Dynamic Price Discovery in a Divergent Expectations Environment

Jacob Paroush
Professor of Economics
Bar Ilan University, Ramat-Gan, Israel
and
Zicklin School of Business
Baruch College, CUNY

Robert A. Schwartz
Professor of Finance
Zicklin School of Business
Baruch College, CUNY

Avner Wolf
Professor of Finance
Zicklin School of Business
Baruch College, CUNY

Draft: May 17, 2005
Please do not quote without permission of authors

Dynamic Price Discovery in a Divergent Expectations Environment*

Jacob Paroush, Robert A. Schwartz, and Avner Wolf

Abstract

We use theoretical modeling and simulation to analyze price discovery for equity shares in a divergent expectations environment where *k* percent of participants have a high evaluation, (1-*k*) have a low evaluation, and *k* is not known by agents before they come to the market to trade. Participants can change their evaluations when, from the order flow they observe, they infer each other's evaluations. We show that, in this environment, price discovery is a path dependent process that leads to multiple equilibria. The analysis yields important implications for price volatility, technical analysis, behavioral finance, and market structure.

Key words: Price discovery, divergent expectations, Polya Process, adaptive valuation, multiple equilibria, path dependency

***** We thank Steve Wunsch for his key inputs into this paper. Bill Abrams, Turan Bali, Paul Davis, Ozgur Demirtas, Sonali Hazarika, and Henri Waelbroeck also made helpful suggestions, for which we are grateful. We also appreciate the comments received at the Graduate Center Seminar at the City University of New York. The simulation was professionally produced by Greg Sipress, to whom we are most thankful. We also thank Dror Parnes for his assistance in programming the initial version of the simulation model.

Introduction

A financial market is an arena where the diverse trading desires of a set of participants are harmonized as orders are placed, prices established, and trades made. As trading progresses, orders are specified based on participants' desires to trade and their expectations of the prices at which trades will be made. The realized transaction prices, in turn, are based on the orders that have been submitted. This interplay between expected and realized prices is of particular interest when traders have diverging expectations concerning a security and adjust their evaluations on the basis of the revealed behavior of others (i.e., when participants have adaptive evaluations).

We view the participants in such an environment as operating within the context of a network. The paper focuses on the dynamic process of price discovery in the network environment characterized by divergent expectations and adaptive valuations. It fits into a more general Polya-type model characterized by increasing returns, path dependency, and multiple equilibria.[1] For instance, consider a set of participants who sequentially express their preferences for one of two alternatives, A and B. With increasing returns, as the percentage of the population that selects one of the two options over the other increases, more participants who would otherwise have selected the relatively unsuccessful option switch and choose the increasingly successful alternative. Consequently, the early chance arrival of more participants selecting A can result in A becoming dominant in the market, or the early chance arrival of more participants selecting B can result in B becoming dominant in the market. This means that the process is path dependent and that multiple equilibria exist.

"A vs. B" can represent a spectrum of alternatives such as different technologies (e.g., VCRs vs. Beta Max), different locations for an industry (e.g., Silicone Valley vs. elsewhere), or different valuations for a stock (e.g., \$45 vs. \$55).[2] We focus on the latter. Just as chance, early events can result in VCRs taking the market away from Beta Max and in high tech firms clustering around Silicone Valley instead of elsewhere, a stock's price can settle closer to 45 or closer to 55 simply because of the chance arrival of early events. In the case of technology, increasing returns are explained by economies of standardization. With regard to location, increasing returns are attributable to the

---

[1] See Arthur (1994).

[2] The VCR vs Beta Max and Silicone Valley examples were suggested by Arthur (1994).

economic benefits of spatial clustering for firms in the same industry. For share price, increasing returns are due to the validation provided by more participants agreeing on a price.[3] Our analysis of increasing returns and path dependency suggests that short-term (e.g., intra-day) price volatility is greater when, all else constant, participant expectations are more divergent (this would not be the case in a frictionless trading environment). This short-term volatility analysis has implications for longer run price behavior. For instance, the genesis of the origins of longer term bubbles and crashes could lie, in part at least, in a price discovery process that starts on an event-to-event basis. In addition, our analysis interfaces with various issues concerning technical analysis, behavioral finance, and market structure.

Much financial modeling is based on the assumption that participants have homogeneous expectations. While homogeneity is commonly assumed for purposes of tractability, it is also widely accepted as being reasonable (what one rational person would conclude from an information set, all rational people should conclude). We alternatively relax the stringent assumption of homogeneity and allow for divergent expectations.[4] We do so in light of the enormous size and complexity of the information set that characterizes real world equity markets, the imprecision of the tools available to analyze share values, the production of private information by agents, the observation that analyst recommendations commonly differ, and in recognition of the prevalence of short selling in the marketplace.[5]

The divergence of expectations has major implications for price discovery. Price discovery, which is relatively trivial under homogeneous expectations, becomes a

---

[3] The use of a daily volume weighted average price (VWAP) as a performance benchmark suggests the importance institutional participants place on trading at a validated price.

[4] See Miller (1977) and Harrison and Kreps (1978) for earlier discussions of heterogeneous expectations in a static context. More recently, Scheinkman and Xiong (2003) have presented a dynamic model that explicitly addresses heterogeneous expectations.

[5] While it is outside the scope of this paper to infer the range of individual assessments, a back of the envelope calculation suggests that it could be quite large. To illustrate, consider a simple dividend discount model, and assume that in one year a firm will start paying a dividend of $1.35 a year and that its equity cost of capital is 10.00%. Further assume that one analyst expects a 7.00% annual growth rate, that a second analyst expects a 7.545% growth rate and that, accordingly, the first analyst values the shares at $45 and the second analyst values the shares at $55. The example shows how a disparity of 55 basis points for the growth rate can translate into a $10 disparity for share value in the price range we have considered. With respect to price, the difference is very large, but with respect to a growth rate, the difference is very small. Is any analyst able to assess a growth rate with 55 basis points precision? Alternatively viewed, a firm with annual earnings of $2 will have a share value of $45 if its price/earnings multiple is 22.5, and a share value of $55 if its P/E multiple is 27.5, a 22% difference. Can a P/E multiple be assessed with 22% precision?

complex, dynamic process when the homogeneity assumption is relaxed. The dynamic complexity itself can produce what might be referred to as "price discovery noise." Formal theoretical analyses of price discovery in a divergent expectations environment are sparse, however. In an earlier paper, Ho, Schwartz and Whitcomb (HSW, 1985), by modeling price determination under transaction costs and transaction price uncertainty, showed that a theoretically desirable equilibrium will generally not be obtained.[6] Handa and Schwartz (1996) analyzed the price paths that can be traced out in an HSW environment. Neither of these papers took account of divergent expectations, however, and, consequently, their insights into price discovery were limited.[7]

Our analysis builds on the divergent expectations, quote setting model developed by Handa, Schwartz and Tiwari (HST, 2003), who relate bid and ask prices to the proportion of bullish vs bearish traders.[8] HST provide a simplified setting where *k* percent of participants assign a relatively high valuation to shares (V_H), and 1-*k* assign a relatively low valuation to shares (V_L). They show how bid and ask prices are set in a continuous limit order book market when V_H, V_L, and *k* are all common knowledge. But, in reality, *k* is not observable until participants come forward and trade. Therefore, we relax the assumption that *k* is known. In our formulation, *k* discovery and price discovery are synonymous. The dynamic price discovery process may be viewed as starting at any daily opening or, more generally, following any news event, either overnight or within the trading day.

Once divergent expectations are allowed for, the way is opened for investors to have adaptive valuations. That is, an agent may change his or her valuation upon knowing the valuation of others. In our stylized model, we take the revealed proportion of buyers and sellers (HST's *k*) to be the conduit through which participants communicate their valuations to each other. Specifically, as more buyers (sellers) arrive at the market, the aggregate mood across all participants becomes more bullish (bearish). In our formulation, the only news that occurs as trading progresses is information about

---

[6] In HSW (1985), two conditions must be satisfied to achieve an equilbirum price: participants must be symmetrically distributed with respect to their desires to hold shares of a risky asset, and their expectations of the market clearing price must be accurate.

[7] Others have followed a substantially different approach that considers price discovery when order flow is distributed over multiple fragmented but integrated markets. See Yan and Zvot (2004) for a recent discussion and further references.

[8] Our formulation draws from HST (2003) and is related to Foucault (1999) who also assumes that investors' share valuations can differ.

the mood of the market as reflected by the proportion of participants who have revealed themselves to be buyers or sellers.

An analysis of the divergent expectations, adaptive valuations environment yields insights that are not readily derived from models based on the assumption of homogeneous expectations. If participants submit their orders to a single price call auction, it is trivial to show that price is V_H if $k > 0.5$, that price is V_L if $k < 0.5$, and that it is indeterminant in the range from V_L to V_H if $k = 0.5$.[9] Consequently, a small change of k in the neighborhood of 0.5 can have a sizable effect on price. In the HST continuous limit order book framework that we use in the current paper, price depends on k throughout the range from V_L to V_H. HST have shown that, as k goes to an extreme of either 0 or 1, the bid and the ask quotes go to V_L or to V_H, respectively and, as they do, that the spread tightens (the spread is at a maximum when k is 0.5). We show, when k is not known and participants have adaptive valuations, that a security's price is particularly sensitive to "small" early events (whether the first arriving participants are bullish or bearish) and, consequently, that price discovery is a multiple equilibria, path dependent process.

Our analysis of divergent expectations contrasts in the following way with the asymmetric information models of, e.g., Copeland and Galai (1983) and Glosten and Milgrom (1985). Under asymmetric information, one group of participants has information that the others do not possess and, when trading with the uninformed, the informed profit from their information. Expectations are homogeneous within each of the two groups (the informed and the uninformed). After the information of the informed becomes common knowledge, expectations are homogenous across all participants.

In contrast, our analysis is structured to be neutral between the group with the relatively high asset valuation (V_H) and the group with the relatively low valuation (V_L). All participants share identical initial estimates of k, say k_0. As trading progresses, the movement of price between V_H and V_L depends on the chance sequence of the order flow (particularly the early order arrivals). Throughout the trading session, it is not possible to identify one group as being better informed than the other. Only the relatively distant future will reveal which group has assessed the information set more accurately.

---

[9] As we show in Section II.A., the solution differs with adaptive valuations.

Our paper is organized as follows. Section I sets forth the behavioristic model. In so doing, we summarize HST (2003), present our behavioral assumptions, our assumptions about individual choice, the probabilistic assumptions, and the analysis of group dynamics. Section II then focuses on price equilibria. Section III presents the simulation analysis we have used to examine the multiple equilibria yielded by our formulation, and the results of our simulation tests. We discuss two extensions of the model in Section IV and, in Section V, consider various implications of the analysis. Section VI is a brief summary.

I. Behavioristic Model

A. The HST Framework

Our analysis builds on the HST analytic framework. Their model uses the following simplifying assumptions.

- Participants arrive in random order at a continuous trading, limit order book market to buy or to sell shares of one risky asset. All orders are of the same size (e.g., one round lot).[10] Each participant, upon arriving, views the limit order book and either enters a limit order or transacts by market order against a counterpart limit order. The quotes establish the market's best bid and offer and thus the bid-ask spread.

- Participants are divided into two categories with respect to share valuation. In each category, all participants are identical with regard to their share valuation. The relatively bullish participants assign a value of V_H, and the relatively bearish participants assign a value of V_L. V_H is the highest (reservation) price at which the bullish participants will buy shares, and V_L is the lowest (reservation) price at which bearish participants will sell shares.[11] The parameter *k* denotes the percentage of the population with the reservation price V_H, and (1-*k*) is the percentage with the reservation price V_L. Thus the parameter *k* represents the

[10] A limit order that has been placed on the book remains alive only until the next order arrives, at which point the limit order either executes or is cancelled. As an extension, HST also model the case where a limit order can remain on the book for two events.

[11] It is implicitly assumed that there are no restrictions or borrowing costs on short selling.

aggregate mood of the market based on the share valuations of the participants. As *k* goes to either of its extreme values of 0 or 1, the expectations of the market become more homogeneous.

- Investors place their orders with respect to their knowledge of V_H, V_L, and *k*.

- The value of *k* is known with certainty by all participants (i.e., *k* is public knowledge).

The HST model allows for information change to occur, and the quotes set in their environment adjust for this possibility. The values of V_H, V_L, and *k* are taken to be known by a participant at the time when he or she submits an order to the market. Price is determinant in the HST formulation even though *k* percent of the participants have a reservation value of V_H and 1-*k* have a reservation value of V_L. Each participant benefits from trading when having traded at a better price than his or her reservation value (V_H for buyers and V_L for sellers). Just how aggressive participants are with regard to price depends on V_H, V_L, and *k*. Equilibrium quotes can be found (with the offer less than V_H and the bid greater than V_L) because participants are willing to risk non-execution for the chance of realizing a transaction at a more favorable price.

B. Our Background Environment

With only a few exceptions, our environment is the same as that of HST. Most importantly, we relax the assumption that *k* is initially known by all participants and that it is a constant. We take *k* to be unknown following any news event or any halt in the continuous market (such as the overnight close). With an unobservable, non-constant *k*, we track individual quote-placement decisions as expectations of *k* change with the progression of events (the sequential arrival of new participants at the market). The advantage of this approach is that, with it, we are able to address the issue of price discovery directly. We also allow participants to have adaptive evaluations (specifically, as an increasing proportion reveal themselves to have V_H valuations, an increasing number who initially valued shares at V_L change their valuation to V_H, and vice versa as an increasing proportion reveal themselves to have V_L valuations). Our analysis shows that price discovery is a path dependent, multiple equilibria process. This insight yields important implications concerning price volatility and technical analysis. Further, our formulation builds a bridge between traditional microstructure and behavioral economics.

We further differ from HST in two relatively inconsequential ways. First, we take V_H and V_L to be constant parameters throughout a trading session and, in so doing, do not allow for an information event to occur. Second, at each event i, we do not consider an individual's decision of whether to place a limit order (to buy or to sell), or to execute by market order against a limit order that has already been posted. Rather, we simply focus on the optimal buy and sell limit orders that would be placed given expectations concerning k at each i^{th} event, and then represent price as the mid-point of the bid-ask spread. In the simulation model described in Section III of this paper, we use the HST formulas to solve for the optimal values of the bid and offer quotes, given participants' expectation of k.

C. Expectations Formation

As noted, the major point of departure for our current analysis from that of HST is that we relax the assumption that the true value of k is initially known. We make the more realistic assumption that everybody forms an expectation of k, and we take this expectation to be the same for all participants. As trading progresses with the arrival of orders, the common expectation of k changes. We let k_0 denote the initial expectation of k and assume that, for all participants, the expectation of k is revised with the succession of events indexed by i according to the equation,

$$k_i = k_0(1-\lambda_i) + \hat{k}_i \lambda_i \tag{1}$$

where

k_i is the current expected value of k at i,

i is an index on the order of events,

k_0 is the initial expected value of k (i.e., before any events have occurred),

$\hat{k}_i$ is the observed value of k at i, and

$(1-\lambda)$ is a weight that denotes the relative importance in the expectation formation of the initial expectation, k_0, $0 \le \lambda_i \le 1$ and $\lambda_{i-1} < \lambda_i$

The variable k_i is the expectation of k at event i based on initial expectations (k_0) and the actual proportion of participants over the previous i-1 events who, by their orders, have revealed themselves to have the valuation V_H. In our model, the factor λ_i can be interpreted as the speed of adaptation to current information regarding the mood of the market. The model's dynamic has another facet, to wit, the value of λ_i is increasing with the progression of trades, which means that a decreasing weight is given to the initial expectation, k_0, and an increasing weight is assigned to the value of k that is observed based on the orders that have actually been placed.

Two terms on the right hand side of equation (1) are i dependent: $\hat{k}_i$ and λ_i. With regard to $\hat{k}_i$, we take the sequence in which participants arrive at the market to be exogenously determined and, in equation (1), $\hat{k}_i$ reflects the proportional number of buyers and sellers who have actually arrived up to event i. Note that, in this formulation, because the sequence is exogenously determined, the sequence in which buyers and sellers arrive has no information content. The term λ_i increasing in i reflects the fact that, with the progression of events, a participant gives an increasing weight to the *observed proportion* of buyers, and a decreasing weight to his or her *initial expectation*, k_0.

Equation (1) may be further understood by writing it as

$$k_i - k_0 = \lambda_i(\hat{k}_i - k_0) \qquad ; \lambda_i \text{ is increasing in i,} \qquad (2)$$

and contrasting it with the adaptive expectations model that is standard in economics,

$$k_i - k_{i-1} = \lambda(\hat{k}_i - k_{i-1}) \qquad ; \lambda \text{ is a constant} \qquad (3)$$

In both models, in forming expectations, decreasing weight is given to k_0 as k_i evolves. In the standard economics formulation, weight is also implicitly given to all previous observations of $\hat{k}_i$; as discussed above, this is not required in our formulation because we take the sequence of order arrival to be exogenous.[12] Further, in the standard

[12] The previous points can be explicitly shown as follows. Equation (3) can be written as

$$k_i = (1-\lambda)\, k_{i-1} + \lambda\, \hat{k}_i \qquad (3a)$$

Since equation (3a) holds for every i, one can derive by using repeated substitution the equivalent expression, economics formulation, the speed of adaptation, λ, is a constant, whereas in our formulation, the speed of adaptation, λ_i, is increasing in i. That is, we assume that, the larger the number of events participants observe, the more importance they place on their observations relative to their initial anticipations.

D. Individual Behavior within the Group

We have thus far presented the assumptions that set forth our basic analytic framework. Within that framework, individuals' expectations on *k* depend on the aggregate mood of the market as summarized by the observed variable, $\hat{k}_i$. As $\hat{k}_i$ evolves, not only do expectations change but, with adaptive valuations, the underlying population distribution of individuals between V_H and V_L also changes. We next discuss the individual's decision of whether to be a V_H participant or a V_L participant when he or she comes to the market.

The individual's specific choice of the security's value, V_H or V_L, is a function of three arguments: the information set (Ω) concerning the state of the world that is common for all participants; the assessment of other participants as reflected by the variable, $\hat{k}_i$; and initial expectation, k_0. We allow the specific form of the functional relationship to be unique to an individual. More formally,

$$V_{ji} = f_j(\Omega, \hat{k}_i, k_0) \tag{4}$$

where j is an index for an individual; i is, as defined above, the numeric count of events; and, for every individual and event, V_{ji} has the value, V_H or V_L. Individuals come to the market as either buyers (if they assign the value V_H), or as sellers if they assign the value

---

$$k_i = (1-\lambda)^i\, k_0 + \lambda \sum_{t=0}^{i-1} (1-\lambda)^t\, \hat{k}_{i-t} \tag{3b}$$

From (3b), it is clear that, first, k_i is not only a function of k_0 and $\hat{k}_i$, but also of all the intermediate values, $\hat{k}_1, \hat{k}_2, \ldots, \hat{k}_{i-1}$. Second, since $\lambda \sum_{t=0}^{i-1} (1-\lambda)^t = 1 - (1-\lambda)^i$

if all the intermediate values are identical to $\hat{k}_i$, then (3b) can be written as

$$k_i = (1-\lambda)^i\, k_0 + [1 - (1-\lambda)^i]\, \hat{k}_i \tag{3c}$$

Thus it is obvious that the weight of k_0 is decreasing and that that of $\hat{k}_i$ is increasing, as in our adaptive expectations equation (1).

V_L). We assume that there exists a reference point (a critical value) that determines whether an individual is a buyer or a seller. We specify the function f_j in equation (4) in the following way:

$$V_{ji} = f_j(\Omega, \hat{k}_i, k_0) = \lfloor V_H \delta_{ji} + V_L(1-\delta_{ji}) \rfloor \qquad (5)$$

where the δ_{ji} are indicator functions that take the values of zero or one such that

$$\delta_{ji} = \begin{cases} 1 \text{ if } k_i \geq k_j^*(\Omega) \\ 0 \text{ if } k_i < k_j^*(\Omega) \end{cases}$$

and k_i is determined by equation (1).

Note that the critical value, k_j* is unique to individual j even with Ω being common knowledge. Whether participant j is a buyer or a seller is determined by equation (5). In other words, j would be a buyer (have a V_H valuation) at event i if, based upon observing $\hat{k}_i$, his or her expectations k_i exceed his or her threshold value, k_j^*, and would be a seller otherwise. The probability of a buyer coming to the market at event i is determined by the proportion of participants for whom, at event i, $k_i \geq k_j^*$

E. Group Dynamics

With regard to the individual valuations, when approaching the market, each individual, j, does so as either a buyer or a seller, depending upon the relationship between k_i and k_j* as described by equation (5) that applies at the point of the agent's arrival. We assume that the arrival sequence of agents is exogenous, that they share the same information set (Ω), that they have the same initial value of k which is k_0, observe the same $\hat{k}_i$, and use the same decision rule, to wit the adaptive valuation equation (1). Thus, agents differ only with respect to their critical point, k_j*. We assume that k_j* is uniformly distributed between 0 and 1, and treat the arrival process as equivalent to a random sampling of k_j* with replacement. As will be seen below, the uniformity assumption is of critical importance for our analysis.

The term k_j* reflects the confidence of the j^{th} participant in his or her own assessment of the fundamental information set (Ω). For k_j* = 1.0, the j^{th} participant is a bear regardless of the proportion of others who assess shares at V_L. At the other end of the spectrum, for $k_j* = 0.0$, the j[th] participant is a bull regardless of the proportion of others who assess shares at V_H. Halfway between, for $k_j* = 0.5$, the j[th] participant has no independent assessment at all but simply goes with the majority.

In our structure, the population proportion of buyers (sellers) can be interpreted as the probability that the next arriving participant will be a buyer (seller). With adaptive valuations, this probability morphs with the progression of events. At each i[th] event, the probability that the next participant will be a buyer (seller) depends, as discussed in Sub-section D above, on the relationship between k_i and the distribution of the k_j*. Concurrently, at each i[th] event, participant expectations concerning arrival probabilities are characterized by k_i. A special case of particular interest exists when the actual (population) probability of the next arriving participant being a buyer equals the expected probability, k_i. This equality is achieved if and only if the k_j* are uniformly distributed between 0 and 1, as we have assumed them to be.[13] This case is of particular interest because of the insights it yields into the price formation process, as will be discussed below.

Note that trending exists in our formulation for all k_i not equal to 0.5. For instance, for k_i equal to 0.6, the probability that the next arriving participant will be a buyer is 0.6 and, if a buyer does arrive, price will rise, k_{i+1} will be greater than 0.6, and so on, implying a trend. For our case, the trend should not lead to profitable trading opportunities because, being apparent to all participants, the quotes will be sufficiently wide to nullify the profitability of any trading strategy designed to exploit the trend.

As noted, the population probability of future events is in harmony with the probabilities displayed by past events if the k_j* are uniformly distributed over the domain 0 and 1, as we have assumed them to be. With harmony between actual and expected arrival rates, our price formation process is consistent with a standard Polya process.[14]

---

[13] If the expected probability rises from some k_i, to some $k_i + \Delta$, the actual probability can maintain equality with the expected probability if and only if an identical percentage of the k_j* are in the range k_i, to $k_i + \Delta$ for any value of k_i and Δ. For this condition to be satisfied, the k_j* must be uniformly distributed.

[14] Arthur (1994, pp. 36) describes the Polya process as follows. "Think of an urn of infinite capacity to which are added balls of two possible colors – red and white, say. Starting with one red and one white ball in the urn, add a ball each time indefinitely, according to the rule: Choose a ball in the urn at random and replace it; if it is red, add a red; if it is white, add a white. Obviously this process has increments that are path-dependent – at any time the probability that the next ball added is red exactly equals the proportion red... Polya proved that... in a scheme like this the proportion of red balls does tend to a limit X, and with probability 1. But X is a random variable uniformly distributed between 0 and 1."

The key feature of the Polya process is that it leads to multiple, path dependent outcomes. In our formulation, the outcome of ultimate interest is price. We turn to the issue of price equilibria in the next section.

II. Price Equilibria

We now use the behavioristic model set forth in the previous section to analyze price determination in a market comprised of a sufficiently large number of participants who are each seeking to buy or to sell one share of a stock. We first consider a simple environment where all of the participants arrive at the market simultaneously. We next turn to a continuous trading environment where participants do not have a memory (do not have prior beliefs about k_0), and then to the more behaviorially realistic continuous trading environment where participants have a memory (give some weight to their prior beliefs about k_0).

A. A Simple Solution

A simple, simultaneous solution would apply if all participants were to arrive at the market at the same time and trade in a single price call auction. This special case is of interest because, unlike the continuous market solution described below, it leads to a unique market clearing value.

In this market, any j^{th} participant would wish to buy one share at any price greater than $V_L + k_j^*(V_H - V_L)$, and would wish to sell one share otherwise. As the price rises in the range from V_L to V_H, it crosses an increasing number of critical price values from below and, in so doing, leads to an increasing number of participants being buyers rather than sellers (and vice versa as the price falls). When the participants all meet simultaneously in a single price call auction, all that each need reveal to the market is his or her critical value of price.[15] For the market to clear, we must have an equal number of buyers and sellers; accordingly, the clearing price equals the median value of the price distribution, $V_L + \bar{k}^* (V_H - V_L)$, where $\bar{k}^*$ is the median of k_j^*. For k_j^* uniformly distributed between 0 and 1 (as we have assumed them to be), the expected value of k_j^* is

[15] Specifically, the participant would submit a stop loss limit order to buy at any price equal to or greater than the price implied by k_j^* up to a value of V_H, or to sell at any price less than the price implied by k_j^* down to a value of V_L.

0.5, so the expected equilibrium value would retain the midpoint value, $(V_H + V_L)/2$. However, because of the inverted buy and sell responses to price, it is an unstable equilibrium.[16] The instability of the system leads us to consider next the process of price formation in a continuous trading environment where, in a sense, the counterpart to instability is multiple equilibria.

B. The Continuous Market Without Memory

The simple solution shown above is a product of the simultaneity of participant interactions and of the way the call auction market poses a question to a participant, “at what prices would you be a buyer, and at what prices would you be a seller?” Both the question and the solution change when we move to a continuous trading environment. In the continuous environment, the market presents bid and offer quotes to a participant and asks the question, “at these prices, are you a buyer or a seller?” Given the value of $\hat{k}_i$ at any i^{th} event in the continuous market, the bid (ask) of a buyer (seller) can be determined following HST (2003) and, in contrast with the call auction environment, standard limit orders (not stop loss orders) are posted at these values.[17]

We first analyze equilibria in the continuous market without a memory. To do so, we consider a special case where $\lambda_i = 1$ for all i, and where a bid and an ask quote have been posted at the start of the trading session with the mid-point of the spread halfway between V_H and V_L. With this configuration, our formulation conforms exactly to a standard Polya process. We demonstrate this in the following points:

- The standard Polya process starts with one red ball and one white ball in an urn. The red ball represents the buyer who has placed the bid, and the white ball represents the seller who has placed the offer.

[16] Ho, Schwartz, and Whitcomb (1985) also showed that, under certain stylized conditions, a call auction environment can lead to accentuated price volatility. Nevertheless, when investors have multiple valuations rather than two, and when adaptive valuations are not conditioned solely on price, a call auction can actually help to control volatility because it is a price discovery procedure.

[17] In the continuous environment, the seller knows exactly how low an offer must be for a buyer to take it if a buyer arrives, a buyer knows exactly how high a bid must be for it to be hit by a seller if a seller arrives, and neither the buyer nor the seller will give a contra the opportunity to trade at a better price. In contrast, market clearing values are not known ex ante in a call auction and, as noted in footnote 14, participants use stop loss limit orders in our call auction environment.

- For the standard Polya process, the initial probability of picking a red ball is 0.5. Likewise, from equation (1) and given that the k^*_j are uniformly distributed, the probability that the next arriving participant in our process will be a buyer is 0.5.

- For the standard Polya process, the probability of a red ball being picked at any i^{th} draw equals the proportion of red balls in the urn as of the i^{th} event. In our model, the probability of the arriving participant being a buyer at any i^{th} event equals the proportion of previous participants who have revealed themselves to be buyers through the preceding events.

- For the standard Polya process, a single ball that has the same color as the ball just drawn is added to the urn at each event as a device to change the probabilities for the next draw. In our model, the valuation of one more participant is revealed at each event, and this addition changes $\hat{k}_i$ in an identical fashion.

- In the standard Polya process, balls are drawn from and added to an urn of infinite capacity that increases in size by one ball at each event; in the process, the proportion of red balls changes and the probability of drawing a red ball at the next event changes identically. In our process, new participants are drawn from an underlying (infinite) population and the number of revealed participants increases by one at each event; in the process, invoking the assumption that the k^*_j are uniformly distributed, the probability of the next arriving participant being a buyer exactly equals the proportion of buyers observed in the preceding i-1 events.

- The Polya process requires that the draw of a red ball be matched with the addition of a red ball to the urn. Our process operates equivalently in that the arrival of a buyer in the market is matched with a proportionate transformation of participants in the underlying population from being sellers to being buyers.

Recognizing that the special case of our formulation is identical to a standard Polya process, we can invoke two important results from Polya. First, each "play of the game" results in its own path, and each path converges on its own equilibrium outcome as the number of events becomes sufficiently large. That is, Polya is a path dependent process that displays a multiplicity of possible asymptotic outcomes. Second, Polya (1931) has shown that, for the standard process described above, each path converges on a value, let us call it X, and that X is a random variable uniformly distributed in the range 0 to 1. In our formulation, the values are the probabilities of the next participant being a buyer. Using HST, these probabilities (*k*) are readily translated into shares values. The Polya result is that *k* converges on a value that is uniformly distributed over the range 0 to 1, while price, which following HST (2003) is a non-linear transformation of *k*, converges on a value that has a U-shaped distribution over the range V_L to V_H. A numerical illustration that shows the non-linearity of the price, *k* relationship is presented in an appendix to this paper.

C. The Continuous Market With Memory

In our more general model, λ_i is not one for all i. The more general case is of interest because it reflects the reality of human beings (unlike red and white balls) having memories of prior events, that beliefs based on past events affect current expectations, and that learning occurs with the progression of events. As learning progresses, λ_i increases and approaches one asymptotically.

A simple, intuitive explanation will show the effect of memory (i.e., λ<1) on the values that prices converge on. The uniform distribution of prices that characterizes the simple Polya process (the continuous market without memory) is attributable to the set of probabilities attached to the nodes at event (i-1) and to the set of transitional probabilities from the nodes at event (i-1) to the nodes at the i^{th} event.[18] With a memory, initial expectations affect the transitional probabilities and, by extension, the probabilities

[18] The uniform distribution can be shown intuitively as follows. There is one buyer and one seller at the start of the trading session. As of the i^{th} event, let B_i be the total number of participants who will have revealed themselves to be buyers, and S_i be the total number who have revealed themselves to be sellers. Letting p be probability, we then have, for i=1,

$p(B_1=2, S_1=1) = p(B_1=1, S_1=2) = \frac{1}{2}$

and, for i=2,

$p(B_2=3, S_2=1) = p(B_1=2)\, p(B_2=3 \mid B_1=2) = (\frac{1}{2})(\frac{2}{3}) = \frac{1}{3}$

$p(B_2=1, S_2=3) = p(S_1=2)\, p(S_2=3 \mid S_1=2) = (\frac{1}{2})(\frac{2}{3}) = \frac{1}{3}$

$p(B_2=2, S_2=2) = p(B_1=2)\, p(B_2=2 \mid B_1=2) + p(S_2=2, S_2=2) = p(S_1=2,)\, p(S_2=2 \mid S_1=2) = 2(\frac{1}{2})(\frac{1}{3}) = \frac{1}{3}$

Generalizing, the distribution at each i^{th} event is uniform in that we have i+1 nodes at each i^{th} event and the probability for each is 1/(i+1).

attached to the nodes at event (i-1). Whatever the value of k_0, the paths with $\lambda < 1$ tend to revert back to k_0 whenever they diverge from it. This reversion tendency (which is absent from pure Polya) causes the distribution of k_i to be unimodal in the neighborhood of k_0. The location of the distribution depends on k_0, and the tightness of the distribution around its mode depends on the rate at which λ_i is increasing in i (the slower the rate, the more protracted is the influence of k_0 and hence the tighter is the distribution around k_0).

III. Simulation Analysis

We next describe the simulations we have used to assess further our model of price discovery. Of particular interest is the sensitivity of the values of *k* and of price at convergence to two key parameters: participants' initial expectations of *k* (k_0), and the duration of memory (I, a controllable parameter that will be explained below). Also of interest are the number of events that each simulation run must extend to in order to achieve convergence, and the number of replications that are necessary to obtain the end of run distributions for *k* and for price that our formulation predicts. An assessment of each of these factors has implications for short period (e.g., intra-day) price volatility.

A. Simulation Structure

The simulation structure is based on the behavioristic model presented in the introduction. There are two types of investors in the simulation: 'bears' who evaluate shares at $V_L = \$45$, and 'bulls' who evaluate shares at $V_H = \$55$. These two valuations set the minimum and maximum values that quotes and transaction prices can take in the simulation. The two investor types trade with each other. The $V_L = \$45$ participants are sellers (they establish the offer prices), and the $V_H = \$55$ participants are buyers (they establish the bid prices). The bids and offers are obtained using a simplified version of the HST (2003) pricing model that excludes the possibility of information change. The requisite HST variables are the two share valuations (\$45 and \$55) and *k,* the proportion of participants who are buyers.[19] The simplified HST equations are

$$B^* = \lambda V_L + (1-\lambda) V_H \qquad (6)$$

$$A^* = \mu V_H + (1-\mu) V_L \qquad (7)$$

where

[19] For simplicity, our discussion refers to the buyer side of the market; the seller side is analogous.

$$\lambda = \frac{1-k}{1-k+k^2}$$

$$\mu = \frac{k}{1-k+k^2}$$

Each simulation run starts with a bid, an offer, and a participant count of two (one buyer and one seller). The opening bid and offer are set using the simplified HST model, the two share valuations, and an initial value of k (k_o) which is a controllable parameter. The simulation progresses to the first event, then to the second, and so on, with the sequential arrival of new participants, each of whom is either a buyer or a seller. Two things happen at each event: (1) a buyer or a seller arrives and the observed proportion of buyers is accordingly updated, and (2) based on this revised proportion, new quotes are established for event i+1.

At each event i, k_i is the expected probability that a buyer will arrive at event i +1. Following equation (1), k_i is a weighted combination of k_o and the proportion of buyers who have arrived during the previous i-1 events.

In the formulation developed in Section I.E., we note that a special case of particular interest exists when the actual probability that the next participant will be a buyer equals the expected probability. We treat this special case in the simulation by letting k_i be both the *actual probability* (the probability that we use to determine whether the next arrival is a buyer), and participants' *common expectation* that the next arrival will be a buyer (the probability that we use, along with the evaluations of \$45 and \$55, to set the bid and ask quotes at the i^{th} event).[20] Price at the i^{th} event is represented by the mid-point of the quotes at the i^{th} event.

The term k_i is given by equation (1). In equation (1), λ_i establishes the relative importance, at each i^{th} event, of participants' initial expectations of the proportion of buyers (k_o) and the proportion of buyers actually observed ($\hat{k}_i$) through the i-1 events. In the simulation, a modified Ogive is used to describe how λ_i increases with i.[21] The modification is that λ_i is assigned an initial value of zero and a maximum value of one

[20] Implicit in this treatment is that participants have adaptive valuations, as discussed in the introduction.

[21] The Ogive (the cumulative of a normal distribution) has been widely used to describe the dissemination of, e.g., news and communicable diseases.

that it attains at some event I. I and k_o are the two controllable parameters in the simulation. A lower value of I reflects participants having a "shorter memory" of events preceding the current trading session (i.e., the events upon which their initial expectation, k_o, is formed). A case of particular interest is when $\lambda_i = 1$ at I=1. In this case, prior memory effects only the location of the initial bid and offer prices and, with $k_o = 0.5$, the simulation follows a pure Polya process. As discussed in Section I, for this case, the distribution of end of run values of k is expected to be uniform, and the distribution of end of run values of price is expected to be U-shaped.

B. Simulation Results

Figure 1 shows the price paths for ten simulation runs, each of which comprises 250 events and has parameter settings of $k_0 = 0.5$ and I = 50. Because $k_0 = 0.5$, each of the ten price paths starts at \$50. The paths generally fan out to higher and lower prices with the succession of events, and the "fanning out" extends through roughly the first 80 events. After this point, each newly arriving participant has little impact on $\hat{k}_i$ and hence on price but, nevertheless, minor fluctuations in price persist. By the 250th event, each of the paths appears to have converged on an acceptably stable value, P_{250}. To be cautious, we run our simulations for 800 events and analyze the distributions of k_{800} and P_{800}.

Price discovery, as we have modeled it, indeed appears to be a noisy process. Supplementary tests not shown here gave essentially confusing results with 200 replications, noisy pictures with 5,000 replications, and acceptably clear pictures with 8,000 replications. Accordingly, we ran 8,000 replications of each simulation to obtain our means and standard deviations of k_{800} and P_{800}.

The distributions of k_{800} and P_{800} are shown in Figures 2 – 4 for three different combinations of k_0 and I. Each figure has two frames, one for k_{800} and the other for P_{800}. The range for k_{800} is 0 to 1, and for P_{800} is V_L to V_H. Each range is broken into 21 equal subdivisions (buckets), with 1 indicating the smallest and 21 the largest. The 21 buckets are displayed on the horizontal axis. The vertical axis shows the frequency (the percentage of the replications) with which the values of k_{800} or of P_{800} fell in each of the 21 buckets.

Figure 2 presents the results for the simulations run with $k_0 = 0.5$ and I = 1 (no memory). This is the pure Polya process. As expected, the distribution of k_{800} conforms (with some variability) to the uniform, and the distribution of P_{800} is U-shaped. Figure 3 shows the results for $k_0 = 0.5$ and I = 50 (memory). As expected, with memory, the distribution of k_{800} is unimodal and symmetrical around its modal value at the middle of the range, bucket 11. The price distribution shown in Figure 3B is also centered on bucket 11, with the U-shaped pattern seen in Figure 2B shrunk to the three middle buckets, 10, 11, and 12.[22] Interestingly, the distribution of k_{800} is tighter around its mode than the distribution of P_{800}.[23] Figure 4 shows the results for $k_0 = 0.6$ and I = 50 (memory). As expected, the distributions of both k_{800} and P_{800} are now unimodal, the modal values are to the right of the middle of the range (bucket 11), and the distributions are skewed to the left. As in Figure 3, the tails of the P_{800} distribution are somewhat fatter than the tails of the k_{800} distribution (i.e., while the k_{800} distribution extends from bucket 5 – 20, the P_{800} distribution extends from bucket 2 – 21), which is consistent with the transformation from k_{800} to P_{800} being non-linear.[24]

Table 1 shows the means and normalized standard deviations for k_{800} and P_{800} for a larger set of simulation runs that includes three values of memory (I = 1, 25, and 50) and five values of k_0 ranging from 0.2 to 0.8. The standard deviations of k_{800} and P_{800} are normalized by dividing each by each variable's range (1 for k_{800} and 10 for P_{800}). The results for I = 1 (no memory) show that changes in the initial value of k_0 have virtually no effect on either the mean or standard deviation of either k_{800} or P_{800}, although the means for both do increase slightly as k_0 rises from 0.2 to 0.8. The reason for these minimal effects is that, because each run always starts with one buyer and one seller, by the next event we must have either two buyers and one seller, or one buyer and two sellers. Hence, one of these two values takes over for all values of k_0 and, with I = 1, k_0 itself has no further influence. On the low end, it makes virtually no difference whether we start with k_0 equal 0.2 or 0.4; on the high end, it makes virtually no difference whether we start with k_0 equal to 0.6 or 0.8.

---

[22] For larger values of I, the price distribution becomes unimodal and centered on bucket 11.

[23] In supplemental tests that are not shown here which were run for values of I substantially less than 50 (e.g., I = 30) with $k_0 = 0.5$, the price distribution remained centered on bucket 6 but remained U-shape in the immediate neighborhood of bucket 6 (i.e., the distribution had local maxima at buckets 5 and 7).

[24] Tests run with $k_0 = 0.4$ and I = 50 (not reported here) yielded results that were mirror image to those shown in Table 4.

A more interesting pattern emerges when memory plays a protracted role. For I equals 25 and 50, the means of both k_{800} and P_{800} increase sharply with increases in k_0 (the increases are greater for I = 50). For instance, for $k_0 = 0.8$, the mean of k_{800} is .507, 0.973 and 0.983 for I equal to 1, 25 and 50, respectively, and the mean of P_{800} is 50.050, 53.050 and 53.097 for I equal to 1, 25 and 50, respectively. The standard deviations of k_{800} and of P_{800} are consistently lower for I = 25 and I = 50 than for I = 1.

An interesting contrast exists between the normalized standard deviations of k_{800} and P_{800}. For I = 1, all of the normalized standard deviations are roughly 10% higher for P_{800} than for k_{800}. For I = 25 and I = 50, the normalized standard deviations are roughly 25% higher for P_{800} than for k_{800} when $k_0 = 0.5$ but, for the other four values of k_0, the normalized standard deviation of $k_{800,}$ is the higher of the two. This translates into the distribution of the standard deviation of k_{800} with respect to k_0 being U-shaped, while the distribution of the standard deviation of P_{800} with respect to k_0 is an inverted U.[25]

The co-existence of the U-shaped pattern for k_{800} and the inverted U-shaped pattern for P_{800} can be attributed to the following. The distributions of k_{800} and P_{800} are both skewed to the left (right) when k_0 is greater than (less than) 0.5. Skewness itself increases the standard deviations of both variables, and accounts for the distribution of the standard deviation of k_{800} with respect to k_0 being U-shaped. The distribution of P_{800}, however, also reflects a second reality: price is a non-linear transformation of k. It can easily be shown that any deviation (in either direction) of k_{800} from the middle of its range is associated with a value of P_{800} that deviates further from the middle of its range. For any value of memory, the distribution of P_{800} clusters more closely around its mean, the further k_0 deviates from 0.5; for any value of k_0, the distribution of P_{800} clusters more closely around its mean the greater the value of I. Now, if the simulation starts in the middle of its range ($k_0 = 0.5$ and $P_0 = \$50$) and there is no memory (I = 1), P_{800} can go to either of its bounds with equal probability, and it clusters closer to both of its bounds than does k_{800}. That is why, as seen in Figure 2, the distribution of P_{800} is bi-modal with the modes at the extremes (Figure 2B), while the distribution of k_{800} is essentially uniform (Figure 2A). If memory is protracted, the bi-modality of P_{800} disappears as k_0 deviates in either direction from 0.5, and the distribution of P_{800} accordingly becomes less dispersed.

[25] The U-shaped and inverted U-Shaped distributions apply over the range of values displayed in Table 1 (i.e., k_0 equals 0.2 to 0.8). Results not shown here indicate that the standard deviation of k_{800} also starts to decrease as k_0 approaches either of its boundaries (0 or 1). Clearly, for k_0 in the close neighborhood of either 0 or 1, the bounds act as absorbing barriers and the variability of both k_{800} and P_{800} shrink.

Apparently, as k_0 deviates in either direction from 0.5, the diminution of the standard deviation accentuation that is attributable to price being a non-linear transformation of k is stronger than the increase in the standard deviation that is attributable to skewness. Accordingly, the distribution of the standard deviation of P_{800} with respect to k_0 is an inverted U.

These simulation results underscore three characteristics of price discovery volatility. First, the process itself is noisy, which explains why we used 800 events and 8,000 replications to establish the patterns we are looking for. Second, as I, the duration of memory, decreases, the distribution of k_{800} flattens out around its modal values, and the distribution of P_{800} becomes U-shaped, which means that price discovery gets more volatile, as is seen in Table 1. For instance, for $k_0 = 0.5$, the normalized standard deviation for P_{800} is .171, .177, and 3.23, for I equal to 50, 25, and 1, respectively. Third, for a meaningful value of I (e.g., 25 and 50), the standard deviation of P_{800} shrinks as k_0 is placed closer to either of its extremes, 0 or 1, as is seen in Table 1. For instance, for I = 25, the normalized standard deviation of P_{800} is .177 for k_0 equal to 0.50; it decreases to .170 and .126 for k_0 equal to 0.40 and 0.20, respectively, and it decreases to .170 and .128 for k_0 equal to 0.6 and 0.8, respectively . The second and third observations have an interesting implication for the relationship between the divergence of expectations and price discovery volatility. To see this, note that, V_H and V_L given, expectations become more divergent as I goes to zero (as memory plays a diminished role) and/or as k_0 goes to 0.5. It follows that a greater divergence in expectations leads to greater price instability in brief intervals of time.

iV. Model Extensions

We have analyzed the dynamic behavior of price in a highly stylized divergent expectations model that places all participants in one of only two groups: those with a high valuation, V_H, and those with a low valuation, V_L. Further, the variable k (the proportion of participants with the valuation V_H) is the sole mechanism through which participants revise their valuations in response to the information they receive about the valuations of others (as observed k increases, some participants switch their valuations from V_L to V_H and, as observed k decreases, some participants switch their valuations from V_H to V_L). In this section, we consider the implications of relaxing these conditions.

A. Multiple Valuations

The assumption that all participants fit into two groups with respect to their valuations of a security is sufficient to show the existence of the dynamic price behavior that can occur in a divergent expectations environment, and dealing with two groups only keeps the analysis simple. However, the "two groups" assumption is limiting: it restricts price to the range, V_L to V_H. To extend beyond this range, multiple valuations must be allowed. Just as an increase in observed *k* can result in some bearish participants shifting their valuations from V_L to V_H, we should also allow for the possibility of some bullish participants increasing their valuations to a new $V'_H > V_H$. That is, some V_H participants, upon observing a preponderance of other V_H participants, may themselves become more bullish and, in so doing, move to the yet higher valuation, V'_H. Similarly, some V_L participants, upon observing a preponderance of other V_L participants, may become more bearish and move to the yet lower valuation, V'_L.

The multiple valuations extension has the pleasing feature of symmetry: bulls and bears are both free to increase their valuations upon observing the presence of more buyers in the market, or to decrease their valuations upon observing more sellers. Of greater importance, multiple valuations suggest the possibility of wider price swings. Further, it opens the possibility of explosive price movements occurring if the response of the bulls to a higher *k* (which must now be interpreted as a vector) is greater than the response of the bears, or if the response of the bears to a lower observed *k* is greater than the response of the bulls.

B. Alternative Valuation Signals

With regard to information transmission, we have taken the valuation signal at any point in a trading session to be *k,* the observed proportion of buyers and sellers who have arrived at the market up to that point. Because all orders are assumed to be of identical size, *k* is also the proportion of volume that is buy-triggered. Participants, however, may respond to other reflections of a market's collective assessment. They may, for instance, attribute importance to the length of time that price has stayed above (or below) a certain level, to the fact that a stock's price has "broken out" and set a new high (or low), to the movement of other share prices or indexes, or to a myriad of other factors. To the extent that participants focus on a valuation signal of longer duration (e.g., greater than a day), price discovery might play out over considerably longer periods than our focus on *k,* an inherently intra-day variable, might suggest.

V. Implications

We have shown that an information set, Ω, which is public knowledge, translates into a unique market clearing value if all participants submit their orders to buy or to sell one share of a stock to a call auction that batches all orders together for simultaneous execution at a single price. Our call auction clearing price, however, is an unstable equilibrium. Consequently, it is not surprising that, in a continuous trading environment, price discovery is a path dependent, multiple equilibria process. This insight yields several implications. The first that we consider is the relationship between dynamic price discovery and trading volume.[26]

Trading occurs in our formulation because participants have divergent expectations and thus differing trading motives. In this context, volume is maximized in the call auction environment because, in this environment, the clearing price is given by $V_L + \bar{k}^* (V_H - V_L)$, where $\bar{k}^*$ is the median of k_j^*. At the median price value, half of the participants are buyers, half are sellers, and all of them trade. This is not the case in the continuous market where some participants place limit orders that expire unexecuted. In the context of our model, inaccuracies in price discovery will result in lower execution rates and thus lower trading volume. Further analysis of the relationship between these two variables would be desirable, but is outside the scope of the current paper.

It is well documented that short-period (e.g., intra-day) price volatility is accentuated vis-à-vis longer-period (e.g., one-week) volatility.[27] The accentuation can be attributed, in part, to standard microstructure factors such as bid-ask spreads and market impact effects. We suggest that the accentuation may in large part be a reflection of dynamic price discovery.

In our path dependent, multiple equilibria environment, a stable equilibrium price is converged on only as a succession of events becomes sufficiently large, and the

---

[26] It has been known in the literature that trading volume is linked to volatility. See Karpoff (1986) for an insightful analysis. Our formulation suggests that the two variables may be linked in the context of the price discovery process as it evolves dynamically in a divergent expectations environment.

[27] See e.g., Ozenbas, Schwartz and Wood (2002) for recent evidence and further references.

specific value that is converged on is path dependent. The implication for volatility follows, not just from the volatility inherent in any particular path, but also from the price variation across the different paths at any i^{th} event.[28] The link between this cross-path volatility and inter-temporal volatility is established by letting price discovery be a repetitive process. With repetition of the process and in the absence of news, price will fluctuate simply because a different path is followed in each repetition. This differs from Handa, Schwartz and Tiwari (2003) where k is a known constant and the bid and ask quotes are themselves constant in the absence of news. In the absence of news, the only source of volatility in HST is the bounce between the quotes.

While our model structure suggests that price discovery accentuates predominantly short-period (e.g., intra-day) volatility, our model extensions raise the possibility that the impacts are more protracted. That is, the joint effect of multiple divergent expectations and multiple indicators of market sentiment (as discussed in the previous section) could result in a process where price keeps rising (falling) for an extended period of time to ever higher (lower) levels. If so, the price discovery process that we examine in this paper may be linked to longer run swings that have been characterized as bubbles and crashes.

Our study of the price discovery process sheds a different light on technical analysis. The standard view is that technical analysis simply uses historic price changes to forecast future price changes. A somewhat different interpretation is that technical analysis may be used to assess whether or not a current price level is sustainable. When faced with high intra-day volatility, a participant may want to trade at a price that has been validated in the sense that a substantial number of other orders have been filled at that price.[29] The widespread institutional investor practice of VWAP (volume weighted average price) trading may be similarly justified. In this context, technical analysis may be viewed as one approach to assessing whether or not a transaction can be validated as a

---

[28] Note, however, that the volatility inherent in any particular path is muted in our formulation because we assume that: participants are divided into only two groups, the proportion $\hat{k}_i$ is the sole indicator of the market's "mood," k^*_j is uniformly distributed across participants, and there is no external information change. A relaxation of any of these assumptions could imply higher volatility along any given path.

[29] Consider a portfolio manager who would be willing to purchase shares at a price up to \$50 but who is reluctant buy at \$48 because he or she expects that, in the current trading session, price will likely dip to \$47 in the absence of news. Buying at \$48 and then seeing price quickly go to \$47 in the absence of news not only means that an opportunity to realize a higher expected return has been missed, but that the buyside trader who placed the order will regret his decision. See Paroush and Venezia (1979) for discussion and further references.

reasonable purchase or sale. For instance, persistent trading at a price (or within a trading range) tends to validate that price (or trading range). On the other hand, penetration of a support (or a resistance) level may be a signal that participants are indeed more bullish (bearish) and, therefore, that price is not apt to revert to a previous trading range.[30]

The divergent expectations, adaptive valuations model suggests that the price level converged on following a sufficiently lengthy sequence of events is sensitive to early arriving orders, but that it is insensitive to late arriving orders. This suggests that a participant who is in a position to assess the relative sensitivity of a market to a current buy/sell imbalance might, as a consequence of path dependency and multiple equilibria, be able to profitably game the market. For instance, a participant might buy (or sell) early in the price discovery process (with relatively large market impact), and then unwind the position toward the end of the process (with relatively little market impact). Such an agent will closely resemble a momentum player. The difference between a momentum player and a manipulator is that the former simply attempts to exploit a trend, while the latter seeks to create one. Neither momentum trading nor manipulative behavior could be profitable, however, in the strict form of our model because the bid-ask spread could easily be widened so as to prevent it.

Our dynamic price discovery process conforms to a standard Polya process, and a special case of our model fits the standard Polyia process exactly. The Polyia process itself is depicted as a sequence of draws, from an urn, of red and white balls, with the proportion of red and white balls (and thus the probability of a red ball or a white ball being picked) changing as the sequence of events progresses. In the price discovery application, the red and white balls represent human participants (buyers and sellers); consequently, the price discovery process reflects behavioral realities that the (mechanical) Polya process does not. In brief, the interfaces with behavioral issues include: the importance of prior expectations (the tenacity with which participants give continuing weight to their initial expectations, k_0); the signals of other participant valuations that an individual may take account of (in our model, the signal is given by $\hat{k}_i$); and, for a signal, the trigger points that will cause an individual to change his or her valuation (in our model, the k^*_j, which we take to be uniformly distributed across

[30] In a frictionless world characterized by instantaneous price discovery, technical analysis would have no role to play.

individuals). A more general behavioral issue that requires further investigation concerns how individuals cope with enormous and enormously complex information sets, and the way in which they make decisions when incapable of measuring input variables with precision.

Our analysis underscores the importance of market structure. We have analyzed price discovery as the product of a network. The scope and quality of a network very much depends on the structural environment within which it develops and operates. Market structure affects both the generation of information (e.g., quotes, transaction prices, and volume) and the dissemination of information (e.g., pre- and post-trade transparency, and the consolidation of quotes and prices). An array of market structure issues should be considered in light of the price discovery process, including the rules that determine how orders are brought together and translated into trades, the consolidation of markets, and market transparency.

VI. Conclusion

We have assumed a set of participants in an equity market who have divergent expectations and adaptive valuations. In contrast, standard formulations in financial economics (e.g. the capital asset pricing model) assume homogeneous expectations. The homogeneous expectations assumption is not only an extremely important simplification for many applications, it is also thought by many to be realistic: information is objective and, based on it, all rational agents should form identical expectations.

Nevertheless, the assumption does not appear to match reality. Analyst recommendations typically differ. Recent evidence suggests that markets are commonly two-sided (i.e., that, for presumably non-liquidity related reasons, some customers are seeking to buy shares at about the same time that others are looking to sell).[31] Short selling is prevalent. Clearly, information pertaining to the equity markets is enormous in both size and complexity, and it is not subject to precise assessment. Consequently, many agents produce what is commonly referred to as "private information" that results in their forming divergent expectations.

The divergent expectations assumption yields useful insight into issues pertaining to price discovery. Our analysis of price formation in a divergent expectations

[31] See Sarkar, Schwartz and Wolf (2005).

environment builds upon Handa, Schwartz and Tiwari's (2003) quote setting model where *k* percent of participants have a high valuation (V_H), 1-*k* percent have a low valuation (V_L), and V_H, V_L, and *k* are common knowledge. We depart in a significant manner from HST in that we relax the assumption that *k* is known. We take *k* to be revealed as orders are sequentially submitted to a continuous order book market. In our setting, price discovery and *k* discovery are analogous.

Participants in our model have adaptive valuations. That is, we assume that an agent's assessment of share value (V_H or V_L) is a function not only of his or her own analysis, but also of the assessments of others. This interaction may be interpreted as agents responding to signals of each others' assessments; it may also reflect participants' desires to "go with the herd," and to trade at "validated" prices.

The dynamic price discovery process, like path dependent, multiple equilibria processes in general, is sensitive to early events. This sensitivity suggests the possibility of gaming. That is, a manipulator may realize excess profits by making a purchase or a sale early in the discovery process (when prices are relatively responsive to the arrival of a buyer or a seller), and then unwinding the position later in the process (when prices are less responsive to the arrival of a buyer or a seller). Gaming in our model, however, can easily be defended against by setting bid-ask spreads that are sufficiently wide.

An important property of the divergent expectations, adaptive valuations model is that it treats price discovery as taking place in a network environment characterized by path dependency and multiple equilibria. The perspective obtained by viewing a market as this kind of a network may yield fresh insights into a variety of behavioral and market structure issues pertaining to how orders are submitted and translated into trades and transaction prices. Both the production of information (through order placement and trading) and the dissemination of information (through intermarket linkages, consolidation and transparency) determine the efficiency with which prices are discovered. These are complex issues that we leave for future research.

References

Arthur, W. Brian, *Increasing Returns and Path Dependency in the Economy,* University of Michigan Press, 1994.

Copeland, Thomas and Dan Galai, "Information Effects on the Bid-Ask Spread," *Journal of Finance*, 38, 1983, pp.1457-1469.

Foucault, Thierry, "Order Flow Composition and Trading Costs in a Dynamic Order Driven Market," *Journal of Financial Markets*, 2, 1999, 99-134.

Glosten, Lawrence and Paul Milgram, "Bid,Ask and Transaction Prices in a Specialist Market With Heterogeneously Informed Traders, *Journal of Financial Economics*, 14, 1985, pp. 71-100.

Handa, Puneet and Robert A. Schwartz, "Dynamic Price Discovery," *Review of Quantitative Finance and Accounting*, 1996, pp. 5 - 28.

Handa, Puneet, Robert A. Schwartz, and Ashish Tiwari, "Quote Setting and Price Formation in an Order Driven Market," with Puneet Handa and Ashish Tiwari, *Journal of Financial Markets,* 6, 2003, pp. 461-489.

Harrison, Michael and Kreps, David, "Speculative investor behavior in stock market with heterogeneous expectations," *Quarterly Journal of Economics*, 1978, 92, 323-336.

Ho, Thomas S.Y., Robert A. Schwartz, and David K. Whitcomb, "The Trading Decision and Market Clearing Under Transaction Price Uncertainty," *Journal of Finance,* 1985, pp. 21 - 42.

Karpoff, Jonathan, "A Theory of Trading Volume," *The Journal of Finance,* 1986, no 41, pp. 1060-1088.

Miller, Edward, "Risk, Uncertainty and divergence of opinion," *Journal of Finance,* 1977, 32, 1151-1168

Ozenbas, Deniz., Robert A. Schwartz and Robert A. Wood, "Volatility in U.S. and European Equity Markets: An Assessment of Market Quality," *International Finance,* Volume 5 Number 3, 2002, pp. 437-461.

Paroush, Jacob and Itzhak Venezia, "On the Theory of the Competitive Firm with a Utility Defined on Profits and Regret" *European Economic Review,* 1979. pp. 193-202.

Polya, G. "Sur Quelques Points de La Theorie des Probabilites." Ann. Inst. H. Poincare. 1:117-61.

Sarkar, Asani, Robert A. Schwartz. and Avner Wolf, "On the Existence and Nature of Two-Sided Markets," working paper, 2005

Scheinkman, Jose and Wei, Xiong, "Overconfidence and Speculative Bubbles," *Journal of Political Economy* 111, 2003, pp.1183-1219.

Ser-Huang Poon, Clive W. J. Granger, "Forecasting Volatility in Financial Markets: A Review" *Journal of Economic Literature,* Jun 2003, Volume. 41, Iss2; pp. 478-540.

Yan, Bingcheng and Eric Zivot, "The Dynamics of Price Discovery," working paper, University of Washington, October 2004.

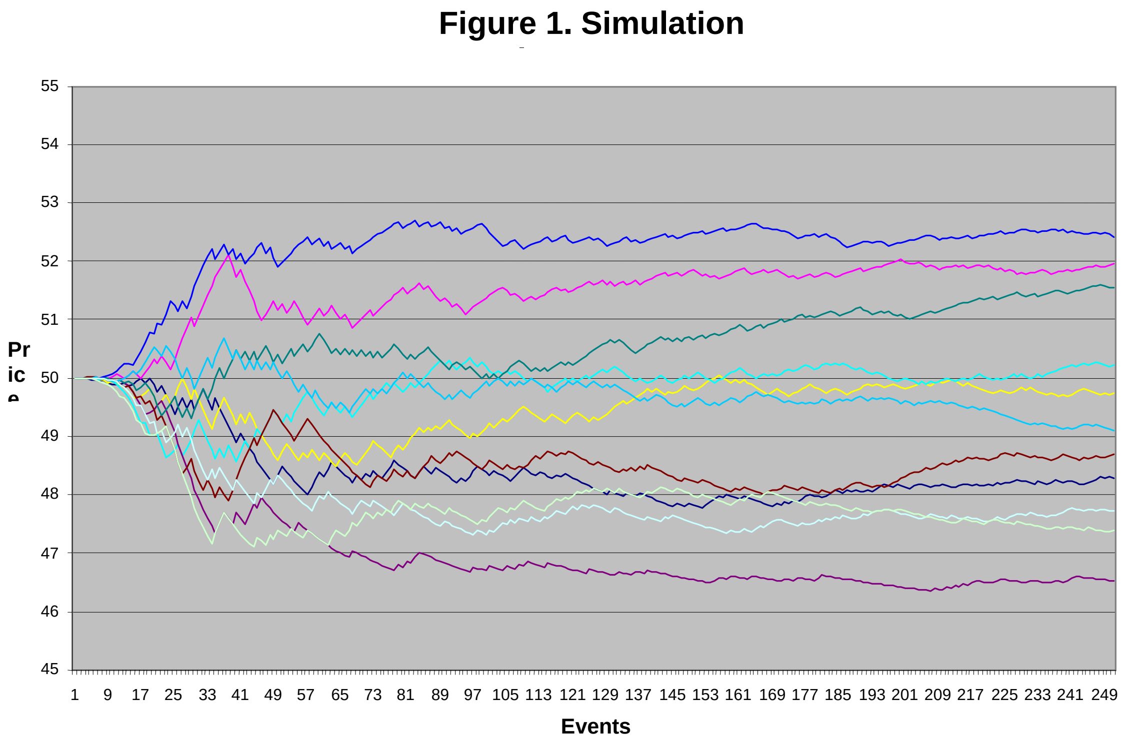

Figure 1. Typical Price Paths. This figure shows ten simulation price paths that were each traced out through 250 events (each event is the arrival of either a buyer whose share valuation is $55 or a seller whose share valuation is $45) when the initial expectation of the proportion of participants who are buyers (k_0) is 0.5, and the influence of k_0 extends, with diminishing strength, up to but not beyond the 50th event (i.e., participants memory extends to event I = 50). At each event, bid and offer prices are solved for and price is taken to be the mid-point of the bid-ask spread. As price and the revealed proportion of participants who are buyers evolve, participant expectations of the proportion who are buyers change along with the proportion who actually are buyers, i.e., participants have adaptive evaluations. This process results in multiple, path dependent equilibria (i.e., by event 200, price has tended to stabilize on 10 different values).

Figure 2A. Distribution of k_{800}

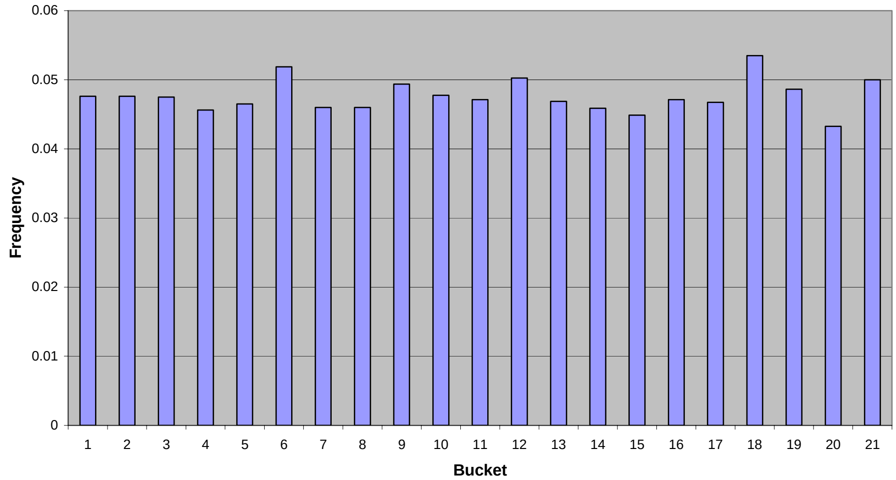

Figure 2B. Distribution of P_{800}

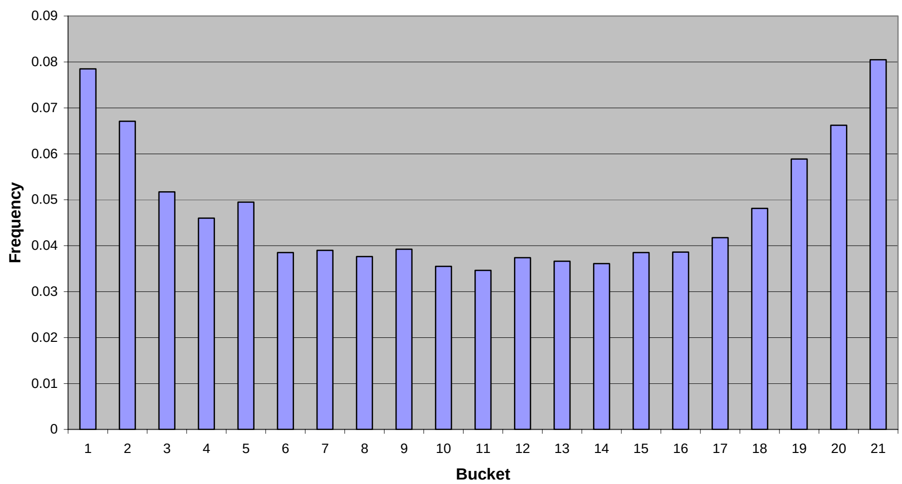

Figure 2. Distribution of k_{800} and P_{800}, with k_0 = 0.5 and I = 1 (No Memory). This figure shows the distribution of the proportion of participants who are buyers (buyer value shares at \$55 and sellers value shares at \$45) and the distribution of prices at the 800th event, based on 8,000 replications of the simulation when the initial expectation of the proportion of participants who are buyers (k_0) is 0.5, with the influence of k_0 not extending beyond the first event (i.e., participants have no memory, or I = 1). At the 800th event, bid and offer prices are solved for and price is taken to be the mid-point of the bid-ask spread. The range for k_{800} is 0 to 1, and for P_{800} is \$45 to \$55. Each range is broken into 11 equal subdivisions (buckets), with 1 indicating the smallest and 11 the largest. The 11 buckets are displayed on the horizontal axis. The vertical axis shows the frequency (the percentage of the replications) with which the values of k_{800} or of P_{800} fall in each of the eleven buckets. k_0 = 0.5 and I = 1 results in a pure Polya process. Accordingly, the distribution of k_{800} is expected to be uniform and, with price being a non-linear transformation of k, the distribution of P_{800} is expected to be U-shaped.

Figure 3A. Distribution of k_{800}

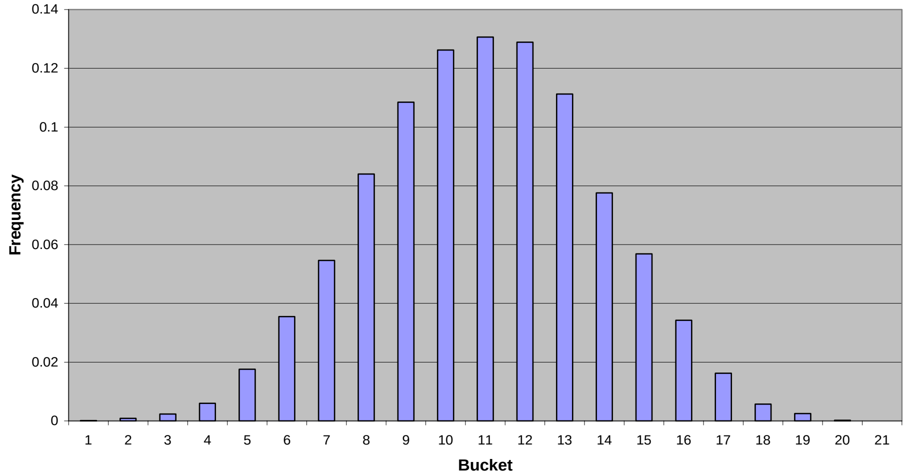

Figure 3B. Distribution of P_{800}

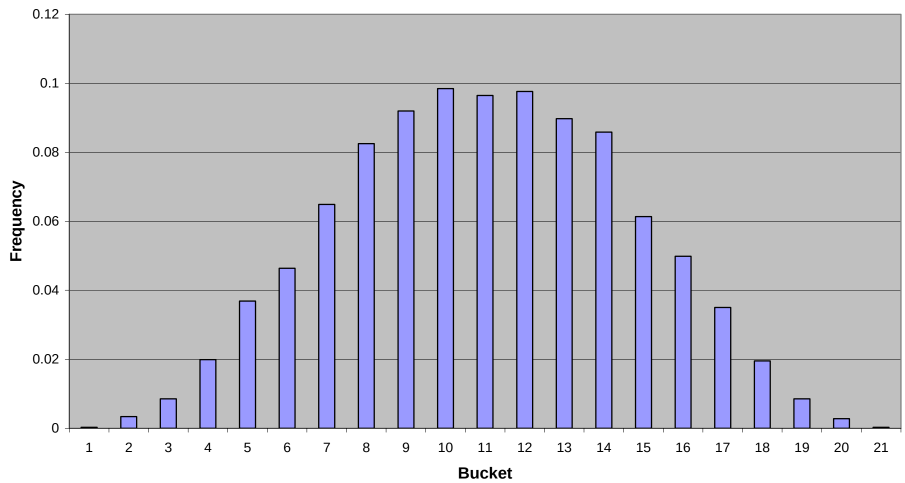

Figure 3. Distribution of k_{800} and P_{800}, with k_0 = 0.5 and I = 50 (Memory). This figure shows the distribution of the proportion of participants who are buyers (buyer value shares at \$55 and sellers value shares at \$45) and the distribution of prices at the 800th event, based on 8,000 replications of the simulation when the initial expectation of the proportion of participants who are buyers (k_0) is 0.5, but the influence of k_0 does not extend beyond the 50th event (i.e., participants have memory with I = 50). At the 800th event, bid and offer prices are solved for and price is taken to be the mid-point of the bid-ask spread. The range for k_{800} is 0 to 1, and for P_{800} is \$45 to \$55. Each range is broken into 11 equal subdivisions (buckets), with 1 indicating the smallest and 11 the largest. The 11 buckets are displayed on the horizontal axis. The vertical axis shows the frequency (the percentage of the replications) with which the values of k_{800} or of P_{800} fall in each of the eleven buckets. With I = 50, the process is not a pure Polya, and the distributions of both k_{800} and P_{800} are expected to be unimodal; with k_0 = 0.5, the mode for both k_{800} and P_{800} is expected to be at the center of the range, i.e., bucket 6.

Figure 4A. Distribution of k_{800}

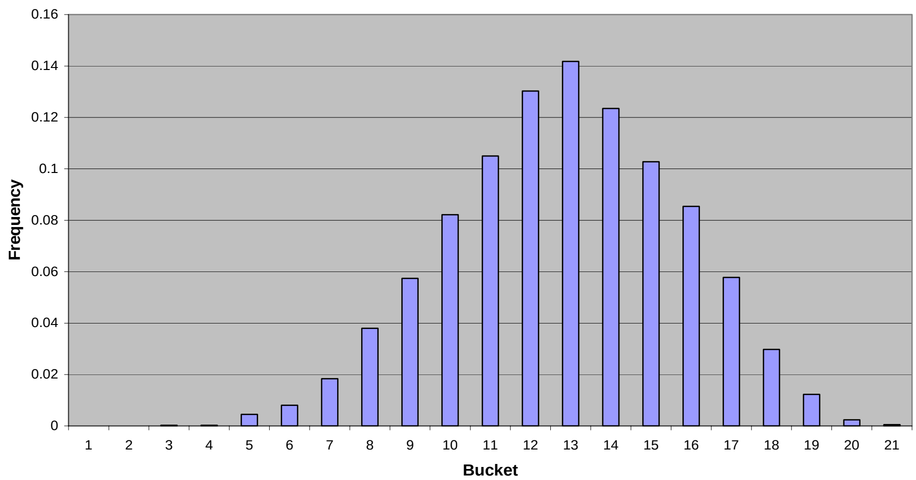

Figure 4B. Distribution of P_{800}

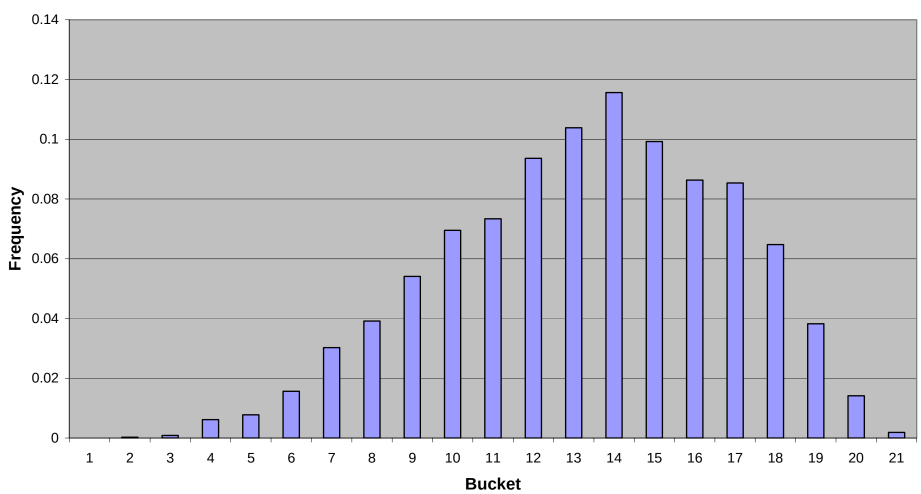

Figure 4. Distribution of k_{800} and P_{800}, with k_0 = 0.6 and I = 50 (Memory). This figure shows the distribution of the proportion of participants who are buyers (buyer value shares at \$55 and sellers value shares at \$45) and the distribution of prices at the 800th event, based on 8,000 replications of the simulation when the initial expectation of the proportion of participants who are buyers (k_0) is 0.6, but the influence of k_0 does not extend beyond the 50th event (i.e., participants have memory with I = 50). At the 800th event, bid and offer prices are solved for and price is taken to be the mid-point of the bid-ask spread. The range for k_{800} is 0 to 1, and for P_{800} is \$45 to \$55. Each range is broken into 11 equal subdivisions (buckets), with 1 indicating the smallest and 11 the largest. The 11 buckets are displayed on the horizontal axis. The vertical axis shows the frequency (the percentage of the replications) with which the values of k_{800} or of P_{800} fall in each of the eleven buckets. With I = 50, the process is not a pure Polya, and the distributions of both k_{800} and P_{800} are expected to be unimodal; with k_0 = 0.6, the mode is expected to be on the upper part of the range (above bucket 6) and the distributions of k_{800} and P_{800} are expected to be skewed to the left.

Table I: Mean and Standard Deviation of the k_{800} and P_{800} Distributions.

		k_{800}		P_{800}	
		Mean	**St. Dev. ÷ Range**	**Mean**	**St. Dev. ÷ Range**
I = 1	$k_0 = 0.2$	0.498	0.288	49.979	0.323
	$k_0 = 0.4$	0.491	0.290	49.945	0.325
	$k_0 = 0.5$	0.501	0.290	50.013	0.323
	$k_0 = 0.6$	0.505	0.288	50.019	0.322
	$k_0 = 0.8$	0.507	0.289	50.050	0.324
I = 25	$k_0 = 0.2$	0.026	0.247	46.984	0.126
	$k_0 = 0.4$	0.271	0.199	48.935	0.170
	$k_0 = 0.5$	0.498	0.141	50.017	0.177
	$k_0 = 0.6$	0.735	0.204	51.053	0.170
	$k_0 = 0.8$	0.973	0.245	53.050	0.128
I = 50	$k_0 = 0.2$	0.018	0.248	46.928	0.118
	$k_0 = 0.4$	0.255	0.206	48.915	0.164
	$k_0 = 0.5$	0.498	0.137	49.989	0.171
	$k_0 = 0.6$	0.740	0.203	51.094	0.166
	$k_0 = 0.8$	0.983	0.247	53.097	0.117

Table I: Mean and Standard Deviation of the k_{800} and P_{800} Distributions. This figure shows the mean and normalized standard deviation of k_{800} (the proportion at the 800th event of participants who are buyers, where buyers value shares at \$55 and sellers value shares at \$45) and the mean and normalized standard deviation of P_{800} (prices at the 800th event), based on 8,000 replications of the simulation when the initial expectation of the proportion of participants who are buyers (k_0) has five alternative values ranging from 0.2 to 0.8, for three different values for memory, I = 1, 25, and 50, where memory refers to the influence of k_0 on subsequent expectations (I = 1 means that participants have no memory after the first event). At the 800th event, bid and offer prices are solved for and price is taken to be the mid-point of the bid-ask spread. The standard deviations of k_{800} and P_{800} are normalized by dividing each by each variable's range (1 for k_{800} and 10 for P_{800}).

Appendix

The appendix presents a numerical illustration that shows the non-linearity of the price, *k* relationship. As we have noted in the text, this non-linearity explains two different observations. (1) As displayed in Figure 2, as *k* discovery and price discovery proceed in the pure Polya environment, *k* converges on values that are uniformly distributed over the range 0 to 1, while price converges on values that have a U-shaped distribution over the range V_H and V_L. (2) For the simulation results shown in Table 1, the distribution of the standard deviation of k_{800} with respect to k_0 is U-shaped, while the distribution of the standard deviation of P_{800} with respect to k_0 is an inverted U.

In our simulation analysis, price is represented by the mid-point of the bid-ask spread. Normalizing V_L and V_H to 0 and 1 respectively and using the simplified HST (2003) equations (6) and (7) in this text gives:

k	**Price**	**Price/*k***
.5000	.5000	1.0000
.6000	.6316	1.0527
.7000	.7532	1.0760
.7500	.8077	1.0769
.8000	.8571	1.0714
.9000	.9396	1.0440
.9900	.9949	1.0049

The above illustrates that when *k* is in the middle of its range, price is in the middle of its range (i.e., k = P = 0.5); that when 0.5 < k < 1 price is greater than *k*, and the ratio of price to *k* is at a maximum at $k = 0.75$; and that as *k* goes to its bound of 1, so too does price, and $k \approx$ P. The relationships for k <. 5 are symmetrical. Thus, as *k* moves linearly from 0.5 to either of its bounds, price follows a bow-shaped path and the *k*, price relationship is non-linear.