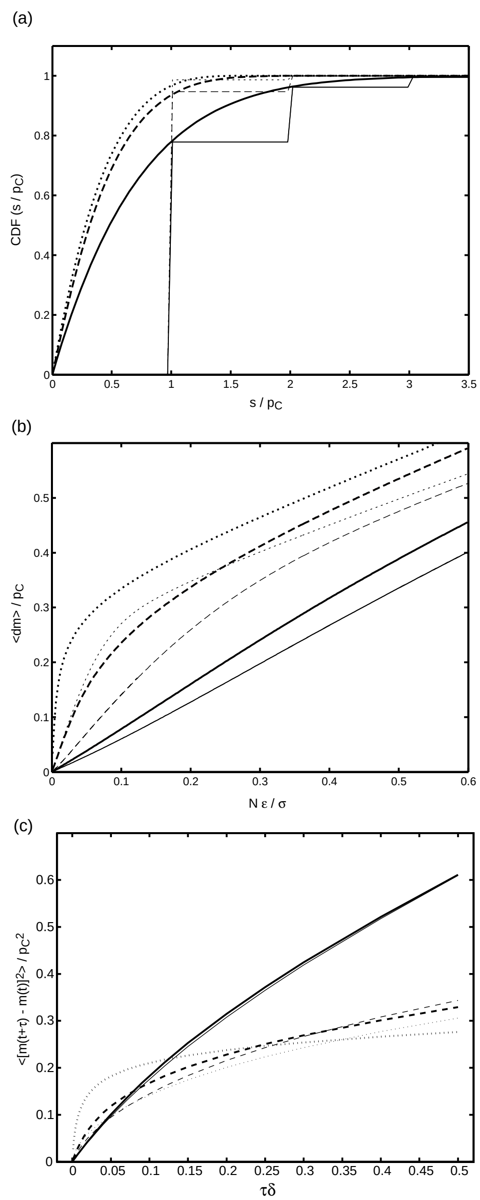

FIGURE 14 Dependence of market properties on tick size. Heavy lines are $dp/p_c \to 0$; light lines are $dp/p_c = 1$. Cases correspond to figure 3, with $\epsilon = 0.2$ (solid), $\epsilon = 0.02$ (dash), $\epsilon = 0.002$ (dot). (a) is the cumulative distribution function for the nondimensionalized spread. (b) is instantaneous nondimensionalized price impact, (c) is diffusion of the nondimensionalized midpoint shift, corresponding to figure 11.

Similarly, in a midpoint-centered frame,

$$n_m(p,t) \equiv n(p - m(t), t) . \tag{5}$$

In bid-centered coordinates, the order-placement rate density is the constant α, and the mean-decay rate in bin p is $\delta \langle n_b(p) \rangle$, where angle brackets denote interchangeably either time or ensemble average. In addition, sell limit orders are removed by the placement of buy market orders, at the rate $\mu/2$. The fact that the number of orders placed must equal the number removed implies that

$$\frac{\mu}{2} = \sum_{p=b+dp}^{\infty} (\alpha dp - \delta \langle n_b(p) \rangle) . \tag{6}$$

This relationship is somewhat more complicated in midpoint-centered coordinates, since orders placed below the midpoint induce a shift in the center of the coordinate system that place them above the new midpoint. This occurs whenever $0 \leq p \leq s/2$, where s is the spread. Thus the average additional deposition rate in midpoint-centered coordinates is $\alpha\langle s/2 \rangle$.

$$\frac{\mu}{2} = \alpha \frac{\langle s \rangle}{2} + \sum_{p=b+dp}^{\infty} (\alpha dp - \delta \langle n_m(p) \rangle) . \tag{7}$$

Equations (6) and (7) are exact constraints on the mean-order depths, which will be respected as well by the approximate solutions below.

3.2 ORDER-DENSITY MASTER EQUATION

In this section we give an overview of a treatment based on a master equation. Instantaneous order-book configurations are one-dimensional *profiles*, which evolve stochastically under order placement and removal, as well as shifts in the origin of the comoving coordinate system when there is a change in the best quotes. The number of such profiles is too large to index tractably, but if the resulting fluctuations of the number at each price are uncorrelated, the statistical properties of the limit order book can be approximately described by the density $\pi(n,p,t)$, which gives the probability of finding n orders at price p at time t. This satisfies the conservation law

$$\sum_n \pi(n,p,t) = 1, \ \forall p, t . \tag{8}$$

The approximation of uncorrelated fluctuations is never satisfied in the bid-centered frame, but under appropriate conditions, discussed later, it is sometimes satisfied in the midpoint-centered frame. Therefore, n will denote n_m in the remainder of this subsection.

The master equation describing the flow of probability from individual placement, expiration, and execution events, as well as coordinate shifts, is straightforward to write down as

$$\begin{aligned}\frac{\partial}{\partial t}\pi(n,p) = {} & \frac{\alpha(p)\,dp}{\sigma}\left[\pi(n-\sigma,p)-\pi(n,p)\right] \\ & + \frac{\delta}{\sigma}\left[(n+\sigma)\,\pi(n+\sigma,p)-n\pi(n,p)\right] \\ & + \frac{\mu(p)}{2\sigma}\left[\pi(n+\sigma,p)-\pi(n,p)\right] \\ & + \sum_{\Delta p} P_{+}(\Delta p)\left[\pi(n,p-\Delta p)-\pi(n,p)\right] \\ & + \sum_{\Delta p} P_{-}(\Delta p)\left[\pi(n,p+\Delta p)-\pi(n,p)\right].\end{aligned} \tag{9}$$

We are assuming time increments are sufficiently small that the time difference is well approximated by a continuous derivative, and have neglected to write the variable t that appears in every term on the right side. Here $P_{\pm}(\Delta p)$ are the rate densities for upward and downward shifts of the frame by Δp. They will be assumed equal for simplicity, and must be found self-consistently with the solution for mean n. To do this we assume that the shift events are otherwise uncorrelated with placement and removal events. In comoving coordinates the order placement rates are now functions of price, $\mu(p)$ and $\alpha(p)$. $\mu(p)$ represents the average rate at which market orders remove limit orders at price p. Similarly, the average rate of limit order deposition $\alpha(p)$ is affected by the fact that the best bid and ask prices are often changing. Far from the midpoint the deposition rate is unaffected, so that $\alpha(\infty) = \alpha$. These functions must be solved self-consistently.

For the analytic treatment we assume that all orders are of the same size (σ shares). In the limit $dp \rightarrow 0$ the number of orders within any given price bin of width dp is either zero or one. A solution for the mean-depth profile $\langle n(p)\rangle$ can be obtained by multiplying eq. (9) by n, summing over n, and setting the time dependence to zero. Differences of π at adjacent prices are replaced by derivatives with respect to p, and the finite-Δp shifts are expanded in Taylor's series. Two transport parameters are defined: a diffusivity

$$D \equiv \sum_{\Delta p} P(\Delta p)\,\Delta p^{2}, \tag{10}$$

and a mean-absolute price shift

$$\langle \Delta p\rangle \equiv \frac{\sum_{\Delta p} P(\Delta p)\,\Delta p}{\sum_{\Delta p} P(\Delta p)}. \tag{11}$$

It is convenient to express the mean-share density nondimensionally, defining

$$\psi\left(\hat{p}\right) \equiv \frac{1}{\epsilon} \frac{\delta \left\langle n\left(p\right)\right\rangle}{\alpha\left(\infty\right) dp}, \tag{12}$$

and similarly to introduce nondimensionalized transport parameters $\beta \equiv D/(p_C^2 \delta)$ and $\langle \Delta \hat{p} \rangle \equiv \langle \Delta p \rangle / p_C$.

Under the mean-field approximation of independent fluctuations, it is convenient to think of buy market orders as being deposited at logarithmic price $p = 0$, and moving to the right until they are annihilated by a sell limit order. (Recall that by definition in midpoint coordinates there are never any limit orders stored at $p < 0$.) Using non-dimensional price coordinates $\hat{p}$, the fraction of market orders surviving to price p, which is by definition also the cumulative distribution function for $\hat{s}/2$, has a simple expression in terms of the mean density:

$$\frac{\mu\left(\hat{p}\right)}{\mu\left(0\right)} = \Pr\left(\hat{s}/2 \geq \hat{p}\right) = \exp\left(-\int_0^{\hat{p}} d\hat{p}' \psi\left(\hat{p}'\right)\right) \equiv \varphi\left(\hat{p}\right) . \tag{13}$$

Equation (13), together with the excess order-deposition relation in the midpoint-centered frame $\alpha\left(\hat{p}\right)/\alpha\left(\hat{\infty}\right) = 1 + \Pr\left(\hat{s}/2 \geq \hat{p}\right)$ (explained above), may be used to reduce the static first-moment (mean) solution of eq. (9) to

$$1 + \varphi = -\left[\frac{d\varphi}{d\hat{p}} + \epsilon\left(1 - \beta\frac{d^2}{d\hat{p}^2}\right)\frac{d\log\varphi}{d\hat{p}}\right]. \tag{14}$$

Direct integration over prices recovers the order-flow conservation (7), but now in nondimensional coordinates.

$$\int_0^{\infty} d\hat{p}\left(1 - \epsilon\psi\right) = 1 - \frac{\langle \hat{s} \rangle}{2}, \tag{15}$$

as long as $\psi \equiv 0$ at $p < 0$ and $\psi \to 1/\epsilon$ for $\hat{p} \to \infty$.

The simple fact that stored limit orders can never cross the midpoint can be used to derive an additional condition that yields a unique solution. While we cannot enforce this condition microscopically, we can at least enforce it on average. Consider the bin at $\hat{p} = 0$. The flux of orders from above is the gradient in order density times the rate of positive shifts $\langle \Delta \hat{p} \rangle$. In contrast, there is no flux of orders from below (since there are no orders below); the rate of removal is proportional to the order density. This argument can be made more formally (see Smith et al. [26]), leading to the condition

$$\langle \Delta \hat{p} \rangle \left.\frac{d\psi}{d\hat{p}}\right|_0 - \psi\left(0\right) \approx 0\,. \tag{16}$$

The self-consistently determined parameters evaluate to

$$\beta = \frac{4}{\epsilon}\int_0^{\infty} d\Delta\hat{p}\left(\Delta\hat{p}\right)^2 \varphi\left(\Delta\hat{p}\right), \tag{17}$$

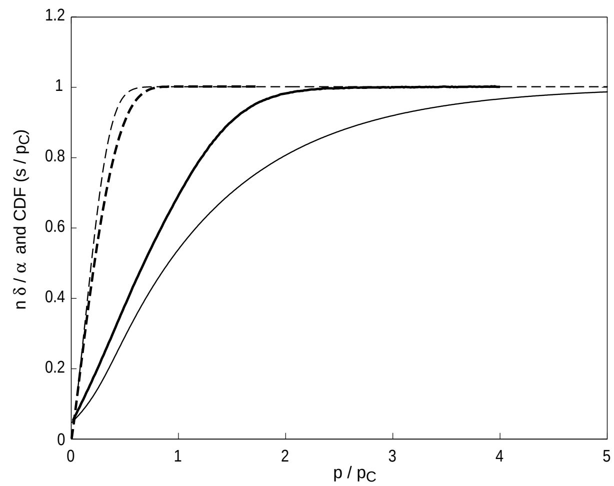

FIGURE 15 Fit of self-consistent solution to simulation results for the midpoint-centered frame, $\epsilon = 0.02$. Thin solid line is the analytic mean-number density; thick solid line is simulation. Thin dashed line is analytic $\Pr(\hat{s}/2 \leq \hat{p})$; thick dashed line is simulation.

and

$$\langle \Delta\hat{p} \rangle = \frac{\int_0^\infty d\Delta\hat{p}\ (\Delta\hat{p})\,\varphi\,(\Delta\hat{p})}{\int_0^\infty d\Delta\hat{p}\,\varphi\,(\Delta\hat{p})\,.} \tag{18}$$

Simultaneous solution of eqs. (14) and (16)–(18) produces the density and cumulative spread distribution shown in figure 3.2.

3.3 INDEPENDENT INTERVAL APPROXIMATION

An alternative to considering the share depth at $dp \to 0$ is to consider the set of price intervals x_i between orders, which in this limit are ensured to be sparse. x_0 is defined to be the spread, negative i index intervals between buy limit orders (bids), and positive i index intervals between offers. (For simplicity in defining the model, new orders are excluded from price bins containing existing orders. The resulting corrections vanish as $dp \to 0$.)

The instantaneous intervals change stochastically, by splitting when orders are added, or by joining when they are removed. For the spread, the processes and their rates are as follows:

1. $x_0 \to x_0 + x_1$, rate $(\delta + \mu/2\sigma)$ (ask removed by cancellation or market order).
2. $x_0 \to x_0 + x_{-1}$, rate $(\delta + \mu/2\sigma)$ (symmetric process for bid removal).

3. $x_0 \to x' \in (1, x_0 - 1)$, when a new offer is placed, rate $\alpha dp/\sigma$ at each unoccupied position of width dp.
4. $x_0 \to x' \in (1, x_0 - 1)$ when new bid is placed, again rate $\alpha dp/\sigma$ per bin.
5. Rate for events leaving x_0 unchanged is therefore $1 - 2\delta - \mu/\sigma - 2\alpha dp\,(x_0 - 1)/\sigma$.

To simplify notation in what follows σ will be set to one without loss of generality.

The expected $x_0\,(t + dt)$, given definite $x_i\,(t)$, is then

$$\begin{aligned}\langle x_0\,(t + dt)\rangle = x_0\,(t)\,[1 - 2\delta - \mu_0 - 2\alpha\,(x_0 - 1)] \\ + (x_0 + x_1)\left(\delta + \frac{\mu}{2}\right) + (x_0 + x_{-1})\left(\delta + \frac{\mu}{2}\right) \\ + (\alpha_0 dp) x_0\,(x_0 - 1)\ .\end{aligned} \tag{19}$$

The average of eq. (19) would generate a recursion for $\langle x_0 \rangle$ from $\langle x_{\pm 1} \rangle$ if we could evaluate $\left\langle x_0^2 \right\rangle$. The mean-field approximation for independent intervals is to assume some relation $\left\langle x_0^2 \right\rangle$ by $a\langle x_0 \rangle^2$, with a to be determined self-consistently. Making this assumption, and abusing the notation by letting x_i denote the mean value of stationary solutions (and no longer the instances), gives

$$\left(\delta + \frac{\mu}{2} s\right)(x_1 + x_{-1}) = a\alpha dp x_0\,(x_0 - 1)\,. \tag{20}$$

For sparse orders, the relation of mean interval to mean density depends on the fluctuation spectrum, but it is qualitatively $x_i \approx 1/\left\langle n\left(\sum_{j=0}^{i-1} x_j dp\right)\right\rangle$, becoming exact at large i. Thus, corresponding to the density solution, it is convenient to nondimensionalize the x_i as

$$\hat{x}_i \equiv \epsilon \frac{\alpha}{\delta} x_i dp = \frac{x_i dp}{p_C} \approx \frac{1}{\psi\left(\sum_{j=0}^{i-1} \hat{x}_j\right)}\ , \tag{21}$$

placing eq. (20) in the form

$$(1 + \epsilon)\,(\hat{x}_1 + \hat{x}_{-1}) = a\hat{x}_0\,(\hat{x}_0 - d\hat{p})\,. \tag{22}$$

The same sequence of steps may be followed for all x_k at $k \geq 1$, to yield the nondimensional recursion relations

$$(1 + k\epsilon)\,\hat{x}_k = \frac{a}{2}\hat{x}_{k-1}\,(\hat{x}_{k-1} - d\hat{p}) + \hat{x}_{k-1}\sum_{i=0}^{k-2}(\hat{x}_i - d\hat{p})\ , \tag{23}$$

which may then be solved numerically, given a convergence condition on $k \to \infty$ and the assumption of symmetry $x_i = x_{-i}$.

For asymptotically constant order placement rate density α, $\hat{x}_k$ must converge to some value $\hat{x}_\infty$ at large k. Taking $\hat{x}_{k+1} \to \hat{x}_k$ in eq. (23), it follows that $\hat{x}_\infty = \epsilon + d\hat{p}$, providing the convergence condition that constrains $\hat{x}_0$ in numerical evaluation, and agreeing with $\psi\,(\infty) \to 1/\epsilon$ at $dp \to 0$.

The partial sum in eq. (23) may also be evaluated in the same limit, by dividing by $\hat{x}_\infty$, to give

$$(1 + k\epsilon) = \frac{a}{2}\left(\hat{x}_\infty - d\hat{p}\right) + \sum_{i=0}^{k-2}\left(\hat{x}_i - d\hat{p}\right) . \tag{24}$$

Expressing $k\epsilon$ as a partial sum over $\hat{x}_\infty - d\hat{p}$, it follows that

$$1 + (1 - \frac{a}{2})\epsilon = S_\infty, \tag{25}$$

where $S_\infty \equiv \sum_{i=0}^{\infty}\left(\hat{x}_i - \hat{x}_\infty\right)$.

S_∞ may be interpreted in terms of global order flow. The rate of decay of the $k+1$ orders in the price range $\sum_{i=0}^{k} x_i$ is $\delta\,(k+1)$, while the rate of market order executions is $\mu/2$. These must balance the rate of market order additions, which is $(\alpha dp)\sum_{i=0}^{k} x_i$ in the bid-centered frame (where the addition rate is not expressed directly in terms of the fluctuation spectrum of x_0). Thus, re-expressing eq. (6),

$$\frac{\mu}{2\sigma} + \delta\,(k+1) = \frac{\alpha dp}{\sigma}\sum_{i=0}^{k} x_i\,, \tag{26}$$

which nondimensionalized gives

$$1 = S_\infty\,. \tag{27}$$

Matching eq. (27) to eq. (25) gives the self-consistency condition $a = 2$, which would be exact if the x_k were all independent and exponentially distributed. The actual distribution obtained from simulations is approximately exponential for large k, though it is more nearly Gaussian for x_0, and has some transitional form for small k. Even so, the correspondence with the bid-centered density profile is qualitatively good across a broad range of ϵ, as shown in figure 3.3.

3.4 UTILITY OF ANALYTIC RESULTS AND AGREEMENT WITH SIMULATION

Both of the foregoing solution methods make simplifying assumptions about fluctuations: the master equation assumes that depth fluctuations are independent, while the independent interval approximation assumes that higher order interval moments have a fixed relation to mean values. Comparison to simulations has shown that both forms of approximation are good for parameter ranges $\epsilon \gtrsim 0.1$, but that they lead to progressively larger quantitative errors for smaller ϵ, so that only qualitative features of the density profile or probability distribution for the spread remain correctly predicted for $\epsilon \lesssim 0.001$.

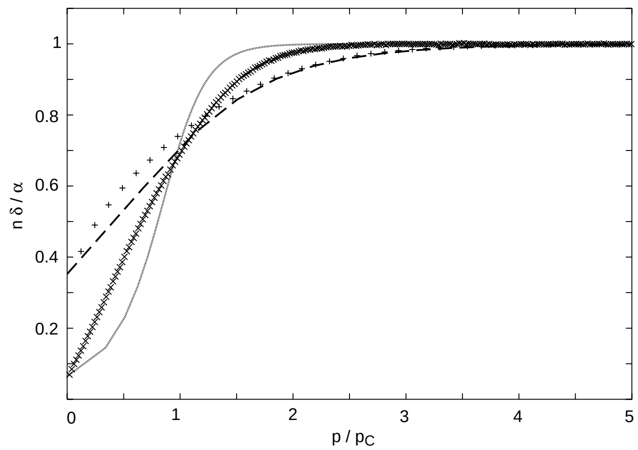

FIGURE 16 Bid-centered density profiles from Monte Carlo simulation (markers) and the Independent Interval Approximation (lines). Pluses and dash line are for $\epsilon = 0.2$, while crosses and dotted line are for $\epsilon = 0.02$.

Nevertheless, the analytic methods correctly capture the progression of the profile from concave-everywhere, to inflected, as ϵ decreases (see fig. 3.3). Qualitatively, this result allows us to understand the progression of the market impact from nearly linear, to sublinear-power dependence on order size, and shows that this relation is recovered in very different mean-field treatments, thus is not dependent at leading order, on precise properties of fluctuations.

The ϵ dependence of the mean profile, combined with the global conservation laws, also gives insight into the nature of autocorrelation of the midprice movement, and relates it to the impact through the mean order-book profile. The quantity S_∞ represents the area, in the nondimensionalized coordinates of figure 3.3, between the mean profile and a constant function with value unity. The conservation law (27) implies that this area is in fact independent of order-flow parameters. With the qualitative behavior of the book just noted, a larger market order rate (smaller ϵ) produces a lower profile near the bid, thus requiring that the profile more rapidly asymptote to one for $p/p_C \gtrsim 1$. The sparser profile near the bid indicates larger or more frequent steps in the random walk, due to shifts in the ask (and so, by symmetry, also in the bid), leading to a short-time diffusivity that should increase with decreasing ϵ. However, the exponent with which the profile asymptotes to unity at large p/p_C is inversely proportional to the late-time diffusivity, as it functions in simple diffusion models [5] indicating that this quantity must decrease with decreasing ϵ. The two are quantitatively related by the constraint (27), and in fact are shown to scale with inverse powers of $\epsilon^{1/2}$ in simulations. The physical interpretation of this relation is that, while more market orders lead to a sparser interior profile and more rapid initial price diffusion, an even greater fraction of the early steps is reversed by negatively autocorrelated later steps, reflecting the relatively greater immobility of the deeper book at large prices.

Perhaps the most quantitatively successful aspect of the analytic treatment, though, is that it motivates the nondimensionalization of the problem, by showing relatively invariant defining equations and qualitative solutions, in appropriately chosen coordinates. It shows relatively easily the existence of the continuum limit for tick size $dp/p_C \to 0$, and concurrently the nonexistence of a regular limit for order granularity $\sigma/N_C \equiv \epsilon \to 0$. It motivates parameter ranges for simulation studies, by showing that the regions of most rapid qualitative change occur over the range $\epsilon \in 0.001 - 0.1$, and gives some qualitative meaning to the order-flow values in real markets, in terms of the sensitivity to change where they occur.

4 CONCLUDING REMARKS

4.1 ONGOING WORK ON EMPIRICAL VALIDATION

This model predicts many different aspects of markets. To test these predictions quantitatively it is necesssary to measure order flow rates, which are not available in most data sets. It is nonetheless possible to compare some of the qualitative predictions of the model to those of data. For example, in a recent careful study Lillo et al. have carefully measured the price impact function for 1000 stocks traded on the New York Stock Exchange [20]. They find a price impact function that is quite concave. It does not appear to follow any simple functional form, such as a power law or logarithm, but increases roughly as the 0.5 power for small orders and the 0.2 power for larger orders. This is roughly the behavior our model generates for small values of ϵ, e.g., $\epsilon \approx 0.001$.

Members of our group are also working to test this model using data from the London Stock Exchange [11]. We have chosen this data set because it contains every order and every cancellation, which makes it possible to measure all the parameters of the model directly. It is also possible to reconstruct the order book and measure all the statistical properties we have studied in this chapter. Our empirical work so far shows that, despite its crude approximations, many of the predictions of the model are quite good. In particular, for a preliminary set of nine stocks the model explains 70% of the variance of the mean daily spread. Also, when plotted using the nondimensional coordinates defined here, the price impact function for the nine different stocks collapses rather well onto a single function. There are also some discrepancies; for example, the collapse seems to be independent of ϵ. If we somewhat arbitrarily choose $\epsilon \approx 0.001$, we get a good fit to the data. The shape of the price impact function is strikingly similar to that observed for the NYSE.

We believe that the discrepancies between the predictions of our model and the data can be dealt with by using a more sophisticated model of order flow. We summarize some of the planned improvements in the following subsection.

4.2 FUTURE ENHANCEMENTS

As we have mentioned above, the zero intelligence, IID order flow model should be regarded as just a starting point from which to add more complex behaviors. We are considering several enhancements to the order flow process whose effects we intend to discuss in future papers. Some of the enhancements include:

- *Trending of order flow.* We have demonstrated that IID order flow necessarily leads to non-IID prices. The converse is also true: Non-IID order flow is necessary for IID prices. In particular, the order flow must contain trends, i.e., if order flow has recently been skewed toward buying, it is more likely to continue to be skewed toward buying. If we assume perfect market efficiency, in the sense that prices are a random walk, this implies that there must be trends in order flow.
- *Power law placement of limit prices.* For both the London Stock Exchange and the Paris Bourse, the distribution of the limit price relative to the best bid or ask appears to decay as a power law [5, 29]. Our investigations of this show that this can have an important effect. Exponents larger than one result in order books with finite numbers of orders. In this case, depending on other parameters, there is a finite probability that a single market order can clear the entire book [26].
- *Power law or log-normal order size distribution.* A real order placement process has an order size distribution that appears to be roughly like a log-normal distribution with a power law tail [21]. This has important effects on the fluctuations in liquidity.
- *Non-Poisson order cancellation process.* When considered in real time, order placement cancellation does not appear to be Poisson [7]. However, this may not be a bad approximation in event time rather than real time.
- *Conditional order placement.* Agents may conditionally place larger market orders when the book is deeper, causing the market impact function to grow more slowly. We intend to measure this effect and incorporate it into our model.
- *Feedback between order flow and prices.* In reality, there are feedbacks between order flow and price movementss beyond the feedback in the reference point for limit order placement built into this model. This can induce bursts of trading, causing order flow rates to speed up or slow down, and give rise to clustered volatility.

The last item is just one of many examples of how one can surely improve the model by making order flow conditional on available information. However, we believe it is important to first gain an understanding of the properties of simple unconditional models, and then build on this foundation to get a fuller understanding of the problem.

4.3 COMPARISON TO STANDARD MODELS BASED ON VALUATION AND INFORMATION ARRIVAL

In the spirit of Gode and Sunder [17], we assume a simple, zero-intelligence model of agent behavior and show that the market institution exerts considerable power in shaping the properties of prices. While not disputing that agent behavior might be important, our model suggests that, at least on the short time scale, many of the properties of the market are dictated by the market institution, and, in particular, the need to store supply and demand to facilitate trading. Our model is stochastic and fully dynamic, and makes predictions that go beyond the realm of experimental economics, giving quantitative predictions about the fundamental properties of a real market. We have developed what were previously conceptual toy models in the physics literature into a model with testable explanatory power.

This raises questions about the comparison to standard models based on the response of valuations to news. The idea that news might drive changes in order flow rates is compatible with our model. That is, news can drive changes in order flow, which in turn cause the best bid or ask price to change. But notice that in our model there are no assumptions about valuations. Instead, everything depends on order flow rates. For example, the diffusion rate of prices increases as the $5/2$ power of market order flow rate, and thus volatility, which depends on the square root of the diffusion rate, increases as the $5/4$ power. Of course, order flow rates can respond to information; an increase in market order rate indicates added impatience, which might be driven by changes in valuation. But a change in long-term valuation could equally well cause an increase in limit order flow rate, which *decreases* volatility. Valuation *per se* does not determine whether volatility will increase or decrease. Our model says that volatility does not depend directly on valuations, but rather on the urgency with which they are felt, and the need for immediacy in responding to them.

Understanding the shape of the price impact function was one of the motivations that originally set this project into motion. The price impact function is closely related to supply and demand functions, which have been central aspects of economic theory since the 19th century. Our model suggests that the shape of price impact functions in modern markets is significantly influenced not so much by strategic thinking as by an economic fundamental: The need to store supply and demand in order to provide liquidity. *A priori* it is surprising that this requirement alone may be sufficient to dictate at least the broad outlines of the price impact curve.

Our model offers a “divide and conquer” strategy to understanding fundamental problems in economics. Rather than trying to ground our approach directly on assumptions of utility, we break the problem into two parts. We provide an understanding of how the statistical properties of prices respond to order flow rates, and leave the problem open of how order flow rates depend on more fundamental assumptions about information and utility. Order flow rates have the significant advantage that, unlike information, utility, or the cognitive powers of an agent, they are directly measurable. We hope that breaking the problem into two pieces will greatly simplify the problem of understanding markets.

ACKNOWLEDGMENTS

We would like to thank the McKinsey Corporation, Credit Suisse First Boston, the McDonnel Foundation, Bob Maxfield, and Bill Miller for supporting this research. We would also like to thank Paolo Patelli, R. Rajesh, Spyros Skouras, and Ilija Zovko for helpful discussions.

REFERENCES

[1] Bachelier, L. "Théorie de la spéculation." (Thesis) **Annales Scientifiques de Ecole Normale Superieure**, III–17, pp. 21–86, 1900. English translation in *Random Character of Stock Prices*, edited by P. H. Cootner. Cambridge, MA: MIT Press, 1964.

[2] Bak, P., M. Paczuski, and M. Shubik. "Price Variations in a Stock Market with Many Agents." *Physica A* **246** (1997): 430.

[3] Bollerslev, T., I. Domowitz, and J. Wang. "Order Flow and the Bid-Ask Spread: An Empirical Probability Model of Screen-Based Trading." *J. Econ. Dynamics & Control* **21** (1997): 1471.

[4] Bouchaud, J.-P., and R. Cont. "A Langevin Approach to Stock Market Fluctuations and Crashes." *European Phys. J. B* **6** (1998): 543.

[5] Bouchaud, J. P., M. Mézard, and M. Potters. "Statistical Properties of Stock Order Books: Empirical Results and Models." *Quant. Fin.* **2** (2002): 251.

[6] Bridgman, P. W. *Dimensional Analysis*. New Haven, CT: Yale University Press, 1922.

[7] Challet, D., and R. Stinchcombe. "Analyzing and Modeling 1+1d Markets." *Physica A* **300** (2001): 285.

[8] Challet, D., and R. Stinchcombe. "Exclusion Particle Models of Limit Order Financial Markets." *Physica A* **300** (2001): 285.

[9] Chiarella, C., and G. Iori. "A Simulation Analysis of the Microstructure of Double Auction Markets." *Quant. Fin.* **2** (2002): 346.

[10] Cohen, K. J., R. M. Conroy, and S. F. Maier. "Order Flow and the Quality of the Market." In *Market Making and the Changing Structure of the Securities Industry*, edited by Y. Amihud, T. Ho, and R. Schwartz. Lexington, MA: Lexington Books, 1985.

[11] Daniels, Marcus, J. Doyne Farmer, László Gillemot, Paolo Patelli, D. Eric Smith, and Ilija Zovko. "Predicting the Spread, Volatility, and Price Impact Based on Order Flow: The London Stock Exchange." Preliminary manuscript.

[12] Daniels, M., J. D. Farmer, G. Iori, and D. E. Smith. "How Storing Supply and Demand Affects Price Diffusion." 2001. ⟨http://xxx.lanl.gov/cond-mat/0112422⟩.

[13] Domowitz, I., and Jianxin Wang. "Auctions as Algorithms." *J. Econ. Dynamics & Control* **18** (1994): 29.

[14] Eliezer, D., and I. I. Kogan. "Scaling Laws for the Market Microstructure of the Interdealer Broker Markets." ⟨http://xxx.lanl.gov/cond-mat/9808240⟩.

[15] Farmer, J. D. "Market Force, Ecology, and Evolution." *J. Ind. & Corporate Change* **11(5)** (2002): 895–953.

[16] Farmer, J. D. "Slippage 1996." Technical Report, Prediction Company, Santa Fe, NM, 1996. ⟨http://www.predict.com⟩

[17] Gode, D., and S. Sunder. "Allocative Efficiency of Markets with Zero Intelligence Traders: Markets as a Partial Substitute for Individual Rationality." *J. Pol. Econ.* **101** (1993): 119.

[18] Hausman, J. A., and A. W. Lo. "An Ordered Probit Analysis of Transaction Stock Prices." *J. Fin. Econ.* **31** (1992): 319–379.

[19] Kempf, A., and O. Korn. "Market Depth and Order Size." *J. Fin. Markets* **2** (1999): 29.

[20] Lillo, F., J. D. Farmer, and R. Mantegna. "Single Curve Collapse of the Price Impact Function." *Nature* (2002): to appear.

[21] Maslov, S., and M. Mills. "Price Fluctuations from the Order Book Perspective—Empirical Facts and a Simple Model." *Physica A* **299** (2001): 234.

[22] Maslov, S. "Simple Model of a Limit Order-Driven Market." *Physica A* **278** (2000): 571.

[23] Mendelson, H. "Market Behavior in a Clearing House." *Econometrica* **50** (1982): 1505–1524.

[24] Plerou, V., P. Gopikrishnan, X. Gabaix, and H. E. Stanley. "Quantifying Stock Price Response to Demand Fluctuations." *Phys. Rev. E* **66** (2002): 027104.

[25] Slanina, F. "Mean-Field Approximation for a Limit Order Driven Market Model." *Phys. Rev. E* **64** (2001): 056136.

[26] Smith, Eric, J. Doyne Farmer, László Gillemot, and Supriya Krishnamurthy. "Statistical Theory of the Continuous Double Auction." ⟨http://xxx.lanl.gov/cond-mat/0210475⟩.

[27] Tang, L.-H., and G.-S. Tian. " Reaction-Diffusion-Branching Models of Stock Price Fluctuations." *Physica A* **264** (1999): 543.

[28] Torre, N. *BARRA Market Impact Model Handbook*. BARRA Inc, Berkeley CA, 1997. ⟨http://www.barra.com⟩.

[29] Zovko, I. I., and J. D. Farmer. "The Power of Patience: A Behavioral Regularity in Limit Order Placement." *Quant. Fin.* **2** (2002): 387.