Forecasts of Future Prices, Unbiased Markets, and "Martingale" Models

Benoit Mandelbrot

Journal of Business, Volume 39, Issue 1, Part 2: Supplement on Security Pricing (Jan., 1966), 242-255.

FORECASTS OF FUTURE PRICES, UNBIASED MARKETS, AND "MARTINGALE" MODELS*

BENOIT MANDELBROT†

I. INTRODUCTION AND SUMMARY OF EARLIER INVESTIGATIONS

THE behavior of speculative prices has always been a subject of extreme interest. In most past work, including [6], the emphasis has been on the statistical behavior of price series themselves. The present paper will attempt to relate the behavior of prices to more fundamental economic "triggering" quantities. This effort will constitute a simplified but detailed application of certain ideas current in economic theory, concerning the roles of anticipation and of expected utility.

My findings will depend upon both the behavior of the underlying "triggering" variable and the relationship between the "triggering" variable and price. It is possible to develop models where the price series follows a pure random walk. On the other hand, it is also possible to conceive of models where successive price changes are *dependent* so that prices do not follow a pure random walk, but where the nature of the dependence is such that it cannot be used to increase expected profits. In the terminology of probability, this is expressed by calling price a "martingale." Before exploring these intriguing possibilities, however, it is appropriate to begin with a brief review of the current state of affairs in the field.

In examining prices alone, one automatically chooses to consider all other quantities as being unknown, and their effects on the development of the price series $Z(t)$ as being random. The stochastic mechanism that will generate the future values of Z may, however, depend upon its past and present values. Insofar as the prices of securities or commodities are concerned, the strength of this dependence has long been of concern to market analysts and certain academic economists, and remarkably contradictory conclusions have evolved.

Among the market analysts, the technicians claim that a speculator can considerably improve his prospects of gain by interpreting correctly certain telltale "patterns" that a skilled eye can help him extract from the records of the past. This naturally implies that the future development of $Z(t)$ is greatly, though not exclusively, influenced by its past, and also that different traders, concentrating upon different portions of the past record, should make different estimates of the future price $Z(t + T)$.

Some academic economists, on the other hand, like to emphasize that, even if successive price changes were generated by tossing a fair coin, there would be

* An earlier version, entitled "Speculative Prices and a 'Martingale' Stochastic Model," was privately circulated in February, 1964, and was presented December 29, 1964, to the Annual Winter Meetings of the Econometric Society, held in Chicago. I am greatly indebted to Professor Eugene F. Fama, with whom I repeatedly discussed the draft of this paper, and who helped clarify the presentation of my ideas on this subject; the present text also incorporates changes suggested by the comments of Professor Charles Ying.

† Research staff member, IBM Watson Research Center, Yorktown Heights, New York, and Institute Lecturer, Massachusetts Institute of Technology, Cambridge, Massachusetts.

spurious "patterns" in the price series. One should therefore expect that more elaborate probabilistic generating mechanisms could account for some other patterns as well, and possibly even for all patterns. As a result, the basic attitude of economists is that the significance of any pattern must always be evaluated in the light of some stochastic model. Economists also tend to be skeptical of systematic trading schemes (see, e.g., Sec. VI of [6], and its continuation in [4]).

The study of such stochastic models of price behavior is not at all new, since many of the fundamental techniques of random processes happen to have been first considered in the context of economics, in 1900, by Louis Bachelier [1]. However, it is only recently that one has begun to feel the influence of [1] upon the theory of price variation and upon its empirical verification. (See [2] for a useful collection of references.)

Bachelier conceived several models, each with a different level of complexity. His most general and least developed model is that the present price is an unbiased estimator of the price at any moment in the future. Bachelier's second-level model asserts that, whichever t and T, the increment $Z(t+T) - Z(t)$ is independent of the values of Z up to and including time t; this assumption is best referred to as the "random walk." As to the third level, which was the only level to be fully developed, it asserts that $Z(t+T) - Z(t)$ is a Gaussian random variable with zero mean and a variance proportional to T ("Brownian motion").

This Gaussian model is, however, clearly contradicted by the facts [6]. Thus, from the viewpoint of the study of each speculative time series as if it were alone in the world, it is natural to proceed to independence of successive price increments combined with a distribution other than the Gaussian. In searching for a generalization, it would be desirable to preserve in part a basic "self-similarity" property of the normal law: If one divides by $T^{-1/2}$ the increment of a Gaussian process over a time increment T, one obtains an expression $T^{-1/2}[Z(t+T) - Z(t)]$ that has a distribution independent of T. The generalization of this property was discovered by Paul Lévy, who showed that, if $0 < \alpha < 2$ and the homogeneity exponent $\frac{1}{2}$ is replaced by the larger exponent $1/\alpha$, one obtains a family of probability laws now called "stable Paretian." They have the property that their population variance is infinite; this feature was felt by some to be shocking, but it does not make them "improper" in describing natural phenomena (see [5]).

In my paper, "The Variation of Certain Speculative Prices" [6], it is demonstrated that a model of independent stable Paretian price increments accounts surprisingly well for many properties of the extremely long price series that are provided by market records. This work, and its continuation and elaboration by Fama [3], has encouraged me to strive for an even better model.

The marginal distribution of price changes will be Paretian, but now the increments need not be independent, as was assumed by the "random-walk" assumption of [6].[1] The sample variation of price will exhibit a variety of striking "patterns," but these could be of no benefit to the trader, on the average. Another feature of the price series to be described is that they are generated by an explicit economic model.

[1] To be perfectly honest, an assumption of independence will creep in by the back door, through the hypotheses that will be made concerning the intervals between changes of weather. It would be easy to make less specific probabilistic assumptions, but very hard to carry out their implications.

The stochastic processes $Z(t)$ to be examined are special cases of "martingales," a concept that will now be defined: Let t, $t + T$, and t_i^0 designate, respectively, the present instant of time, a future instant, and an arbitrary set of past instants. $Z(t)$ will be a martingale if $E[Z(t + T)$, given the values of $Z(t)$ and of all the $Z(t_i^0)] = (t)$.

One particular feature of this definition is that one has $E[Z(t + T)$, given the value of $Z(t)] = Z(t)$. However, *much more* is implied in the martingale equality, namely, the statistical independence between future anticipations and past values of Z: Thus, to define a martingale, one may begin by postulating that it is possible to speak of a single value for $E[Z(t + T) \mid Z(t)]$, without having to specify by which *past* values this expectation is conditioned. In a later stage, one will add the postulate that $E[Z(t + T) \mid Z(t)] = Z(t)$. This two-stage definition should underline the central role that martingales are likely to play in the problem to which the present work is devoted: that of the usefulness of a knowledge of past prices for purposes of forecasting.

It should also be stressed that the *distribution* of $Z(t + T)$, conditioned by known values of $Z(t)$ and of the $Z(t_i^0)$, may very well depend upon the past values $Z(t_i^0)$: the expectation alone is unaffected by the $Z(t_i^0)$.

The application of martingales to price behavior gives meaning to the loose idea that prices are somehow "unbiased." That idea goes back at least to Bachelier, in whose mind "unbiasedness" meant that price determination in active speculative markets is governed by a linear utility function.[2]

Interest in martingales among pure probabilists is such that an immense variety of martingale processes has been described. If we dealt with a single economic series, namely the price, the choice among all this wealth of possibilities could only be directed by purely mathematical criteria—which are a notoriously poor guide. Hence, the present step beyond the random walk was only undertaken within the context of a "fundamental analysis," in which the price attempts to follow something that can be described as "value": That is, the present price $Z(t)$ is a function of past prices, and of the past and present values of the exogenous trigger $Y(t)$. In the present paper, the process generating value will be such that, as T increases, the expectation of $Y(t + T)$ will tend fairly rapidly toward a limit. If that limit is taken as the present price $Z(t)$, one will achieve two things: (1) price and value will occasionally coincide; (2) price will be generated by a martingale stochastic model in which the present $Z(t)$ is an unbiased estimator of $Z(t + T)$; moreover, for large enough values of T, $Z(t)$ is an unbiased estimator of $Y(t + T)$.

If, however, the process generating Y had other properties, the forecast future value $E[Y(t + \text{infinity})]$ need not be a martingale. An example to the contrary is given in Section IIG. Therefore, the fact that forecasting of the value leads to

[2] Let us now consider, however, some non-linear function F of the price; the expectation of $F[Z(t + T)]$ will *not* be in general equal to the present value $F[Z(t)]$. This means that, if our speculator's private utility function is not linear in Z, playing on Z may be advantageous or disadvantageous for him. Moreover, individual speculators need not be led by the same utility as the market considered as a whole: They may, e.g., either seek or avoid a large dispersion of possible future prices $Z(t + T)$. Even in the case of a martingale, an increasingly detailed knowledge of the past may be useful for such purposes.

Similarly, if log Z is a martingale, playing on Z will be advantageous to speculators having a linear utility function. The fact that unbiasedness is linked to a choice of scale for Z is well known to mathematical statisticians.

a martingale in the prices tells us something about the structure of the value as well as the structure of the market mechanism. If forecasted value does not follow a martingale, prices could follow a martingale only if they do not follow value.

The above considerations are linked with the oft-raised question of whether one can divide the speculators into several successive groups, as follows: the members of the first group know only the present and past values of Z; the members of the next group also know the present and past values of the single series Y, and know how the price will depend upon the variation of Y; the members of further groups also know the temporal evolution of various series that contribute to Y, and again know how these series affect the price. In the model of Y to which the present paper will be devoted, there is no advantage, on the average, in knowing anything beyond the present $Z(t)$.[3]

II. THE FORECASTING FUNCTION OF EXCHANGE MARKETS AND THE PERSISTENCE OF PRICE MOVES FOR AGRICULTURAL COMMODITIES

A. THE PROBLEM

The present section will be devoted to the series of *equilibrium prices* for an agricultural commodity. Consideration of fluctuations around this series, due to temporal scatter of supply and demand, will be postponed to Section III. Here, the price $Z(t)$ will be equal to the expected value of the future crop, which in turn only depends upon past and future weather, according to the following rules: (1) weather can only be good, bad, or indifferent; (2) one is only interested in deviations of the price from some "norm," so that it is possible to neglect the price effects of indifferent weather; (3) when there were g good days and b bad days between the moments t' and t'' within the growing season, the size of crop will have increased by an amount proportional to $b - g$; (4) under the conditions of rule (3), the "value" $Y(t)$ of a unit quantity of the crop will have decreased by an amount proportional to $b - g$. The final rule (5) runs as follows: At any instant t, there is a single price of a unit quantity for future delivery, equal to

$$\lim_{T\to\infty} E[Y(t+T)].$$

Units will be assumed so chosen that price will increase by 1¢ when the ultimate expected value $Y(t)$ increases by the effect upon the crop of one day's bad weather. These rules are very simplified,[4] and they do not even take into account the effect upon future prices of the portions of past crops that are kept in storage. The total problem is so complex, however, that it

[3] Martingales are naturally closely related to other techniques of time-series analysis that involve conditional expectations, such as regression theory, correlation theory, and spectral representation. In particular, if $Z(t)$ is a martingale, its derivative is spectrally "white" in the sense that the covariance $C(\tau)$ between $Z'(t)$ and $Z'(t+\tau)$ vanishes if $r \neq 0$. The expected value of the sample spectral density of $Z'(t)$ will therefore be constant, that is, independent of frequency. If a market can associate such a series $Z(t)$ to the exogenous $Y(t)$, this market can be called a "whitener" of the derivative $Y'(t)$. However, one must keep in mind that spectral methods are concerned with measuring correlation rather than dependence. Spectral whiteness expresses lack of correlation, but it is *not* synonymous with full independence, except when the joint distributions of prices at different times are Gaussian, which is surely not the case for the examples I constructed for this paper. In fact, whiteness is even weaker than the martingale identity.

[4] It is really acknowledged that they would have been made much more realistic if they referred to log Z instead of Z, and similarly to the logarithms of other quantities. This transformation was avoided, however, in order to avoid burdening the notation. The interested reader will readily make the transformation by himself.

is best to begin by following up each of its aspects separately.

Our rational forecast of Y naturally depends upon the weather forecast, i.e., upon the past of Y, the probability distribution of the lengths of the weather runs, and the rules of dependence between the lengths of successive runs.

The crudest assumption is to suppose that the lengths of the runs of good, bad, or indifferent weather are ruled by statistically independent exponential variables—as is the case if weather on successive days is determined by independent random events. Then the future discounted knowing the past is exactly the same as the future discounted *not* knowing the past; in particular, if good and bad days are equally probable, the discounting of the future will not change the prices based upon the present crop size. This means that the process ruling the variation of $Z(t)$ is the simplest random walk, with equal probabilities for an increase or decrease of price by 1¢.

Our "intuition" about the discounting of the future is of course based upon this case. But it is not necessary that the random variable U, designating the length of a good or bad run, be exponentially distributed. In all other cases, some degree of forecasting will be possible, so that the price will be influenced by the known structure of the process ruling the weather. The extent of this influence will depend upon the conditional distribution of the random variable U, when it is known that $U \geq h$. The following subsection will therefore be devoted to a discussion of this problem.

B. THE DISTRIBUTION OF RANDOM VARIABLES CONDITIONED BY TRUNCATION

Exponential random variables.—To begin with, let us note that the impossibility of forecasting in the exponential case can be restated as being an aspect of the following observation: Let U be the exponential random variable for which $P(u) = Pr(U \geq u) = \exp(-bu)$, and let $U(h)$ designate the conditioned random variable: U, conditioned by $U \geq h > 0$. Bayes's theorem then yields the following results: If $u < h$, one has $Pr[U(h) \geq u] = 1$; if $u > h$, one has

$$Pr[U(h) \geq u] = Pr(U \geq u \mid U \geq h) = \frac{\exp(-bu)}{\exp(-bh)} = \exp[-b(u-h)].$$

This means that $U(h) - h$ is a random variable independent of h, but having a mean value $1/b$ determined by the original scale of the unconditioned U.

Hyperbolic random variables.—Assume now that U follows the hyperbolic law, made familiar by Pareto's income distribution, which depends upon the two positive parameters σ and α as follows: If $u < \sigma$, one has $Pr(U \geq u) = 1$; if $u > \sigma$, one has $Pr(U \geq u) = (u/\sigma)^{-\alpha}$. In the present case, Bayes's theorem yields the following results: If $h < \sigma$, one has $Pr[U(h) \geq u] = Pr(U \geq u)$; if $\sigma < u < h$, one has $Pr[U(h) \geq u] = 1$; finally, if $\sigma < h < u$, one has

$$Pr[U(h) \geq u] = Pr(U \geq u \mid U \geq h) = \frac{(u/\sigma)^{-\alpha}}{(h/\sigma)^{-\alpha}} = (u/h)^{-\alpha}.$$

It is clear that the various typical values of $U(h)$, such as all quantiles or the expectation, are proportional to h. For example, $hq^{-1/\alpha}$ gives the value of $U(h)$ that is exceeded with the probability q. As to the mean of $U(h)$, it is finite only if $\alpha > 1$; in that case, one has

$$E[U(h)] = \int_h^\infty \alpha h^\alpha u^{-\alpha} du = \frac{\alpha h}{(\alpha - 1)}.$$

Thus, $E[U(h) - h] = h/(\alpha - 1) = E[U(h)]/\alpha$, which is greater or smaller than h according to whether α is smaller or greater than 2; if $\alpha = 2$, $E[U(h) - h] = h$.

As to the marginal probability that $h < U < h + dh$, knowing that $h < U$, it is equal to $\alpha h^{-(\alpha+1)}dh/h^{-\alpha} = \alpha dh/h$, which decreases with h.[5]

An important property of the present conditioned or truncated variable $U(h)$ is that it is *scale-free* in the sense that its distribution does not depend upon the original scale factor σ. One may also say that the original hyperbolic law is self-similar. Conversely, this property characterizes the hyperbolic law,[6] and is very systematically exploited in my studies of various empirical time series and spatial patterns. In particular, it turns out that runs whose duration is hyperbolically distributed provide a very reasonable approximation to the "trend" component of a number of meteorological time series, and this is, of course, the motivation of their use in the present context.

C. PRICES BASED UPON A FORECAST CROP SIZE

With the above preliminaries in mind, let us resume the crop-forecasting problem raised in Section IIA, assuming that the lengths of successive weather runs are statistically independent[7] random variables following Pareto's hyperbolic law. It is clear that a knowledge of the past now becomes useful in predicting the future. The results become especially simple if one modifies the process slightly to assume that weather alternates between "passive runs" of indifferent behavior, and "active runs" when it can be good or bad with equal probabilities. Then, as long as one is anywhere within a "passive run," prices will be unaffected by the number of indifferent days in the past. But if there have been h good or bad days in the immediate past, the same weather is likely to continue for a further period which has $h/(\alpha - 1)$ as its mean.[8]

[5] In order to fully assess the above findings, it is good to contrast them with the result valid in the Gaussian case. As a simplified intermediate case, consider the random variable U for which $Pr(U \geq u > 0) = \exp(-bu^2)$. Then the arguments developed above show that, for $u > h$, one has

$$\begin{aligned} Pr[U(h) \geq u] &= Pr(U \geq u \mid U \geq h) \\ &= \exp[-b(u^2 - h^2)] \\ &= \exp[-b(u + h)(u - h)]. \end{aligned}$$

It follows that all the typical values of $U(h) - h$, such as the expected value or the quantiles, are smaller than the mean and the quantiles of an auxiliary exponential variable $W^0(h)$ such that $Pr[W^0(h) \geq w] = \exp(-2hb\,w)$; this shows that the mean of $U(h) - h$ is smaller than $1/2hb$, and it therefore tends to zero as h tends to infinity.

Things are very similar in the Gaussian case, but the algebra is complicated and need not be given here.

[6] It definitely necessitates that the ratio $Pr(X \geq u)/Pr(X \geq h)$ be the same when X is the original variable U or the variable U divided by any positive number k. For this condition to be satisfied, the function $P(u) = Pr(U \geq u)$ must satisfy $P(u)/P(h) = P(ku)/P(kh)$. Let $R = \log P$ be considered as a function of $v = \log u$; the above requirement can then be written as

$$R(v) - R(v^0) = R(v + \log k) - R(v^0 + \log k).$$

This means that $R = \log P$ must be a linear function of $v = \log u$, which is, of course, the definition of the hyperbolic law through doubly logarithmic paper, in the manner of Pareto.

[7] See n. 1.

[8] Things are actually slightly more involved: A positive hyperbolic random variable must indeed have a minimum value $\sigma > 0$; therefore, after an active run has started, its expected future length jumps up to $\sigma/(\alpha - 1)$ and stays there until the actual run length has exceeded σ. Such a fairly spurious jump will also appear in the exponential case if good weather could not follow bad weather, and conversely. One can in fact modify the process so as to eliminate this jump in all cases, but this would greatly complicate the formulas for a small profit.

It is also interesting to derive the forecast value of $E[Y(t + T) - Y(t)]$, when T is finite and the instant

Recall then that the crop growth due to one day of good weather decreases the price by 1¢. A good day following h other good days will then decrease the price by the amount $[1 + 1/(\alpha - 1)]$¢ $= [\alpha/(\alpha - 1)]$¢, in which the $1/(\alpha - 1)$¢ portion is due to revised future prospects. But, when good weather finally turns to "indifferent," the price will go *up* by $h/(\alpha - 1)$, to compensate for unfulfilled fears of expected future bounties. It should be noted that $h/(\alpha - 1)$ *is not* a linear function of the known past values of Y. This implies that the best linear forecast is not optimal.

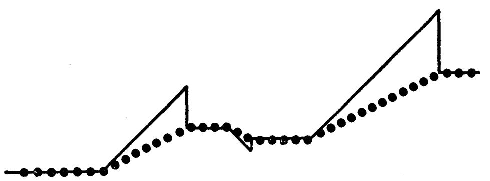

FIG. 1.—For both the dotted and the bold lines, the abscissa is time. For the dotted line, the ordinate is $Y(t)$; for the bold line, the ordinate is $Z(t)$.

As a result, the record of the prices of our commodity will appear as a random alternation of three kinds of period, to be designated as "flat," "convex," and "concave," and defined as follows: During flat periods, prices will vary little and "aimlessly." During concave periods, prices will go up by small equal amounts every day, except that on the last day of the period they will fall by a fixed proportion of their total rise within the whole period. Precisely the opposite behavior will hold for convex periods.

---

t is the hth instant of a bad weather run: One readily finds that

$$(1/h)E[\,Y(t+T) - Y(t)\,]$$
$$= \alpha(\alpha - 1)^{-1}[1 - (1 + T/h)^{1-\alpha}] - 1 .$$

This shows that the convergence of $E[Y(t + T)]$ to its asymptote is fast when h is small, and slow when h is large.

Examples of these three kinds of periods have been shown on the process illustrated by Figure 1. If a run of good weather is interrupted by a single indifferent day, the pattern of prices will be made up of a "slow fall, a rapid rise, a slow fall, and a rapid rise." Up to a small day's move, the point of arrival will be the same as if there had been no indifferent day in between; but that single day will "break" the expectations sufficiently to prevent prices from falling as low as they would have done in its absence.

When one nears the end of the growing season, the above forecasts should naturally be modified to avoid discounting the weather beyond the harvest. If the necessary correction is applied, one will find that the final price will precisely correspond to the crop size, determined by the difference between the number of days of good and bad weather. However, the corresponding corrections will not be examined here.

D. THE MARTINGALE PROPERTY OF FORECASTED PRICES

The random series $Y(t)$ is *not* a martingale. To prove this fact, it suffices to exhibit one set of past values of Y for which the martingale property is not verified. We shall show that the conditioned expectation,

$$E[Y(t + 1) - Y(t), \text{ knowing the number } h \text{ of past good days}],$$

is non-vanishing.

PROOF: $Y(t + 1) - Y(t) = 0$ if and only if the run of good weather breaks today, an event of probability $[h^{-\alpha} - (h + 1)^{-\alpha}]/h^{-\alpha} \sim \alpha/h$. Otherwise, $Y(t + 1) - Y(t) = -1$. Thus, the expectation of $Y(t + 1) - Y(t)$ equals the probability that $Y(t + 1) - Y(t) = -1$, which is $1 - \alpha/h$, a non-vanishing function of the past weather (whose history is fully represented for our present purposes by the duration of the current good weather run).

On the contrary, *the price series* Z(t) *is a martingale*. To begin with, let us assume that h is known and evaluate the increment

$$E[Z(t+1)\text{, given the value of } Z(t) \text{ and given that the number of preceding good days was exactly } h] - Z(t)\,.$$

Let h be sufficiently large to avoid the difficulties due to existence of a lower bound to the duration of a weather run. Then, if today's weather continues good, price will go down by an amount equal to $\alpha/(\alpha-1)$; we saw that this event has a probability equal to $(h+1)^{-\alpha}/h^{-\alpha} \sim 1 - \alpha/h$. Suppose, on the contrary, that the current day is indifferent; this fact alone implies that the good-weather run is over and that the advance discounting of the effect of future weather was unwarranted. As a result, the price will climb up abruptly by an amount equal to $h/(\alpha - 1)$; this event has a probability equal to α/h. The expected price change is thus approximately

$$\left(1-\frac{\alpha}{h}\right)\frac{\alpha}{\alpha-1}-\frac{\alpha}{h}\,\frac{h}{\alpha-1}=\frac{1}{h}\,\alpha^2(\alpha-1)\,,$$

which is approximately zero. The rigorous derivation of the expected price change is more involved, but its result is simpler, namely: the expected price change is exactly (and not just approximately) equal to zero.

Let us now take account of the fact that one's actual knowledge of the past is usually not represented by the value of h but by some past values $Z(t_i^0)$ of $Z(t)$. The number of good days in the current run is then a random variable H, and $D(h) = Pr\ (H < h)$ is a function determined by the $Z(t_i^0)$. It follows that

$$E[Z(t+1)\text{, given the value of } Z(t) \text{ and the past prices } Z(t_i^0)] = \int dD(h)E[Z(t+1)\text{, given the value of } Z(t) \text{ and the value of } h] = \int dD(h)Z(t) = Z(t)\,.$$

One shows very similarly that $E[Z(t+T)] = Z(t)$ when T exceeds 1, so that Z(t) *is indeed a martingale process*.

Variance of Z(t + 1) − Z(t).—If the number of preceding good days was exactly h, this variance is equal to

$$\left(1-\frac{\alpha}{h}\right)\left(\frac{\alpha}{\alpha-1}\right)^2+\frac{\alpha}{h}\left(\frac{h}{\alpha-1}\right)^2,$$

which becomes proportional to h when h is large. If h is not a known number, but is generated by a random variable H, conditioned by some known past prices $Z(t_i^0)$, the variance of $E(t+1) - Z(t)$ is proportional to $E(H)$.

(If no past price is known, and $1 < \alpha < 2$, one can show that $E(H)$ is infinite, and one falls back upon the infinite-variance property of the random walk with stable Paretian increments—my original model [6].)

Comment.—We have now, in a sense, reached the climax of this story, and it may be good to comment again upon some observations made in the Introduction. If the price Z were generated by a random walk, then, whichever measure of risk has been adopted, no knowledge of the past should influence one's estimate of the risks involved in trading in Z. If, on the contrary, Z is generated by the present martingale, then the only "risk" that is not influenced by the past is constituted by the expectation of Z. A martingale is thus a "fair game." But, as h increases, so do the expected deviations from the expectation of $Z(t+1)$ and so do all other measures of "risk." This was to be expected, since, as h increases, so does the relative contribution to Z of anticipated changes in Y. Clearly, all risk-seeking and risk-avoiding traders will want to know how the market value of a crop is divided between its present value and the changes anticipated before harvest time!

Note also the following: If risk-avoiders exceed risk-seekers in influence on the market, the martingale equality should be replaced by $E[Z(t+T)] - Z(t) > 0$, the difference increasing with the variance of $Z(t+T)$. This would have as a consequence that prices would increase in time, on the average, especially during the periods of high variance. However, this "tendency toward price increase" would be of significance only for traders who seek risk more than does the average trader on the market.

E. THE DISTRIBUTION OF PRICE CHANGES

This distribution is symmetric, and it will suffice to derive it when $\Delta Z = Z(t+1) - Z(t)$ is positive or zero. It will be useful to designate by W' the mean duration of an indifferent weather run, and by W'' the mean duration of a good or bad weather run. Moreover, it will be assumed to simplify that W' and W'' are both large when measured in days.

The most significant price changes are those that satisfy $\Delta Z > \alpha/(\alpha - 1)$. They occur only on the last days of good weather runs, so that their total probability is $1/2\ (W' + W'')$. Their precise distribution is obtained by simply rescaling the law ruling the duration of good weather runs. One has therefore: For $z > \alpha(\alpha - 1)$,

$$Pr(\Delta Z \geq z) = \left[\frac{z+1}{\alpha/(\alpha-1)}\right]^{-\alpha} \times \frac{1}{2(W'+W'')}\ .$$

Next, consider the probability that $Z' = 0$. This event occurs when t is anywhere within a run of indifferent weather, so that its probability is $W'/(W' + W'')$.

Finally, $\Delta Z = \alpha/(\alpha - 1)$ when the instant t is within a bad weather run but is not the last instant in that run. This event has the probability $(W'' - 1)/2(W' + W'')$.

The over-all distribution of daily price change is thus a "bell" with two Paretian tails. It is shaped very much like a stable Paretian law; in this sense, the present model may be considered as providing a further elaboration of the process first proposed in my [6].

It is now safe to mention that the martingale property of forecast prices holds independently of the distribution $P(u)$ of bad weather runs, as long as the runs are statistically independent. However, any non-hyperbolic form for $P(u)$ would predict a marginal distribution of price change that is in conflict with the evidence brought forth in [6].

F. A MORE INVOLVED AGRICULTURAL COMMODITY

Although still very crude, the preceding model seems more realistic than could have been expected. A further touch may be added by taking account of the possibility of crop destruction by a natural calamity, such as hail. I have found that at least some among natural calamities have Paretian distributions. The extent of such calamities is presumably known only gradually, and they may therefore give rise to "patterns" similar to those we have studied above. The main interest of a mixture of several exogenous variables is, however, that it is unrealistic to believe that there is proportionality between the distribution of large price changes and that of the time intervals between them. Such a proportionality holds in the case of the single trigger $Y(t)$, but not in the case of many triggers.

G. BEST LINEAR FORECASTS CANNOT BE USED TO DEFINE PRICES

I shall state without proof some results that can be omitted without interrupting the continuity of the present work. Let us suppose that, instead of being ruled by the process $Y(t)$ that we have described, the value is ruled by a process $Y^*(t)$ with the following properties: $\Delta Y^*(t) = Y(t+1) - Y(t)$ is a stationary Gaussian process whose covariance function is equal to that of $\Delta Y(t)$. Then the best extrapolate $E^*[Y(^*t + T)$, knowing $Y^*(t)$ for $s \leq t]$ is linear. It is also identical to the best linear extrapolate of $Y^*(t + T)$. As $T \rightarrow \infty$, this extrapolate tends to infinity and therefore it cannot be used to define a price series $Z(t)$.

The above example suffices to show that, in order that price follow a martingale process, it does not suffice that price be based upon a forecast of value.

III. PERSISTENCE OF PRICE MOVES FOR INDUSTRIAL SECURITIES

A. FIRST APPROXIMATION

The arguments of Section II can be directly translated into terms of "fundamental analysis" of security prices. Suppose, indeed, that the market value of a corporation is equal to the expected value of its future size X, computed taking account of current and past values of its size $Y(t)$. If the rules of growth are of the form that we shall presently describe, it is meaningful to specify "the" expected future size by a single number, independent of the moment in the future to which one refers, and independent of the elements of the past history available for forecasting. The resulting theory is again greatly simplified (note also the omission of all reference to current yield).

Our rules of growth are such that the lengths of periods of growth and decline are random, independent, and Paretian. Then, the longer a company has grown straight up, the more the outsiders should justifiably expect that it will grow in the future. Its market value Z should therefore justifiably increase by the multiple $1 + 1/(\alpha - 1)$ of any additional growth actually observed for $Y(t)$. If, however, the growth of Y ever stops sufficiently for all to know it, one should observe a "break of confidence" and a fall of Z justifiably equal to the fraction $1/\alpha$ of the immediately preceding rise. If the growth of Y is stopped by "breathing spells," the growth of Z will have a sawtooth pattern. If a long growth period of Y, ending on a breathing spell, is modified by adding an additional intermediate breathing spell, the ultimate value of the company would be unchanged. But a single big tooth of Z would be replaced by two teeth, neither of which attains equally dizzy heights. In the absence of breathing spells, price can go up and up, until presumably the discounted future growth would have made this corporation bigger than the whole economy of its country, necessitating corrections that will not be examined in this paper.

Most of the further developments of this model would be very similar to those relative to the commodity examined in Section IIC. There is, however, a difference in that, if α is small, the expected length of the further growth period may be so long that one may need to discount the future growth at some non-vanishing rate.

B. SECOND APPROXIMATION

Let us now examine the case of an industrial security whose fundamental value $X(t)$ follows a process of independent increments: either Bachelier's process of independent Gaussian increments, or the process in which the increments are stable Paretian [6]. In both models, the rate of change of X may sometimes be very rapid; in the latter model it may even be instantaneous. But it will be assumed that the market only follows $X(t)$ through a smoothed-off form $Y(t)$ for which the maximum rate of change is fairly large but finite. (In some cases, the establishment of an upper bound u^* to the changes of Y may be the consequence of deliberate attempts to insure "market continuity.")

In order to avoid mathematical complications, let time be discrete,[9] and the maximum rate of change u^* be known. It is clear that, whenever the market observes that $Y(t) - Y(t-1) < u^*$, it will be certain that there was no smoothing off at time t and that $X(t) = Y(t)$. If u^* is large enough, the equality $X = Y$ will hold for most values of t. Thus the market price $Z(t)$ will be equal most of the time to the fundamental value $X(t)$. Every so often, however, one will reach a point of time where $Y(t) - Y(t-1) = u^*$, a circumstance that may be due to any change $X(t) - X(t-1) \geq u^*$. At such instants, the value of $X(t) - X(t-1)$ will be greater than the observed value of $Y(t) - Y(t-1)$, and its conditional distribution will be given by the arguments of Section IIB; it will therefore critically depend upon the distribution of $X(t) - X(t-1)$.

If the distribution of the increments of X is Gaussian, and u^* is large, the distribution of $X(t) - X(t-1)$, assuming that it is at least equal to u^*, will be very much clustered near u^* and so will also the distribution of $X(t+1) - X(t-1)$. There will therefore be a probability very close to 1 that $X(t+1) - X(t-1)$ be smaller than $2u^*$ and $X(t+1) - X(t)$ smaller than u^*. As a result, $Y(t+1)$ will equal $X(t+1)$ and $Z(t+1)$ will be matched to $X(t+1)$. In other words, the mismatch between Z and Y will be small and short-lived in most cases.

Suppose now that the distribution of ΔX has two Paretian tails with $\alpha < 2$. If $Y(t) - Y(t-1) = u^*$, while $Y(t-1) - Y(t-2) < u^*$, one knows that $X(t) - X(t-1) > u^*$ and has a conditional expectation independent of the scale of the original process and equal to $\alpha u^*/(\alpha - 1)$. The market price increment $Z(t) - Z(t-1)$ should therefore amplify, by the factor $\alpha/(\alpha - 1)$, the increment $Y(t) - Y(t-1)$ of the smoothed-off fundamental value. Now proceed to time $t+1$ and distinguish two cases: If $Y(t+1) - Y(t) < u^*$, the market will know that $X(t+1) = Y(t+1)$. Then the price $Z(t+1)$ will equal $X(t+1) = Y(t+1)$, thus falling from the inflated anticipation equal to $X(t-1) + \alpha u^*/(\alpha - 1)$. But if $Y(t+1) - Y(t) = u^*$, the market will know that $X(t+1) - X(t-1) = Y(t-1) + 2u^*$. It follows that the conditioned distribution of the difference $X(t) - X(t-1)$ will be very close to a Paretian law truncated to values greater than $2u^*$. Thus the expected value of $X(t+1)$—which is also the

[9] Continuous-time processes with independent increments were considered in [6]. In the stable Paretian case, one finds that $X(t)$ is significantly discontinuous in the sense that, if it changes greatly during a unit time increment, this change is mostly performed in one big step somewhere within that time. Therefore, the distribution of large jumps and that of large changes over finite time increments are practically identical. In the Gaussian case, on the contrary, the interpolated process is continuous. More generally, whenever the process $X(t)$ is interpolable to continuous time and its increments have a finite variance, there is a great difference between the distributions of its jumps (if any) and of its changes over finite time increments. This shows that the case of infinite variance—which in practice means the Pareto case—is the only one for which the restriction to discrete time is not severe at all.

market price $Z(t+1)$—will be equal to $Z(t-1) + 2u^*\alpha/(\alpha - 1) = Z(t) + u^*\alpha/(\alpha - 1)$.

After $Y(t)$ has gone up, n times in succession, in steps equal to u^*, the value of $Z(t+n-1) - Z(t-1)$ will approximately equal $nu^*\alpha/(\alpha-1)$. Eventually, however, one will reach a value of n such that $Y(t+n-1) - Y(t-1) < nu^*$, which implies that $X(t+n-1) - X(t-1) < nu^*$ and the market price $Z(t+n-1)$ will then crash down to $X(t+n-1)$, losing in one swoop all its excessive growth.

As the size of the original jump of X increases, the number of time intervals involved in smoothing also increases, and correction terms must be added.

Let us now discuss qualitatively the case where the value of the threshold u^* is random. Then, after the market observes a change of $Y(t)$, it will question it to determine whether it is a fully completed change of fundamental conditions, equal to a change of $X(t)$, or the beginning of a large change. In the first case, the motion need not "persist," but it will persist in the second case. This naturally involves a test of statistical significance: A few changes of Y in the same direction may well "pass" as final moves, but a long run of rises should be interpreted as due to a "smoothed-off" large move. Thus, the following more complicated pattern will replace the gradual rise followed by fast fall that was earlier observed: The first few changes of Z will equal the changes of Y, then Z will jump to such a value that its increase from the beginning of the rise equals $\alpha/(\alpha-1)$ times the increase of Y; whenever the rise of Y stops, Z will fall to Y; whenever the rise of Y falters, and then resumes, Z will fall to Y and then jump up again.

In a further generalization, one may consider the case where large changes of Y are gradually transmitted with probability q and very rapidly in other cases. Then the distribution of the changes of Z will be a mixture of the distribution obtained in the previous argument and of the original distribution of changes of Y; however, the Paretian character is preserved in such a mixture. See [5].

C. MORE COMPLEX ECONOMIC MODELS

It is natural to examine the case when there is more than one "tracking" mechanism of the kind examined so far. It may for example happen that $Z(t)$ attempts to predict the future behavior of a smoothed-out form Y of $X(t)$, while $X(t)$ itself attempts to predict the future behavior of some other function $X^*(t)$. This would lead to zigzags larger than those observed so far. Therefore, for the sake of stability, it will be very important in every case to ascertain whether the driving function $Y(t)$ is a smoothed-off fundamental economic quantity or is already influenced by forecasting speculation.

Suppose now that two functions $Z_1(t)$ and $Z_2(t)$ attempt to track each other, with lags in each case. The zigzags will get increasingly amplified, as in the divergent case of the cobweb phenomenon. All this hints at the difficulty of studying in detail the process by which "the market is made" through the interactions among a large number of traders; it also underlines the necessity of making a detailed study of the function of "insuring the continuity of the market" that is assigned to the specialist.

IV. ADDITIONAL COMMENTS

A. THE VALUATION OF OIL FIELDS

The use that was made of the results of Section IIB may also be illustrated by the example of oil fields in a new country.

One knows "intuitively" that there

is a high probability that the total oil reserves in this country are very small; but, if it turns out to be oil rich, its reserves would be very large. This means that the a priori distribution of the reserves X is likely to have a big "head" near $x = 0$ and a long "tail"; indeed, that distribution is hyperbolic. Let us now consider a forecaster who only knows the partial value, $Y(t)$, of the recognized reserves at time t. As long as the reserves have not been completely explored, their expected market value $Z(t)$ should be equal to $\alpha Y(t)/(\alpha - 1)$: The luckier the explorers have been in the past, the more they should be encouraged in digging or the more they should expect to have to pay for digging rights in the immediate neighborhood of a recognized source. Eventually, however, $Y(t)$ will reach X and it will then cease to increase; at this very point, $Z(t)$ will tumble down to $Y(t)$, thus decreasing by $Y(t)/(\alpha - 1)$.

If the distribution of X had been exponential, $Z(t)$ would always exceed $Y(t)$ by an amount independent of $Y(t)$ and equal to the market value of entirely unexplored territory. If $Y(t)$ had been a truncated Gaussian variable, the premium for expected future findings would have rapidly decreased with $1/Y(t)$.

It would be interesting to study actual forecasts in the light of the triple alternative that I have sketched. But the example of oil fields was mainly brought in to further demonstrate how the variation of prices can be affected by certain *unavoidable* delays in the "transmission of information about the physical world."

B. FACETIOUS APPLICATIONS

The simple mathematical result of Section IIB may be used to construct various facetious "paradoxes" related to the concept of expectation; with some trepidations, two examples will now be presented.

First paradox: Why does anyone who stops young stop in the middle of a promising career? According to certain results of A. Lotka, the distribution of the number of scientific papers due to any single author is Paretian with exponent 2. Assume, therefore, for the sake of argument, that a "scientific career" is something the duration of which is unpredictable and hence can be considered as being random and Paretian with exponent 2: This would mean that most people have a very short career but a few have very long careers indeed. Then, however long a career has already proceeded, it should be expected to continue for an equal additional time span. Since all careers eventually stop, they may be considered as breaking off half of their promise. The only way of avoiding such apparent disappointment is to live to be so old that age corrections must be added in the computation of the expected career.

The parable of the receding shore. "Once upon a time, there was a country called the Land of Ten Thousand Lakes, and those landmarks were affectionately known to some of its inhabitants as Biggest, Second Biggest, . . . , rth Biggest, etc., down to 10,000th Biggest. The widest was a sea 100 miles across, the width of the rth Biggest was $100/\sqrt{r}$, so that the smallest had a width of only 1 mile. But each lake was always covered with a haze that made it impossible to see across and thus identify its width. One would, of course, find out if one could discover its name; but the land was poorly mapped and poorly marked, and had few inhabitants whom the traveler could ask for instructions. The people of that land were expert at measuring distances, however; they also knew all about the computation of averages and were great believers in expected utilities. They knew, therefore, that as one of them stood on an unknown shore, he had before him a stretch of water of expected width equal to 2 miles. He could very well travel 1 mile to reach the center of "the" expected lake; but he could never go beyond this point! Suppose, indeed, that he succeeded in sailing forth to a new total distance just short of $100/\sqrt{r}$ miles. In the meantime, the other shore would have "moved on," to a new mean distance from him equal precisely to $100/\sqrt{r}$. It is clear therefore that those Lakes were ruled by Spirits who would never let them be crossed by a stranger. However far the traveler might sail, the Spirits would spread the lake ever farther, and the stranger would always remain right in the middle of water; his boldness should eventually be punished by death, but all travelers were eventually reprieved by a special grace."

REFERENCES

1. Bachelier, Louis. "Théorie de la spéculation" (doctoral dissertation in mathematics, University of Paris, March 29, 1900), *Annales de l'Ecole Normale Supérieure*, Ser. 3, XVII (1900), 21–86. An English translation appears as pp. 17–75 of P. H. Cootner (ed.) [2].
2. Cootner, Paul H. (ed.). *The Random Character of Stock Market Prices*. Cambridge, Mass.: M.I.T. Press, 1964. (This book constitutes a very complete collection of reprints with many references relevant to the topic of the present essay.)
3. Fama, Eugene F. "The Behavior of Stock Market Prices," *Journal of Business*, XXXVIII (January, 1965), 34–105.
4. Fama, Eugene F., and Blume, Marshall E. "Filter Rules and Stock Market Trading," *Journal of Business*, Vol. XXXIX, No. 1, Part II (Supplement, January, 1966).
5. Mandelbrot, B. "New Methods in Statistical Economics," *Journal of Political Economy*, LXXI (1963), 421–40.
6. ———. "The Variation of Certain Speculative Prices," *Journal of Business*, XXXVI (1963), 392–417, reprinted as pp. 307–32 in P. H. Cootner (ed.) [2].