Financial Analysts Journal
Volume 66 • Number 5
©2010 CFA Institute

Simulating Security Markets in Dynamic and Equilibrium Modes

Bruce I. Jacobs, Kenneth N. Levy, CFA, and Harry M. Markowitz

An asynchronous discrete-time model run in "dynamic mode" can model the effects on market prices of changes in strategies, leverage, and regulations, or the effects of different return estimation procedures and different trading rules. Run in "equilibrium mode," it can be used to arrive at equilibrium expected returns.

Analysts in the natural sciences, manufacturing, logistics, and warfare frequently rely on asynchronous discrete-time models. Such models are less well known in finance, where analysts have tended to rely on continuous-time models.[1] But in finance, as in these other areas, discrete-time models offer several important advantages over continuous-time models.

Both continuous-time and discrete-time models are a form of dynamic model. Dynamic models allow one to represent the evolution of a system, such as a financial market, over time. In continuous-time dynamic models, the system changes continuously over time; in discrete-time dynamic models, an internal system clock advances in discrete increments. Discrete-time models can be further classified into synchronous and asynchronous models. Synchronous discrete-time models use system clocks that advance by fixed increments, such as a day or a year, with the status of the system updated at each increment. Asynchronous discrete-time models use system clocks that advance from one event to the next, whereby the time intervals between events are typically not constant. Asynchronous models can provide a more realistic representation of markets than synchronous models.[2]

The most commonly used dynamic models in finance assume that security prices follow a continuous-time process. This process is frequently assumed to be random and is modeled as a Brownian motion or as a function of a Brownian motion. A major advantage of continuous-time models is that some of them can be solved explicitly, which allows one to evaluate investment strategies analytically. Most familiar, perhaps, are option-pricing models that can be solved given a fixed-price process for the underlying security.

Most discrete-time models cannot be solved analytically; large and detailed asynchronous models that attempt to model complex systems require computer simulation. Asynchronous discrete-time models, however, can provide insights not available from purely analytical procedures. They can be used to examine the mechanisms behind price movements and can thus be used to test the effects on security prices of such real-world events as changes in investors' strategies, modifications in overall leverage, and switches in regulatory regimes.

Consider how continuous-time and discrete-time models deal with so-called liquidity black holes. On Black Monday, 19 October 1987, liquidity disappeared from the market as large numbers of investors all attempted to sell at the same time. Similar black holes developed in connection with the collapse of the hedge fund Long-Term Capital Management in 1998 and, more recently, during the 2008–09 credit crisis (see, e.g., Jacobs 1999, 2004, 2009). In these and other, less extreme cases, the price process was not fixed. Continuous-time modeling that assumes the contrary may provide misleading results.

Discrete-time models are better suited than continuous-time models to deal with changes in underlying parameters. For example, Kim and Markowitz (1989) presented an asynchronous discrete-time simulation in which investors are assumed to be either rebalancers or portfolio insurers. The former tend to sell as market prices rise and buy as market prices fall, whereas the latter buy as prices rise and sell as prices fall. Kim and Markowitz showed that the behavior of the market changes radically as the proportions of the two kinds of investors vary.

Bruce I. Jacobs and Kenneth N. Levy, CFA, are principals at Jacobs Levy Equity Management, Florham Park, New Jersey. Harry M. Markowitz is president of Harry Markowitz Company, San Diego.

In the five years leading up to the 1987 market crash, more and more investors became portfolio insurers. Perhaps a model that was able to incorporate the actual strategies of market participants would have been able to anticipate that the trend-exacerbating activities of these portfolio insurers would lead to a market crash.

More generally, in order to make meaningful statements about anticipated market activities, one needs a model that can represent the strategies of different kinds of participants, as well as changes in the composition of those participants. Asynchronous simulation is well suited to this purpose.

Asynchronous simulation can also be used to test hypotheses about the current behaviors of market participants. For example, consider the controversy over whether investors are predominantly rational beings who optimize some given utility function or not-so-rational beings who are swayed by fads, fashions, and other cognitive biases. Surely, the market includes a mixture of these investors and others with still different investment patterns. A detailed asynchronous simulation can tell us whether given combinations of investor behaviors lead to price patterns that resemble observed market behavior. A continuous-time model that starts by assuming some random price process cannot do this.

Finally, asynchronous models can be used to arrive at equilibrium expected returns for a variety of realistic financial markets without requiring the kinds of unrealistic assumptions, such as unconstrained leverage, that some analytical models require (see, e.g., Black and Litterman 1992).

Simulation Overview

When simulating any system, one is free to choose how the system is represented. We find it convenient to use five basic types of entities to represent a market: securities, statisticians, investors, portfolio analysts, and traders.[3] Our simulation determines the prices and trading volumes of securities endogenously. Simulated statisticians provide return estimates, variances, and covariances. Ideal portfolio weights are determined by portfolio analysts who use the inputs from statisticians and investors' risk-aversion parameters and portfolio constraints. Prices and volume arise as traders seek to complete the desired trades to move investors' current portfolios toward their ideal portfolios. All investors are mean–variance investors who seek to maximize their utility: $U = E - KV$, where, for each investor, E is the portfolio's expected return, K is the investor's risk-aversion parameter, and V is the variance of the portfolio's return. **Exhibit 1** describes the simulator's entities and their relationships in more detail.

Statisticians can use two types of return estimation procedures: *HIST* and *RPS_C*. With *HIST*, statisticians use historical security price data to estimate returns. With *RPS_C*, they form return-per-share estimates by dividing a given expected constant dollar gain by the security's current price.

Our asynchronous discrete-time simulator can operate in two modes. When the objective is to model the evolution of certain time-varying quantities—in this case, market prices and volumes—the simulator operates in the dynamic analysis (DA) mode. When the objective is to use the simulator as an iterative parameter estimator to find the values of such parameters as equilibrium-implied expected returns for securities—on the basis of the composition of the market portfolio and the preferences of market participants—the simulator operates in the capital market equilibrium (CME) mode. In the following sections, we describe our findings in these two modes.

Dynamic Analysis

In the DA mode, we modeled the reaction of security prices and volumes to different scenarios, environments, and policies. In particular, we examined how prices and volumes react to changes in (1) the initial random seeds that determine individual investor initial wealth and cash flows over the simulation, (2) the proportion of entities that use various methods of estimating expected returns, and (3) trading and anchoring rules.

Different Initial Random Seeds. For this analysis, we used a base case with the following parameters:

- A 4,000-day run (corresponding to 16 2/3 simulated years)
- 16 securities, each with a starting price of $200
- Four *RPS_C* statisticians
- Four *HIST* statisticians
- 4,000 investors who relied on *RPS_C* statisticians (i.e., each *RPS_C* statistician provided estimates to 1,000 investors)
- 200 investors who relied on *HIST* statisticians (i.e., each *HIST* statistician provided estimates to 50 investors)

Table 1 summarizes information about the base case.[4] We ran three simulations (Runs A, B, and C) that differed only with respect to their initial random seeds.

Figure 1 plots the equally weighted market indices for the three simulation runs. In each run, the index starts at 200 (because in each run, the initial price of each security was set at $200) and ends between 200 and 300 after 4,000 simulated days. In between, each index exhibits considerable short-term noise and a few longer movements.

Exhibit 1. Entities Used in the Asynchronous Discrete-Time Simulation

Securities	Securities are the market's fundamental tradable assets. Their prices and the volumes traded are determined endogenously by the simulation.
Statisticians	Statisticians provide estimates of securities' expected returns, variances, and covariances. The number of statisticians can be controlled by the user. The user can also control the method that statisticians use to form return estimates and the frequency and number of observations they use in making their estimates. The simulation provides for two methods of return estimation. In the *HIST* method, statisticians use historical security price data to form their estimates. In the *RPS_C* ("return per share constant") method, statisticians estimate expected returns by dividing a given expected constant dollar gain by current price. Suppose, for example, that a particular security is expected to return $10 per share over the next year; if the price of the security is currently $200, the *RPS_C* statistician estimates expected return to be $10/$200, or 5 percent.
Investors	Investors hold portfolios of securities, and all want to maximize their utility, which is defined as expected portfolio return minus a risk-aversion parameter times portfolio variance (subject to budget and leverage constraints). Individual investors with similar characteristics are grouped together into investor templates. The user can control the number of investor templates, the number of investors stamped from each template, how frequently the investors reoptimize, which portfolio analyst the investors use, which trader the investors use, the risk aversion of the investors, and the amount of leverage the investors can use. Investors within a particular template differ from one another only insofar as their wealth and cash flow are concerned. The initial wealth of each investor is drawn from a lognormal distribution, the mean and standard deviation of which depend only on the template to which the investor belongs. In addition, each investor can experience a random cash flow (deposit or withdrawal) at each portfolio reoptimization; the user can determine the probability of such a cash flow, as well as the parameters relating to the probability distribution of the size of the deposit or withdrawal.[a]
Portfolio analysts	Portfolio analysts link statisticians and investors. Each portfolio analyst uses the expected return, variance, and covariance estimates created by a specific statistician and, from these estimates, produces an ideal, or target, portfolio for a given investor template. The user can control parameters relating to analysts, including how many analysts there are, which statistician each analyst uses, and such information as the maximum leverage that each analyst can use in creating target portfolios and whether the analyst's investors can sell securities short.[b]
Traders	Traders buy and sell securities for investors. Traders enter bids and offers into a central limit order book and attempt to move each investor's portfolio in a straight line (in portfolio space) from its current composition to its target composition. For example, if the investor's current portfolio vector is x_0 and the target portfolio vector is x_1, then traders move the portfolio along the trajectory defined by $(1 - h)x_0 + hx_1$ for $0 \leq h \leq 1$. Traders are subject to controls that limit the extent of trading used to move toward an investor's target portfolio and by trading rules used in setting bid and offer prices. Traders observe these rules to ensure that investors' orders are executed in a timely but cost-effective manner. The rules also determine how the traders form their initial bids and offers, how they vary the bids and offers over time, and how they determine price limits beyond which they would not execute trades. One trader, known as a liquidation trader, liquidates positions when any investor's portfolio violates maintenance margin requirements. This trader's parameters can be set to be very aggressive, thus ensuring that all necessary trades are made, virtually irrespective of the cost. When bid and offer prices for the same security overlap, a trade occurs. The price at which this trade happens is regarded as the current price of the security until another trade occurs. Thus, prices are determined endogenously by trades.[c]

[a]Each investor starts with his or her initial wealth invested in a portfolio that contains all securities in equal weights. In practice, the first few months of a simulation often show greater trading volume and price variability as investors move to their desired portfolios.
[b]For portfolio optimization with long and short positions, see Jacobs, Levy, and Markowitz (2005, 2006).
[c]Because prices—and thus returns—are determined endogenously, the return distributions need not be normal. Levy and Markowitz (1979), however, provided motivation for using mean–variance preferences as an approximation for direct expected utility maximization even when returns are not normal.

In each of the three runs, the market started by moving sharply upward and then declined. Afterward, each of the three markets recovered and each experienced its own particular ups and downs. Although each of the simulated markets had its own unique movements, just as actual markets do, all the markets were relatively stable. As we will show later in this section, that is not always the case.

Figure 2 shows the 90-day moving averages of total daily volumes for Runs A, B, and C. A common complaint about many analytically tractable market models is that they often imply little volume and sometimes none. This complaint does not apply to our simulator, at least with respect to the parameter settings in our examples. Throughout the 16 2/3 years of each simulation, none of the three runs had a 90-day moving average volume of less than 200,000 shares a day. For the 16 securities, this averages to at least 12,500 shares a day. This amount seems to be a reasonable volume for a market with 4,200 investors.

One possible concern is the apparent downward slope of all three volume curves. To ascertain whether the model gives large volumes at first and

Table 1. Summary of Investor Template Parameters for the Base Case

Investor Template No.	Risk-Aversion Parameter	No. of Investors	Frequency of Reoptimization	Return Estimation Procedure
0	3	1,000	Daily	*RPS_C*
1	2	1,000	Monthly	*RPS_C*
2	2	1,000	Quarterly	*RPS_C*
3	4	1,000	Daily	*RPS_C*
4	3	50	Daily	*HIST*
5	2	50	Monthly	*HIST*
6	2	50	Quarterly	*HIST*
7	4	50	Daily	*HIST*

Notes: This table shows the parameters used in the base case simulation, against which others are compared. The base case uses eight investor templates, of which four are of the *RPS_C* type and four are of the *HIST* type. Each *RPS_C* investor template is used by 1,000 investors, and each *HIST* template is used by 50 investors. The table also shows each investor's risk aversion and frequency of reoptimization.

Figure 1. Equally Weighted Market Indices for Three DA Mode Simulation Runs with Identical Parameters but Different Initial Seeds

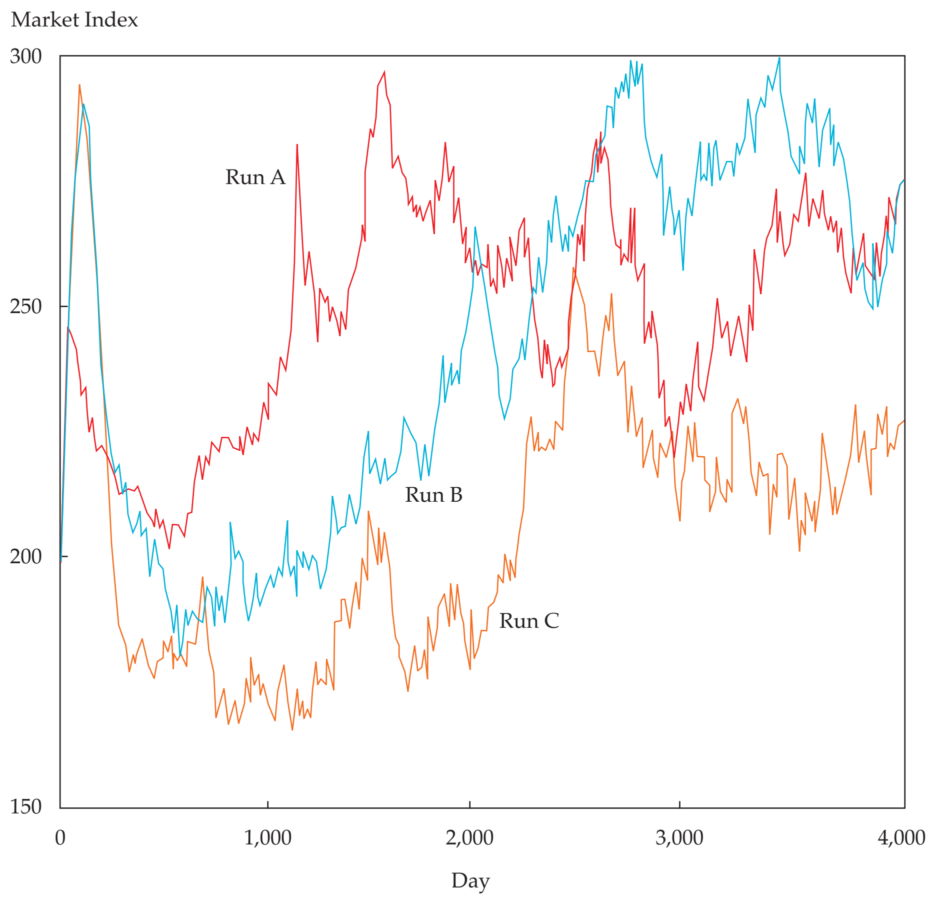

then reduces to little or no volume, we continued several runs for much longer periods and found that the model maintains a reasonable volume level.

In summary, we found, as expected, that changes in the initial random seeds merely gave rise to idiosyncratic differences in security prices and volumes. We were thus assured that the simulated markets were insensitive, in the aggregate, to random seed changes.

Different Ratios of Momentum to Value Investors. If all market participants use historical information to form their expectations about future security returns but use different aspects of history

Figure 2. Total Trading Volume 90-Day Moving Averages for Three DA Mode Simulation Runs

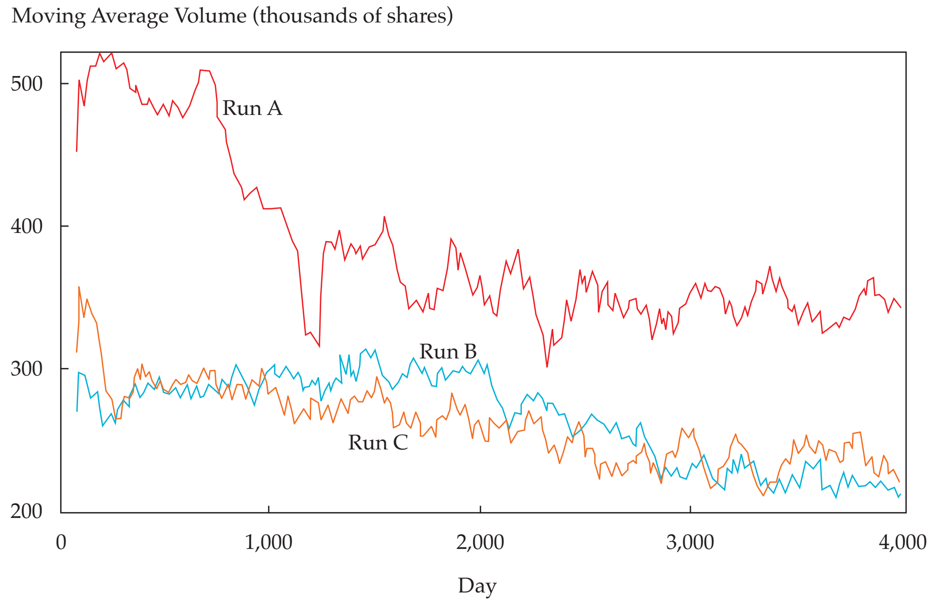

(e.g., different frequencies of observation), then one plausible theory is that an active but stable market will emerge because of sufficient diversity among the participants.

We tested this theory by assuming that all statisticians used securities' historical returns to estimate expected returns, variances, and covariances (i.e., all statisticians were *HIST* statisticians). We varied the frequency of return observation (daily, monthly, or quarterly) and the number of periods (days, months, or quarters) used in estimating returns.[5] We found that when all statisticians were *HIST* statisticians, the price of at least one of the securities invariably grew exponentially over time.

In effect, all investors were momentum investors, and they destabilized the market by introducing positive feedback. When a security's price increased, all statisticians raised the expected return estimate of that security, all portfolio analysts wanted more of it in their target portfolios, and all investors attempted to buy it. Traders' purchase orders pushed the security's price up, further increasing demand for it. With the positive feedback inherent in momentum investing, prices exploded. Conversely, when a security's price declined, all investors sold, which put further downward pressure on it.

Another possible theory is that a stable and active market will emerge if, in addition to *HIST* statisticians, there are other statisticians who base their estimates on fixed or constant future dollar returns per share (i.e., *RPS_C* statisticians). With *RPS_C* statisticians (for a given security with a constant dollar return per share), expected return declines as share price increases. This is the opposite of *HIST* statisticians, who raise return expectations as prices rise. Thus, investors who rely on *RPS_C* statisticians will tend to buy more (less) of a security as its price declines (increases).

We categorized investors as momentum investors if they used portfolio analysts who obtain their estimates from *HIST* statisticians, and we categorized investors as value investors if they used portfolio analysts who obtain their estimates from *RPS_C* statisticians. Momentum investors can destabilize markets because their actions tend to prolong and exacerbate trends. Value investors stabilize markets; as prices rise, these investors tend to sell, and as prices fall, they tend to buy.

To see how differing proportions of value investors and momentum investors affect overall market behavior, we used the previously described base case but varied the number of momentum investors in the simulation. First, we increased the number of momentum investors so that there was 1 momentum investor for every 10 value investors (400 momentum investors versus 4,000 value investors). Prices became more volatile, but markets were not explosive.

The results changed dramatically when we further increased the number of momentum investors so that there was 1 momentum investor for every 5 value investors (800 momentum investors versus 4,000 value investors). **Figure 3** plots the daily prices for this case. Here, the price of Security

Figure 3. Security Prices vs. Time for Simulation with One Momentum Investor per Five Value Investors in DA Mode

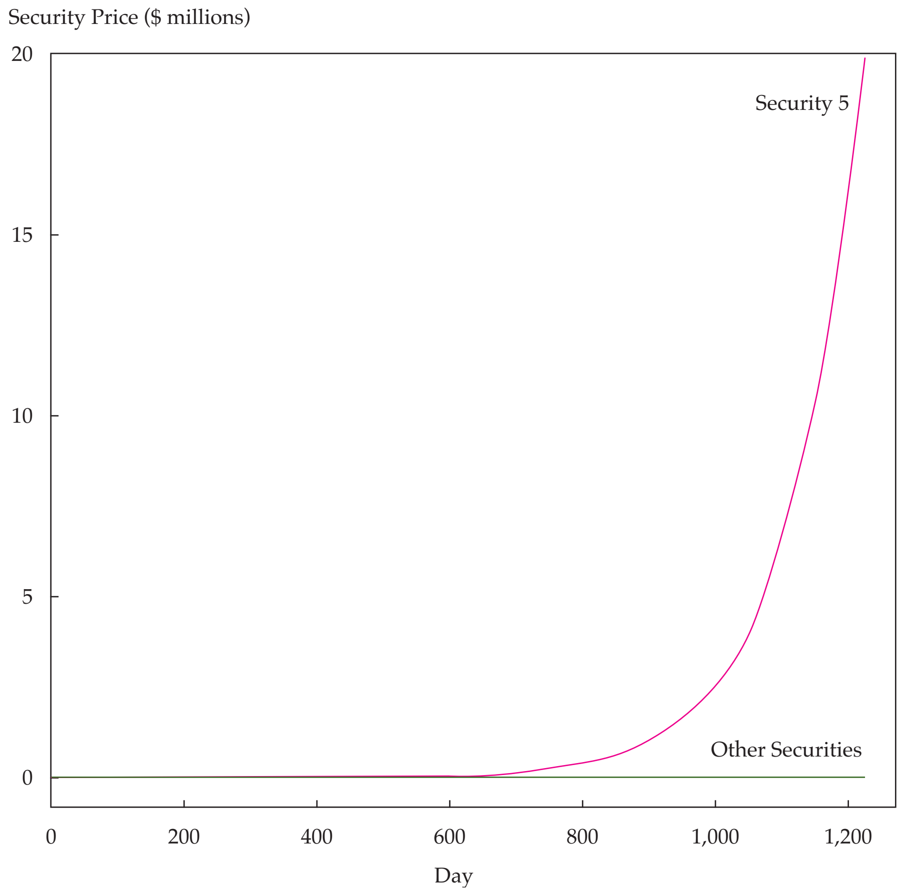

5 reaches $20 million after Day 1,200, and the simulation terminates. The prices of the other 15 securities are also plotted in Figure 3, but they are so small in comparison to that of Security 5 that they all appear as a single horizontal line near $0.

Adding even more momentum investors caused prices to explode even faster. In the explosive cases, share volume started at about the same level as in the nonexplosive cases but dropped rapidly (though not monotonically) to much lower levels. Note, however, that the dollar volume of trading did not diminish in the explosive cases.

That the market explodes if all investors estimate expected returns by historical average returns may seem worthy of little note because it is a negative finding. It tells us how the market should *not* be modeled rather than how it should be modeled. Some, however, would consider the use of historical returns to be the obvious natural initial hypothesis. This use has a long history in financial theory. For example, Sharpe (1963) used historical average returns to estimate expected returns, and Michaud (1998) argued for the superiority of a resampled frontier based on historical averages. Even those who argue for Bayes-Stein estimation methods use historical averages in their null hypothesis.

Proponents of Bayes-Stein estimators or resampled frontiers argue that their methods produce better decisions than mean–variance methods that use historical averages for estimating expected returns. But they do not argue that if all investors used methods based only on historical averages, markets would self-destruct. To our knowledge, such an argument is not the implication of any prior simulation or analytical market model. It is the result of our more literal description of the estimation, optimization, and trading process. It suggests that a stable, active market requires a predominance of value investors. If there are too many momentum investors who buy and sell on the basis of historical prices, the market becomes unstable and prices of some securities either explode or plunge to zero.

Trading and Anchoring Rules. In real markets, security analysis, portfolio construction, and trading occur asynchronously at discrete-time intervals, with security analysts or statisticians providing their estimates to portfolio analysts, who provide target portfolios to traders, who execute trades to move investors' portfolios toward their target portfolios. In the time between security analysis and trading, changes in the prices of securities may affect trading goals. Traders need rules to guide them in deciding whether securities whose prices have moved should still be traded and, if so, at what prices to place revised orders. To provide a realistic representation of a market, an asynchronous discrete-time simulator needs these types of rules.

■ *Trading rules*. Our trading rules specify how bid and offer prices for an investor's orders are to be formed and adjusted over time. They also specify the conditions under which bids or offers must be withdrawn. We describe a trading rule for offers; an analogous rule applies to bids.

According to the offer trading rule, each trader sets its own initial offer price, P_O, as follows:

$$P_O = \alpha_S + \beta_S P_C, \tag{1}$$

where α_S is the particular trader's user-specified "sell alpha," β_S is the trader's "sell beta," and P_C is the current price of the security. For example, assume that the security most recently traded at \$200. Then, a trader with an α_S of \$1 and a β_S of 1 would submit a limit order to sell (i.e., set its own initial offer price) at \$201.

If a trade is not executed at or above this price within a specified period, the trader adjusts the alpha and beta as follows:

$$\alpha_S \leftarrow \alpha_S + \Delta\alpha_S \text{ and } \beta_S \leftarrow \beta_S + \Delta\beta_S, \tag{2}$$

where $\Delta\alpha_S$ and $\Delta\beta_S$ are the trader's user-defined "alpha increment" and "beta increment." The trader alters the limit order to sell the security at the revised price computed from Equation 1 by using the adjusted alpha and beta values. If a trade is still not executed within a specified period, the trader readjusts the alpha and beta to produce a new value for its offer price. If a trade has not been executed after a specified number of adjustments, the trader withdraws the sell order.

Under these trading rules, when traders review their unexecuted orders, they will either cancel their buy (sell) orders or increase (decrease) their own bid (offer) prices. The bids will thus tend to march upward toward the offers, and the offers will tend to march downward toward the bids (until a trade occurs, a bid or offer expires because of time limitations, or the investor reoptimizes). The arrival of new orders may exacerbate these tendencies. For example, if a preponderance of buy orders is entering the book, sequential transactions might result in a rising best offer and in transactions occurring at higher and higher prices. The arrival of sell orders might have the opposite effect.

■ *Anchoring rules*. In some situations, price changes can be extremely rapid. For example, in the base case, we found that orders are generated every day for as many as 2,100 investors who reoptimize daily, plus some for the 2,100 investors who reoptimize monthly or quarterly. If these thousands of participants want to purchase a particular security, they could issue a series of new buy orders, each of which establishes a slightly higher bid price. These orders will notch the price up repeatedly; a security that sold for \$100 at the beginning of the day might sell for thousands of dollars at the end of the day.

The underlying cause of these rapidly changing prices is that simulated traders, subject exclusively to the trading rules previously discussed, consider only the current situation. They do not recall, for example, that a security whose best offer is in the thousands of dollars may have sold for \$100 earlier in the day. In contrast, human traders have a good sense of how much to pay for a security. If prices deviate too much from recent levels, human traders will either revise or withdraw their bids and offers. We refer to this sense as "anchoring." Computer algorithms do not automatically incorporate this sense of anchoring. Without it, simulated traders can drive a security's price to unrealistic levels.[6] In our simulations that used only these trading rules, sometimes with as little as one momentum investor for every 20 value investors, the market was explosive. Even with all value investors and no momentum investors, market behavior was sometimes unstable.

To approximate a human trader's sense of how much realistically to pay for a security, we devised and incorporated anchoring rules. The simulated trader follows the trading (i.e., price adjustment) rules described in Equation 1 and Equation 2 until the adjustment conflicts with the anchoring rules, at which point the anchoring rules govern. That is, the anchoring rules become effective when the security's price deviates too far from its recent level. The user sets parameters that define "recent" and what prices are considered "too far" from the recent level. We describe an anchoring rule for offers; an analogous rule applies to bids.

The offer anchoring rule enables the user to set the minimum offer price as follows:

$$P_{O,\min} = P_R - cP_L, \tag{3}$$

where P_R is a recent price (which the user can specify to be the average price or the minimum price of the security over a specified number of recent days), c is a user-specified parameter, and P_L is the average recent price or the standard deviation of recent prices (as defined by the user). Thus, for example, the anchoring rule could specify that the trader will not sell at a price more than 10 percent below the 20-day moving average price. This structure permits flexibility in modeling trader behavior by allowing the trader to cancel orders when security prices "move away." It also helps reduce the problem of price explosions and implosions.

Our experience with trading rules confirmed another advantage of asynchronous discrete-time simulation. Our initial micro-hypotheses about trading behavior led to macro-behavior that did not resemble a real-world market. Our simulation results indicated that something was wrong with our assumptions and led us to form better ones.

Capital Market Equilibrium

We have described the use of an asynchronous discrete-time simulator in its dynamic analysis mode. Simulation, however, is not limited to dynamic analysis. Simulation can be used to estimate parameters of realistic representations of systems that are too complicated for analytic approaches to handle. This section gives an example of how an asynchronous discrete-time simulator can be used as a parameter estimator. In particular, we describe an alternative mode of operation—the capital market equilibrium (CME) mode—that can be used to find equilibrium expected returns that are consistent with a given covariance structure and market portfolio. It allows the user to find equilibrium expected returns for any of the great variety of markets that can be simulated, results that would be impossible to find by using closed-form analytic techniques.

Black and Litterman (BL 1992) suggested a procedure to find equilibrium expected returns that are consistent with a given covariance matrix and a specified market portfolio. The BL procedure operates on the assumption that investors live in a capital asset pricing model (CAPM) world—that is, they can either borrow all they want at the risk-free rate or sell short without limit and use the proceeds to buy long, subject only to a budget constraint that all their holdings sum to 100 percent. Under these assumptions, the BL procedure uses "reverse optimization" to compute equilibrium expected returns from the given covariance matrix and the specified market portfolio.

The BL procedure for estimating expected returns has the following inputs: a covariance matrix, percentages of the market portfolio invested in various securities, and views about expected returns for some, all, or none of the securities. If the user supplies no views, the BL procedure produces CME expected return estimates. These estimates are expected returns that would clear the market if investors are essentially unconstrained and can borrow all they want at the risk-free rate.

Under the BL assumptions, the Tobin (1958) separation theorem applies, and all investor portfolios lie on the straight capital market line (CML). Portfolios on the CML consist of various combinations of the riskless security and the *same* portfolio of risky securities. In reality, contrary to the assumptions of the BL procedure, investors are constrained and cannot borrow all they want. Because of these constraints, investor portfolios do not all lie on the CML. Instead, they lie on the curved efficient frontier at positions determined by investor risk aversions, and the compositions of the portfolios of risky securities differ from investor to investor. In such cases, the market portfolio may not even be efficient (see Markowitz 2005).

With a simulator used as a parameter estimator, equilibrium expected returns are not subject to the BL procedure's unrealistic assumptions regarding constraints and borrowing. That is, in the CME mode, the simulator seeks equilibrium expected returns for markets in which the CAPM assumptions do not necessarily hold. It allows users to solve for expected returns for markets with real-world constraints, including those in which investors can neither borrow without restriction nor short without restriction. In other words, it can be used to seek equilibrium expected returns for any market that can be simulated. This statement is subject to two caveats. (1) Not all such markets are consistent with equilibrium solutions. (2) We have not explored the convergence properties of the simulator for all such markets.

In the following sections, we describe how a simulator can compute CME expected returns and present a sample case that illustrates the speed with which a simulator converges toward the market-clearing expected returns.

Expected Return Estimation Method. The objective is to find CME expected security returns. Recall that in the DA mode, the statistician provides securities' expected returns. In the CME mode, securities' expected returns are instead determined by the following iterative adjustment procedure: If the weight of a security in the simulated market portfolio is above the specified (or target) market portfolio weight, its expected return is lowered. If the weight of a security in the current market portfolio is below the target weight, its expected return is raised. In response to the expected return changes, investors change their portfolios in such a way that the aggregate of all investors' portfolios converges toward target market portfolio weights.

To implement this adjustment procedure, we introduce four non-negative parameters (a_0,a_1, b_0,b_1). For the i-th security, the iterative adjustment proceeds as follows: Let δ_i be the difference between the weight of the security in the simulated market portfolio, w_i^m (computed from the aggregate of investors' holdings), and the weight of the security in the target portfolio, w_i^t; that is,

$$\delta_i = w_i^m - w_i^t.$$

If the simulated market weight is close enough to the target weight such that

$$|\delta_i| < a_0 + a_1 w_i^t,$$

no action is taken. Otherwise, if δ_i is positive (negative), the simulator subtracts (adds) $b_0 + b_1 w_i^t$ from (to) the current return estimate for security i. That is,

$$r_i \leftarrow r_i - \text{sgn}(\delta_i)(b_0 + b_1 w_i^t).$$

The parameters a_0 and a_1 thus define a tolerance band around the target weights, outside of which adjustment of the return estimates is deemed necessary, and the parameters b_0 and b_1 define the degree to which the estimates should be adjusted. Because this return-estimating procedure is homogeneous, the overall level of the estimates must be fixed arbitrarily. Our method is to adjust expected returns so as to seek an equilibrium solution in which the average expected return is the same as its initial value. Therefore, the market weights after convergence may differ from the target weights. Nevertheless, the relative market weights of all securities (excluding cash) are the same as their relative target weights.

Case Study. To create a realistic representation of market participants' holdings, we created 10 investor templates of 10 investors each that would place representative portfolios on various parts of the efficient frontier and not just on the CML.[7] With such placement, the BL assumptions are no longer satisfied. Therefore, the BL procedure would not provide correct equilibrium expected returns.

All investors of a given template choose the portfolio that maximizes $U = E - KV$ for their given risk aversion, K. In this case, we let K = 0.5, 1.0, 2.0, . . . , 9.0 for the 10 investor templates, respectively. The templates are otherwise identical. We created eight securities plus cash and stipulated that the specified market portfolio comprise 8 percent in each of Securities 0 through 5, 4 percent in Securities 6 and 7, and the remainder in cash. We assigned arbitrary initial values to the return estimates, but their mean was 10 percent.[8]

We used a common constant covariance matrix. In other words, all entities shared the same estimated covariance matrix, and this covariance matrix remained constant over time.

Table 2 shows the results at the end of the run. Each row of the table represents the equilibrium portfolio weights (rounded to integer percentage points) for a particular investor template, and each column represents each investor template's holdings of a particular security. The table shows that, in equilibrium, Investor Template 0 (K = 0.5) holds a portfolio with 29 percent in Security 1, 28 percent in Securities 3 and 5, 10 percent in Security 7, and 5 percent in cash. Investor Template 0 holds none of Securities 0, 2, 4, or 6.

Table 2. Investor Template Portfolios at the End of the CME Run

Investor Template No.	Risk-Aversion Parameter	S0	S1	S2	S3	S4	S5	S6	S7	Cash
0	0.5	0	29	0	28	0	28	0	10	5
1	1	9	20	9	20	9	20	0	11	3
2	2	17	11	17	11	17	11	7	6	3
3	3	16	7	16	7	16	7	9	4	17
4	4	12	5	12	5	12	5	7	3	39
5	5	10	4	10	4	10	4	6	3	50
6	6	8	4	8	4	8	4	5	2	58
7	7	7	3	7	3	7	3	4	2	64
8	8	6	3	6	3	6	3	3	2	69
9	9	5	2	5	2	5	2	3	1	73

Notes: This table shows the equilibrium portfolios for each of the representative investor templates (0 through 9) at the end of the CME run. The second column shows each investor template's risk aversion. The third through eleventh columns give the (rounded) percentage that each investor holds in each of the eight securities (S0 through S7) and cash.

In contrast, Investor Template 9 ($K = 9$) holds a portfolio with positive positions in all securities. Thus, the investors in this simulation are on different segments of the efficient frontier. Their portfolios cannot be constructed by using a convex combination of a market portfolio and cash. That is, the portfolios do not all lie on the CML, and the assumptions underlying the BL procedure do not hold.

Figure 4 illustrates the convergence of the simulated market portfolio toward its target weights. The weights, plotted monthly, converge within approximately 300 iterations or 300 days.[9] Securities 0 through 5 each converge to approximately 9 percent of the market. Securities 6 and 7 each converge to approximately 4.5 percent of the market. Cash converges to an average of about 36 percent of the market. Note that the ratio of the weights of each security to their sum (excluding cash) is the same as that of the target portfolio, as previously described.

Figure 5 shows the evolution of the return estimates. Consistent with the convergence of the market portfolio, the estimates converge within approximately 300 iterations or days. For Securities 0 through 7, the final estimates are 9.3 percent, 11.0 percent, 9.3 percent, 11.0 percent, 9.3 percent, 11.0 percent, 8.9 percent, and 10.3 percent, respectively. The average of these estimates is 10 percent, the same as the average of the initial estimates, as required by the estimate anchoring rule.

Figure 4. Convergence of Weights in CME Mode

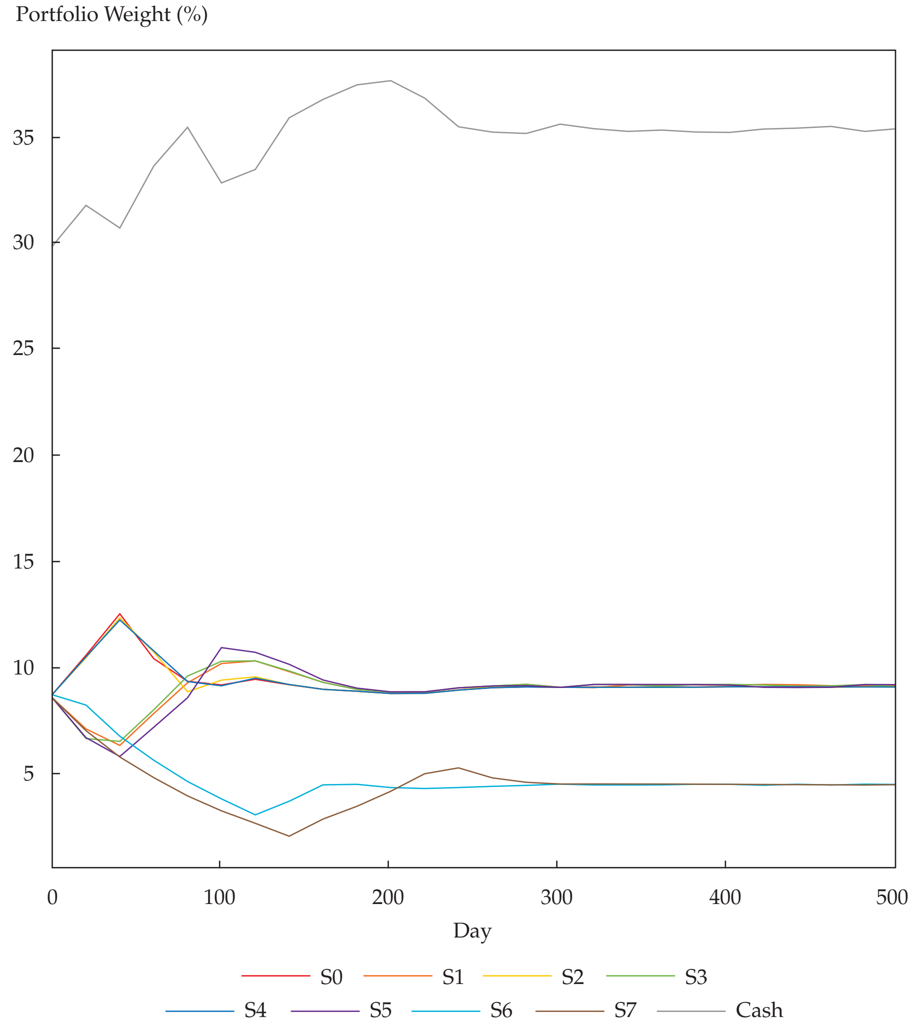

Notes: The objective of the simulator in the CME mode is to arrive at market-clearing equilibrium expected returns under realistic conditions. The simulator does so by iteratively adjusting the expected returns and, consequently, each investor's desired portfolio weights until the aggregate of all investors' weights equals the target market portfolio. This figure demonstrates the algorithmic convergence of the aggregate weights toward the target market weights. In this CME mode example, the investors' weights converge toward target market weights within approximately 300 iterations (with one iteration per day).

Figure 5. Evolution of Return Estimates in CME Mode

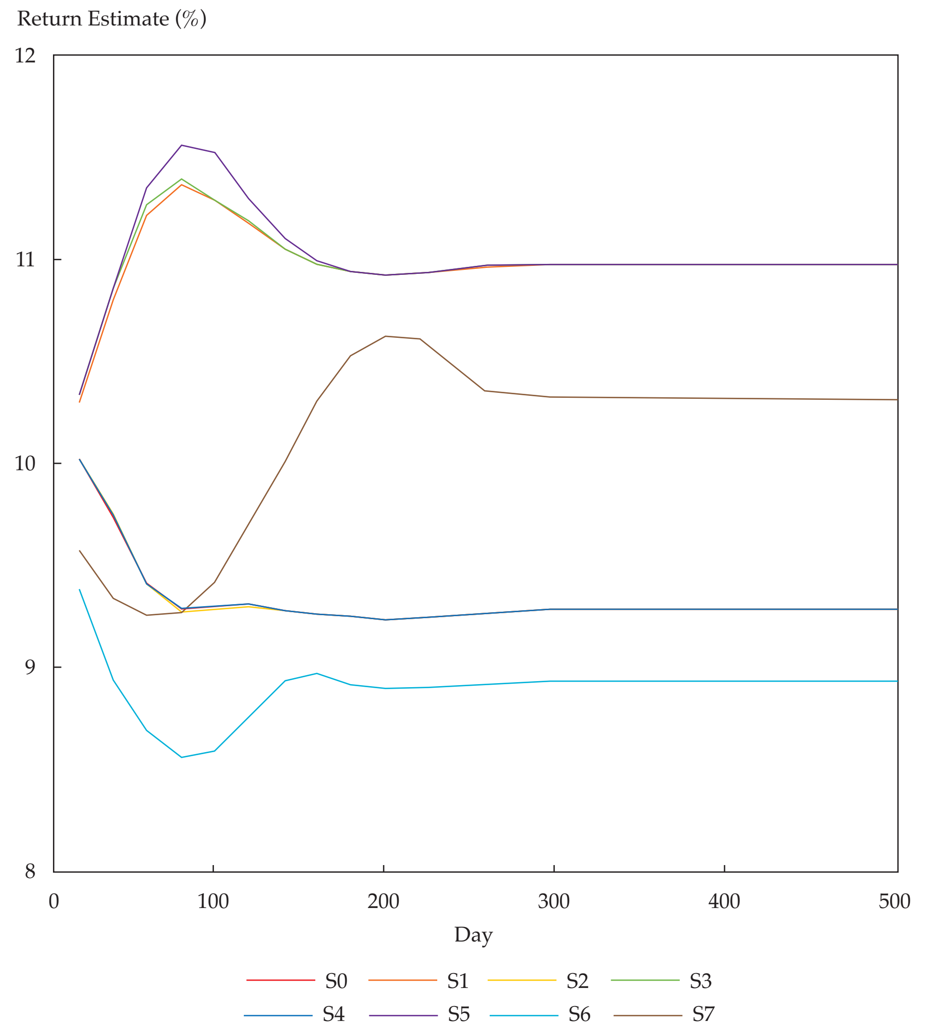

The estimated equilibrium expected returns at the end of the run are the final results in the CME mode. These returns are consistent with the given market portfolio and the given covariance matrix. Furthermore, they derive from a simulation with realistic assumptions regarding limits on investors' ability to borrow.

Conclusion

With asynchronous discrete-time simulation models, researchers can create realistic dynamic models of the market. These realistic models can help test the effects on securities' prices of such real-world events as changes in investment strategy, regulatory policy, levels of passive portfolio management, leverage, capital gains taxation, and circuit breakers. In addition, the impact of such institutional structures as minimum tick sizes and the use of crossing networks can be investigated.

With the use of an asynchronous discrete-time simulator, we showed that a relatively small proportion of momentum investors can destabilize a market. When the ratio of momentum investors to value investors is low, market prices fluctuate but do not become unstable. As this ratio increases, price volatility increases. When the ratio is large enough, the price of at least one security explodes.

We also showed that traders need anchoring rules. Without anchoring rules, traders who want to buy a security may drive the security's price up to an unrealistically high level; alternatively, traders who want to sell a security may drive the security's price down to an unrealistically low level.

Finally, we showed that asynchronous discrete-time simulation can be used to compute CME returns (i.e., implied equilibrium expected returns) under realistic constraints on borrowing and short selling that would make the same problem analytically intractable.

This article qualifies for 1 CE credit.

Notes

1. Merton (1990) surveyed continuous-time models in finance.
2. We believe that asynchronous discrete-time models are superior to synchronous discrete-time models (e.g., the microscopic simulation model of Levy, Levy, and Solomon 2000), which assume that a market equilibrium price is computed each period from the demand and supply curves of all investors on the basis of optimizing or behavioral considerations similar to those that can be incorporated into an asynchronous model. But a world in which prices are set by equilibrium calculations based on all investors' demand and supply curves is different from the real world, in which investors can enter and leave at any time and may or may not find other investors waiting to trade with them.
3. See Jacobs, Levy, and Markowitz (2004). Our asynchronous discrete-time simulator that uses these entities is publicly available on the Jacobs Levy Equity Management website (www.jlem.com/jlmsim).
4. A full set of base case parameters for the DA mode is included in documentation available at www.jlem.com/jlmsim. Levy and Levy (1996) showed the importance of heterogeneity with respect to sample length. Accordingly, one *HIST* statistician used 90 monthly observations for the mean calculation and 250 daily observations for the covariance calculation. Another *HIST* statistician used 60 monthly observations for the mean calculation and 60 monthly observations for the covariance calculation.
5. We defined a month as equal to 20 days.
6. The "flash crash" of 6 May 2010 illustrates the importance of anchoring rules. On that day, the Dow Jones Industrial Average plunged 1,000 points and then rebounded in a matter of minutes. Some computerized trading algorithms that lacked anchoring rules sold stocks like Accenture down to a level of 1 cent before prices rebounded.
7. Assigning only a small number of representative investors to each template is feasible in this case and saves computation time.
8. Further details concerning the input parameters we used for this case study can be found at www.jlem.com/jlmsim.
9. Although the chosen calendar interval was arbitrary, we equated each iteration to one day.

References

Black, Fischer, and Robert Litterman. 1992. "Global Portfolio Optimization." *Financial Analysts Journal*, vol. 48, no. 5 (September/ October):28–43.

Jacobs, Bruce I. 1999. *Capital Ideas and Market Realities: Option Replication, Investor Behavior, and Stock Market Crashes.* Malden, MA: Blackwell.

———. 2004. "Risk Avoidance and Market Fragility." *Financial Analysts Journal*, vol. 60, no. 1 (January/February):26–30.

———. 2009. "Tumbling Tower of Babel: Subprime Securitization and the Credit Crisis." *Financial Analysts Journal*, vol. 65, no. 2 (March/April):17–30.

Jacobs, Bruce I., Kenneth N. Levy, and Harry M. Markowitz. 2004. "Financial Market Simulation." *Journal of Portfolio Management*, vol. 30, no. 5 (30th Anniversary):142–152.

———. 2005. "Portfolio Optimization with Factors, Scenarios, and Realistic Short Positions." *Operations Research*, vol. 53, no. 4 (July/August):586–599.

———. 2006. "Trimability and Fast Optimization of Long–Short Portfolios." *Financial Analysts Journal*, vol. 62, no. 2 (March/April):36–46.

Kim, Gew-rae, and Harry M. Markowitz. 1989. "Investment Rules, Margin, and Market Volatility." *Journal of Portfolio Management*, vol. 16, no. 1 (Fall):45–52.

Levy, Haim, and Harry M. Markowitz. 1979. "Approximating Expected Utility by a Function of Mean and Variance." *American Economic Review*, vol. 69, no. 3 (June):308–317.

Levy, Moshe, and Haim Levy. 1996. "The Danger of Assuming Homogeneous Expectations." *Financial Analysts Journal*, vol. 52, no. 3 (May/June):65–70.

Levy, Moshe, Haim Levy, and Sorin Solomon. 2000. *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena*. San Diego: Academic Press.

Markowitz, Harry M. 2005. "Market Efficiency: A Theoretical Distinction and So What?" *Financial Analysts Journal*, vol. 61, no. 5 (September/October):17–30.

Merton, Robert C. 1990. *Continuous-Time Finance.* Cambridge, MA: Blackwell.

Michaud, Richard O. 1998. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Cambridge, MA: Harvard Business School Press.

Sharpe, William F. 1963. "A Simplified Model for Portfolio Analysis." *Management Science*, vol. 9, no. 2 (January):277–293.

Tobin, James. 1958. "Liquidity Preference as Behavior Towards Risk." *Review of Economic Studies*, vol. 25, no. 2 (February):65–86.