How speculation can explain the equity premium

Glenn Shafer

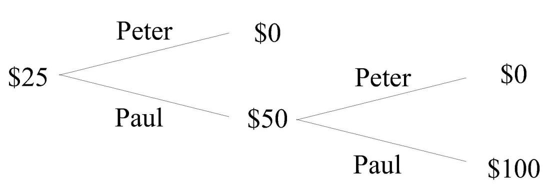

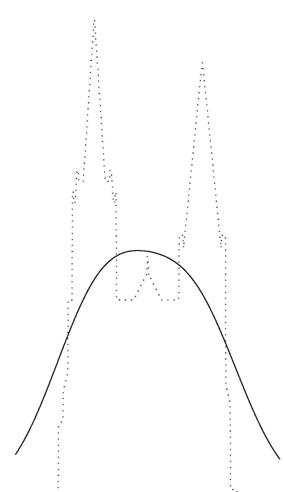

The Game-Theoretic Probability and Finance Project

Working Paper #47

November 9, 2016

Project web site:
http://www.probabilityandfinance.com

Abstract

When measured over decades in countries that have been relatively stable, returns from stocks have been substantially better than returns from bonds. This is often attributed to investors' risk aversion: stocks are thought to be riskier than bonds, and so investors will pay less for an expected return from stocks than for the same expected return from bonds.

The game-theoretic probability-free theory of finance advanced in our 2001 book *Probability and Finance: It's Only a Game* suggests an alternative explanation, which attributes the equity premium to speculation. This game-theoretic explanation does better than the explanation from risk aversion in accounting for the magnitude of the premium.

Contents

1 Introduction

The game-theoretic probability-free theory of finance advanced in our 2001 book *Probability and Finance: It's Only a Game* and in subsequent working papers (see Section 8) suggests that the equity premium can be explained by speculation. Speculation is involved in three ways:

1. Speculation causes volatility. This is widely accepted by traders and experts in option pricing.

2. Speculation removes opportunities for low-risk profit in the market, resulting in a market that is efficient, in the sense that an investor can very rarely do better than hold all tradable assets in proportion to their capitalization.

3. Assuming that the market is efficient in this way, and assuming that an index that holds all assets in proportion to their capitalization is tradable, speculation forces this index to appreciate in proportion to the square of its volatility.

This game-theoretic explanation does better than the explanation from risk aversion in accounting for the magnitude of the premium.

In a series of working papers posted at `www.probabilityandfinance.com/articles/index.html`, we have provided a rigorous mathematical elaboration of Point 3. Other working papers in the same series explain the larger framework of game-theoretic probability and finance. The most relevant of these working papers are listed in Section 8, where they are numbered with the prefix GTP (game-theoretic probability). See especially GTP38 and GTP44.

This note does not give mathematical proofs. Instead, it summarizes the game-theoretic explanation of the equity premium and touches on its implications and on the empirical questions it raises.

2 What causes volatility?

It is widely accepted among finance practitioners that volatility is primarily caused by speculation. Perhaps the most authoritative statement to this effect is by John Hull in his widely used textbook, *Options, Futures, and Other Derivatives* [12]. The following passage appears on page 329 of the ninth edition:

> What Causes Volatility?
>
> It is natural to assume that the volatility of a stock is caused by new information reaching the market. This new information causes people to revise their opinions about the value of the stock. The price of the stock changes and volatility results. This view of what causes volatility is not supported by research. With several years of daily stock price data, researchers can calculate:

> 1. The variance of stock returns between the close of trading on one day and the close of trading on the next day when there are no intervening nontrading days.
> 2. The variance of the stock price returns between the close of trading on Friday and the close of trading on Monday.
>
> The second of these is the variance of returns over a 3-day period. The first is a variance over a 1-day period. We might reasonably expect the second variance to be three times as great as the first variance. Fama (1965), French (1980), and French and Roll (1986) show that this is not the case. These three research studies estimate the second variance to be, respectively, 22%, 19%, and 10.7% higher than the first variance.
>
> At this stage one might be tempted to argue that these results are explained by more news reaching the market when the market is open for trading. But research by Roll (1984) does not support this explanation. Roll looked at the prices of orange juice futures. By far the most important news for orange juice futures prices is news about the weather and this is equally likely to arrive at any time. When Roll did a similar analysis to that just described for stocks, he found that the second (Friday-to-Monday) variance for orange juice futures is only 1.54 times the first variance.
>
> The only reasonable conclusion is that volatility is to a large extent caused by trading itself. (Traders usually have no difficulty accepting this conclusion.)

Hull's view of the matter is also supported by instances in which addition of hours to the trading day or days to the trading week has increased volatility.

Trading not based on new information is often called *noise trading*, and possible causes have been studied by a number of authors. Several studies have concluded that short-term institutional investors are noise traders [19, 24, 2, 3, 25]. Some authors have suggested that there is excessive speculation when there are differences in opinion between agents and there are short sale constraints on stocks [18, 1, 22].

3 What is an efficient market?

According to the definition introduced by Eugene F. Fama in 1965 [5], a market is efficient if its prices reflect all available information. The financial economists who accept this definition further assume that the available information determines a probability distribution for future payoffs of the assets priced by the market, and even that the future payoffs will objectively obey that distribution. Unfortunately, as acknowledged by Fama and everyone else who has studied the topic, we cannot directly test the hypothesis that a particular market is efficient. We can only test a model that specifies the probability distribution of future payoffs, the risk preferences of investors, and the way in which these elements interact to determine equilibrium prices.[1]

The game-theoretic probability-free theory of finance offers a simpler and more testable definition of market efficiency. According to this definition, a market is efficient if no one can devise a strategy for trading over an extended future time period that will multiply the capital it risks by a large factor. This can easily be tested: define a trading strategy and implement it. If you multiply the money you risk (this means all the money you risk, your own money and money from anyone else foolish enough to lend it to you) by 1000, you can reject the hypothesis of efficiency with as much confidence as you would have in rejecting a hypothesis that failed a statistical test with pre-specified significance level equal to 0.001.[2]

Why should a market be efficient in this sense? Because of speculation. Suppose many traders are at work, trying every conceivable strategy that might multiply the capital it risks by a large factor. Each strategy creates a demand for certain positions. If successful, the strategy will be played with more and more capital, and hence demand for these positions will push up their prices until the strategy becomes ineffective. In the limit, there should be no gain from shifting capital from one asset to another, and hence no portfolio should beat holding all assets in the market in proportion to their capitalization.

What is “the market” in this picture? It is a community of institutions in which ready trading is possible, and the assets in this market are all the assets that can be readily traded. In today’s world, this means all the assets that can be traded from minute to minute or perhaps even from microsecond to microsecond. For a trader working in a bank in New York City or in a trading center in New Jersey, the market might include several electronic markets, and the assets will include stocks of large corporations and perhaps some derivatives and foreign currencies. Real estate, privately held corporations, bonds and other forms of cash, and even many small-cap stocks would not be included. Some short-term bonds can be traded readily (the trader will be using a money-market account to do his trading), but the supply of these bonds is elastic, and an money-market account does not represent ownership of a productive asset.

3.1 The efficient index hypothesis (EIH)

The efficient market hypothesis motivated by this picture says that you cannot multiply the capital you risk by a large factor relative to an index defined by the total value of all the readily tradable assets. To make it a simple slogan: you cannot beat this index. In order to distinguish this hypothesis from Fama’s efficient market hypothesis, we call it the *efficient index hypothesis* (EIH).

[1] There is a vast literature on this topic, but Fama’s Nobel Prize lecture [6] provides a brief summary of the main issues.

[2] In the spirit of Blaise Pascal’s solution of the problem of points, game-theoretic probability uses this idea to define probability: the probability of an event is the least amount of money one must risk in order to get a unit amount if the event happens; see GTP32.

The EIH does not rely on any assumptions about the risk attitudes of investors. Instead, it locates the cause and meaning of efficiency in speculation. This may be increasingly appropriate in an age when high-frequency trading and other computerized strategies dominate the market.

In order to make the EIH plausible and derive interesting theoretical implications, we need an index that can itself be readily traded. In practice, a New York financial institution might have a difficulty computing an index that represents all the assets in which it can readily trade and would surely have difficulty trading in such index with negligible transaction costs.[3] So we must approximate our ideal index with something that can be at least approximately traded with small transaction costs. An obvious choice for a New York institution is the S&P 500 index.

3.2 Mathematical implications of the EIH

Suppose for simplicity that the risk-free interest rate is zero,[4] and that speculators can borrow and invest at this zero rate. Suppose that trading is so frequent that a continuous-time mathematical model is appropriate. Finally, suppose that we measure time by (cumulative) relative quadratic variation rather than by calendar and clock.

The *relative quadratic variation* of a continuous path at clock time t is the sum of the squared returns from time 0 to time t (i.e., the square of the *volatility*), considered in the limit as the length of the interval over which returns are measured tends to zero.[5] The EIH implies that this limit exists. As measured by relative quadratic variation, time flows faster when trading is unusually intense and volatility is high.

Suppose s is time as measured by the relative quadratic variation of the index, and let S_s be the value of the index at time s. Suppose further, for simplicity, that $S_0 = 1$. Then as explained in Section 4 of GTP44, the EIH implies that as s increases the trajectory S_s will look as if

$$S_s = e^{s/2+W_s}, \tag{1}$$

where W_s is a Brownian motion.[6] When W_s is Brownian motion, the random variable W_s for a fixed value of s is normally distributed with mean zero and

[3] The theory discussed in the following sections assumes zero transaction costs. But transactions costs can be brought into the theory, and the practical significance of apparent departures from efficiency can be measured by the level of transaction costs that would explain them; see GTP23.

[4] This is equivalent to assuming that all assets in the market, including the index I, are measured as multiples of the value of a bank account that earns the risk-free rate continuously.

[5] In finance theory, the observed sum of the squared returns over some period of time is often called the *realized variance*; its square root is the observed *volatility* for that unit of time. Here we are considering a realized path—the path that I's price actually takes. But we are considering an idealized continuous-time version of this realized path, not a discrete-time path that can actually be observed. So instead of using the practical-sounding term *realized variance*, we have borrowed *quadratic variation* from probability theory. We add *relative* because a return is a relative rather than an absolute price change.

[6] This is a special case of geometric Brownian motion, the probability model assumed by

standard deviation $\sqrt{s}$. When s is very large, $\sqrt{s}$ and W_s will be negligible in comparison with $s/2$, and hence we will have $S_s \approx e^{s/2}$, or

$$\ln S_s \approx \frac{s}{2}. \tag{2}$$

In words: after a long period of time, the natural logarithm of the index will be approximately equal to half the relative quadratic variation.

What do we mean when we say that the EIH implies (1)? We mean that for every property of Brownian motion, there is a trading strategy that will multiply the capital it risks by a large factor relative to the index (by an infinite factor in the idealized continuous-time picture we use in GTP44 and GTP45) if the trajectory $\ln S_s - s/2$ does not have the property. To the extent that we believe the EIH—i.e., to the extent that we think such trading opportunities have already been fully exploited, we will not expect the property to fail.

In this paper, we will not review the general proof given in GTP44 and GTP45. But in the next section, we will sketch the proof of the particular property (2). In other words, we will describe a trading strategy that will multiply the capital it risks by a large factor relative to the index if (2) fails. This strategy is very simple—so simple that we can be confident that it is being so massively implemented by algorithmic traders that prices will have adjusted to prevent its success and hence assure the validity of (2).

4 The equity premium

The equity premium is the amount by which the average return from stocks exceeds the risk-free interest rate. Its size is widely considered puzzling.

Here we briefly review the puzzle and summarize how it can be resolved using the EIH. Following the working papers we are summarizing, we use an idealized continuous-time model. (Only GTP1 used discrete time.) The purpose of working in continuous time is to produce a clean theory, without distracting approximations, that can be compared to the established probabilistic continuous-time theory that uses geometric Brownian motion and other Itô processes.

4.1 Why the premium is a puzzle

The empirical study by Mehra and Prescott reported in Table 2 of [15] estimates the premium over the period 1889–2005 as 6.36%. Other empirical studies have produced similar estimates. As first noted by Mehra and Prescott in the late 1970s, these numbers are too large to be explained by risk attitudes in the context of standard probabilistic assumptions. Standard theory suggests that the premium should be about 1% ([16], page 146). A number of behavioral

the Black-Scholes formula for pricing options. If a price S_t follows geometric Brownian motion with drift μ and volatility σ, then $S_t = \exp(\mu t - (\sigma^2/2)t + \sigma W_t)$, where W_t is a Brownian motion, and this reduces to $S_t = \exp(t/2 + W_t)$ when $\mu = \sigma = 1$. So the EIH is consistent with the Black-Scholes theory when time is measured in a way that makes the volatility appear constant.

models and other explanations have been proposed to account for the difference, but none has gained general acceptance [17].

4.2 The premium implied by the EIH

Denote the value of the index I at clock time t by I_t, and assume for simplicity that $I_0 = 1$. Denote the relative quadratic variation at time t by Σ_t. Then (2) can be written

$$\ln I_t \approx \frac{\Sigma_t}{2}. \tag{3}$$

This tells us that if I is efficient, variation in its value must be accompanied by commensurate growth relative to the risk-free rate.

Most empirical studies of the equity premium, including those just mentioned, consider the average returns of a market index I rather than its growth. The continuous-time counterpart of the average return is the sum of the returns from time 0 to time t, considered in the limit as the length of the time interval over which returns are measured tends to zero. Denote this quantity by M_t. Then

$$\mathrm{M}_t = \ln I_t + \frac{\Sigma_t}{2}. \tag{4}$$

This relation, which is Lemma 3.2 in GTP44, expresses the difference between the geometric and arithmetic cumulation of returns. Combining (3) and (4), we obtain

$$\mathrm{M}_t \approx \Sigma_t. \tag{5}$$

In words: if the market index I is efficient, then its average return approximates the square of its cumulative volatility. Since we are using an account that earns the risk-free rate as I's numéraire, this average return is the equity premium.

The annualized volatility of the S&P 500 is approximately 20% ([14], page 8). Squaring this, we obtain an equity premium of 4%. This is closer to the empirical estimates than the 1% obtained from standard theory, and GTP38 (Section 4) shows that it is within (5)'s anticipated error of approximation.

4.3 The trading strategy that implies the premium

Section 5 of GTP44 discusses a relatively simple trading strategy that multiplies the capital it risks by a lot if (5) is violated.

To make your capital grow even faster than I if I is growing faster than our theory predicts—i.e., if M_t is substantially greater than Σ_t, you invest more than you have in I—i.e., you invest all you have and borrow money (at the risk-free rate) to invest even more. Suppose you do this for K rounds of trading, borrowing a small fraction ϵ of your current capital each time so that you always have

$$(1+\epsilon) \times (\text{current capital}) \tag{6}$$

invested in I. Let m_k be I's return on round k. Then the value of I will be multiplied by $1 + m_k$ on that round, while the value of your capital will be multiplied by $1 + (1+\epsilon)m_k$. So your capital will increase (or decrease) relative to I by the factor

$$\frac{1 + (1+\epsilon)m_k}{1 + m_k}.$$

Using Taylor's series for the logarithm, we obtain the approximation

$$\ln \frac{1 + (1+\epsilon)m_k}{1 + m_k} \approx \epsilon m_k - \epsilon m_k^2 - \frac{\epsilon^2}{2} m_k^2.$$

So over K rounds, your capital will grow relative to I by a factor whose logarithm is approximately

$$\begin{aligned} \epsilon \sum_{k=1}^{K} m_k - \epsilon \sum_{k=1}^{K} m_k^2 - \frac{\epsilon^2}{2} \sum_{k=1}^{K} m_k^2 &= \epsilon \mathrm{M}_t - \epsilon \Sigma_t - \frac{\epsilon^2}{2} \Sigma_t \\ &= \epsilon(\mathrm{M}_t - \Sigma_t) - \frac{\epsilon^2}{2} \Sigma_t, \end{aligned} \tag{7}$$

where t is the time required for the K rounds. If you continue until Σ_t is so large that $\epsilon \Sigma_t$ is a substantial even though ϵ is small, and M_t exceeds Σ_t substantially, then (7) will be substantial; you will have multiplied your capital by a large factor relative to I. (For example, suppose $\epsilon = 0.01$ but $\epsilon \Sigma_t = 3$, and suppose M_t is 50% greater than Σ_t. Then (7) is approximately 1.5, meaning that you have multiplied your capital relative to I by $e^{1.5} \approx 4.5$.)

To similarly make money if I grows too slowly—i.e., if M_t is substantially less than Σ_t, you can take ϵ in the preceding argument to be a small negative number. In other words, you keep a small fixed fraction of your capital in the risk-free bond on each round, investing the rest in I.

You can implement the two strategies simultaneously: put half your initial capital on one of them and half on the other. So you have a strategy that will multiply its initial capital substantially relative to I unless $\mathrm{M}_t \approx \Sigma_t$. (We promised a strategy that multiplies *the capital it risks*, so you need to implement the strategy just sketched in a way that risks no more than its initial capital. You can do this by stopping the strategy if its capital gets close to zero. In GTP44 we rely on the assumption that the price path is continuous to make sure we can stop in time. Weaker assumptions can also be accommodated.)

4.4 Implications of the premium

This probability-free explanation of the equity premium raises questions about the impact of the market on the wider economy.

The market represents one portion of society's productive assets—the portion that is so liquid that it is volatile in the short term. If the level of this volatility is driven by institutional and technological factors not directly related to the productivity of these assets, and this in turn drives changes in the valuation of these assets relative to cash[7] and hence relative to society's other productive assets, can this create durable imbalances in valuation?

Consider an extended period in which the publicly traded and very liquid portion of the economy is exceptionally productive, but volatility is relatively low, so that speculation keeps the value of I from increasing relative to the other assets in the economy to the extent justified by economic fundamentals. What will happen? We might conjecture that entrepreneurs will avoid the market: start-ups may remain privately held, some public corporations may go private, and even successful large corporations may not have the means or incentives to acquire start-ups. If these tendencies do not correct the imbalance, the prices of at which productive assets outside the liquid market can be sold might fall.

On the other side, if the publicly traded and very liquid portion of the economy is not as productive relative as the rest of the economy, but volatility is very high, so that speculation forces an increase in the value of I not justified by economic fundamentals, we might see a greater tendency for companies to go and stay public or to be acquired by public companies. Perhaps assets outside the liquid market would rise in value.

We can also ask whether there are mechanisms that might make I's volatility adjust when I's value is too high or too low relative to the valuation of other assets in the economy. In other words, can volatility somehow be made endogenous to the picture?

Of course, our theory and these speculations assume a lot about liquidity. Our argument for I's being efficient requires that capital be shifted easily, with negligible transaction costs, within the universe of assets I represents and not so easily to other assets aside from the risk-free bond. The strategies that force (5) further require that the index itself can be traded with negligible transaction costs.

We are also relying on the assumption that changes in value will not be too abrupt. In our mathematical picture, the price path I_t is continuous, so that the speculator can always avoid a catastrophic loss by liquidating his exposure when his capital hits some small level close to zero. As a practical matter, the liquidity needed for this may fail when the market falls abruptly. Such crashes may therefore achieve corrections that would not be possible otherwise.

5 The game-theoretic CAPM

As explained in GTP44, the strategies just described, when they mix I and another easily traded asset S instead of mixing I and cash, lead to the conclusion that

$$\mathrm{M}_t^S \approx \Sigma_t^{S,I}, \tag{8}$$

[7]We continue to assume that traders have accounts that enable them to invest and borrow at the risk-free rate. So "cash" refers to the unit of value of an account that earns the risk-free rate.

where M_t^S is S's *summed returns* and $\Sigma_t^{S,I}$ is the *relative covariation* of S and I:

$$\mathrm{M}_t^S := \sum_{k=1}^{K} m_k^S \quad \text{and} \quad \Sigma_t^{S,I} := \sum_{k=1}^{K} m_k^S m_k^I,$$

where t is the time it takes to complete K rounds and m_k^S and m_k^I are the returns for S and I, respectively. The approximate relationship (8) is a simple form of the game-theoretic Capital Asset Pricing Model (CAPM) that we first discussed in GTP1 in 2001. The approximation (5), $\mathrm{M}_t^I \approx \Sigma_t^I$, is the special case of (8) where $S = I$.

Our game-theoretic CAPM can be also thought of as a prediction from the CAPM of standard finance theory, which assumes that the returns m^S and m^I have a joint probability distribution and relates their theoretical expected values by a formula that reduces, when the risk-free interest rate is zero, to

$$\mathbf{E}(m^S) = \frac{\mathbf{Cov}(m^S, m^I)}{\mathbf{Var}(m^I)} \mathbf{E}(m^I).$$

This theoretical relationship predicts an approximate equality when the theoretical expected values and the theoretical covariance and variance are replaced by the corresponding idealized realized values. Making this replacement, we obtain

$$\mathrm{M}_t^S \approx \frac{\Sigma_t^{S,I}}{\Sigma_t^I} \mathrm{M}_t^I.$$

By (5), Σ_t^I and M_t^I approximately cancel each other out, and we obtain (8).

6 Need for empirical work

The theory we have summarized raises a number of questions that merit empirical study.

6.1 Attaining efficiency

In Section 3, we argued that the efficiency of a market index can emerge from speculation. This might be confirmed by simulation studies. We are not aware of such studies having been carried out. They could help us understand the frequency of trading and level of liquidity required for the continuous-time theory to become accurate.

6.2 Measuring the equity premium

It would be useful to revisit the data on the size of the equity premium using the definitions and viewpoint of our theory.

As we noted in Section 4, our theory uses the cumulative volatility rather than clock time as its time scale. At a practical level, this means that our measures of volatility, average return, and covariance are based on Lebesgue rather than Riemann sums. Instead of fixing a small amount dt of clock time and measuring changes over successive periods of this duration, we fix a small squared increment $(dI)^2$ as the unit of time over which we measure returns, covariances, etc. Taking this into account might make some difference in the empirical calculations.

In addition to studying the extent to which the observed equity premium matches the predictions of our continuous-time theory, we should also look at how the predictions are affected by realistic assumptions about transaction costs. On the other side, we might add to the model traders who play more aggressive strategies that might enhance the effect of the strategies we have described.

Our emphasis on speculation and liquidity may also suggest different choices for the market index I. In 1977, Richard Roll pointed out that tests of the standard CAPM should look at a market index that includes all the many assets that could be considered by an investor interested in balancing current and future consumption [20]. As Roll argued, the S&P 500, which includes only assets traded on the NYSE and NASDAQ, hardly meets this criterion, and it seems impossible to define an index that does meet it. A corresponding criticism of our theory is that we should consider an index that includes all the many assets, now including foreign currencies and ETFs, that have sufficient liquidity to be used by the trading strategies that our theory considers. The S&P 500 does not meet this criterion, but it is conceivable that by investigating the trading actually being conducted by hedge funds, banks, and other large traders, we could construct an index that does meet it.

6.3 Testing the game-theoretic CAPM

In their 2004 review of the standard CAPM [7], Eugene F. Fama and Kenneth R. French concluded that empirical data did not accord with it well enough to justify the use made of it by finance professionals. It would surely be worthwhile to make a similar study of the degree to which the game-theoretic CAPM and our game-theoretic theory of the equity premium accord with data.

As we have seen, the game-theoretic prediction concerning the equity premium is a special case of the game-theoretic CAPM. (The approximation (5) is a special case of the approximation (8).) The general argument for the game-theoretic CAPM seems to involve more approximation, however, and it is not unreasonable to expect that it will fit the data poorly, or at least that it will require time horizons longer than could be of interest to finance professionals (see [13]). It would nevertheless be of interest to understand the degree to which the empirical results lie within the ranges that the game-theoretic approach predicts when realistic transaction costs are taken into account.

Perhaps of relevance are the working papers by John R. Graham and Campbell R. Harvey [10, 11]. These papers document the correlation between CFO's predictions of the equity premium over the next ten years and the estimates of future volatility implied by the VIX.

7 Acknowledgements

This paper was prepared for a seminar at the Rutgers Business School. The authors would like to thank the participants in the seminar for helpful comments and suggestions. Frank McIntyre and Darius Palia have been especially helpful.

8 Relevant GTP papers

At `http://www.probabilityandfinance.com/articles/index.html`.

GTP1. The Game-Theoretic Capital Asset Pricing Model. Vladimir Vovk and Glenn Shafer. March 2002. First posted in November 2001. Journal version: *International Journal of Approximate Reasoning*, 49:175–197, 2008. This paper derives a discrete-time version of the game-theoretic CAPM, using Riemannian sums.

GTP2. Game-theoretic capital asset pricing in continuous time. Vladimir Vovk and Glenn Shafer. December 2001. This paper translates the ideas of GTP1 into continuous time using nonstandard analysis.

GTP5. A game-theoretic explanation of the $\sqrt{dt}$ effect. Vladimir Vovk and Glenn Shafer. January 2003. This paper gives a game-theoretic explanation of the empirical observation that price series have quadratic variation—i.e., that the sum of squared changes in price tends to be proportional to the length of time.

GTP23. Testing lead-lag effects under game-theoretic efficient market hypotheses. Wei Wu and Glenn Shafer November 2007. Game-theoretic efficient market hypotheses identify the same lead-lag anomalies as the conventional approach: statistical significance for the autocorrelations of small-cap portfolios and equal-weighted indices, as well as for the ability of other portfolios to lead them. Because the game-theoretic approach bases statistical significance directly on trading strategies, it allows us to measure the degree of market friction needed to account for this statistical significance. The authors find that market frictions provide adequate explanation.

GTP28. Continuous-time trading and the emergence of probability. Vladimir Vovk. May 2015. First posted in April 2009. Journal version: *Finance and Stochastics*, 16:561–609, 2012. `arXiv:0904.4364v4` [math.PR]. This article considers an idealized financial security with continuous price path, without making any stochastic assumptions. It is shown that typical price paths possess quadratic variation. When time is replaced by the quadratic variation process, the price path becomes Brownian motion. This is similar to the celebrated Dubins-Schwarz theorem, except that the probabilities (constituting the Wiener measure) emerge instead of being postulated.

GTP32. How to base probability theory on perfect-information games. Glenn Shafer, Vladimir Vovk, and Roman Chychyla. December 2009. This paper reviews the basics of game-theoretic probability.

GTP38. The efficient index hypothesis and its implications in the BSM model. Vladimir Vovk. October 2011. First posted in September 2011. `arXiv:1109.2327v1` [q-fin.GN]. This article assumes the Black-Scholes model, which says that the value of a traded security follows geometric Brownian motion, adds the efficient index hypothesis, and derives the predictions concerning the equity premium that are derived in GTP44 without initially assuming the Black-Scholes model.

GTP39. The Capital Asset Pricing Model as a corollary of the Black-Scholes model. Vladimir Vovk. September 2011. `arXiv:1109.5144v1` [q-fin.PM]. This article similarly obtains the game-theoretic CAPM beginning with the Black-Scholes model.

GTP43. Getting rich quick with the Axiom of Choice. Vladimir Vovk. May 2016. `arXiv:1604.00596v2` [q-fin.MF]. The axiom of choice, a nonconstructive axiom that is almost universally used both by pure and applied mathematicians, is known to lead to some paradoxical conclusions. It implies, for example, that a spherical ball can be decomposed into a finite number of pieces, which can be recombined to form a number of balls of the same size. The condition of measurability for events and variables, which is used in standard probability theory, allows us to avoid such paradoxes. This condition was not needed in the discrete-time game-theoretic probability studied in [23]. As this paper shows, it is needed in continuous-time game-theoretic probability, because the axiom of choice implies, paradoxically, that knowledge of a continuous path up to a particular time T almost always allows one to predict its values for some short period after T.

GTP44. A probability-free and continuous-time explanation of the equity premium and CAPM. Vladimir Vovk and Glenn Shafer. July 2016. The principal results of this article are summarized in these notes.

GTP45. Basics of a probability-free theory of continuous martingales. Vladimir Vovk and Glenn Shafer. July 2016. The results of GTP44 are presented as special cases of very general results in game-theoretic continuous-time probability.

Other references

[1] Patrick Bolton, José Scheinkman, and Wei Xiong. Executive compensation and short-termist behaviour in speculative markets. *Review of Economic Studies*, 73:577–610, 2006. 2

[2] Brian J. Bushee and Christopher F. Noe. Corporate disclosure practices, institutional investors, and stock return volatility. *Journal of Accounting Research*, 38:171–202, 2000. 2

[3] Patrick Dennis and Dean Strickland. Who blinks in volatile markets, individuals or institutions? *Journal of Finance*, 57:1923–1949, 2002. 2

[4] Eugene F. Fama. The behavior of stock-market prices. *Journal of Business*, 38:34–105, 1965. 2

[5] Eugene F. Fama. Random walks in stock market prices. *Financial Analysts Journal*, pages 55–59, September/October 1965. 2

[6] Eugene F. Fama. Two pillars of asset pricing, 2013. Nobel Prize lecture, `www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2013/fama-lecture.pdf`. 3

[7] Eugene F. Fama and Kenneth R. French. The Capital Asset Pricing Model: Theory and evidence. *Journal of Economic Perspectives*, 18:25–46, 2004. 10

[8] Kenneth R. French. Stock returns and the weekend effect. *Journal of Financial Economics*, 8:55–69, 1980. 2

[9] Kenneth R. French and Richard Roll. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics*, 17:5–26, 1986. 2

[10] John R. Graham and Campbell R. Harvey. The equity risk premium in January 2007: Evidence from the Global CFO Outlook Survey, 2007. Working Paper. SSRN 959703. 10

[11] John R. Graham and Campbell R. Harvey. The equity risk premium in 2014, 2014. Working Paper. SSRN 2422008. 10

[12] John C. Hull. *Options, Futures, and Other Derivatives.* Prentice-Hall, Upper Saddle River, NJ, ninth edition, 2015. 1

[13] Philippe Jorion and William N. Goetzmann. Global stock markets in the twentieth century. *The Journal of Finance*, 54:953–980, 1999. 10

[14] Rajnish Mehra. The equity premium puzzle: a review. *Foundations and Trends in Finance*, 2:1–81, 2006. 6

[15] Rajnish Mehra and Edward C. Prescott. The equity premium: ABCs. Chapter 1 of [17], pp. 1–36. 5

[16] Rajnish Mehra and Edward C. Prescott. The equity premium: a puzzle. *Journal of Monetary Economics*, 15:145–161, 1985. 5

[17] Rajnish Mehra and Edward C. Prescott, editors. *Handbook of the Equity Risk Premium.* Elsevier, Amsterdam, 2008. 6, 13

[18] Edward M. Miller. Risk, uncertainty, and divergence of opinion. *Journal of Finance*, 32:1151–1168, 1977. 2

[19] Michael Porter. *Capital choices: Changing the way America invests in industry.* Council on Competitiveness and Harvard Business School, Boston, 1992. 2

[20] Richard Roll. A critique of the asset pricing theory's tests. Part I: On past and potential testability of the theory. *Journal of Financial Economics*, 4:129–176, 1977. 10

[21] Richard Roll. Orange juice and weather. *American Economic Review*, 74:861–880, 1984. 2

[22] José A. Scheinkman and Wei Xiong. Overconfidence and speculative bubbles. *Journal of Political Economy*, 111:1183–1220, 2003. 2

[23] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001. 12

[24] Richard W. Sias. Volatility and the institutional investor. *Financial Analysts Journal*, 52:13–20, 1996. 2

[25] Yexiao Xu and Burton Malkiel. Investigating the behavior of idiosyncratic volatility. *Journal of Business*, 76, 2003. 2