STOCHASTIC MARKET EFFICIENCY

OLE PETERS
AND ALEXANDER ADAMOU

"NEITHER A BORROWER NOR A LENDER BE"

Hamlet Act 1, Scene 3, 75

It is argued that the simple trading strategy of leveraging or deleveraging an investment in the market portfolio cannot outperform the market. Such stochastic market efficiency places strong constraints on the possible stochastic properties of the market. Historical data confirm the hypothesis.

KEYWORDS: Stochastic market efficiency.

In Sec. 1 we summarize a few key properties of geometric Brownian motion that were pointed out in (Peters, 2010). We indicate the main elements of the analogy that is often drawn between this and the dynamics of markets. Section 2 introduces the concept of stochastic efficiency, namely the hypothesis that the properties of price fluctuations in real markets are strongly constrained by stability and efficiency arguments so as to make investments of leverage 1 optimal. This hypothesis is motivated by the considerations in Sec. 1 but goes beyond the simple model discussed there. Section 3 constitutes the main body of the study, where the hypothesis is tested empirically using data from American stock markets. The arguments leading to the hypothesis are neither specific to the mathematical model nor to American stock markets. We choose the model because it is analytically tractable and American stock markets because they have been well observed for a sufficiently long time. The arguments are robust enough to yield insights into other assets, such as housing or indeed national economies or the global economy. Section 4 summarizes the rationale and main results of the study.

1. MATHEMATICAL BACKGROUND

Notation: The present study uses three different levels of realism. To avoid tedious nomenclature and confusion between these, we use three different superscripts:

1. Superscript ^{m} refers to the mathematical toy model used to motivate and guide our investigations.
2. Superscript ^{s} refers to simulations performed using market data to test our main hypothesis.
3. Superscript ^{r} refers to vaguely corresponding quantities in the context real people's behavior.

Toy model: The variable $x(t)$ is said to undergo geometric Brownian motion if

$$(1) \qquad dx(t) = x(t)(\mu^{\mathrm{m}} dt + \sigma^{\mathrm{m}} dW),$$

where dW is a Wiener process. Owing to the analogy with market dynamics, where $x(t)$ is likened to the price of an asset or the value a market index, we call μ^{m} the expected return (see below) and σ^{m} the volatility.

Equation (1) is non-stationary, in the sense that $P_x(x,t)$, the distribution of $x(t)$, does not converge to a probability distribution in the limit $t \to \infty$. Rather, $\lim_{t\to\infty} P_x(x,t) = 0$ for any value of x. This implies that the process is non-ergodic in the sense of (Grimmet and Stirzaker, 2001) Ch. 9. Therefore, even if both quantities are meaningful, the time average of an observable arising from this process need not be identical to the ensemble average (over all possible realizations of the process) of the same observable.

Following (Peters, 2010) it is possible to compute the ensemble average, $\langle g \rangle$, and the time average, denoted by $\overline{g}$, of the exponential growth rate of $x(t)$ using the estimator

$$(2) \qquad \widehat{g^{\mathrm{m}}}(T,N) \equiv \frac{1}{T} \ln \left(\frac{1}{N} \sum_i^N \left(\frac{\int_0^T dx_i(t)}{x_i(0)} \right) \right),$$

where i indexes realizations of the process described by (Eq. 1). Taking the limit $N \to \infty$ for finite T defines the ensemble average, while taking the limit $T \to \infty$ for $N = 1$ defines the time average. This procedure yields clear interpretations of two well-known characteristics of geometric Brownian motion:

$$(3) \qquad \langle g^{\mathrm{m}} \rangle \equiv \lim_{N\to\infty} \widehat{g^{\mathrm{m}}}(T,N) = \mu^{\mathrm{m}}$$

shows that the ensemble-average growth rate is equal to the expected return; and

$$(4) \qquad \overline{g^{\mathrm{m}}} \equiv \lim_{T\to\infty} \widehat{g^{\mathrm{m}}}(T,1) = \mu^{\mathrm{m}} - \frac{(\sigma^{\mathrm{m}})^2}{2} = \frac{\langle d\ln(x) \rangle}{dt}$$

shows that the time-average growth rate is equal to the rate of change of the ensemble average of $\ln(x(t))$, which is obtained by applying Itô's formula to (Eq. 1).

Referring to the analogy with stock markets, it was pointed out in (Peters, 2010) that an investor should be more concerned about the time average of a single realization of the process (his investment) than the average over many realizations (conceptually equivalent to multiple copies of his investment in parallel universes).

We introduce another parameter, l^{m}, to (Eq. 1) in order to extend the market analogy to leveraged investments. We split the expected return into a so-called

"riskless" part, $\mu_{\text{riskless}}^{\text{m}}$, (to represent an asset with zero volatility) and a so-called "excess" part, $\mu_{\text{excess}}^{\text{m}}$, that is added in proportion to the amplitude of the fluctuations $l^{\text{m}}\sigma^{\text{m}}$. Thus (Eq. 1) becomes

$$(5) \qquad dx_l = x_l((\mu_{\text{riskless}}^{\text{m}} + l^{\text{m}}\mu_{\text{excess}}^{\text{m}})dt + l^{\text{m}}\sigma^{\text{m}}dW).$$

The case $l^{\text{m}} = 0$ results in exponential growth with rate $\mu_{\text{riskless}}^{\text{m}}$, and the case $l^{\text{m}} = 1$ is equivalent to (Eq. 1), where $\mu^{\text{m}} = \mu_{\text{riskless}}^{\text{m}} + \mu_{\text{excess}}^{\text{m}}$.

Leverage $l^{\text{m}} < 0$ reflects short-selling; $0 \le l^{\text{m}} \le 1$ reflects part of the investor's equity being invested in the market and part kept safe paying a return of $\mu_{\text{riskless}}^{\text{m}}$; and $l^{\text{m}} > 1$ reflects what is commonly referred to as leveraging, *i.e.* an investment in the market that exceeds the investor's equity and includes borrowed funds. The volatility in $x_l(t)$, reflecting the investor's equity, is $l^{\text{m}}\sigma^{\text{m}}$, proportional to the leverage, and the expected return of $x_l(t)$ is $\mu_{\text{riskless}}^{\text{m}} + l^{\text{m}}\mu_{\text{excess}}^{\text{m}}$, reflecting a safe interest rate and the excess expected return of the market added in proportion to the leverage. Thus leveraging causes both the excess return and the fluctuations to increase linearly.

Equation (5) has the leveraged ensemble-average growth rate

$$(6) \qquad \langle g_l^{\text{m}} \rangle = \mu_{\text{riskless}}^{\text{m}} + l^{\text{m}}\mu_{\text{excess}}^{\text{m}}$$

and the leveraged time-average growth rate

$$(7) \qquad \overline{g_l^{\text{m}}} = \mu_{\text{riskless}}^{\text{m}} + l^{\text{m}}\mu_{\text{excess}}^{\text{m}} - \frac{(l^{\text{m}}\sigma^{\text{m}})^2}{2}.$$

Crucially, (Eq. 7), unlike (Eq. 6), is non-monotonic in l^{m} and establishes the existence of an objectively optimal leverage,

$$(8) \qquad l_{\text{opt}}^{\text{m}} = \frac{\mu_{\text{excess}}^{\text{m}}}{(\sigma^{\text{m}})^2},$$

which maximizes $\overline{g_l^{\text{m}}}$.

Equation (8) implies that unless $l_{\text{opt}}^{\text{m}} = \mu_{\text{excess}}^{\text{m}}/(\sigma^{\text{m}})^2 = 1$ it is possible to choose l^{m} in (Eq. 5) such that $x_l(t)$ (reflecting the dynamics of equity in a leveraged investment) consistently outperforms $x(t)$ of (Eq. 1) (reflecting the market portfolio). Due to the non-linear effects of fluctuations in the multiplicative process (Eq. 1) or (Eq. 5) it is mathematically possible to outperform a rising market by keeping a fraction of one's equity in a savings account.

In reality, the outcome of an investment held for some time is given by the time-average growth of the (real) investment over that time period. The ensemble-average growth rate is *a priori* irrelevant in practice; it is a means of conceptualizing randomness (Peters, 2011; Gell-Mann and Hartle, 2007; Gell-Mann and Lloyd, 2004; Ehrenfest and Ehrenfest, 1912). Attempting to optimize (Eq. 6) leads to the recommendation of maximizing l^{m} (or $-l^{\text{m}}$). But (Eq. 7) shows that this would lead to a negatively diverging time-average growth rate. Thus,

if (Eq. 6) is falsely believed to reflect the quantity an investor should optimize, and l^{m} is interpreted as the leverage used in the investment, the investor will be led to exceed (positively or negatively) the leverage that would truly be most beneficial – worse, this excess is likely to ruin the investor.

There is a long history of the struggle to make sense of the misleading recommendations derived from (Eq. 6), starting with N. Bernoulli, see (de Montmort, 1713), and carrying on to (Markowitz, 1952, 1991), see (Peters, 2011). The optimal leverage for (Eq. 5) was identified at least as early as (Kelly Jr., 1956), albeit in a somewhat different context and form, but has been used seemingly little by investors since then. The present study is concerned with the dynamic properties of the optimal simulated leverage derived from time series of real markets.

2. STOCHASTIC EFFICIENCY

The efficient market hypothesis (in its weakest form) (Bachelier, 1900; Fama, 1965) states that the price of an asset traded in an efficient market reflects all the information publicly available about the asset. The corollary is that it is impossible for a market participant, without access to privileged ("insider") information, consistently to achieve returns exceeding the time-average growth rate of the market ("to beat the market") by choosing at which prices to buy and sell assets. We shall refer to this concept as "price efficiency".

We propose another kind of market efficiency, which we call "stochastic efficiency": it is impossible for a market participant without privileged information to beat the market by choosing *how much* to invest in an asset or portfolio of assets, *i.e.* by choosing his leverage. Simple strategies such as borrowing money to invest, $l^{\mathrm{r}} > 1$, and keeping some money in the bank, $l^{\mathrm{r}} < 1$, should not yield consistent market outperformance ("no leverage-arbitrage"). This reasoning was used in (Peters, 2010) to hypothesize that real markets self-organize such that

$$(9) \qquad l^{\mathrm{r}}_{\mathrm{opt}} = 1$$

is an attractive point for their stochastic properties (represented by $\mu^{\mathrm{m}}_{\mathrm{riskless}}, \mu^{\mathrm{m}}_{\mathrm{excess}}$ and σ^{m} in the model).

The hypothesis we are about to test is motivated by the model (Eq. 5) and its properties (Eq. 6), (Eq. 7) and (Eq. 8) in the sense that this model motivates the existence of an optimal leverage. But it is by no means derived from the model, as the hypothesis requires the dynamic adjustment, or self-organization, of the stochastic properties of the system themselves. One would have to think of $\mu^{\mathrm{m}}_{\mathrm{riskless}}, \mu^{\mathrm{m}}_{\mathrm{excess}}$ and σ^{m} as slowly (compared to the fluctuations) varying functions of time, related to one another as well as to l^{m} through a dynamic which has $l^{\mathrm{m}}_{\mathrm{opt}} = \mu^{\mathrm{m}}_{\mathrm{excess}}/(\sigma^{\mathrm{m}})^2 = 1$ as an attractor. We do not devise a quantitative model of this here, but discuss qualitative features which such a model may require. The reader is referred to the work by others who have proposed and investigated quantitative models, see (Geanakoplos, 2010; Thurner et al., 2009) and references therein.

Although inspired by a mathematical toy model, the hypothesis in (Eq. 9) does not rest on model-specific properties. Crucial for it are the identification of the time-average growth rate as the practically relevant measure for deciding levels of investment and the consequent establishment of an optimal leverage, about which economic arguments may be framed.

Stochastic efficiency is a tantalizing concept. It posits that the market has a different quality of knowledge than implied by price efficiency. Price efficiency is essentially a static concept, as it states that prices coincide with some form of value, but it does not specify the fast dynamics and fluctuations necessary to equilibrate to these values. Stochastic efficiency, on the other hand, constrains price *dynamics* and predicts properties of fluctuations.

To argue convincingly for stochastic efficiency we have to elucidate those aspects of the dynamics that enforce it. For stochastic efficiency to manifest itself, there needs to be negative feedback similar to the feedback between prices and supply-demand imbalances. In addition to the "no leverage-arbitrage" argument above, the following dynamic arguments show that both $l_{\text{opt}}^{\text{r}} = 1$ and $l_{\text{opt}}^{\text{r}} = 0$ are particularly attractive values, and that the interval $0 \leq l_{\text{opt}}^{\text{r}} \leq 1$ constitutes a potentially stable regime, whereas values outside it are unstable:

1. **Leverage feedbacks:**

    (a) If $l_{\text{opt}}^{\text{r}} > 1$, investors will eventually borrow money to invest. Highly leveraged investments are liable to margin calls and tend to increase volatility (Geanakoplos, 2010). In addition, high leverage will lead to price increases and (through an outflow from deposit accounts) to an increase in yields on safe bonds $\mu_{\text{riskless}}^{\text{r}}$. Both effects reduce $\mu_{\text{excess}}^{\text{r}}$. Both the rise in σ^{r} and the fall in $\mu_{\text{excess}}^{\text{r}}$ imply that $l_{\text{opt}}^{\text{r}}$ decreases.

    (b) If $0 < l_{\text{opt}}^{\text{r}} < 1$, there can be no margin calls on optimally leveraged investments, leading to lower volatility, and as l^{r} decreases, *i.e.* investors withdraw from the market, prices fall, yields on safe bonds decrease, and $\mu_{\text{excess}}^{\text{r}}$ increases. Thus the pressures from case (a) for l^{r} to decrease are relaxed and l^{r} is free to increase.

    (c) If $l_{\text{opt}}^{\text{r}} < 0$ investors will eventually borrow stock to short-sell. Highly negatively leveraged investments are liable to margin calls and tend to increase volatility. High negative leverage will lead to price decreases and (through an inflow into deposit accounts) to a decrease in yields on safe bonds $\mu_{\text{riskless}}^{\text{r}}$, making $\mu_{\text{excess}}^{\text{r}}$ less negative. Both the rise in σ^{r} and the increase in $\mu_{\text{excess}}^{\text{r}}$ imply that $l_{\text{opt}}^{\text{r}}$ becomes less negative.

2. **Global stability:** It is difficult to envisage globally stable economies existing with optimal leverage outside the interval $0 \leq l_{\text{opt}}^{\text{r}} \leq 1$ because:

    (a) If $l_{\text{opt}}^{\text{r}} > 1$ everyone should invest in the market more than he owns. This is not possible because the funds to be invested must be provided by someone.

(b) If $l^{\mathrm{r}}_{\mathrm{opt}} < 0$ everyone should sell more market shares than he owns. This is not possible because the assets to be sold must be provided by someone.

Thus the range $0 \leq l^{\mathrm{r}}_{\mathrm{opt}} \leq 1$ is special in not being globally unstable. We believe the above to be the main drivers behind stochastic efficiency. There are additional effects, however, which reinforce it.

1. **Economic paralysis:** In an economy with $l^{\mathrm{r}}_{\mathrm{opt}} \leq 0$ there is no incentive to invest, which may result in no productive economic activity.
2. **Covered short-selling:** An investment with $l^{\mathrm{r}} < 0$ is punished by the costs of borrowing stock to short-sell, *i.e.* covered as opposed to naked short-selling, which makes $l^{\mathrm{r}}_{\mathrm{opt}} = 0$ special.
3. **Risk premiums:** The interest received by a depositor is typically less than the interest paid by a borrower. Therefore, an investment with $l^{\mathrm{r}} < 1$ is punished by low deposit interest rates and an investment with $l^{\mathrm{r}} > 1$ is punished by high borrowing costs. This reinforces $l^{\mathrm{r}}_{\mathrm{opt}} = 1$ as an attractive point.
4. **Transaction costs:** The costs of buying and selling assets (fees, market spreads, etc.) punish any strategy that requires trading. Holding an investment of constant leverage generally requires trading to rebalance the ratio of assets to equity. The two exceptions are investments with $l^{\mathrm{r}} = 0$ and $l^{\mathrm{r}} = 1$.

Following these considerations we arrive at a refined hypothesis: on sufficiently long time scales $l^{\mathrm{r}}_{\mathrm{opt}} = 1$ is a strong attractor (which we refer to as "strong" stochastic efficiency). Deviations from this attractor are likely to be confined to the interval $0 \leq l^{\mathrm{r}}_{\mathrm{opt}} \leq 1$ ("weak"), whose end points are sticky. In the following we submit this hypothesis to an empirical test using market data.

3. TESTS OF STOCHASTIC EFFICIENCY IN HISTORICAL DATA

We test the stochastic efficiency hypothesis by simulating leveraged investments in the Standard & Poor's index of 500 leading U.S. companies (S&P500) using historical data of the daily returns of the index over the last 55 years.

3.1. *Data sets*

The data used in this study are publicly available from the Federal Reserve Economic Data (FRED) website, hosted by the Federal Reserve Bank of St. Louis at `http://research.stlouisfed.org/fred2`. We use the daily closing prices, adjusted for dividends and stock splits, of the S&P500 (FRED time series: SP500) from 4^{th} August 1955 to 17^{th} November 2010. Additionally we use the daily effective federal funds rate (FRED time series: DFF) and the daily bank prime loan rate (FRED time series: DPRIME) over the same period. The estimate of optimal leverage using these data is generous, as the S&P500 represents a well diversified portfolio of large and successful companies, and – since bankrupt

companies (such as Enron) are replaced – is positively affected by survivorship bias. All studies were repeated for Dow Jones Industrial Average and NASDAQ data with essentially identical results (not shown).

3.2. *Simulations*

An investment of constant leverage over a given time period, or "window," is simulated as follows. At the start of the first day we assume unit equity, comprising holdings of l^s in the risky asset (S&P500) and cash deposits of $1 - l^s$. At the end of the day the values of these holdings and deposits are updated according to the historical market returns and interest rates, after which the portfolio is rebalanced, *i.e.* the holdings in the risky asset are adjusted so that their ratio to the total equity is again l. On non-trading days the return of the market is zero, whereas deposits continue to return interest payments, which leads to an unrealistic but negligible rebalancing on those days. The investment proceeds in this fashion until the final day of the window, at which point the final equity is recorded. If at any time the total equity falls below zero, the investment is declared bankrupt and the simulation stops, *i.e.* we do not allow recovery from negative equity. The procedure is then repeated for different leverages, and the optimal simulated leverage, l^s_{opt}, is the leverage for which the final equity is maximized (found using a golden section search algorithm (Press et al., 2002), Chap. 9).

Figure 1 shows as a function of leverage simulated returns for an investment over the largest possible window, namely the entire time series. The four curves in the figure correspond to four sets of assumptions about interest rates and transaction costs, mentioned as additional effects in Sec. 2. We list these in order of increasing complexity and resemblance of actual conditions and practices in financial markets:

- **Simulation 1** (red line in Fig. 1) is the simple case, where the effective federal funds rate is applied to all cash, whether deposited or borrowed. No costs are incurred for short-selling ($l^s < 1$), akin to naked short-selling, wherefore market returns apply to negative holdings in the market exactly as they apply to positive holdings. Transaction costs are neglected. This results in a smooth curve.
- **Simulation 2** (yellow line) is like the first case, but federal interest rates are paid on short positions, corresponding to fees for borrowed stock. This penalization of negative holdings in the market introduces a discontinuity in the derivative, or kink, at $l^s = 0$.
- **Simulation 3** (green line) is like the second case, but now federal interest rates are received on cash deposits, whereas prime interest rates are paid on borrowed funds or stock. This resembles the effect of risk premiums and introduces a kink at $l^s = 1$.
- **Simulation 4** (blue line), the complex case, is like the third case, but whenever the portfolio is rebalanced a loss in equity of 0.2% of the value
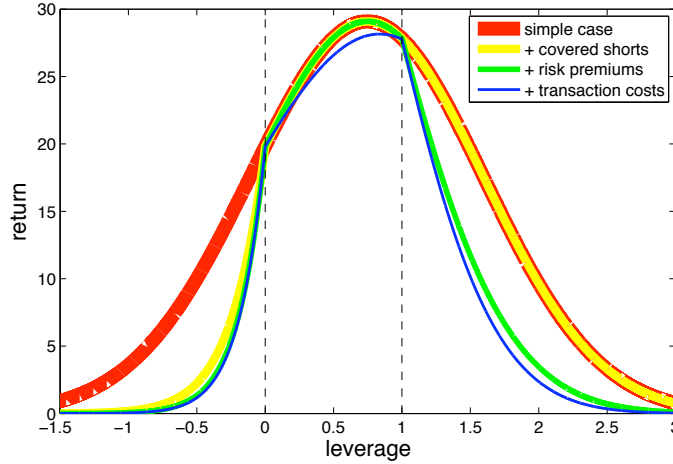
FIGURE 1.— Total return from a constant-leverage investment in the S&P500, starting 4^{th} August 1955 and ending 17^{th} November 2010, as a function of leverage. For descriptions of the simulations, see text.
Red line: Simulation 1.
Yellow line: Simulation 2.
Green line: Simulation 3.
Blue line: Simulation 4.

of the assets traded is incurred. This resembles transaction costs.

As discussed in Sec. 2, risk premiums, covered short-selling and transaction costs tend to penalize investments with leverages other than 0 or 1. This is reflected in the simulation results by the kinks described above and visible in Fig. 1.

For many time windows the discontinuity in the derivative of the return-leverage curve is accompanied by a change in sign of the derivative, making the point a global maximum and fixing $l_{\text{opt}}^{\text{s}}$. Certainly, this corresponds to a real effect observed in real markets. However, even without these effects, the red line shows that $0 \le l_{\text{opt}}^{\text{s}} \le 1$. This, being the simplest case with the fewest assumptions and approximations, despite also being the least realistic, provides the strongest support for our hypothesis.

3.3. *The entire time series*

The return-leverage curve for an investment window spanning the entire time series over the last 55 years shows an optimal leverage of $l_{\text{opt}}^{\text{s}} = 0.75$ for the simple case (simulation 1) and $l_{\text{opt}}^{\text{s}} = 0.84$ for the complex case (simulation 4).

We will discuss in Sec. 3.4 in how far this confirms the hypothesis.

The time-average growth rate (Eq. 7) of the specific model (Eq. 5) is parabolic in l^{m}. We show in Fig. 2 the simulated time-average growth rate for the simple case, which is the logarithm of the simulated return, shown in Fig. 1, divided by the window length. Given the known deficiencies of the geometric Brownian motion model, the quality of a parabolic fit (black dashed line) is remarkable. The parameters of the fitted parabola are a fit of the model (Eq. 7) and can be taken as meaningful definitions in the simulations of the riskless return, $\mu^{\mathrm{s}}_{\mathrm{riskless}}$, the excess return, $\mu^{\mathrm{s}}_{\mathrm{excess}}$, and the volatility, σ^{s}, for the S&P500 over the last 55 years. A least-squares fit estimates these parameters as $\mu^{\mathrm{s}}_{\mathrm{riskless}} = 5.4\%$ $p.a.$, $\mu^{\mathrm{s}}_{\mathrm{excess}} = 1.9\%$ $p.a.$, and $\sigma^{\mathrm{s}} = 15.8\%$ per square root of one year. We used a one-parameter fit, fixing the co-ordinates of the maximum first, and then fitting on the range $-7 \leq l^{\mathrm{s}} \leq 3$. The ensemble-average annual outperformance of the S&P500 over federal deposits over this period amounts to a mere 1.9%. Due to the wealth-depleting effect of the substantial volatility, which manifests itself in the model as the $-(l^{\mathrm{m}}\sigma^{\mathrm{m}})^2/2$ term in (Eq. 7), this small excess is insufficient for a full investment ($l^{\mathrm{s}} = 1$) to be optimal. Full investment does outperform federal deposits, albeit barely, by 0.6% $p.a.$.

3.4. *Shorter time scales*

Over the full time series our simulations yield an optimal leverage between zero and one. How significant a corroboration of our hypothesis is this? Does it rule out $l^{\mathrm{r}}_{\mathrm{opt}} = 1$? Even assuming that $l^{\mathrm{r}}_{\mathrm{opt}}$ is attracted to a particular value, we expect random deviations from it to increase as the investment window gets shorter, since returns over shorter windows are more heavily influenced by noise. To take an extreme example, in the simple simulation of a daily rebalanced portfolio, the observed optimal leverage over a window of one day can take one of two values: $l^{\mathrm{s}}_{\mathrm{opt}} = +\infty$ if the market beats the federal funds rate on that day, $l^{\mathrm{s}}_{\mathrm{opt}} = -\infty$ if federal funds beat the market. Indeed, the magnitude of the observed optimal leverage will be infinite for any window over which the daily returns are either all greater than, or all less than, the federal funds rate. This is unlikely for windows of months or years but occurs commonly over windows of days or weeks. The longest run of consecutive up-moves relative to the federal funds rate in the S&P500 was 14 trading days from 26[th] March 1971 to 15[th] April 1971, and the longest draw-down relative to the federal funds rate was 12 trading days from 22[nd] April 1966 to 9[th] May 1966 and from 10[th] November 1969 to 25[th] November 1969. Even without this divergence, shorter windows are more likely to result in larger positive and negative optimal leverages because relative fluctuations are larger over shorter time scales.

We quantify this idea in the model: solving (Eq. 5) yields the following estimate for the time-average growth rate after a finite time, T:

$$(10) \qquad \overline{g^{\mathrm{m}}_l}(T, N = 1) = \mu^{\mathrm{m}}_{\mathrm{riskless}} + l^{\mathrm{m}}\mu^{\mathrm{m}}_{\mathrm{excess}} - \frac{(l^{\mathrm{m}}\sigma^{\mathrm{m}})^2}{2} + \frac{l^{\mathrm{m}}\sigma^{\mathrm{m}}W(T)}{T}.$$
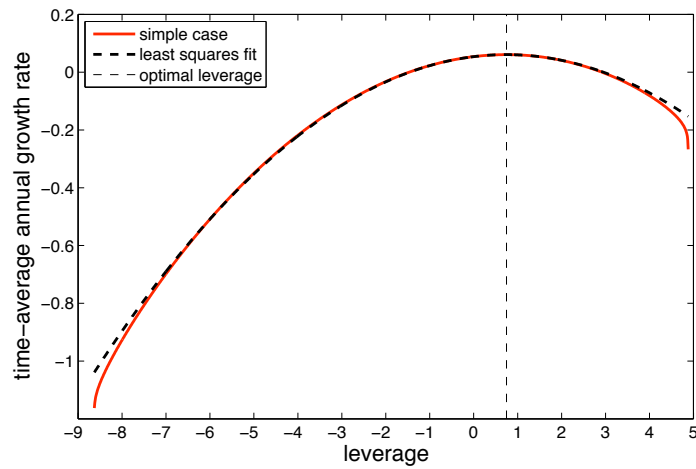
FIGURE 2.— Computed time-average growth rates closely follow a parabola as a function of leverage. The deviation from parabolic form for extreme leverages is due to crashes and sudden recoveries. A 20.47% daily drop as occurred on 19 October 1987 leads to bankruptcy (an infinitely negative return) in our simulations for leverage $l^s >= 4.89$, and a 11.58% as occurred on 13 October 2008 leads to bankruptcy for $l^s < -8.64$, showing the well-known asymmetry between negative and positive extreme events.

Maximizing this generates an estimate for the optimal leverage over a window of length T,

$$(11) \qquad \widehat{l_{\mathrm{opt}}^{\mathrm{m}}}(T, N = 1) = l_{\mathrm{opt}}^{\mathrm{m}} + \frac{W(T)}{\sigma^{\mathrm{m}} T}.$$

Thus, in the model optimal leverage for finite-time windows is normally distributed with mean $l_{\mathrm{opt}}^{\mathrm{m}}$ and standard deviation

$$(12) \qquad \mathrm{stdev}(\widehat{l_{\mathrm{opt}}^{\mathrm{m}}}(T, N = 1)) = \frac{1}{\sigma^{\mathrm{m}} T^{1/2}}.$$

To assess the significance of our finding $l_{\mathrm{opt}}^{\mathrm{s}} = 0.75$ (or 0.84 for the complex simulation), we estimate the time scale at which the standard deviation of $\widehat{l_{\mathrm{opt}}^{\mathrm{s}}}$ is 1. Only for time windows on this scale or larger can the assertion that $l_{\mathrm{opt}}^{\mathrm{s}}$ is confined to a range of size one be significantly corroborated by a single measurement (the window of the entire record). Substituting the computed volatility $\sigma^{\mathrm{s}} = 15.8\%$ per square-root of one year, as estimated in Sec. 3.3, for the model volatility σ^{m} would fix this time scale at about 40 years. Since we do not trust the specific form of the model, however, we empirically test the relation suggested by (Eq. 12).

Likening $\widehat{l_{\mathrm{opt}}^{\mathrm{m}}}(T, N = 1)$ to observed optimal leverage over a window of size T, we investigate how well the model predicts the fluctuations in $l_{\mathrm{opt}}^{\mathrm{s}}(T)$. We compile histograms of $l_{\mathrm{opt}}^{\mathrm{s}}(T)$ by moving windows of size T across the record and compare the standard deviation of $l_{\mathrm{opt}}^{\mathrm{s}}(T)$ found in these histograms to the standard deviation of $\widehat{l_{\mathrm{opt}}^{\mathrm{m}}}(T, N = 1)$. Even for windows considerably shorter than the entire record (months or a few years), the standard deviations of the corresponding histograms are considered meaningful, and the relation (Eq. 12) can be tested.

Figure 3 shows, on double-logarithmic scales, the standard deviation of $l_{\mathrm{opt}}^{\mathrm{s}}$ against the window length for the simple case. Remarkable agreement is found with the model-specific prediction in (Eq. 12). We note that, for shorter time scales, the standard deviation is slightly higher than predicted. For longer time scales, the standard deviation drops below the prediction. This is because for window lengths approaching the entire period under study, the number of independent windows in the sample is small, and hence the standard deviation of the sample is depressed.

In Fig. 4 the diminishing fluctuations in $l_{\mathrm{opt}}^{\mathrm{s}}$ are illustrated as follows: for every day the optimal leverage is computed for the longest available window (the window starting on 4^{th} August 1955) for the simple and complex cases. As time passes and the statistics improve, the optimal leverage is seen to be consistent with an approach to $l_{\mathrm{opt}}^{\mathrm{s}} = 1$. About $1/3$ of the measurements lie outside the one-standard deviation band, as would be the case in the model. No measurements lie outside the two-standard deviations band, whereas this would occur about 5% of the time in the model. Of course the period investigated could
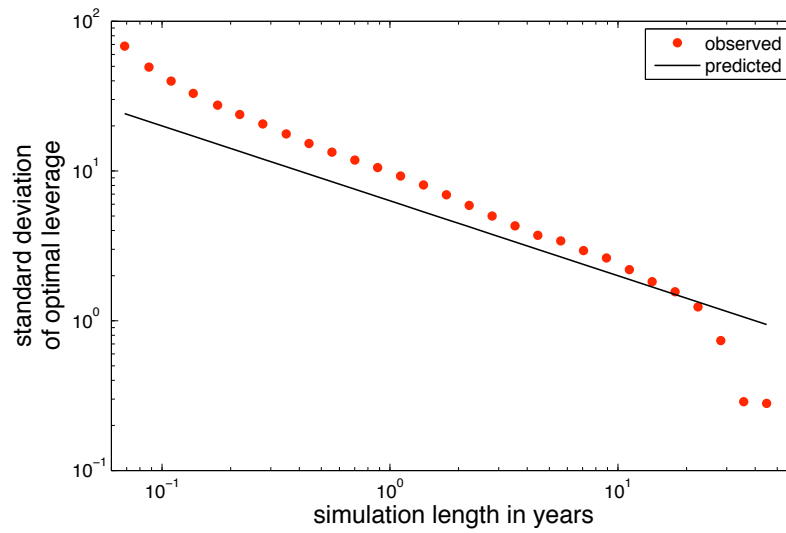
FIGURE 3.— The standard deviation of $l_{\text{opt}}^{\text{s}}$ in the simple case (symbols) as a function of window length can be predicted based on the specific model (Eq. 1) (straight line), using the parameters found in Sec. 3.3. Prediction and observations yield an estimate of the window length necessary to make meaningful statements of the type "$l_{\text{opt}}^{\text{s}}$ is within $0 \leq l_{\text{opt}}^{\text{s}} \leq 1$".
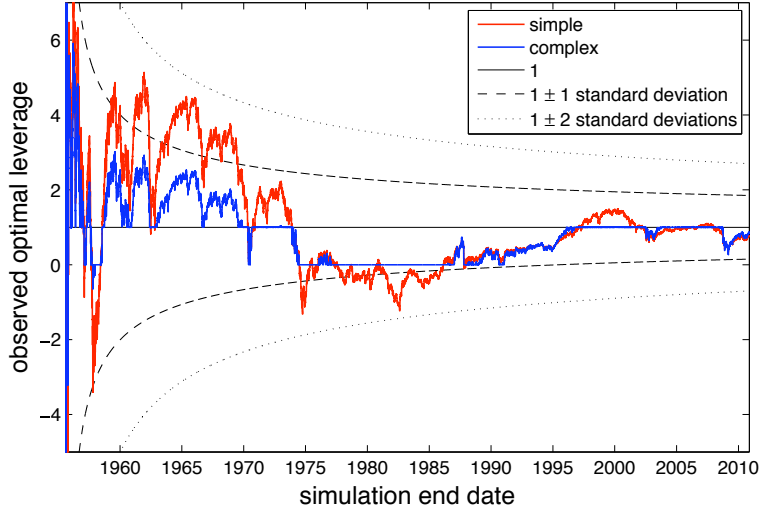
FIGURE 4.—

Daily optimal leverages are computed for an expanding window, starting on 4^{th} August 1955. Also shown are the one- and two-standard deviation envelopes about $l_{\text{opt}}^{\text{s}} = 1$, based on the estimate $\sigma^{\text{s}} = 15.8\%$ per square-root of one year in Sec. 3.3.

be atypical, but we attribute the lack of large deviations to the inadequacy of the the model (Eq. 1). The large fluctuations in daily closing prices, whose likelihoods are underestimated by (Eq. 1), prohibit very large values of l^{s}.

Figure 4 illustrates the convergence of $l_{\text{opt}}^{\text{s}} \to 1$ over time, but provides no information regarding the typicality of the time series. Further insight into the dynamics of $l_{\text{opt}}^{\text{s}}$ can be gained by examining time series for fixed window lengths. Figure 5 (a) shows $l_{\text{opt}}^{\text{s}}$ for the simple case for windows ranging from 5 years to 40 years as a function of the end date of the window. Figure 5 (b) shows the same for the complex case, which we claim is a more realistic simulation of market conditions and practices. From the strong fluctuations over short time scales emerges attractive behavior consistent with both the strong and weak forms of the stochastic efficiency hypothesis. The effects of the stickiness of the points $l_{\text{opt}}^{\text{s}} = 0$ and $l_{\text{opt}}^{\text{s}} = 1$ in the complex model are clearly visible and lend additional support to both forms of the hypothesis as it applies to real markets. In particular, over the last decade or so optimal leverage for the 20- and 40-year periods remained close to unity. This may be seen as an evolution of the market system towards strong stochastic efficiency over the course of the last half century.
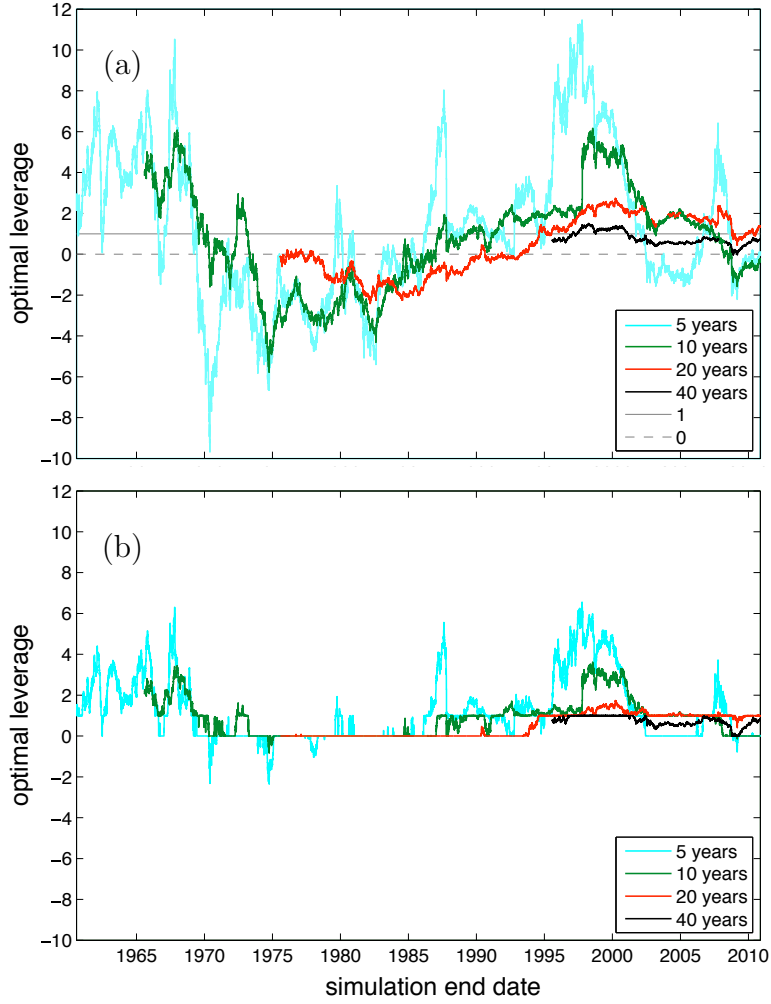
FIGURE 5.—

(a) In simulation 1, see text, observed optimal leverage fluctuates strongly on short time scales but appears to converge to $l_{\mathrm{opt}}^{\mathrm{s}} = 1$ on long time scales, which constitutes the central result of the study.

(b) In simulation 4, see text, the kinks in Fig. 1 ensure that $l_{\mathrm{opt}}^{\mathrm{s}} = 0$ and $l_{\mathrm{opt}}^{\mathrm{s}} = 1$ are often found exactly. The 40-year simulation supports the strong stochastic efficiency hypothesis, $l_{\mathrm{opt}}^{\mathrm{r}} = 1$, with a dip to $l_{\mathrm{opt}}^{\mathrm{s}} = 0$ only during the credit crunch of 2008.

4. DISCUSSION

Nothing in nature, including Brown's pollen (Mazo, 2002), can truly follow Brownian motion, whether geometric or not. Nor is anything in nature knowably faithfully described by any mathematical expression (Rényi, 1967). However, just as the dynamics of Brown's pollen, in the appropriate regime, have some properties in common with a Wiener noise, so the dynamics of share prices have some properties in common with geometric Brownian motion. Specifically, the daily excess returns for the markets investigated – like the returns in geometric Brownian motion – are sometimes positive and sometimes negative. For any time-window that includes both positive and negative daily excess returns, regardless of their distribution, a well-defined optimal leverage exists in our simulations, Sec. 3.2. We have empirically investigated the properties of such optimal leverages.

Stability arguments, which do not depend on the specific form of the distribution of returns and go beyond the model of geometric Brownian motion, led us to the quantitative prediction that on sufficiently long time scales real optimal leverage will be between $0 \leq l^{\mathrm{r}}_{\mathrm{opt}} \leq 1$.

We used specific properties of geometric Brownian motion to estimate the time necessary to obtain a meaningful empirical test of this prediction. Over short time scales the fluctuations in $l^{\mathrm{s}}_{\mathrm{opt}}$ are too large for us to have confidence, to the required precision of order 1, in a single measurement of $l^{\mathrm{s}}_{\mathrm{opt}}$ (for the longest time window only one measurement is available). The model predicts a time scale of $\frac{1}{(\sigma^{\mathrm{m}})^2} \approx 40$ years and we confirmed this estimate in the scaling of the standard deviation of $l^{\mathrm{s}}_{\mathrm{opt}}$. We therefore consider our main finding $l^{\mathrm{s}}_{\mathrm{opt}} = 0.75$ (complex case: 0.84) for the longest possible window of 55 years a significant corroboration of the hypothesis $0 \leq l^{\mathrm{r}}_{\mathrm{opt}} \leq 1$. Both ends 0 and 1 are clearly special due to the kinks in Fig. 1. The economic paralysis argument suggests that $l^{\mathrm{r}}_{\mathrm{opt}} = 1$ is a stronger attractor than $l^{\mathrm{r}}_{\mathrm{opt}} = 0$, and our observations support this argument, especially the 40-year window in Fig. 5 (a) and (b) .

The existence of optimal leverage is important conceptually and its value and the stability arguments associated with it are of practical significance. While these arguments do not preclude special conditions under which it is optimal to invest more than one's equity or to short-sell an asset, they give a fundamental scale to leverage in general. In other words, if it appears that optimal leverage is outside the band $0 \leq l^{\mathrm{r}}_{\mathrm{opt}} \leq 1$, then a special reason – such as insider knowledge or a tax-incentive – for this violation of stochastic efficiency must exist. Artificially maintaining such conditions will lead to instabilities. Consider housing: many societies consider it desirable for an individual to be able to purchase a home whose price exceeds his equity without having to take reckless risks. Indeed, some societies perversely consider rampant inflation in the cost of housing, unlike other basic utilities, a social good. Without carefully designed restrictions on speculative home purchases, policies which aim to achieve the corresponding market conditions, *i.e.* $l^{\mathrm{r}}_{\mathrm{opt}} > 1$, will defeat their purpose and create an invest-

ment bubble followed by a crash.

Ordinary price efficiency provides constraints which underlie many theoretical and practical considerations. However, it is not "accountable" in the sense of the word used by Popper (1982) Chap. I.2, who demanded that a "theory will have to account for the imprecision of the prediction"[1]. While empirical studies have contested price efficiency (Farmer and Geanakoplos, 2008; Brody et al., 2010), the lack of accountability of the concept of price efficiency makes it difficult to evaluate any such apparent counter-examples. Stochastic efficiency, in contrast, is accountable in that it predicts its own imprecision, (Eq. 12), and the degree of its validity can be meaningfully and objectively tested. This is particularly important given the complexity of the systems involved.

We emphasize that our work is in no way meant to advocate or evaluate constant-leverage or any other investment strategies. Stochastic efficiency is a fundamental organizing principle for the stochastic properties of markets. The simulations in this study are an empirical test of this fundamental principle.

It has been argued that central banks, focusing their attention on interest rates, pay insufficient attention to leverage (Geanakoplos, 2010). Arguing in the context of the model, a strong link between the two is (Eq. 8): reducing the risk-free interest rate $\mu_{\mathrm{riskless}}^{\mathrm{m}}$ (something we liken to the rate at which governments lend) increases optimal leverage because, assuming that overall expected returns μ^{m} do not change, it implicitly increases $\mu_{\mathrm{excess}}^{\mathrm{m}}$ (it creates an incentive to invest rather than save), which tends to lead to an eventual increase in real leverage. We agree with the criticism in (Geanakoplos, 2010). Effecting an increase (decrease) in real leverage through a decrease (increase) in $\mu_{\mathrm{riskless}}^{\mathrm{r}}$ is rather indirect. This appears problematic given how sensitive $l_{\mathrm{opt}}^{\mathrm{m}}$ is to $\mu_{\mathrm{riskless}}^{\mathrm{m}}$ – an increase of $\mu_{\mathrm{riskless}}^{\mathrm{m}}$ by $\mu_{\mathrm{excess}}^{\mathrm{m}}$ (estimated at about 1.9% over the last 55 years) sets $l_{\mathrm{opt}}^{\mathrm{m}}$ to zero, removing any incentive to invest. Conversely, a decrease of $\mu_{\mathrm{riskless}}^{\mathrm{m}}$ by $\mu_{\mathrm{excess}}^{\mathrm{m}}$ doubles $l_{\mathrm{opt}}^{\mathrm{m}}$.

On a global scale we can argue that real leverage can only be exactly 1. This reflects the fact that there is nothing outside the global economy, or that the economy is fully invested in itself. Given that this is the case, reversing the argument, what would be the optimal behavior? In other words, we ask: given that we are forced to use a certain leverage as a global community (namely $l^{\mathrm{r}} = 1$), what kinds of risks should we take to make this the optimal leverage? As a global community, no action should be taken for which $(\sigma^{\mathrm{r}})^2 > \mu_{\mathrm{excess}}^{\mathrm{r}}$, as this would imply $l_{\mathrm{opt}}^{\mathrm{r}} < 1$, whereas $l^{\mathrm{r}} = 1$ always. A caricature interpretation of this conceptually dubious statement is this: if the risks we create through allowing global warming or nuclear proliferation are so great that we should all partly move to another (safe) planet, then we should not take those risks.

We do not consider our arguments specific to financial markets. They are relevant also to other regularly traded assets and commodities, related even to

---

[1]Popper does not refer to stochastic theories in this discussion. To apply his arguments to our case, we replace "precision in the initial conditions" in his Chap. I.3 by "window length". Both concepts quantify the information available about the system.

such basic needs as food and shelter, such as the price of wheat or apartments in Manhattan. They are relevant to macroeconomic decisions. Indeed, the same type of dynamics – multiplicative growth with fluctuations – is at work in many other systems. Equation (1) is commonly used to describe the growth of populations in ecology (Lewontin and Cohen, 1969) or the early spread of a disease in epidemiology (Daley and Gani, 1999).

We have argued that $0 \leq l_{\mathrm{opt}}^{\mathrm{r}} \leq 1$ is a natural attractor for an economic or market system, with the end points of the interval being sticky. Worryingly in the aftermath of the financial crisis of 2008, we note that the sticky $l_{\mathrm{opt}}^{\mathrm{r}} = 0$ may correspond to a Depression: in this case there is no incentive to invest and to take risks. The aim of economic policy may be viewed as creating conditions where $l_{\mathrm{opt}}^{\mathrm{r}}$ for the entire economy is within $0 \leq l_{\mathrm{opt}}^{\mathrm{r}} \leq 1$ and close to the upper bound.

5. ACKNOWLEDGMENTS

Department of Mathematics
and Grantham Institute
Imperial College London
180 Queen's Gate
London SW7 2AZ
ole@santafe.edu

and

Seasure Shipping Ltd.
10 Barley Mow Passage
London W4 4PH
alex.adamou@gmail.com

REFERENCES

Bachelier, L. J. B. A., 1900: *Le jeu, la chance, et le hasard*. Paris: Gauthier-Villars.

Brody, D., J. Brody, B. K. Meister, and M. F. Parry, 2010: Outsider trading. *arXiv:1003.0764v2*.

Daley, D. J. and J. M. Gani, 1999: *Epidemic modeling: An introduction*. Cambridge University Press.

de Montmort, P. R., 1713: *Essay d'analyse sur les jeux de hazard*. 2d ed., Jacque Quillau, Paris. Reprinted by the American Mathematical Society, 2006.

Ehrenfest, P. and T. Ehrenfest, 1912: *Begriffliche Grundlagen der statistischen Auffassung in der Mechanik. Translation "The conceptual foundations of the statistical approach in mechanics" by J. Moravcsik, Cornell University Press, 1959*. Teubner, Leibzig.

Fama, E., 1965: The behavior of stock-market prices. *J. Business*, **38 (1)**, 34–105.

Farmer, D. and J. Geanakoplos, 2008: The virtues and vices of equilibrium in the future of financial economics. *Complexity*, **14 (3)**, 11–38, .

Geanakoplos, J., 2010: Solving the present crisis and managing the leverage cycle. *Federal Reserve Bank of New York Economic Policy Review*, 101–131.

Gell-Mann, M. and J. B. Hartle, 2007: Quasiclassical coarse graining and thermodynamic entropy. *Phys. Rev. A*, **76**, 022 104 (16).

Gell-Mann, M. and S. Lloyd, 2004: Effective complexity. *Nonextensive entropy*, M. Gell-Mann and C. Tsallis, Eds., Oxford University Press, 387–398.

Grimmet, G. and D. Stirzaker, 2001: *Probability and Random Processes*. 3d ed., Oxford University Press.

Kelly Jr., J. L., 1956: A new interpretation of information rate. *Bell Sys. Tech. J.*, **35 (4)**.

Lewontin, R. C. and D. Cohen, 1969: On population growth in a randomly varying environment. *Proc. Nat. Ac. Sci.*, **62 (4)**, 1056–1060.

Markowitz, H., 1952: Portfolio selection. *J. Fin.*, **1**, 77–91.

Markowitz, H. M., 1991: *Portfolio Selection*. 2d ed., Blackwell Publishers Inc.

Mazo, R. M., 2002: *Brownian Motion*. Oxford University Press.

Peters, O., 2010: Optimal leverage from non-ergodicity. *Quant. Fin., doi:10.1080/14697688.2010.513338*.

Peters, O., 2011: The time resolution of the St. Petersburg paradox. *To appear in Phil. Trans. R. Soc. A, arXiv:1011.4404v1*.

Popper, K., 1982: *The open universe: An argument for Indeterminism*. Routledge.

Press, W. H., B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, 2002: *Numerical Recipes in C: The Art of Scientific Computing*. 2d ed., Cambridge University Press, New York.

Rényi, A., 1967: A Socratic dialogue on mathematics. *18 unconventional dialogues on the nature of mathematics*, R. Hersh, Ed., Springer (2006), Chap. 1.

Thurner, S., D. Farmer, and J. Geanakoplos, 2009: Leverage causes fat tails and clustered volatility. *arXiv:0908.1555v2*.