Underbetting the Market: Behavioral and Bayesian Insights on the Equity Premium Puzzle

Andrew Grant* David Johnstone†

Abstract

Conventional asset pricing models presume investors "know" the true probability distribution of returns and invest accordingly. In contrast, this paper explores how rational investors under uncertainty may deliberately "underbet" their beliefs. Drawing on prospect theory, financial decision theory, and the mathematical literature on gambling, we show how perceived downside risks lead investors to overweight small-probability adverse outcomes and reduce their allocations below normative "optimal" levels. Using simulations and established theory, we demonstrate that even modest doubts about the reliability of return forecasts cause substantial reductions in market exposure. These findings offer a unified explanation for conservative investment behavior and help resolve the equity premium puzzle: endogenous underinvestment leads to high observed returns, even in the absence of irrationality. When payoff risk is partly exogenous, such uncertainty aversion results in equilibrium outcomes with elevated equity premia and limited exposure to ruin.

Keywords: equity premium puzzle, Kelly criterion, parameter uncertainty, prospect theory, underinvestment, Bayesian portfolio choice, gambling theory

JEL Classification: G11, G12, D81, C63

*University of Sydney Business School. Email: andrew.grant@sydney.edu.au

†University of Sydney Business School. Email: david.johnstone@sydney.edu.au

1 Introduction

Investors in the stock market have imperfect information and must assess probabilities subjectively. The purpose of this paper is to develop a simple model of how investors might allow heuristically for uncertain probability beliefs and how their approach affects (i) their investment choices and results, and (ii) stock market characteristics including the price implied rate of return.

There is a widely accepted argument in finance suggesting that investors set asset prices specifically to earn whatever is their required return on the asset relative to its risk (or, moreover, relative to how it affects portfolio risk). For example, if they were betting on a coin toss and required a 25% return, they would "price" a bet that pays \$1 or nothing at no more than 40 cents, thus obtaining expected return $0.5/0.4 = 1.25$. Moving from the casino to the stock market, the "equity premium puzzle" points at the long history of high average returns and asks why investors demand such high returns relative to the much lower returns from bonds. Drawing on both behavioral finance and mathematical models used in gambling, we propose a simple model of why and how investors might do that.

In the simplest statistical models of investment, probabilities are "physical" and investors are like casino owners who know the physical probabilities of their roulette wheels. However, there is no such physical reality in the stock market. Stock market investing is more like betting on horses, or on events like tomorrow's weather, rather than roulette. Horse races and the weather can be seen in terms of probabilities but not with any known objective distributions. Hence, investors have no choice but to rely on personal, subjective and usually tenuous assessments of payoff distributions. Ultimately, they price assets nervously using whatever information seems relevant and allowing however they can for their own known limitations.

When thinking probabilistically, investors use the same laws of probability as we all learn in classical frequentist statistics, but they do not see the stock market as a stationary physical process. Tomorrow is not a random draw from the stock market's empirical history.[1] Worse still, no single day's outcome, up or down, can be said to "have" a probability or probability distribution.[2]

By recognizing that stock returns have different distributions under different conditions with different probabilities, investors can allow for subjective considerations that rationally cannot be ignored and for circumstances where objective historical returns data is of little relevance. There are no objective probabilities.[3] All that exists objectively is the period's result - observable after the event. This is a problem of "parameter uncertainty." Stock returns are modelled as probability distributions but as distributions with unknown parameters θ, and investors seek returns that compensate for their own uncertainty not only about the returns but about their parameters θ. Accepting the Bayesian rather than frequentist view[4] of markets and investment, we attribute the high returns on stocks to investors' pragmatic ways of accounting for their lack of confidence in θ (in their own probability forecasts), that is, in their recognition of parameter risk, "unknown

[1]Barry (1978) suggested that any newly arising threat of political and economic shock can make models based on "objective data" (historical stock returns) virtually irrelevant.

[2]One of the deepest arguments in Bayesian theory is that "probability does not exist" and cannot be regarded as having objective reality (de Finetti, 1974, p.x).

[3]Modern Bayesian portfolio admit subjective returns distributions based partly on historical data and partly on current circumstances affecting future returns (e.g., Black and Litterman, 1992). The frequentist simplification of "plugging in" assumed parameter values, as if they are known to be true, is heavily criticized (e.g., Coles and Loewenstein 1988; Bauder, Bodnar, Parolya and Schmid, 2021; Rachev, Hsu, Bagasheva and Fabozzi, 2008; Lewellen and Shanken, 2002; Avramov and Zhou, 2010).

[4]In a straight frequentist approach, past market data is used to generate the maximum likelihood estimates of the unknown parameters and these are plugged into the optimization model as if they are known true parameters.

unknowns" and even the threat of black swans. Ultimately, in this mindset, investors react by investing and buying stocks conservatively, thus choosing to "underbet" whatever they sense as the true or likely beliefs.

2 Underbetting

There are multiple strands of the decision-making literature concluding basically in unison that decision makers do, and perhaps even should, invest as if they overweight the small probability of a very bad outcome. Perhaps the strongest analytical explanation for this form of self-restraint exists in the mathematical literature on gambling, which, although highly developed, is not widely known in finance. In finance and financial economics, similarly analytical work known as cumulative prospect theory (CPT) presents investment and probability models where investors (i) instinctively overweight losses relative to gains while also (ii) biasing personal probability assessments to inflate the probabilities of the most unlikely events. Both theoretical approaches suggest that investors are inclined to "underbet their own beliefs."

2.1 Prospect Theory

In economics and finance, the seminal paper by Benartzi and Thaler (1995) on myopic loss aversion considered a simple problem where the investor is faced with a bet with known positive expected value but also carrying the chance of significant loss on the first and perhaps only ever bet. In a well-known statement of this problem, Samuelson (1963) asked a colleague whether he would accept a 50% chance to win $200 in exchange for a 50% chance to lose $100. The colleague felt not, but countered that he would readily accept an independent series of such wagers. Rejecting Samuelson's normative proof that this combination is irrational,[5] Benartzi and Thaler (1995) showed how the colleague's response occurs naturally (call it "rationally") under prospect theory. First, by weighting losses more heavily than gains (loss aversion), the decision maker can "rationally" reject the single bet while accepting a sufficiently long series of identical bets. Second, by framing investment as "for the long term" and eschewing a more human myopic concern with just each single bet along the series, the investor can accept the first bet and every following bet one at a time, simply because she has fixated mentally on a long horizon, eschewing concern with the single case and relying on what is implied by the "law of averages."

> The longer the investor intends to hold the asset, the more attractive the risky asset will appear, so long as the investment is not evaluated frequently. (Benartzi and Thaler, 2005, p.75)

Thus, investor psychology in regard to (i) how deeply immediate or short term losses "hurt," and (ii) what time horizon the investor has in mind, offers one way to explain the apparently irrationally high market return on equity (the risk premium puzzle). Although the Benartzi and Thaler (2005) explanation holds perfectly well with known objective probabilities, its extension via cumulative prospect theory (Tversky and Kahneman, 1992, Tversky and Wakker, 1995) adds a further dimension. Investors conceive subjective weights that usurp the role of objective probabilities. The documented result is that events with objectively very small probability are understood *as if* their

[5] Pratt and Zeckhauser (1987) proposed a notion of "proper" risk aversion under which investors choosing not to make the single bet will not accept more than one of the same bets.

probability is something near double the objective (e.g., frequency based) probability. Typically, for example, a 1% probability will become mentally 2% or 3% or more. Using the calibrated probability weighting function developed by Tversky and Kahneman (1992, p.312) a probability of 3% is taken as if it is 6%.

Investors who mentally inflate small probabilities of loss are applying the CPT version of Murphy's Law by which the worst is to be expected (Diecidue and Wakker, 2000; Lopes and Oden, 1999). This kind of exaggerated response to downside risk appears aligned with the personality profile of cautious investors, who exhibit greater financial conservatism than gamblers (Jadlow and Mowen, 2010). At the other extreme in the probability spectrum, probabilities near one are discounted and effectively understated. Both of these "cognitive biases" occur naturally under rank dependent utility,[6] thus giving rise to the well-known S-shaped prospect theory weighting function. In response to such distortions, Kale (2006) proposes a growth-optimizing utility function that incorporates downside protection, leading to lower allocations in the risky asset than standard log-utility.

2.2 Gambling Theory

The mathematical theory of gambling (MacLean and Ziemba, 1999; 2006; MacLean, Ziemba and Blazenko, 1992; MacLean, Ziemba and Li, 2005; MacLean, Sanegre, Zhao and Ziemba, 2004) mirrors the log utility branch of utility theory but evolved independently.While mathematically advanced (e.g., Aucamp, 1993), it culminates in simple down-to-earth reasons for why bettors should err conservatively by betting smaller amounts than they would were they to "actually believe their own probability forecasts." The mathematical rationale for "underbetting" is not only that investors are aware of their own commonly inaccurate probability assessments[7] but, just as importantly, that betting "too much" and losing heavily is more damaging to capital growth than betting too little and foregoing some return (Thorp, 2011).

To illustrate, consider a simple extension of the problem posed by Samuelson (1963). The investor considers a bet on a coin toss that costs $\pi = 0.5$ per unit and pays \$1 for heads, and zero for tails. The investor believes – let us say correctly – that the coin is biased and has physical probability $p = 0.6$ of heads. Rather than saying yes or no to the bet, the investor can buy units of the bet by risking dollar amount f of starting wealth W_0 by buying $f\,W_0/\pi$ units. The bet has a positive expected return of $p/\pi - 1 = 0.6/0.5 - 1 = 20\%$, which exceeds the risk-free return factor $R_f = 1.05$, so it appears to offer a favorable investment opportunity. Nonetheless, the investor knows that she must not allocate 100% of her wealth to a bet on heads, because tails would bring immediate bankruptcy. This type of problem was critiqued in the calibration arguments of Rabin (2000), and illustrates the tension between expected value maximization and realistic investor behavior. Rather than rejecting the gamble entirely, we ask instead: what fraction f of initial wealth can she comfortably bet to best take advantage in her own mind of such an agreeable opportunity?

Her initial finance style approach is to simulate the bet mentally and look at the ex ante Sharpe ratio of each possible f. Starting with initial wealth of $W_0 = 1$ and imagining values of f from zero upwards, she finds that her ending wealth W_1 after buying fW_0/π units of the bet fraction

[6] The assumption of rank dependent utility (where the decision weights of outcomes depend on their rank against all possible outcomes) provides a normative basis for the S-shaped weighting function (Diecidue and Wakker, 2001).

[7] Realized events can make probability assessments "look wrong" even if they were well based. We can't know ex post that a probability was "wrong," any given assessment whether apparently right or wrong can lead to action that causes loss.

has mean,

$$E[W_1] = p\,f\,W_0/\pi + (1-f)W_0 R_f$$

variance

$$\mathrm{var}(W_1) = \left(\frac{f\,W_0}{\pi}\right)^2 p(1-p),$$

implied return factor $R = E[W_1]/W_0 = E[W_1]$, and finally portfolio ex ante Sharpe ratio

$$\mathrm{SR} = \frac{E[R]-R_f}{\sqrt{\mathrm{var}(R)}} = \frac{(p-\pi)\,R_f}{\sqrt{p\,(1-p)}}.$$

Unfortunately that approach comes to a dead end because the Sharpe ratio depends only on p, π and R_f and is independent of f. However, taking another approach altogether, the "right" choice of f has been explored at length in the gambling literature from a mathematical yet pragmatic perspective. Under this approach, spawned in casino blackjack betting[8], the gambler chooses f to achieve a personally agreeable compromise between long run expected capital growth and volatility along the growth path. Fitting well with Benartzi and Thaler's (1995) model of how investors overcome their myopic inclinations, the gambler's choice of f is couched as a choice between different long runs. Applied to our problem, the natural upper bound on f is

$$f^{Kelly} = (p - \pi R_f)/(1-\pi R_f),$$

proved simply by maximizing $E[\log W_1]$.[9] This is known as log optimal investment or "Kelly betting" after Kelly (1956).[10] Mathematicians and some finance theorists have argued that f^{Kelly} represents the objective upper bound on sensible f because any higher f brings not only lower expected growth but simultaneously higher volatility, the worst of combinations (Thorp, 2011). The practical conclusion in gambling theory is that there is a domain of f somewhat lower than f^{Kelly} where the gambler achieves a large reduction in the volatility of wealth combined with relatively little reduction in long run growth (MacLean and Ziemba, 1999; 2006). In this domain, the relationship between growth and volatility is disproportionate, prompting the gambler to reduce the size of bets just enough to find a sweet spot that combines acceptably high capital growth with much reduced volatility. Such conservatism or underbetting - by design - is known in the gambling and mathematics literature as "fractional Kelly" and is practised by investing a fixed fraction of the "full Kelly" bet, typically something like "30% Kelly" (MacLean and Ziemba, 1999; 2006; Thorp, 2011).

To illustrate why this strategy has practical support, we have simulated the growth path of an investor making repeated bets on heads at price $\pi = 0.5$ per bet, using $p = 0.6$ as the investor's probability of heads. This thought experiment reveals to the investor how she can expect to grow her wealth if indeed her belief is objectively correct (which she knows is uncertain). Simulation reveals two potential growth paths. One is for the "full Kelly" bettor who bets fixed fraction $f = f^{Kelly} = (p-\pi R_f)/(1-\pi R_f) = 0.1579$ based on $p = 0.6$ and the other is for a "fractional Kelly" bettor who bets fixed fraction $\lambda\, f^{Kelly}$ with $\lambda = 0.3$, giving final $f = 0.0474$.

[8]The archetypal results developed when mathematicians identified positive expected value opportunities in some hands of blackjack (Thorp 1969, 2011) and sought the best way to exploit that occurrence.

[9]Maximizing the expected sum of log returns is equivalent to maximizing expected capital growth.

[10]Kelly betting has sometimes been advocated as obviously optimal because it is the best model asymptotically for the long run. That conclusion was mocked famously by Samuelson, who emphasised that "there are other utility functions" and any tradeoff between return and risk is personal.

Results are generated by assuming that after each bet the investor re-allocates whatever wealth exists (up or down after the previous trial) by betting the same constant fraction of wealth on heads in the next toss, with the rest always held in R_f. The investor's starting wealth is normalized to \$1.

Figure 1 illustrates the simulated distribution of investor wealth after 30 trials under different λ-Kelly strategies. Each panel is based on 100,000 bootstrap simulations. Brighter regions indicate more frequent outcomes. When the investor's belief $p = 0.6$ is accurate (top row), the full-Kelly bettor ($\lambda = 1$) achieves higher upside potential but at the cost of considerable downside risk. In contrast, fractional Kelly betting ($\lambda = 0.3$) produces a narrower, more stable distribution of wealth outcomes. This difference in volatility is especially pronounced when we introduce model misspecification (bottom row), assuming the true success probability is only $p = 0.525$. Under belief error, the full-Kelly investor suffers significantly more dispersion and downside risk, while the conservative fractional Kelly bettor maintains a tighter distribution. Despite their different volatility profiles, the expected log-wealth per trial is quite similar: 1.056 for $\lambda = 0.3$ and 1.062 for $\lambda = 1$. This modest difference highlights the robustness and appeal of underbetting in the face of uncertainty.

Table 1: Probabilities of wealth W_1 outcomes after 30 trials (with wealth rounded up to the nearest integer), assuming the true probability of $p = 0.6$ is correct.

	$\lambda = 0.3$	$\lambda = 1$
$\Pr(W_1 \leq 1)$	0.000	0.021
$\Pr(W_1 \leq 2)$	0.000	0.097
$\Pr(W_1 \leq 3)$	0.022	0.176
$\Pr(W_1 \leq 4)$	0.175	0.284
$\Pr(W_1 > 10)$	0.004	0.291
$\Pr(W_1 > 15)$	0.000	0.094
$\Pr(W_1 > 25)$	0.000	0.043
$\Pr(W_1 > 50)$	0.000	0.006

The similar average growth rates for full and fractional Kelly strategies are easy to explain: the full-Kelly investor experiences both greater upside and greater downside. Extremely good and extremely poor outcomes occur more frequently, tending to offset each other and yielding an average growth rate that lies "near the middle." Table 1 illustrates the respective probabilities of very low and very high final wealth outcomes under $\lambda = 0.3$ and $\lambda = 1$. The downside of the full-Kelly ($\lambda = 1$) rule is its elevated risk of severe losses: in 18% of repeated 30-trial sequences, the investor ends with $W_{30} \leq 3$. In contrast, the fractional Kelly strategy ($\lambda = 0.3$) produces such low outcomes only 2.2% of the time. Notably, both strategies (despite their differences) leave a 28% probability of underperforming the risk-free all-bonds portfolio, which would deliver $W_{30} = R_f^{30} = 1.05^{30} = 4.32$.[11]

Our simulated results align with what has been observed in real-world gambling and investment settings.[12] While the full-Kelly strategy maximizes expected long-run growth, both professional

[11] The all-bonds portfolio assumes a constant monthly return of 5%, compounded over 30 trials.

[12] Gamblers have found that betting just 30–40% of the full Kelly amount results in only a modest reduction in average capital growth, while substantially reducing the volatility of wealth (MacLean, Ziemba, and Blazenko, 1992; MacLean and Ziemba, 1999; MacLean et al., 2004; MacLean, Thorp, and Ziemba, 2010).

gamblers and market investors[13] have found it to be subjectively "too risky" in practice—because large drawdowns are frequent and difficult to recover from within realistic time frames. Even when the investor believes she has accurate probability forecasts, full-Kelly betting can lead to highly volatile outcomes. The rationale for adopting a more conservative strategy becomes even stronger once we allow for the possibility of misjudged probabilities.

To explore this, we extend the original simulation to cases where the investor believes $p = 0.6$ but the true probability is lower—specifically, $p = 0.525$. These results are shown in the lower panels of Figure 1. They illustrate how full-Kelly betting becomes far more dangerous under belief error, with a much wider and more skewed distribution of outcomes. By contrast, the fractional Kelly strategy ($\lambda = 0.3$) continues to deliver a relatively tight and stable distribution, demonstrating its robustness in the face of model uncertainty.

Table 2: Probabilities of wealth W_1 outcomes after 30 trials (with wealth rounded up to the nearest integer), assuming the investor's probability of $p = 0.6$ is incorrect, with true probability of $p = 0.525$.

	$\lambda = 0.3$	$\lambda = 1$
$\Pr(W_1 \leq 1)$	0	0.118
$\Pr(W_1 \leq 2)$	0.0013	0.323
$\Pr(W_1 \leq 3)$	0.117	0.463
$\Pr(W_1 \leq 4)$	0.462	0.606
$\Pr(W_1 > 10)$	0	0.083
$\Pr(W_1 > 15)$	0	0.039
$\Pr(W_1 > 25)$	0	0.006
$\Pr(W_1 > 50)$	0	0.004

The lower panels of Figure 1 make it clear that full-Kelly investors face a substantial risk of poor outcomes when their probability beliefs are overstated. When the investor's belief $p = 0.6$ is correct, there is a 9% chance of finishing with less than \$2 in wealth after 30 trials. However, if the true probability is just $p = 0.525$, that same poor outcome occurs 32% of the time. Table 2 further highlights the risk: with accurate beliefs, the probability of ending below initial wealth ($W_{30} < 1$) is only 2.1%; under belief error, it jumps to 11.8%.

A more conservative benchmark is the risk-free bond portfolio, which yields $W_{30} = R_f^{30} = 4.32$. The full-Kelly investor under belief error ($p = 0.6$, true $p = 0.525$) underperforms this benchmark in 62% of simulations, compared to just 28% when the belief is accurate. These results highlight the asymmetric consequences of overconfidence. Even when the investor's forecast is only slightly wrong, the performance deterioration is severe. By examining these effects through simulation, a cautious investor can anticipate the financial costs of miscalibrated beliefs and act to hedge against them—most effectively by scaling back exposure.[14]

Finally, suppose again that the investor believes 0.6 when the true probability is 0.525. In this case, used merely to illustrate what can happen when beliefs are wrong, the full-Kelly ($\lambda = 1$) investor has expected geometric growth $E[\log(W_{30}/W_0)] = 1.0379$, which is lower than the fractional-

[13] The earliest application of Kelly-style mathematics to the stock market appears in Thorp (1969; 1971).

[14] The fractional Kelly investor also suffers from overstated probabilities, but the consequences are significantly milder.

Kelly $\lambda = 0.3$ investor's expected growth of $E[\log(W_{30}/W_0)] = 1.0490$. Thus, by attempting to maximize growth, the investor obtains a lower growth rate than that achieved by the more conservative investor, while at the same time her wealth is more volatile and follows a much less predictable growth path.

The message from these simulated results is that even when the investor "knows" the true probability distribution of the asset's payoff, there is a pragmatic argument for "underbetting," or for betting much less than the "growth optimal" Kelly fraction. When the gambler's p is sufficiently overstated, the associated full-Kelly investment fraction f^{Kelly} puts the investor on a sure and fast path to financial ruin. Intentional underbetting requires psychological self-control that works against the losses that come from over-optimism, saving the investor from unrecoverable losses. In realistic settings, it is the big losses rather than reduced profits that are not merely psychologically damaging but most likely to prevent acceptable capital growth or survival as a trader.[15]

The mathematical Kelly literature is in some ways re-invented in economics in "the theory of financial Darwinism." The original Kelly result showed that a gambler will outperform and bankrupt all others when they have (i) the correct probabilities and (ii) they bet full-Kelly. This result is explored in economics by Sandroni (2000) and Blume and Easley (2006; 2010) who show that in a competitive trading market the dominant trader is the one whose forecasting accuracy and affinity for risk-taking leads her to act *as if* she is the Kelly ideal (i.e., a full-Kelly bettor possessing true probabilities). Interestingly, that effect can occur, at least approximately, when an over-confident probability forecaster is brought back to near ideal by being so risk averse that even with overconfident beliefs she chooses bet sizes that effectively mimic a full-Kelly investor with true beliefs.

3 Portfolio Choice

Whether we take a CPT view or the professional gamblers' approach, the generally agreed idea is that investors naturally bet *as if* they overstate the small probability of a bad period in the market.

> Descriptively, a pessimistic attitude can result from an irrational belief that unfavorable events tend to happen more often, leading to an unrealistic overweighting of unfavorable likelihoods (Murphy's law). ...The decision maker may decide that unfavorable outcomes are especially important in decision making and therefore should receive more attention than equally likely favorable outcomes. (Diecidue and Wakker, 2001 p.284)

Special attention to the chance of bad outcomes, or to outcomes that rank worst relative to the other possibilities, is natural. Investors will protect themselves against the possibility of the least desirable results, when the simplest precaution is simply to bet less and live to fight another day.

Common investment advice from financial advisors, often meant to encourage investor activity and commissions, almost always rests on the favorable empirical distribution of many years of historical market returns (Malkiel, 1973). Over time the market rises and it seems on that evidence

[15]In an objective probability setting, traders are shown to outperform all others when they have correct probabilities and are full-Kelly bettors. This idea is extended by Blume and Easley (2006; 2010) who show that the dominant trader is the one whose forecasting accuracy and risk taking affinity leads her to act as if she is the kelly ideal. That can occur because an over-confident probability forecaster can be brought back to the ideal if they are sufficiently risk averse or conservative in the amounts of their bets.

alone that holding an investment in the market for at least a few years is close to a guaranteed path to higher wealth. If we take random 10 year periods in US market history, there is only a 3% chance of losing wealth over that time, thus supporting claims that "investing is not gambling" and investors should not let fear override evidence. There are illustrations in many finance textbooks of the staggering growth that would have been obtained by someone putting $1 into US stocks 100 or more years ago. Simple calculations reveal the more remarkable fact that had we known the future time series of market returns in 1926, the best outcome by 2024 would have been achieved by an investor who maintained a constant factor $f = 2.21$ of personal wealth in the market (with the rest in bonds). Thus, for every dollar of current wealth, the investor who borrowed another $1.21 and carried such a constantly levered portfolio, would be the richest of all.[16]

The question that arises is obvious; while it is clear that the stock market represents an undeniable investment opportunity, it is equally obvious that carrying anything like 221% or even 100% of wealth in the market is not a strategy that most investors will naturally adopt. Instead, real world investors might typically carry only 20%-70% of wealth in the market depending on their age and other variables.

Although there is an empirical returns distribution that investors find highly reassuring, investors view every new day, month or year on the stock market as something "out of the future" occurring under new and always unknown conditions. Knowing too well that they don't know what's coming next, a rational investor will build a level of trepidation into their choice of the component f of wealth to put at risk. In a simple model of how an instinctively apprehensive investor might think, we suppose that the investor looks at the run of historical market returns and makes hypothetical calculations of what investment fraction f would "be best" if a bad market outcome is somewhat more probable than history suggests. By making these calculations, the investor can "discover" whether their rational investment fraction is robust to that probability or very sensitive to it. The basis for making such a hypothetical calculation is instinctive and one of simply experimenting. The investor wants to invest but wants to understand how she should do so under the simple provision that the market is not as reliable as long term data suggests.

To make her simple but informative calculations, the investor first identifies the worst 30 (2.5%) monthly returns in US market history, separating these from the remaining 1152 (97.5%). There are 1,182 monthly returns in the data set.[17] These are expressed as factors R_m rather than as percentages. For example, if the market added 1.5% in the given month then R_m is 1.015. Similarly, the risk-free US bank bill rate for the same month is written as factor R_f. Histograms and descriptive statistics for the two sub-samples of market returns are as shown in Figure 3.

Building caution into the analysis, the investor imagines that a Bad month (one of the worst 30) can occur randomly with probability p somewhat greater than the empirical 2.5%. That simple step proves extremely revealing because simulations reveal that it takes only a small upward bias in p to induce a large effect on the investor's choice of f.

In our model, the investor is an expected utility maximizer who acts rationally except that she allows for a higher probability of a Bad month than is implied by empirical market data. That bias can be taken as psychological self-control or as the investor's way to allow for inherent "parameter uncertainty" about p. Apart from an upward bias in p, relative to the empirical $p = 0.025$, the investor acts rationally according to power utility function $u(w) = w^{1-\gamma}/(1-\gamma)$ for wealth w, with

[16] These calculations rest on the returns 1926–2024 provided via the Ken French Data Library. Better results would have been achieved of course by someone with more foresight who could alter fraction f year by year rather than hold it constant throughout.

[17] Obtained from Kenneth French data library.

personal risk aversion γ.[18] To test the effect of upwardly biased p on the chosen f, the investor varies p upwards from the empirical estimate 2.5% by just enough to reveal its strong effect on f. To allow for investors of different risk aversions, we show results for values of γ from 1 (full-Kelly) to 10.

The first step is to simulate the market returns distribution from the empirical returns distribution but to bias that distribution by drawing $100p\,\%$ of observations from the sample of the 30 Bad months and the remaining data from the remaining $1,182 - 30 = 1,152$ "normal" months. The bond rate for any month drawn is the rate observed in that same month. Thus, we pair each monthly market return drawn from the raw data with the corresponding risk-free rate observed in the same month, thus obtaining a simulated pair of returns (R_m, R_f). By bootstrapping a very large sample of such pairs, we can estimate the mean and variance of monthly returns and risk premia for each assumed value of p. The relevant parameters of the simulated returns are shown for different values of p in Table 3.

Table 3: Parameters of the market return R_m if the probability of a bad month is p

p	$E[R_m]$	$E[R_m - R_f]$	$\text{var}(R_m)$	$\text{var}(\log R_m)$
0.030	1.0085	0.0058	0.0028	0.0029
0.045	1.0062	0.0035	0.0032	0.0033
0.050	1.0053	0.0026	0.0033	0.0035
0.060	1.0037	0.0010	0.0035	0.0037

Using the parameter estimates found by simulation, we use the power (CRRA) utility "Merton fraction" to find analytical results for the investor's optimal investment fraction f^* as a function of personal risk aversion γ. Results for the Merton f^* are shown in Table 4, and calculated as follows:

$$f^* = \frac{\log E[R_m] - \log R_f}{\gamma\,\text{var}(\log R_m)} \approx \frac{E[R_m] - R_f}{\gamma\,\text{var}(\log R_m)}. \tag{1}$$

While the Merton formula is operational and convenient, it might be questioned in terms of accuracy because it assumes a lognormal portfolio returns distribution and is an approximation. Campbell and Viceira (2002) provide a thorough explanation of (1) as an approximation.[19] However, to confirm the accuracy of the approximate fraction, we use simulation to find the exact f that would work best for a given investor (one of given risk aversion γ) when the market returns distribution includes hypothetically one of the Bad months with probability p. By simulating the investor's result under given p and given f, it is easy to calculate the expected utility obtained by investing any given fraction f of wealth in the risky asset. With a large enough number of repeats for given p and f, we find the value of f that gives the highest average utility under that pairing.

The simulation generates random pairs (R_m, R_f) from the hypothetical population defined by p. For each pair, the investor who allocates fraction f of initial wealth $w = 1$ to the risky market obtains payoff

$$W = f\,R_m + (1 - f)\,R_f,$$

[18] Levy (2025) provides excellent background on work that finds CRRA preferences descriptively valid. See also his argument for why logically it must be that $\gamma = 1$.

[19] Similarly, the recent paper by Johnstone and Lin (2024) develops a pricing model using Campbell's results and corroborates the accuracy of (1) for situations where the portfolio return is not lognormal by testing against a distribution-free result that is solved numerically.

and hence power utility $u(W) = \frac{W(1-\gamma)}{(1-\gamma)}$. After drawing a sufficiently large sample of wealth outcomes W, the investor has a reliable estimate of expected (i.e., average) utility $\bar{u}(W)$ at each possible f, and can pick out the particular fraction f for which $\bar{u}(W)$ is highest. This numerically estimated "true best f" is then compared with the theoretical approximate best f, that is f^*, from the Merton formula. Note that f^* is not affected by the amount of starting wealth w, as is well known in CRRA theory. In Table 4, we provide results for f^* found using the two methods. The results are very close to one another throughout. It was not essential to calculate the Merton fraction to find a theoretical estimate of f^* but it is reassuring to confirm how the two approaches coincide.

These results reveal a simple but strong finding. Specifically, while the historical returns distribution tells us that in hindsight the investor's optimal investment fraction is extremely high, it takes very little concern with what might happen next on the market to shake that choice and leave the investor opting for a much lower fraction. For the growth optimal ($\gamma = 1$) investor, when we allow for just the observed historical frequency 2.5% of Bad results, the best fraction is $f^* = 2.21$. However, a minimal upward bias in that probability drives this highly risk tolerant investor's f^* very quickly lower. If we admit just a 6% chance of a Bad month, this most risk tolerant growth-seeking investor is brought to allocating only 25% of wealth to the risky market, rather than 221%. Such a large reduction suggests that the optimal investment fraction is highly sensitive to small changes in investor confidence.

Table 4 shows that similarly strong effects occur for those investors with higher risk aversion. These are investors whose risk tolerance is lower instinctively but remain very sensitive to a higher probability of a bad market. For example, if an investor with risk aversion $g = 5$ takes the empirical market returns distribution as true, she will hold 40% of wealth in the market, but if she gives the market a $p = 6\%$ chance of having a Bad month she will hold only 5% in the market.

Table 4: Best investment fraction f^* given investor's risk aversion γ and subjective probability p of a bad market outcome

γ	$p = 0.03$		$p = 0.045$		$p = 0.05$		$p = 0.06$	
	True	Merton	True	Merton	True	Merton	True	Merton
1	2.01	2.03	0.98	1.01	0.72	0.75	0.25	0.27
2	1.04	1.01	0.53	0.51	0.38	0.35	0.13	0.13
3	0.64	0.68	0.34	0.35	0.25	0.25	0.10	0.09
5	0.43	0.41	0.21	0.21	0.15	0.15	0.05	0.05
10	0.19	0.20	0.14	0.10	0.08	0.08	0.00	0.02

4 Discussion

Much of the theory of finance evolved in the 1950s and 1960s at a time when classical statistics was taken for granted. Frequentist concepts of probability dominated and decision makers talked of probabilities as if they were physical and objective. Although it is now acceptable to admit subjectivist Bayesian probability theory and the psychology of human cognitive bias, much conventional finance theory assumes known probability distributions. Think of the famous efficient frontier in portfolio theory, which is derived from a given joint probability distribution of returns. Similarly, when we calculate beta in CAPM we plug in an assumed joint returns or payoff distribution to find the beta coefficient of a given stock.

A problem for investment theory is that investors don't have probabilities that they believe to be "true probabilities." However, when trying to allow for uncertainty, investors think probabilistically while finding heuristic ways to protect themselves from poor personal probability assessments. Note that when we say poor probability assessments, we don't mean inaccurate probabilities in the sense of not being near the "true probability." Instead, we mean probability assessments that lead us to lose money, which is an objectively observable measure and can be scored without ever knowing what the probabilities "really were."[20]

With the aim of making money, and not too often losing large portions of existing wealth, investors can engineer or jolly up their models to preclude the most costly mistakes. In the preceding analysis we have seen three different devices for doing this. In mathematical gambling theory, gamblers have learnt to purposely underbet their own probability beliefs. In prospect theory, loss aversion amplifies the cost of losing a dollar and probability assessments have a bias that drags low probabilities up and high probabilities down, thus tempering overconfidence at both ends of the probability spectrum.

In the third way, illustrated by our simple calculations, investors take account of the empirical probability of a bad market outcome and temper the level of their investment to allow for a higher but still low probability of such a bad turn of events. This is a totally natural defence because it allows profitable investment while virtually eliminating the chance of ruin. Note for example the positions taken by the growth optimal Kelly investor. By "overstating" the probability of a bad market, this hypothetical investor, known to be highly risk tolerant, takes on a conservative and "sensibly risk averse" position. With a very small overstatement of the chance of a bad market, her fraction f^* invested in the market falls from the "full Kelly" fraction of more than 200% to only about 25%.

Such "underinvestment" implies that investors want the market's payoff, but are unwilling to pay the "full price" at the start of the period. Under the principle of market clearing, the opening price of the market reflects the aggregate capital allocation, that is, the average fraction of wealth f^* that investors commit. If investors in aggregate hold a lower f^*, they are implicitly demanding a higher rate of return from the value created by corporations during the period.

This logic helps explain the equity premium puzzle: underinvestment leads to high realized returns. However, as Lewellen and Shanken (2002) point out, there is a potential circularity. If the end-of-period value of the market is determined by investors themselves — through their beliefs and expectations — then the rate of return they demand becomes subjective. To avoid this loop, we note that for many stocks, end-of-period values are largely exogenous. Consider, for example, a gold mining stock whose value depends on the quantity of gold discovered. If the firm uncovers $\$G$ in net asset value during the period, this directly contributes to the stock's terminal price. A cautious investor who underinvested at the start but enjoys this exogenous payoff earns a correspondingly higher return. This same principle applies to firms realizing R&D gains or asset discoveries. Thus, rational underinvestment can generate high average returns when payoffs are exogenous.

Taking this further, if we assume exogenous asset payoffs — as if a stock is a bet with a known payoff distribution — intentional underbetting still arises naturally. Coles and Loewenstein

[20] The same approach is used in the probability scoring rules literature where probabilities are evaluated ex post by functions such as the Brier score that capture how close the probability is to the actual event, observable ex post. Tailored score functions measure the economic success of a decision based on the forecast probability. See e.g., Johnstone, Jose, and Winkler (2011).

(1988) show that investors who recognize parameter uncertainty will perceive asset payoffs as more volatile and unpredictable than under a known distribution. When parameter risk is explicitly incorporated into decision models, it reduces willingness to pay for the same expected payoff — lowering prices and raising expected returns. In this way, the rational Bayesian literature (e.g., Lewellen and Shanken, 2002) aligns with behavioral prospect theory and gambling-based intuition by providing a formal rationale for why investors underbet. It becomes clear that even without emotional bias, a rational investor facing model uncertainty will behave cautiously, resembling the same undercommitment predicted by probability weighting and regret aversion.

5 Conclusion

Investors are told that there is no better time to invest than yesterday.[21] Conventional wisdom stresses "time in the market, not timing the market." Yet the difficulty for many investors lies in making a full commitment, not because they doubt the long-run record of equity markets, but because they worry that tomorrow is not simply a random draw from yesterday's history. Market history may suggest that one should invest every available cent and wait patiently for the nearly guaranteed long-term returns. But this reasoning feels fragile in the face of genuine uncertainty about what comes next.

Investors want to hold stocks but worry about what can go wrong in the future. Their natural defence is to invest less or "underbet". Underbetting places the investor in the market but with limited exposure to unacceptably large losses. There is a clear asymmetry that supports this strategy. It can be explained either by (i) behavioral theory, such as loss aversion and probability weighting under prospect theory, or by (ii) mathematical principles from professional gambling, which show that the risk of ruin or capital loss grows disproportionately as investment size increases.

To assess investor sensitivity to the risk of loss, we run simulations assuming return distributions with a slightly higher probability of a bad market outcome than historical data suggest. The results reveal that investors are quickly deterred by even mild pessimism. Take the most risk-tolerant type of investor: the so-called "Kelly bettor" in gambling theory, or log-optimal investor in finance. Acting solely on the empirical distribution of market returns, such an investor would allocate 100% of wealth to equities — or more than 200% if allowed to borrow at the risk-free rate. However, if we adjust the model to assume that the probability of a randomly drawn month falling in the worst 2.5% of historical outcomes is actually 6%, the optimal Kelly allocation collapses to just 25%. Similarly, for a conventional investor with relative risk aversion $\gamma = 3$, the optimal allocation drops from 64% to just 10%. These findings show how even slight uncertainty about tail probabilities, or mild aversion to the unknown, can dramatically reduce rational investor commitment.

This analysis offers a behavioral and mathematical perspective on the risk premium puzzle. Suppose that, over time, firms generate exogenous value through innovation, exploration, and productivity gains, as has largely been true since the Industrial Revolution. Yet investors may worry that these advances are not indefinitely sustainable. Such doubts — the fear that the future may not resemble the past — lead investors to discount stocks more heavily and demand higher returns, even when fundamentals remain unchanged. The implication is that background risks, including unquantifiable "black swans," can materially affect expected returns. Although a black swan cannot be predicted in form or timing, simply allowing for its possibility with slightly more than zero probability leads investors to reduce valuations and accept a higher equity risk premium,

[21] See especially the arguments in the "buy and hold bible" by Siegel (2022).

not out of irrationality, but out of humility in the face of uncertainty.

References

[1] Aucamp, D. (1993). On the extensive number of plays to achieve superior performance with the geometric mean strategy. *Management Science*. 39: 1163-1172.

[2] Avramov, D. and Zhou, G. (2010) Bayesian portfolio analysis. *Annual Review of Financial Economics*. 2:25-47.

[3] Barry, C.B. (1974) Portfolio analysis under uncertain means, variances and covariances. *Journal of Finance*. 29: 515-522.

[4] Bauder, D., Bodnar, T., Parolya, N. and Schmid, W. (2021) Bayesian mean–variance analysis: optimal portfolio selection under parameter uncertainty. *Quantitative Finance*. 21: 221-242.

[5] Benartzi, S. and Thaler, R.H. (1995) Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*. 110: 73–92.

[6] Black, F. and Litterman, R. (1992) Global portfolio optimization. *Financial Analysts Journal*. 48: 28-43.

[7] Blume, L. and Easley, D. (2006) If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets. *Econometrica.* 74: 929-966.

[8] Blume, L. and Easley, D. (2010) Heterogeneity, selection, and wealth dynamics. *Annual Review of Economics*. 2: 425-450.

[9] Campbell, J.Y. and Viceira, L.M. (2002) *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*. Oxford University Press.

[10] Coles, J., and Loewenstein, U. (1988) Equilibrium pricing and portfolio composition in the presence of uncertain parameters and estimation risk. *Journal of Financial Economics*. 22: 279-303.

[11] de Finetti, B. (1974) *Theory of Probability*. Vols 1 and 2. New York: John Wiley & Sons.

[12] Diecidue, E. and Wakker, P.P. (2001) On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*. 23: 281–298.

[13] Jadlow, J.W. and Mowen, J.C. (2010) Comparing the Traits of Stock Market Investors and Gamblers. *Journal of Behavioral Finance*. 11: 67–81.

[14] Johnstone, D.J., Jose, V.R.R. and Winkler, R.L. (2011) Tailored scoring rules for probabilities. *Decision Analysis*. 8: 256–268.

[15] Johnstone, D.J. and Lin, Y-X. (2024) Why Investors Want Risk. *Decision Analysis*. 22: 14-29.

[16] Kale, J.K. (2006) Growth Optimization with Downside Protection: A New Paradigm for Portfolio Selection. *Journal of Behavioral Finance*. 7: 29–42.

[17] Kelly, J. L. (1956) A new interpretation of information rate. *Bell System Technical Journal*. 35: 917–926.

[18] Levy, M. (2025) Relative risk aversion must be close to 1. *Annals of Operations Research*. 346: 127-135.

[19] Lopes, L.L. and Oden, G.C. (1999).The role of aspiration level in risky choice: a comparison of cumulative prospect theory and SP/A Theory. *Journal of Mathematical Psychology*. 43:286-313.

[20] Lewellen, J. and Shanken, J. (2002) Learning, asset-pricing tests, and market efficiency. *Journal of Finance*. 57: 1113-1145.

[21] Levy, M. (2025) Projects with no cost of capital. *Financial Management*. 54: 177–191.

[22] MacLean, L.C., Sanegre, R., Zhao, Y. and Ziemba, W.T. (2004) Capital growth with security. *Journal of Economic Dynamics and Control*. 28: 937-954.

[23] MacLean, L.C. and Ziemba, W.T. (2006) Capital growth theory and practice. In: Zenios, S.A. and Ziemba, W.T. (eds) *Handbook of Asset and Liability Management, Volume 1: Theory and Methodology*. Amsterdam: North Holland. pp. 429–473.

[24] Maclean, L.C., Thorp, E. and Ziemba, W.T. (2010) Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria. *Quantitative Finance*. 10: 681-687.

[25] MacLean, L.C., Ziemba, W.T. and Blazenko, G. (1992) Growth versus security in dynamic investment analysis. *Management Science*. 38: 1562-1585.

[26] MacLean, L.C. and Ziemba, W.T. (1999) Growth versus security tradeoffs in dynamic investment analysis. *Annals of Operations Research*. 85: 193-225.

[27] Malkiel, B. (1973) *A Random Walk Down Wall Street*, New York: Norton.

[28] Pratt, J.W. and Zeckhauser, R.J. (1987) Proper risk aversion. *Econometrica*. 55: 143–154.

[29] Rabin, M. (2000) Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica*. 68: 1281–1292.

[30] Rachev, S., Hsu, J., Bagasheva, B., and Fabozzi, F. (2008) *Bayesian Methods in Finance*. Wiley.

[31] Sandroni, A. (2000) Do markets favor agents able to make accurate predictions? *Econometrica*. 68: 1303-1342.

[32] Samuelson, P.A. (1963) Risk and uncertainty: a fallacy of large numbers. *Scientia*. 98: 108-113.

[33] Siegel, J.J. (2022) *Stocks for the Long Run : The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies* 6th ed. New York: McGraw-Hill.

[34] Thorp,E. O. (1969) Optimal gambling systems for favorable games.*Review of the International Statistical Institute*. 37: 273–293.

[35] Thorp, E.O. (1971) Portfolio choice and the Kelly criterion. In: *Proceedings of the 1971 Business and Economics Section of the American Statistical Association*, pp.215–224.

[36] Thorp, E.O. (2011) The Kelly criterion in blackjack sports betting, and the stock market. In: MacLean, L.C., Thorp, E.O., Ziemba, W.T. (eds) *The Kelly Capital Growth Investment Crtierion Theory and Practice*. Singapore: World Scientific Publishing. pp. 789–832.

[37] Tversky, A. and Kahneman, D. (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*. 5: 297–323.

[38] Tversky, A. and Wakker, P. (1995) Risk attitudes and decision weights. *Econometrica*. 63: 1255–1280.

Figure 1: Simulated distributions of investor wealth under λ-Kelly strategies over 30 trials. Each panel represents the bootstrap distribution of final wealth after 30 investment rounds, based on 100,000 simulated paths.
Investors bet a fraction λ of the Kelly-optimal amount, where the believed probability of success is $p = 0.6$. In the lower panels, the true success probability is only $p = 0.525$. Lighter regions denote more frequent outcomes.

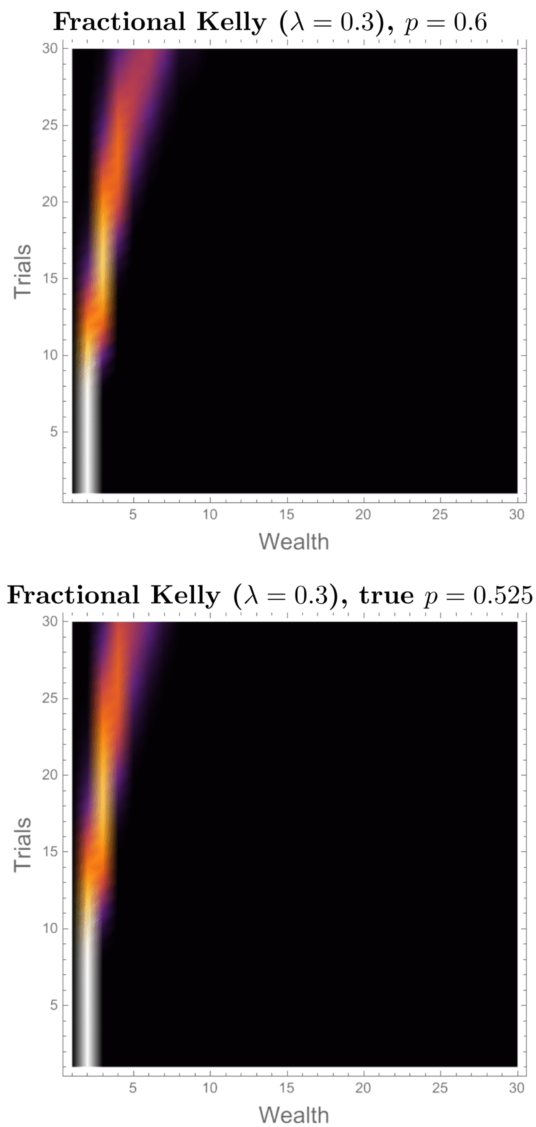

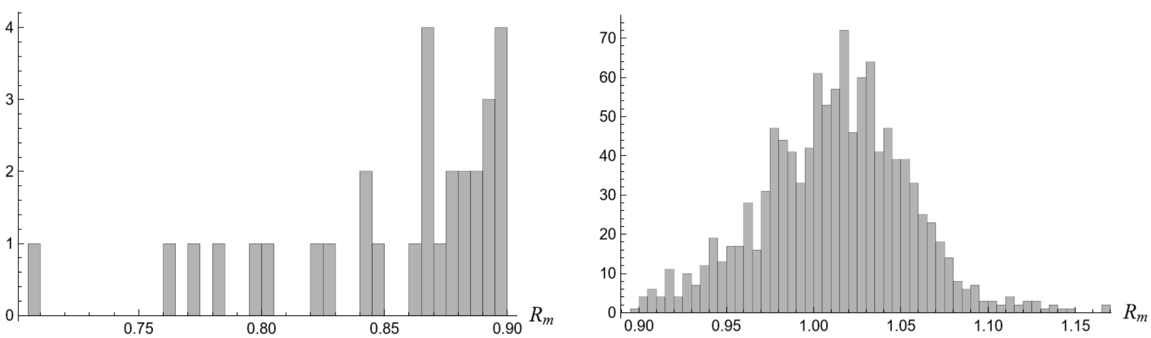

Worst 30 Months
$E[R_m] = 0.853$, standard deviation $= 0.048$

Other Months
$E[R_m] = 1.013$, standard deviation $= 0.045$

Figure 2: Empirical distributions of market returns during worst months vs. all others.