

Investing in Disappearing Anomalies*

Christopher S. Jones
701 Hoffman Hall
Marshall School of Business
University of Southern California
Los Angeles, CA 90089
christopher.jones@marshall.usc.edu

Lukasz Pomorski
Office of Doctoral Studies
Graduate School of Business
University of Chicago
1101 E. 58th Street, Chicago, IL 60637
lpomorsk@gsb.uchicago.edu

First draft: February 2002
This revision: June 2003

*Updated versions of the paper may be found at <http://www-rcf.usc.edu/~christoj>. We are grateful to Doron Avramov, John Cochrane, Eugene Fama, John Heaton, Tobias Moskowitz, Lubos Pastor, Gregor Trojanowski, Pietro Veronesi, and participants of the 2003 MFA Conference, the 2003 LBS Doctoral Conference, and lunchtime workshops at Chicago and USC. All errors remain our own.

Investing in Disappearing Anomalies

Abstract

This paper studies the optimal asset allocation to a security whose anomalous returns are possibly disappearing over time. We apply our framework to two anomalies, positive stock index autocorrelations and the January effect. We summarize the empirical evidence on these phenomena from the point of view of investors with differing beliefs on the behavior of the anomaly. Our results suggest that anomalies decay over time, and that accounting for this decay has a large effect on optimal portfolio weights. Somewhat surprisingly, we find that priors that downweight the historical magnitude of an anomaly can sometimes increase the future allocations to it.

JEL classification: G12, C11.

Keywords: anomalies, Bayesian analysis, asset allocation, January effect, return autocorrelation, structural breaks.

1 Introduction

Since Fama's (1970) seminal work, many challenges to market efficiency have arisen. Notable examples include the size effect (Banz, 1981), the weekend effect (French, 1980), and momentum (Jegadeesh and Titman, 1993). To the extent that these anomalies represent failures of market efficiency rather than the misspecification of a pricing model, it is reasonable to expect their appearances to be finite-lived. Once the market "discovers" a trading opportunity, it should not persist for long unless there are limits to arbitrage. In a review article, Cochrane (1999) observes that the recent decline in the size and value premia "suggests that at least some of the premium the new strategies yielded in the past was due to the fact that they were simply overlooked."

Schwert (2001) demonstrates a lack of robustness across sample periods for a variety of well-known anomalies and attributes at least some of the attenuation in available abnormal returns to the dissemination of academic research findings. Among other results, he finds that funds managed by Dimensional Fund Advisors to exploit the size and value effects have not offered the average returns suggested by the historical sample used to identify these effects.

Not all anomalies seem to disappear as soon as they are discovered. Schwert (2001) shows that the January effect, which measures the tendency of small firms to outperform large firms in the month of January, has endured through the 1990s.¹ Jegadeesh and Titman (2001) find that the magnitude of the momentum effect has been relatively unchanged since their 1993 study.

Identifying those anomalies that disappear and those that do not should be revealing about which ones represent market inefficiencies and which ones are artifacts of a misspecified pricing model or limits to arbitrage. For example, based in part on the persistence of momentum profits over the time since their first study was published, Jegadeesh and Titman (2001) conclude that a behavioral model might offer a more accurate description of asset prices. In comparison, the sudden disappearance of the size and value effects (Schwert, 2001) may make the three-

¹ This is also evident in the returns on the DFA US Small Stock Portfolio. Since DFA introduced it in 1982, it has yielded the average raw return of more than 4% in Januaries, and below 1% in non-January months

factor model of Fama and French (1993) less relevant, with size and value effects perhaps better characterized as truly anomalous but now vanished.²

Besides market inefficiency and model misspecification, data snooping also arises as a potential explanation for apparent anomalies. If an anomaly disappears immediately after it is identified in the academic literature, this lends support to the data snooping hypothesis, although it certainly does not prove it. Data snooping would be an unlikely explanation for an anomaly that did not diminish or disappeared only gradually following its identification.

The time-varying nature of apparent anomalies is perhaps most important to the investor trying to form an investment strategy based on one of them. Asset allocation in this setting is complicated by at least three factors, outlined by Brennan and Xia (2001). The first is that the investor must decide whether or not the anomaly is solely the result of data snooping. The second is that the investor must take into account parameter uncertainty - expected returns are estimated, not known. The third problem, not addressed by Brennan and Xia (2001) or elsewhere, is that anomalies might be expected to disappear.

Investing in such an unstable environment is the topic of this paper. Rather than focus on the explanation of anomalies, we examine the optimal asset allocations of a Bayesian investor who considers the possibility that anomalous past returns might not continue into the future. We examine this question in a framework similar to those of Kandel and Stambaugh (1996), Brennan and Xia (2001), Pastor (2000), and Pastor and Stambaugh (2000). Like all of these papers, we explicitly account for parameter uncertainty in forming optimal portfolios. In addition, following Pastor and Stambaugh (1999), we sometimes consider investors who center their prior beliefs around the parameter values implied by a particular asset pricing model.

The framework we propose here differs from standard models of structural breaks, such as those considered by Andrews (1993) or Pastor and Stambaugh (2001). While these papers assume an instantaneous regime switch, we allow for a prolonged interim period in which the anomalous assets only gradually change their behavior. Like Pastor and Stambaugh (2001),

² The debate on the disappearance of the size and value effects is far from over. Recent theoretical work by Zhang (2003) predicts business-cycle effects in the value premium that should cause it to disappear at times, only to reappear later in full force.

however, the nature of our structural shift is guided by economic reasoning embodied in the combination of model and prior. In their paper, the more probable structural breaks are those that do not radically change the market risk premium or Sharpe ratio. Here, structural breaks are assumed to cause anomalies to disappear and markets to become more efficient. In the interest of simplicity, we do not consider the possibility of multiple breaks.

Our approach accommodates Merton's (1987) view that the time between the discovery of an anomaly and its elimination may be considerable. Merton argues that information published in scientific journals diffuses particularly slowly, as academics and practitioners struggle to determine whether the observation is an artifact of data, a result of market frictions, or a truly attractive investment opportunity that can be taken advantage of. Furthermore, even if money managers are convinced of the trading opportunity, they still face frictions in implementing the strategy on a large scale as they gather data, build models, satisfy prudence requirements, and possibly market the new strategy to their clients. Hence, the anomaly may decay sluggishly over time, especially if it applies to many stocks (e.g., all small stocks) and arbitraging it away requires an orchestrated action of many investors.

While the idea that learning by market participants might reduce market efficiencies over time is not new, few studies have addressed the possibility directly. Watts (1978) conjectures that the shrinking of abnormal reactions to earnings announcements might have been due to learning. More related papers by Mittoo and Thompson (1990) and McQueen and Thorley (1997) examine more directly whether published articles provide a source for learning by the market. The former focuses on publication of academic research related to the size effect, while the latter examines whether analysis published in the *Wall Street Journal* was responsible for the disappearance of a particular anomaly in the pricing of gold mining stocks. Both papers find evidence that anomalies attenuate following publications about them.

In Mittoo and Thompson's (1990) modelling of the size effect, the return on a size-related spread portfolio is assumed to follow a regime-switching model with two regimes.³ They reject

³ The size-related spread portfolio used by Mittoo and Thompson (1990) is the equally-weighted market index minus the value-weighted market index. Because of the greater weight of small stock in the former index and of large stocks in the latter, returns on this spread portfolio should capture the size effect if it is present. Other size-related strategies that involve the largest and smallest firms exclusively might be expected to amplify the

the hypothesis that there is no regime switch in the size effect, finding that the average size effect, around 1% per month prior to the publication of Banz (1981), dropped to zero by the late 1980s.

Our paper differs in the anomalies considered and more generally in its focus on asset allocation. Understanding the asset allocation process, while interesting from our own perspective as potential investors, is also relevant for assessing the likelihood of alternative explanations of underlying anomalies. As Kandel and Stambaugh (1996), Avramov (2002), and Cremers (2002) demonstrate in their analyses of predictive regressions, statistical significance is sometimes a poor measure of economic relevance. If optimal allocations decline to negligible levels after anomalies are discovered, then statistical evidence that continues to support the existence of the anomaly may be discounted as economically trivial. If, on the other hand, optimal allocations remain large for a significant period of time, then explanations that require modifications to standard asset pricing models may gain credence.

As examples of our approach, we consider two of the most puzzling anomalies in empirical finance: short-term index autocorrelation, usually associated with Lo and MacKinlay (1988), and the January effect, identified by Keim (1983) and Reinganum (1983). Both are interesting because they have been the focus of substantial academic debate that is to a large extent unresolved. In addition, both suggest simple investment strategies that have been widely applied in the investment management industry.

Our results demonstrate that these anomalies are likely to be viewed as slowly disappearing, and that as of December 2001, the end of our sample, both continue to imply attractive investment opportunities. Specifically, while autocorrelations in the value-weighted market portfolio are today extremely low, estimates of the equally-weighted portfolio's weekly autocorrelation range from 10-15%. Furthermore, these autocorrelations appear to be available even when excluding firms in the smallest two deciles of market capitalization. The effects on asset allocation are large: given a positive two standard deviation return, an investor might triple his investment in the equally-weighted portfolio; the same investor would short the market follow-effects found.

ing a negative two standard deviation return. As for the January effect, we estimate a decline in January expected abnormal returns from around 9% to 4%. Although reduced, the effects of this smaller expected return remain large: an investor who would normally (in non-Januaries) put about 100% of his assets in the market portfolio would have put, in January 2002, five times his wealth in a size-based spread portfolio.

For both anomalies considered here, accounting for a possible decay is economically important. Allocations to the anomaly are often dramatically different for investors who allow for a decay and those who do not.

A particularly surprising finding is that prior beliefs that appear more strongly in favor of an asset pricing model may perversely lead to greater allocations to the non-benchmark assets. In one case, we find that an investor with a prior belief that annualized CAPM alphas are mean-zero with a 2% standard deviation allocates more to the January portfolio than does a similar investor whose beliefs are diffuse. Intuitively, if alphas are constrained to be small by the prior, and if the data seem to be in violation of the CAPM, then the alphas might have to decay more slowly to be consistent with that data. With the January effect, it turns out that this greater persistence can dominate the lower initial level of alpha implied by the more skeptical prior. At the end of the sample, the investor with the informative prior therefore has a higher predictive mean for future January returns and actually invests more in the January spread portfolio.

In Section 2 of the paper, we proceed with the examination of short-term autocorrelations in stock index returns. This case is somewhat simplified by the fact that it is a study of the index in isolation, so models of asset pricing have no relevance. Section 3 considers the January effect, in which asset allocations are complicated by the existence of multiple assets and possible priors that are motivated by an asset pricing model, which we take to be the CAPM. We conclude in Section 4.

2 Short-horizon autocorrelations in equity index returns

... we learned that over the past decade several investment firms – most notably, Morgan Stanley and D.E. Shaw – have been engaged in high-frequency equity trading strategies specifically designed to take advantage of the kind of patterns we uncovered in 1988.

Lo and MacKinlay (1999)

While short-horizon autocorrelations in individual equity returns had been noticed as far back as Fama (1965), the extremely high short-term autocorrelations of diversified equity indices were generally unknown until Lo and MacKinlay (1988).⁴ Their results imply weekly return autocorrelations as high as 15% for the value-weighted CRSP index and 30% for the corresponding equally-weighted index. While not a violation of market efficiency *per se*, autocorrelations of this magnitude are viewed as anomalous in light of the dominant view in the finance literature, summarized by Ahn et al. (2002), that “time variation in expected returns is not a high-frequency phenomenon.”

While the analysis within Lo and MacKinlay (1988) suggests that index autocorrelations may have begun to decline by the end of their sample, it is difficult to find more recent comparable evidence. Ahn et al. (2002) examine various US index and futures returns since 1982 and find daily autocorrelations ranging from -9% to +22%. Unfortunately, relatively small sample sizes imply that standard errors are often large.

A natural explanation for positive index autocorrelations is the existence of microstructure effects. At least since Fisher (1966), nontrading has been known to induce spurious positive autocorrelation. The fact that equally-weighted indices display more prominent autocorrelation than value-weighted indices supports this explanation, since equally-weighted indices weigh more heavily on less frequently-traded small stocks. While Lo and MacKinlay (1988) and Conrad and Kaul (1989) discount this possibility, more recent analysis by Ahn et al. (2002) favors this explanation. We will set this issue aside temporarily.

⁴ Interestingly, a similar finding was obtained by Hawawini (1980), but it was not the primary focus of his paper and appears to have been generally overlooked.

An alternative explanation is that prices of some stocks are slow to adjust to newly arrived information. Possible reasons for slow price adjustment include information asymmetries (Llorente et al., 2002), behavioral biases (Hong and Stein, 1999), and the investor recognition hypothesis (Merton, 1987, Badrinath et al., 1995, Shapiro, 2002).

Our hypothesis that anomalies disappear, or, more generally, decay over time, may shed light on some of the proposed explanations. For instance, it seems unlikely that behavioral biases inherent in human nature would vanish. As we document below, return autocorrelations have substantially decreased over time, which suggests that this anomaly might not be well-captured by purely behavioral models. On the other hand, the investor recognition hypothesis could be more consistent with autocorrelation decay. Merton (1987) postulates that due to information set-up costs investors will only research and hold a small subset of available investments. If that is true of better informed institutional investors as well, then prices of stocks with high institutional trading might reflect information earlier than prices of other stocks. As discussed in Badrinath et al. (1995) and Sias and Starks (1997), some of the information found by institutional investors will be common for all stocks, inducing the positive cross-autocorrelations found by Lo and MacKinlay (1990).

If recent phenomena such as faster computers, the internet, the proliferation of the financial press, the popularity of business education, the availability of security price data, and the greater dissemination of academic research have led to a decrease in information acquisition and processing costs, then autocorrelation induced by such costs might also be reasonably expected to have declined. An investigation of these possibilities would obviously be difficult, and it is well beyond the scope of the present study.

2.1 Modelling disappearance

The Bayesian investor we model is assumed to have a relatively simple view of how anomalies disappear over time. Specifically, the investor views the anomaly as being in full force up until some time τ , when the anomaly is “discovered”.⁵ Following discovery, the investor be-

⁵ Throughout the paper, we will report τ in terms of calendar dates, although computations are performed by assigning each observation a sequential whole number.

believes that the magnitude of the anomaly disappears at a geometric rate. In the case of return autocorrelations, the investor assumes that the market return R_t is described by

$$R_t - \mu = \rho_0 \delta^{(t-\tau)^+} (R_{t-1} - \mu) + \epsilon_t, \quad (1)$$

where $x^+ = \max\{x, 0\}$ and where $0 \leq \delta \leq 1$. At least initially, we assume $\epsilon_t \sim N(0, \sigma_\epsilon^2)$.

The model implies that returns display a first-order autocorrelation of ρ_0 up until and including date τ . After that time, the autocorrelation disappears at a rate determined by δ . The extreme values $\delta = 0$ and $\delta = 1$ correspond to the cases, respectively, in which the anomaly disappears immediately or does not decay at all. When τ is equal to the last date in the sample, no decay is observed.

Values of δ in between zero and one indicate that the anomaly disappears over time, with values near one indicating persistence.

Were autocorrelation merely the product of data snooping, the true value of ρ_0 would obviously be zero, making δ and τ irrelevant. Given that data snooping has resulted in a positive estimate of ρ_0 over some historical period, however, we should expect the estimate of τ to correspond roughly to the end of that period (possibly close to a publication date) and the estimate of δ to be close to zero. Anomalies that persist well past their “discovery” make the data snooping explanation highly improbable.

It should be noted that our specification is not fully general. There is no way for autocorrelations to increase over time, and if autocorrelations decrease over time ($\delta < 1$) then they will eventually approach zero. In addition, other parameters are not allowed to change over time. An alternative model might allow for instability in multiple parameters and for the autocorrelation to converge to some nonzero value. In the interests of parsimony we do not consider these generalizations.

In light of the well-known heteroskedasticity of weekly market returns, we do pursue a generalization in which volatility is stochastic and that returns, conditional on volatility, are distributed as student-t. This combination of assumptions allows for the variety of departures from normality that Gallant, Hsieh, and Tauchen (1997) find in daily index returns. In this case,

we replace the assumption that $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ with $\epsilon_t \sim t(0, h_t, \nu)$, where

$$\ln h_t = a + b \ln h_{t-1} + cR_{t-1} + \eta_t \quad (2)$$

and $\eta_t \sim N(0, \sigma_\eta^2)$. The parameter b measures the persistence of volatility, while c allows for the possibility of a leverage effect, or a negative correlation between returns and volatilities. ν , the student-t degrees of freedom parameter, measures the degree of leptokurtosis in returns conditional on volatility.

2.2 The investment problem

For simplicity, as in Kandel and Stambaugh (1996), we consider a myopic investor who is able to invest only in the market and a risk-free asset. Similarly to Pastor and Stambaugh (2000), we assume that investors have mean-variance preferences. The mean and the variance of the market return is unknown to the investors and need to be estimated from the data. In our Bayesian specification estimation risk can be accounted for in a fairly straightforward manner. Parameter uncertainty is incorporated in the predictive distribution, obtained by integrating over the posteriors. For a detailed analysis of such investment problems, see Kandel and Stambaugh (1996).

In the case of a single risky asset, at time t the investor forms a predictive mean and standard deviation of R_{t+1} , $\hat{\mu}$ and $\hat{\sigma}$, and maximizes

$$p\hat{\mu} + (1-p)r_f - \frac{A}{2}p^2\hat{\sigma}^2, \quad (3)$$

where p is the fraction of wealth invested in the risky asset. The optimal allocation is

$$p^* = \frac{\hat{\mu} - r_f}{A\hat{\sigma}^2}. \quad (4)$$

In this paper, we assume $A = 2.83$, a value chosen by Pastor and Stambaugh (2000) so that an unconstrained investor would, given their data sample, choose to invest 100% of his portfolio in the market.⁶ In our sample, similar results will obtain approximately. In any case, it is trivial

⁶ Another possibility would be to choose the value of A that would cause some unconstrained investor to fully invest in the market given our own data samples. Since we consider different datasets in Sections 2 and 3, however, this would require a different value of A for each section. For simplicity, we therefore assume Pastor and Stambaugh's (2000) value.

to adjust most of our results for different values of A .

When volatility is stochastic, it is possible that the predictive variance, $\hat{\sigma}^2$, might also account for uncertainty about the current value of the stochastic variance process h_t . We abstract from this possibility, recognizing that the information set available to actual investors is much richer about the current level of volatility. (For example, investors observe returns much more frequently than weekly.) Investors are therefore assumed to know the current value of h_t , though not its future values or the parameters that govern its dynamics.

2.3 Priors

We consider four priors on the parameters of the model. The first and most general, which we call the diffuse prior, is chosen to represent prior ignorance. It is given by

$$p(\mu, \rho_0, \delta, \tau, \sigma_\epsilon) \propto \begin{cases} 1/\sigma_\epsilon & \text{if } 1/1/1962 \leq \tau \leq 12/31/2001 \text{ and } 0 \leq \delta \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

for the constant volatility model and

$$p(\mu, \rho_0, \delta, \tau, a, b, c, \sigma_\eta, \nu) \propto \begin{cases} \lambda \exp(-\lambda\nu)/\sigma_\eta & \text{if } 1/1/1962 \leq \tau \leq 12/31/2001, \\ & 0 \leq \delta \leq 1, \text{ and } 3 \leq \nu \leq 50 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

for the stochastic volatility model. The parameter λ is set to .05, a value that implies a relatively diffuse prior distribution for ν , with a mean of about 18 and a standard deviation of about 12.

Other priors are restricted versions cases of these. An investor who is certain that anomalies do not exist will put all prior mass on $\rho_0 = 0$ (making δ and τ irrelevant). An investor who is dogmatic that anomalies do not disappear puts all prior mass on $\delta = 1$ (making τ irrelevant). We refer to these priors as the “no-anomaly” prior and the “no-decline” or “no-decay” prior. Finally, we consider the case of an investor who believes that the date of discovery was 4/1/1988, which was approximately when Lo and MacKinlay (1988) was published. This investor is said to have a “Lo and MacKinlay” prior.

2.4 Posterior and predictive inference

We estimate model parameters using Gibbs sampling, a Markov chain Monte Carlo approach developed in Geman and Geman (1984). Gibbs sampling is a method of indirectly drawing samples from a multivariate distribution whose density is unknown. Samples are constructed by repeatedly drawing from conditional distributions, which are often easier to obtain than the joint distribution. For a good introduction to this inference technique, see Casella and George (1992).

At least for constant volatility, conditional distributions are obtained for one parameter at a time. For instance, we can rewrite the model as

$$R_t - \rho_0 \delta^{(t-\tau)^+} R_{t-1} = \left(1 - \rho_0 \delta^{(t-\tau)^+}\right) \mu + \epsilon_t. \quad (7)$$

Conditional on all parameters but μ , we may regard this equation as a linear regression without a constant and with a known variance. It is easy to show that the conditional distribution of μ (see Zellner, 1971, for example) is Gaussian. A similar rearrangement can be used to obtain a Gaussian distribution for ρ_0 conditional on the remaining parameters.

For τ and δ , even the conditional distributions are nonstandard. Nevertheless, by Bayes rule we can easily compute the posterior distribution, at least up to a constant of proportionality, using the properties of the Gaussian likelihood function. We draw these parameters, one at a time, in separate blocks of the Gibbs sampler. Drawing δ , for example, is performed using the “gridgy Gibbs” approach of Ritter and Tanner (1992). The $[0, 1]$ interval (the allowable range of δ) is discretized into more than ten thousand discrete points. At each one of these points we compute the posterior density up to an unknown integrating constant using Bayes rule and the known likelihood function, again conditioning on all parameters but δ . This discretized conditional density is then normalized to integrate to one, and a cumulative distribution function is calculated. A random draw of δ is obtained by evaluating the inverse of this c.d.f. at a draw of a uniform $[0, 1]$ random variable. Draws of τ are performed similarly using the natural discreteness of τ to define the grid points.

We allow the Gibbs chain to run for 16,000 iterations and discard the first 1,000 to allow

for the chain to converge. Sample moments computed from the remaining 15,000 draws serve as highly accurate estimates of corresponding posterior moments.

Our approach to estimation of the student-t stochastic volatility model combines the algorithms of Jacquier, Polson, and Rossi (1994) and Geweke (1993). In essence, we augment observed price data with unobserved volatility data. Conditional on the augmented data set, estimation proceeds similarly to method described above, relying on the properties of the Gaussian augmented likelihood. An appendix describes this procedure in more detail.

Predictive means are computed using the law of iterated expectations as

$$\hat{\mu} = \tilde{\text{E}}[E[R_{t+1}|\theta]] = \tilde{\text{E}}\left[\mu + \rho_0\delta^{(t+1-\tau)^+}(R_t - \mu)\right], \quad (8)$$

where a “tilde” denotes a moment estimated by averaging across the 15,000 Gibbs draws and where θ denotes all parameters of the model. For the constant volatility model, predictive variances are calculated using the variance decomposition

$$\hat{\sigma}^2 = \tilde{\text{E}}[\text{Var}(R_{t+1}|\theta)] + \tilde{\text{V\ddot{a}r}}(E[R_{t+1}|\theta]) \quad (9)$$

$$= \tilde{\text{E}}[\sigma_\epsilon^2] + \tilde{\text{V\ddot{a}r}}\left(\mu + \rho_0\delta^{(t+1-\tau)^+}(R_t - \mu)\right) \quad (10)$$

For the stochastic volatility model, predictive variances are instead decomposed as

$$\hat{\sigma}^2 = \tilde{\text{E}}[h_{t+1}|\theta] + \tilde{\text{V\ddot{a}r}}(E[R_{t+1}|\theta]) \quad (11)$$

$$= \tilde{\text{E}}[\exp(a + b \ln h_t + cR_t + \eta_{t+1})] + \tilde{\text{V\ddot{a}r}}\left(\mu + \rho_0\delta^{(t+1-\tau)^+}(R_t - \mu)\right) \quad (12)$$

$$= \tilde{\text{E}}[\exp(a + b \ln h_t + cR_t + .5\sigma_\eta^2)] + \tilde{\text{V\ddot{a}r}}\left(\mu + \rho_0\delta^{(t+1-\tau)^+}(R_t - \mu)\right) \quad (13)$$

where h_t is assumed to be known.

2.5 Data

Following Lo and MacKinlay (1988), we work with weekly returns on two stock market indexes. The data spans the period from July 5, 1962, around the start of the Lo and MacKinlay (1988) sample, to December 26, 2001. We examine weekly returns on the CRSP value and equally-weighted indices of all stocks traded on NYSE or AMEX. Weekly returns are computed by geometrically compounding daily returns from one Wednesday to the following Wednesday. If

a Wednesday return is missing, Thursday's return is used instead; if the Thursday return is also missing, then the Tuesday return is used. If all three returns (Tuesday, Wednesday, and Thursday) are missing, the return for that week is reported as missing. Lo and MacKinlay (1988) report that for their sample less than .5% weeks (out of 1216) are missing. Although our procedure of computing weekly returns is the same as that of Lo and Mackinlay (1988), no week is missing in our data for the period considered in Lo and MacKinlay. This is presumably because errors in the CRSP tapes are corrected over time. The only missing week is from September 11 to 16 of 2001, when trading was suspended due to the events of September 11. Without that observation, our sample consists of 2060 weeks.

2.6 Results

2.6.1 OLS inference

The top panel of Table 1 reports OLS estimates of autocorrelation coefficients for both indexes over different subsamples, along with standard errors calculated under OLS assumptions and using the method of Newey and West (1987). For the pre-Lo and MacKinlay subsample, we obtain essentially the same estimates as they did: autocorrelation coefficients of about 30% for the equally-weighted index and 9% for the value-weighted index. In the whole sample, 1962-2001, the estimates are somewhat lower, at 26% and 4%, respectively. Furthermore, when we use just the post-Lo and MacKinlay subsample (1988-2001), the autocorrelation of the equally-weighted index is only about 19%, while the value-weighted index actually has a negative autocorrelation of -6%. Using Newey-West standard errors with ten lags, neither of these is significantly different from zero.

To assess whether autocorrelation in the equally-weighted index is primarily induced by the smallest stocks, we examine a portfolio of stocks in the largest eight deciles that is chosen to approximately replicate the equally-weighted index. The weights $\{\hat{p}_1, \hat{p}_2, \dots, \hat{p}_8\}$ of this portfolio are obtained by regressing equally-weighted index returns on the returns of the top eight

deciles:

$$R_{ew,t} = p_0 + \sum_{i=1}^8 p_i R_{di,t} + u_t \quad (14)$$

The “decile 1-8 index” is given by $\sum_{i=1}^8 \hat{p}_i R_{di,t}$. Autocorrelations of this index are only slightly diminished from those of the equally-weighted index, indicating that index autocorrelation is not induced solely by the smallest firms.

We also test for a structural break in return autocorrelations between 1987 and 1988, and the results are mixed. The lower panel of Table 1 reports t -statistics for the parameter Δ in the regression

$$R_t = \gamma + \rho R_{t-1} + \Delta R_{t-1} I_{t>1/1/1988} + \epsilon_t, \quad (15)$$

where $I_{t>1/1/1988}$ is an indicator that takes the value 0 before 1988 and 1 thereafter. While statistics are somewhat sensitive to whether OLS or Newey-West standard errors are used, the null hypothesis of no structural break ($\Delta = 0$) can be rejected only in the case of the value-weighted index. For the equally-weighted index and the decile 1-8 index, no rejection can be made.⁷ As in Kandel and Stambaugh (1996), the change in optimal asset allocations may be large even when a frequentist analysis results in no statistical significance. We show below that this is the case here.

2.6.2 Posterior summaries

While the results in Table 1 suggest that weekly autocorrelations may have diminished over time, they are insufficient for knowing the state of return autocorrelations at the end of the sample. We attempt to determine this by estimation of the model in (1). The posterior medians and 90% confidence intervals for the model parameters are exhibited in Table 2. The μ and σ parameters are annualized for easier interpretation.⁸ Figure 1 presents corresponding plots of the posterior distribution of τ , which is computed only under diffuse priors.

⁷ Note that in this specification we fix the break at the time of the publication of Lo and MacKinlay (1988). As we show below, the break likely occurred before that date.

⁸ 90% confidence intervals are computed as the interval between the 5% and 95% quantiles of the posterior distribution.

Panel A of Table 2 was obtained under the assumption of constant volatility. For the equally-weighted index, the medians of ρ_0 are quite large, indicating historical autocorrelations of almost 30%. Under diffuse priors and under the “Lo and MacKinlay” prior, the median δ is close to one, indicating extreme persistence or near-zero decay, and there is little apparent difference between these priors and the one that imposes “no decline.”

For the value-weighted index, results are much different. For both τ fixed (the LM prior) and τ estimated (the diffuse prior), the 95th percentile of the posterior distribution of δ is somewhat below one, indicating the extreme likelihood that value-weighted autocorrelations have declined over time. Furthermore, the speed of this decline is most likely fast, as the medians of δ are below 0.6. For instance, under the diffuse prior, the posterior median of δ implies that in just a month the autocorrelation decreases by about 88%.

Panel B of Table 2 extends the model to the fat-tailed stochastic volatility process proposed in (2). Posteriors of the stochastic volatility parameters a , b , c , and σ_η are fairly typical of those found in the literature. Values of b near one indicate that volatility is highly persistent, while negative values of c are consistent with a leverage effect. Average levels of volatility, implied by a and b , are consistent with the unconditional estimates of Panel A. The parameter ν , which represents the degrees of freedom in the student-t distributed residual, is usually centered between 25 and 35, indicating that stochastic volatility is responsible for most of the leptokurtosis in weekly stock index returns.

Under stochastic volatility, the autocorrelation results for the equally weighted index change somewhat. As shown in Panel B of Table 2, differences in posterior medians of ρ_0 are now somewhat more pronounced, but all three specifications that allow for nonzero ρ_0 continue to imply large values. In addition, under the “Lo and MacKinlay” prior there is slightly more support for values of δ less than one.

For the value-weighted index under stochastic volatility, values of δ remain low, indicating fast decay towards zero. Results for the diffuse prior are particularly interesting because they imply that historical autocorrelations for the value-weighted index might have been as high as 20%.

Because of their asymmetry and multimodality, the posteriors for τ are inadequately described by medians and 90% confidence intervals. Therefore in Figure 1 we display histograms of the posterior distribution of τ for the diffuse prior, the only case in which τ is unknown.

The top left panel displays the posterior of τ for the equally-weighted index under the assumption of constant volatility. In this case, autocorrelation decay most likely began at the end of the 1970s, but τ is estimated quite imprecisely and all values receive substantial support. In particular, one of the posterior modes is very close to the end of the sample, indicating strong support for the possibility that autocorrelation decay did not occur, and that the anomaly remains in full force today. When stochastic volatility is introduced in the top right panel, this late-sample mode in the posterior of τ disappears. The distribution of τ , while still diffuse, now has 96% of its mass prior to the end of 1985.

For the value-weighted index, shown in the bottom two panels, the posterior also displays some bimodality, with one mode around the early 1970s and a lesser mode around the late 1980s. As in the equally-weighted index, the latter mode is reduced substantially when volatility is stochastic, and again a full 96% of the posterior mass of τ lies prior to the end of 1985.

It is notable that in all cases, but especially when volatility is stochastic, there is a high posterior probability that τ substantially predates the Lo and MacKinlay (1988) article, which drew the greatest attention to this anomaly. Thus if we are to attribute the decline in autocorrelations to the high-frequency trading strategies that Lo and MacKinlay (1999) describe in the quote at the beginning of the section, these strategies most likely were well-underway prior to the publication of their original study.

To better understand the timing and magnitude of decay in autocorrelations, the left two columns of Figure 2 plot the posterior median and 90% confidence intervals of $\rho_0 \delta^{(t-\tau)^+}$ as a function of t . For the equally-weighted index, this function is seen to be gradually downward sloping, ending in 2001 at roughly a third to a half of its original level. The decline is more pronounced when volatility is stochastic, which is primarily due to the absence of the second posterior mode of τ that is observed for the constant volatility case.

In comparison, the decline in the value-weighted index autocorrelation is dramatic. By the

mid 1970s, the median plummets to zero under both volatility models. Again, the decay is stronger for the stochastic volatility model.

The rightmost column of Figure 2 contains the year-by-year OLS estimates of return autocorrelations. While it appears that some decline has occurred in both series, the autocorrelations are obscured by significant sampling error.⁹ We conclude that a parametric model, if correctly specified, is likely to offer substantial improvements over the periodic re-estimation of a time-invariant model over short samples.

2.6.3 Asset allocations as of 2001

To gauge the economic significance of these results, we calculate optimal portfolio weights implied by the different prior beliefs. For simplicity, we consider an investor who allocates his wealth between the market and the risk-free instrument. Since autocorrelations imply a lag dependence in expected returns (and in volatility when it is stochastic) we need to specify the return for the week preceding the time the allocation is made. We consider three different cases, +4%, 0%, or -4%, where 4% is roughly a two standard deviation weekly return. In addition, we must specify a risk-free rate. To abstract from the role played by time variation in this rate, we assume that it is 6% per year. Finally, when volatility is stochastic we assume that its lagged value is equal to 2% (on a weekly basis).

Table 3 shows that as of the end of 2001, the lagged return on the equally-weighted index still has a huge effect on optimal portfolio holdings. Under stochastic volatility, for example, an investor with diffuse priors would triple his holdings in the market portfolio following a +4% return, relative to a 0% return. The same investor would take a substantial short position were the lagged market return equal to -4%.

As extreme as these fluctuating allocations are, they are substantially more stable than those of an investor who does not consider the possibility of a decline in autocorrelations. This “no-decline” investor shorts 600% of his wealth in the market following a -4% return and puts over 1,200% of his wealth in the market following a +4% return.

⁹The OLS standard error of each autocorrelation is $1/\sqrt{52} = 0.139$.

Somewhat surprisingly, the lagged return even has an impact on the investor who has a dogmatic belief that no anomaly exists. This is due to the fact that lagged returns impact the investor’s volatility forecast through (2) and (13). While this effect also impacts the allocation for other priors, its magnitude is relatively minor.

Allocations to the value-weighted index are much more stable, with investors allocating much closer to 100% of their wealth to the market portfolio. The sole exception is the “no decline” investor, who maintains a belief that autocorrelations are still present. For all priors, allocations fluctuate more noticeably for the stochastic volatility case, though this is almost entirely due to the impact that lagged returns have on forecasts of volatility, not expected returns.

Another measure of economic significance is the certainty equivalent loss of an investor who is forced to hold a portfolio that is suboptimal given his beliefs. We calculate certainty equivalents similarly to Kandel and Stambaugh (1996): the portfolio weight p optimal for an investor with one prior is evaluated using the predictive moments of another investor with a different prior. The certainty equivalent loss is (in the case of mean-variance utility) the difference between the utility resulting from an investor’s own optimal allocation p^* and the utility from the suboptimal allocation p . It can be interpreted as the riskless return that the investor would be willing to sacrifice to restore his own optimal allocation.

For a given predictive mean and variance, $\hat{\mu}$ and $\hat{\sigma}^2$, following (3) and (4), this loss is therefore given by

$$\left(p^* \hat{\mu} + (1 - p^*) r_f - \frac{A}{2} p^{*2} \hat{\sigma}^2 \right) - \left(p \hat{\mu} + (1 - p) r_f - \frac{A}{2} p^2 \hat{\sigma}^2 \right) \quad (16)$$

Since $\hat{\mu}$ and $\hat{\sigma}^2$ generally depend on lagged returns, we summarize certainty equivalent losses by averaging over all lagged returns in the sample.

Table 4 reports these averaged certainty equivalent losses on an annualized basis. For the equally-weighted index, they are frequently very large, indicating that the ability to invest on short-term trends has substantial value. Under stochastic volatility, an investor with diffuse priors, for instance, would forgo 9% of his wealth annually to choose his own portfolio rather than hold the portfolio that would be chosen by the “no-anomaly” investor. Not surprisingly,

the “no-anomaly” and “no-decline” investors are least satisfied with holding one another’s portfolios, with such suboptimal strategies generating certainty equivalent losses of about 65% per year.

For the value-weighted portfolio, there are few losses associated with holding a suboptimal portfolio. The exception is the “no-decline” investor, who differs from all the others in his belief that a small amount of autocorrelation still exists.

2.6.4 Rolling sample results

To observe how investor beliefs and optimal allocations may have changed through time, we consider a sequence of sample periods all starting in January 1962 and ending in December in the years between 1970 and 2001. Figure 3 plots the rolling sample posterior medians of the “terminal” autocorrelations $\rho_0 \delta^{(T-\tau)^+}$, where T is the last observation of each rolling sample period. Because it is somewhat backwards looking, we do not report results on the “Lo and MacKinlay” prior in this section. Since the “no-anomaly” prior sets autocorrelations equal to zero, there are only two lines per panel in Figure 3.

For the equally-weighted index, median terminal autocorrelations are generally decreasing over time, although the trend is fairly erratic. For the value-weighted index, diffuse priors lead to autocorrelations that have a median of essentially zero after about 1973. For both indexes, autocorrelations are always higher under diffuse priors, which allow for some decay to take place.

The portfolio allocations implied by the different priors are shown in Figures 4 and 5. For the equally-weighted index, shown in Figure 4, allocations following a +4% return generally decline over time, while allocations following a -4% return become less negative. The effects are more noticeable for stochastic volatility. Allocations following a 0% return have also changed, declining at first but rising between 1975 and 2001. The net difference between allocations following a 0% return and a 4% return has therefore fallen more significantly than the top panel alone suggests.

Allocations to the value-weighted index, shown in Figure 5, have changed much more dra-

matically. With the exception of the “no-decline” prior, return-based variation in portfolio weights had almost entirely disappeared by 1975.

2.7 Implications

Taken together, these results provide strong evidence of decay in index return autocorrelations and show that a Bayesian investor would substantially modify his portfolio allocations to account for this possibility. While allocations to the value-weighted portfolio are today fairly constant for the investor with diffuse priors, lagged returns remain important for investment in the equally-weighted index. Furthermore, our results suggest that return autocorrelations were likely “discovered” by the market somewhat before Lo and MacKinlay’s 1988 paper was published.

While our analysis does not directly address the sources of autocorrelation, the observed patterns of decay may prove challenging for some potential explanations. To attribute autocorrelations to data snooping, for instance, one would obviously have to explain why autocorrelations in the equally-weighted index persist today. If microstructure effects are the root cause, as Ahn et al. (2002) argue, then some shift in the 1970s in the magnitude of these effects would seem to be required to explain the sudden disappearance of value-weighted return autocorrelations. While the investor recognition hypothesis could provide an explanation for our results if one believes that information costs declined rapidly over the 1970s, we provide no specific evidence of this possibility.

The differences between the value-weighted and equally-weighted index suggests that transactions costs may play a role in the market’s ability to take advantage of newly discovered anomalies. Since transactions costs are higher for trading small stocks, the equally-weighted autocorrelation was more difficult for a small number of investors to speculate away. This process was likely additionally slowed by the unavailability of the S&P 500 futures contract until 1982 and the Russel 2000 futures contract until 1987.

If transactions costs on the equally-weighted index are large, our asset allocation results likely overstate the usefulness of conditioning portfolio weights on lagged portfolio returns.

Mech (1993), in particular, finds limited evidence that transaction costs are sufficiently high to make trading on this autocorrelation unprofitable. Estimating the level of transactions costs incurred by portfolios that are optimally constructed to minimize these costs is well beyond the scope of this paper. It is notable, however, that the substantial autocorrelations found in the largest eight deciles suggest that these costs may not be exacerbated by illiquidity in the smallest stocks.

Informally, our results demonstrate that short-term trend following continues to be an important part of any asset allocation strategy involving time-varying positions in the equally-weighted index. Significant certainty equivalent gains can be made simply by delaying or accelerating a re-allocation by a single week.

3 The January effect

The January effect persisted for a long time because no one was paying attention to it. Then it became just the talk of everybody.

Robert Shiller

Still, you can't say that anything has changed. The plot just shows that there's been more variability in the last five years or so.

Donald Keim

I think it was all chance to begin with. There are strange things in any body of data.

Eugene Fama¹⁰

Another puzzling phenomenon has been the tendency of small capitalization firms to outperform large capitalization firms in the month of January. Observations related to this anomaly were first made by Rozeff and Kinney (1976), who observed higher returns on an equally-weighted stock index in the month of January, and Banz (1981), who identified a consistent relation between size and risk-adjusted equity returns. These results were refined by Keim (1983) and Reinganum (1983), who showed that January effect and size effect were highly interrelated. In fact, Blume and Stambaugh (1983), after controlling for upward biases in small stocks returns, document that the size effect was significant only in January. For succinctness, we refer to this finding which relates both to firm size and to the month of the year, as “the January effect.”

More recent evidence on the effect is mixed. In an international sample, Hawawini and Keim (1995) find that, except for France and the UK, the size premium is significantly higher in January than the remainder of the year. Unlike Blume and Stambaugh (1983), however, they find that in most countries the premium is positive in other months as well.

In the “post-discovery” domestic sample, conclusions are varied. For instance, Haugen and Jorion (1996) maintain that the January effect has shown no evidence of dissipating, and that

¹⁰ All quotes were taken from “Early January: The Storied Effect on Small-Cap Stocks,” by James H. Smalhout, *Barrons*, December 11, 2000.

“no significant trend portends it’s eventual disappearance.” In contrast, Schwert (2001) documents that the January effect has lessened, but it has not disappeared completely. Specifically, he finds that the magnitude of the effect in the 1990s was roughly half that of the prior two decades.

Potential explanations of the January effect include tax-loss selling (e.g., Reinganum, 1983), the illiquidity of small-cap firms (Reinganum, 1990), and “window dressing” (Haugen and Lakonishok, 1988). While the examination of these hypotheses is outside the scope of this paper, our results may cast some light on the possible validity of these explanations. For example, if tax-loss selling were advocated as the sole explanation of the January effect, and if the January effect were to decline over a particular period, then some change in the tax code around that time would seem to be required for the theory to hold credibly. If there were no changes in the code and a distinct decline in the January effect were observed, we can at least conclude that tax-loss sales are not likely the sole source of the effect.¹¹

In light of the potential instability of the January expected return, the January effect is an interesting phenomenon given the focus of our study. Following most papers on the January effect, we work with monthly returns on the spread portfolio that goes long in a portfolio of small stocks and shorts a portfolio of large stocks. We choose the lowest capitalization decile of the NYSE and AMEX exchanges as the former portfolio and the highest capitalization decile as the latter. Following Reinganum (1983), we consider a sample starting in July 1962. The sample ends in December 2001.

3.1 The investment problem

For simplicity, we again consider a myopic Bayesian investor with mean-variance utility. Now, this investor chooses between two risky assets, the value-weighted market index (NYSE and AMEX) and the size-based spread portfolio described above. In addition, the investor is able to invest in the one-month Treasury bill.

The investor forms his predictive means and variances of the excess market return, $R_{em,t}$,

¹¹ See Grinblatt and Moskowitz (2003) for recent evidence on the effects of tax-loss selling and its potential impact on January returns.

and the spread portfolio return, $R_{\text{spr},t}$, using a simple model:

$$R_{\text{em},t} \sim N(\mu_m, \sigma_m^2) \quad (17)$$

$$R_{\text{spr},t} = \alpha_0 + \alpha_1 I_J(t) \delta^{(t-\tau)^+} + \beta R_{\text{em},t} + \epsilon_t, \quad (18)$$

where $x^+ = \max\{x, 0\}$, $\epsilon_t \sim N(0, \sigma_\epsilon^2)$, and $I_J(t)$ is an indicator that takes the value 1 in Januaries and 0 in all other months.¹²

A January effect can be present only if $\alpha_1 \neq 0$. If it exists, δ measures the speed of its decay and τ measures the last period in which it existed in full force. When $\delta = 1$, the January effect does not decay over time, while if $\delta = 0$ it disappears instantly after time τ . The specification of the spread portfolio return also nests several standard cases. When $\alpha_1 = 0$ the spread portfolio does not ever exhibit any January seasonality, and when $\alpha_0 = \alpha_1 = 0$ the expected return on the spread portfolio conforms with the CAPM.

Assuming that January seasonality only affects the Jensen's alpha coefficient in the spread return equation greatly simplifies the estimation process. A more complete model would obviously allow seasonality in the market's expected return as well, and possibly introduce seasonal effects in all other parameters of the model. We make these implicit simplifying assumptions because it has been suggested that the January effect is limited to small stocks and is not noticeable in the value-weighted market portfolio. Moreover, our sample contains relatively few Januaries, so the estimation of a more complex model is probably unrealistic.

3.2 Priors

We consider several prior distributions in our analysis, motivated in part by the quotes at the beginning of this section. In all cases, the prior on the mean and standard deviation of the market return is flat and independent of the remaining parameters, or

$$p(\mu_m, \sigma_m, \alpha_0, \alpha_1, \beta, \delta, \tau, \sigma_\epsilon) \propto p(\alpha_0, \alpha_1, \beta, \delta, \tau, \sigma_\epsilon) / \sigma_m \quad (19)$$

¹² Any decline in January expected returns should be expected to generate a corresponding positive shock to prices. We considered a more general specification that allowed for a permanent price shock at time τ , but no effect could be inferred and all other results were very similar. We therefore retained the simpler model.

The most general prior (the “Shiller” prior) on the remaining parameters is assumed to be diffuse, or

$$p(\alpha_0, \alpha_1, \beta, \delta, \tau, \sigma_\epsilon) \propto \begin{cases} 1/\sigma_\epsilon & \text{if } 0 \leq \delta \leq 1 \text{ and } 1963 \leq \tau \leq 2000 \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

As before, we also consider several special cases. An investor who believes that anomalies do not exist will put all prior mass on $\alpha_1 = 0$ (the “Fama” prior), while the investor who does not think that anomalies disappear puts all prior mass on $\delta = 1$ (the “Keim” prior). Given the possibility that the motivating quotes misrepresent the true opinions of their sources, and because the quotes do not truly represent ex ante views, we somewhat blandly refer to these three priors as the “diffuse”, “no-anomaly,” and “no-decay” priors.

We also consider the informative, CAPM-based priors for the regression parameters proposed by Pastor and Stambaugh (1999) and Pastor (2000). Namely, the priors on α_0 and α_1 are independent normal with mean zero and standard deviations of either .01 or .02. The prior on β is diffuse and the prior on σ_ϵ is proportional to $1/\sigma_\epsilon$. This prior will shrink α_0 and α_1 towards zero, so that the process governing $R_{\text{SPR},t}$ is closer to the one implied by the CAPM.

3.3 Calculation of the posterior distribution

Because of prior independence, the posterior distribution of μ_m and σ_m can always be computed independently of the other parameters. Using standard results, their posteriors are student-t and inverted gamma, respectively.

Conditional on δ and τ , other model parameters (α_0 , α_1 , β , and σ_ϵ) can be estimated using the standard multiple regression results (see, for instance, Zellner, 1971). Given the diffuse conditional prior

$$p(\alpha_0, \alpha_1, \beta, \sigma_\epsilon | \delta, \tau) \propto 1/\sigma_\epsilon \quad (21)$$

for example, we have the well-known result that the posterior of α_0 , α_1 , and β , conditional on σ_ϵ , is multivariate normal, while the posterior distribution of σ_ϵ is inverse gamma.

In the CAPM-based model, the prior is not conjugate (the normal priors are assumed to be

independent of σ_ϵ) and there is no closed form solution for the posterior. However, draws from the posterior can be easily obtained using the Gibbs sampler.

The last two parameters, δ and τ , conditional on α_0 , α_1 , β , and σ_ϵ , can be estimated using the gridgy Gibbs sampler described in Section 2. To approximate the posterior distribution of δ , its support (the interval $[0, 1]$) is divided into a 1000-point grid.¹³ Next, we iterate the Gibbs sampler 51,000 times. We discard the first thousand draws and retain every tenth draw after that, leaving a sample of 5,000 draws with which we work. Extremely small autocorrelation coefficients of the draws indicate that the Gibbs chain converged and that the sample is highly representative of the true posterior distribution.

Note that in either of the two restricted specifications, no decline and no anomaly, δ and τ do not need to be estimated, which greatly simplifies the analysis.

3.4 Results

3.4.1 MLE estimates

As a preliminary step, we report maximum likelihood estimates for the model above and for a simplified version. Table 5 contains these estimates, along with standard errors.

The top panel contains results for two sample periods for a simplified model in which the January effect is constant. We split our sample in 1981/1982, by which time Booth and Keim (1999) suggest that the size effect had become widely known to investors. Dimensional Fund Advisors small cap 9-10 Fund, for instance, was established in 1982.

The estimates suggest that the January premium, α_1 , has indeed declined over the last few decades. Between the two samples, the estimate of α_1 decreased by about a third. This decrease is not as pronounced as, for instance, Booth and Keim (1999) report, because their sample is shorter and does not incorporate the high January returns of 2000 (10.9%) and 2001 (22.8%). Nonetheless, the estimates indicate clearly that the magnitude of the January effect has decreased over time. Interestingly, the estimates of other parameters also vary between

¹³ Many more gridpoints for δ were needed for return autocorrelations because of the higher frequency that weekly data set. With such high frequency observations, extremely small differences in δ could result in large differences in likelihoods.

the two subsamples (most notably, the point estimate of β changes sign), which may indicate that characteristics of the spread portfolio other than its January return change over time.

Using the whole sample, in the second panel of Table 5 we compute the maximum likelihood estimates for the general model in (18) and the special cases of no anomaly and no decline.¹⁴ The estimates imply that the January premium began to decline in 1976. Given the point estimate of δ , about 57% of the premium had disappeared by the year 2001. Both these parameters are estimated fairly precisely and provide strong evidence that the January anomaly indeed seems to have decayed over time.

To make this point somewhat more formally, we consider a number of criteria provided in the bottom four rows of the table. In addition to the Akaike Information Criterion,

$$AIC = -2 \ln p(\text{data}|\theta_{ML}) + 2d, \tag{22}$$

and the Bayesian Information Criterion,

$$BIC = -2 \log p(\text{data}|\theta_{ML}) + d \ln T \tag{23}$$

we also consider a somewhat ad hoc modification of the BIC,

$$BIC^* = -2 \log p(\text{data}|\theta_{ML}) + d \ln(T/12), \tag{24}$$

where d is the number of parameters and T is the length of the sample. The modified BIC adjusts the sample size to reflect the fact that only one out of twelve observations is actually informative about the January effect. For all three criteria, lower numbers are preferred. Finally, we also report statistics and p-values for likelihood ratio tests of the two restricted models against the general one.

While the “no anomaly” specification is clearly rejected, the distinction between the “no decline” and fully general specifications is less pronounced. Specifically, AIC weakly favors the general specification, while the sample size-dependent BIC favors the no-decline model. When

¹⁴ Estimates for the general model were obtained by maximizing the model likelihood conditional on each possible value of τ , then choosing the τ that yielded the highest likelihood. The parameter δ was constrained to lie in $[0, 1]$. Standard errors for all the parameters were calculated using the outer product of gradients method, at which time we treat τ as continuous. Since $\max\{t - \tau, 0\}$ is non-differentiable, we approximate it as $(t - \tau)\Phi(10(t - \tau))$, where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal.

the sample size is adjusted in the modified BIC, the no-decline model is still favored, but only by a slim margin. Not surprisingly, the likelihood ratio test of this specification is just below its 95% critical value.

As emphasized by Kandel and Stambaugh (1996), it is not clear how an investor should incorporate these results into his portfolio decisions. In particular, he may be inclined to consider the general model even though it is not firmly supported by the *BIC* or the likelihood ratio test. As Kandel and Stambaugh show, a Bayesian investor may find a model important even if frequentist analysis does not suggest statistical significance. We consider this issue in the next section, in which we estimate the three models in the Bayesian framework and calculate the optimal portfolio weights they imply.

3.4.2 Posterior summaries

We initially consider an investor who is deciding, as of December 2001, his asset allocations for January 2002 based on all historical data since July 1962. Posterior medians and 90% confidence intervals are presented in Table 6. In addition to the diffuse prior (20) and the special cases of no decline ($\delta = 1$) and no anomaly ($\alpha_1 = 0$), the table contains results for the “2% CAPM” prior, in which $\alpha_0, \alpha_1 \sim N(0, .02^2)$ and the “1% CAPM” prior, where $\alpha_0, \alpha_1 \sim N(0, .01^2)$.

In general, the posteriors for the diffuse prior and the two CAPM-based priors suggest that any decline in the January effect has been relatively slow, as posterior medians for δ are all above .95. As the prior belief about α_1 becomes tighter around zero, the posterior naturally shrinks towards zero as well. This is not the only effect that is observed as a result of imposing this view, however. In particular, tighter priors for α_0 and α_1 tend to lead to much later estimates for τ . Intuitively, as the initial level of the anomaly is depressed, it must persist for longer to be consistent with the data.

The effects of this shrinkage on the expected January 2002 return,

$$E [R_{\text{spr},2002}] = E \left[\alpha_0 + \alpha_1 \delta^{(T-\tau)^+} + \beta \mu_m \right] \quad (25)$$

are surprising. Initially, as the prior on α_0 and α_1 goes from diffuse to somewhat informative (a standard deviation of .02), $E [R_{\text{spr},2002}]$ actually increases. Only when priors are tightened fur-

ther does this post-sample expected return decrease, although sample information still pushes α_1 well over four standard deviations above its prior mean.

Evidently, by pushing the posterior median of τ from 1976 to 1991, the “2% CAPM” prior reduces the time over which the January effect decays so substantially that it more than offsets the lower initial level of the effect. Although moving to the “1% CAPM” prior moves τ slightly further, to 1995, this change is insufficient to overcome the effects of a smaller value for α_1 . An implication of this result is that a prior that displays skepticism towards anomalies and that generates more in-sample shrinkage may actually lead to more extreme predictive inferences.

The timing of the beginning of the disappearance of the January effect is further explored in the left panel of Figure 6, which displays the posterior histogram of τ computed under the diffuse prior. As was also the case for return autocorrelations, the posterior of τ is multimodal, with the primary posterior mode in 1975 and lesser modes in 1993 and 2000. Thus, the year 1976 again stands out as the most probable year in which the decline of the first decline in the January took place. It is interesting that the first paper on January seasonalities, Rozeff and Kinney (1976), was published in that year, though any link between academic research and market efficiency is at this point purely speculative.

The right panel of Figure 6 provides a visual characterization of the decline in mean January returns over the full sample period. While there is a substantial amount of estimation uncertainty, a downward path is clearly evident in the January effect. As suggested by Table 6, roughly half of the mean January return had decayed by 2001. Extrapolating the trend forward, the effect would appear likely to survive for some time further.

3.4.3 Asset allocations as of 2001

Kandel and Stambaugh (1996) emphasize that weak statistical evidence need not imply economic significance, and the January effect presents another example of this point. Table 7, which displays the optimal allocations to the January spread portfolio and the value-weighted market as of December 31, 2001, shows that allocations differ substantially based on whether or not the investor considers the possibility of disappearance. Specifically, the “no-decline”

investor, who assumes that the January effect is stable, would optimally invest over 1,000% of his wealth in the January spread position. The diffuse prior investor, who allows for decay, invests half that amount.

Consistent with the higher expected returns perceived by the “2% CAPM” investor, the allocation to the January spread is larger for that investor. It is smaller, but still quite large, for the more skeptical “1% CAPM” investor. Finally, the “no-anomaly” investor position in the spread portfolio, while small, is not zero. This is because this investor only considers nonzero values of α_1 to be anomalous and has a slightly positive posterior mean for α_0 .

What is striking about these January allocations is that they remain large despite the likely possibility that they have declined over time. Even with the amount of shrinkage present in the “1% CAPM” prior, the allocation to the January spread portfolio remains at close to 400% of initial wealth. Furthermore, this is not simply an artifact of low risk aversion; the same level of risk aversion underlies the “no-anomaly” investor’s allocation, which is close to zero in the spread portfolio and 100% in the market index.

Certainty equivalent losses again provide an alternative measure of the significance of these results. Since the January effect is only available one month per year, values are expressed on a monthly (non-annualized) basis. The bottom half of Table 7 shows that these losses are small when an investor who allows for some decline (diffuse, 2% CAPM, 1% CAPM) is forced to hold the portfolio of another such investor. For the same investors to hold the “no-decline” portfolio would require substantial compensation, however, on the order of 10-20% of initial wealth. The “no-anomaly” investor is also very reluctant to hold any portfolio but his own.

While some of these numbers may seem high, they are a direct consequence of the differences in predictive moments. As indicated in Table 6, the median January 2002 mean return for the no decline investor is 7.48% per month, while for the no anomaly investor it is 44 times lower, at just .17%.

When one invokes trading costs, obviously high in the case of the spread portfolio, the allocation may remain substantial. For instance, assuming the round trip transaction costs of 2%, the January allocation is still 250% of the investor’s wealth (at the same time, the allocation

to the market rises to almost 70%). In fact, even if we assume that the transaction costs reduce the expected return to the level of the expected excess market return, the allocation to the anomaly is still non-trivial at 60%, while the allocation to the market is about 92%. Hence, even in the presence of substantial trading costs, the January anomaly is still appealing after the investor allows it to decay for some 20 years.

In contrast, the optimal portfolio of the no-decline investor will likely remain extreme even after a trading costs correction. The 2% to 3% round trip costs assumed above will not decrease the allocation to a reasonable level. In the case of 2%, it only decreases to 787% of wealth. To drive the allocation to 100% of wealth, the costs should be 6.72%.

3.4.4 Rolling sample allocations

To observe how inferences and optimal allocations have changed over time, we again perform a rolling sample analysis. All samples begin in July 1962 and ending at the end of years 1972 through 2001. Thus, the shortest sample contains ten Januaries.

Figure 7 graphs the time series of allocations to the spread portfolio and the market portfolio. The allocations implied by diffuse priors are contrasted with those resulting from more restrictive priors. For visual clarity, we do not present the CAPM based allocation in this figure. The year on the horizontal axis indicates the year for which the allocations were made; investors are assumed to decide on their January allocations in December of the previous year.

Although allocations have proven somewhat erratic over time, there is a general downward trend in allocations to the January spread portfolio for the two priors that allow for a January effect. The allocations of the investor who allows for a decline in anomalies vary between 10.03 in 1975 and .28 in 1990. During the same time, allocations to the market index have risen for investors with all three priors. Accounting for decay in the January effect does make a difference in optimal allocations, particularly over the last 15 years. As of the end of 2001, the investor with diffuse priors would invest five times his initial wealth in the January spread portfolio. While high, this allocation is just half that of the investor who assumes no decline has occurred. Just two years before, this difference would have been even more substantial,

with “no-decline” priors implying a similar allocation but diffuse priors suggesting that the optimal allocation be close to zero.

3.5 Implications

Almost 20 years since the publication of Keim (1983) and Reinganum (1983), the posterior mean of the spread portfolio’s January alpha remains a robust 4%, roughly one half of it’s original level. This is even more impressive if one allows for the possibility that this paper’s version of the January spread strategy, namely buying the smallest decile of NYSE and AMEX stocks and selling the largest decile, could likely be improved on. Suggestions by academics and practitioners include forming the portfolio some time prior to the last day of December, holding the portfolio for only two weeks rather than an entire month, and buying only those small stocks that experienced negative returns in the previous year. Although we do not pursue such refinements, we believe that they are likely to exist.

The January spread portfolio remains an important component of the optimal strategies corresponding to most of the beliefs considered here (the only exception being the no-anomaly prior), but the allocations to the January spread portfolio vary substantially across investors. In particular, by the end of our sample the diffuse investor allocates roughly half as much to the anomaly-based strategy as does the “no-decline” investor. Certainty equivalents demonstrate that these differences are economically significant.

While there is considerable imprecision about when the January effect began to disappear, the most likely period is somewhere in the mid-1970s, around the time the first study on January seasonality in stock returns was published by Rozeff and Kinney (1976). Our analysis in no way shows the causality of these events, but it is consistent with the possibility, advocated by Mittoo and Thompson (1990), that the dissemination of academic research may increase market efficiency.

The simplistic view that inefficiencies exist until someone discovers them is, however, contradicted by the fact that this particular anomaly has persisted for so long. While numerous studies have attempted to uncover the sources of the January effect, we believe that plausible

explanations must be consistent with the observed pattern of decay. The window dressing explanation of Haugen and Lakonishok (1988), for example, might be consistent with the average level of the January effect, but the story would seem to require an additional component to make it consistent with the effect's gradual decline. We believe that exploration of such issues would be of interest.

4 Conclusion

We present evidence supporting slow disappearance of both weekly index autocorrelations and the January effect.

Autocorrelations in the value-weighted index have almost certainly declined to a small fraction of their original level. Though there is some uncertainty about when this decline began, the early 1970s appears most probable. In comparison, the equally-weighted index continues to indicate weekly autocorrelations of 10-20% as of the end of 2001. That this value has declined over time is somewhat supported by the data, much more so when volatility is modelled as stochastic. In this case, it is also likely that the decay began sometime in the 1970s and almost certainly by 1985, dates that precede the publication of Lo and MacKinlay (1988), the seminal academic article on the subject.

Frequentist analysis of the January effect reveals marginal evidence against the constancy of the January abnormal return. Bayesian estimates suggest that almost 20 years after the papers by Keim (1983) and Reinganum (1983), the January effect remains at about half its original level. While the starting date of this decay is highly uncertain, the 1970s stands out as most likely, with 1976 being the most probable first year of decline under diffuse priors. While this year predates the papers by Keim and Reinganum, it coincides with the publication of Rozeff and Kinney's first paper on January seasonalities in the stock market.

Accounting for a potential decay can be of substantial economic importance. Following Kandel and Stambaugh (1996), we assess its impact by investigating optimal portfolio allocations. We find in all cases that an investor who allows for the possibility of anomaly decay will

choose portfolio weights that are far different from those chosen by an investor who assumes that the anomaly is constant through time. Moreover, certainty equivalent differences reveal large utility costs for investors required to hold suboptimal portfolios.

We also document a surprising result that shrinking abnormal returns *in sample* towards a benchmark (in our case, the CAPM) may in fact lead to increased allocations to the anomalous asset, as shrinkage can actually cause the *out-of-sample* predictions of abnormal returns to rise. Effectively, a prior that reduces the initial level of an anomaly causes that anomaly to persist for longer, so that its level at the end of the sample might be higher than it would be under a diffuse prior. Hence, skepticism about the existence of tradeable market inefficiencies might be better represented as a belief both about the levels of anomalies *and* their rates of disappearance.

The disappearance of anomalies has important implications for research on their underlying causes. Most obviously, studies using shorter, more recent sample periods may provide little insight on the historical factors that contributed to the anomalous behavior. For example, Ahn et al. (2002) attempt to determine the source of index autocorrelation by examining a variety of value-weighted indexes and their corresponding futures contracts. However, equity index futures were not traded until the 1980s or later, by which time value-weighted (though not equally-weighted) autocorrelations remained at just a tiny fraction of their original levels.

In addition, potential explanations of apparent anomalies should be consistent with the observed patterns of decay. This is likely to be a challenge to most finance theories, in which markets are commonly modelled as single-period or stationary equilibria. In the case of index autocorrelations, explanations involving behavioral biases might be difficult to reconcile with the severe time-variation noticed in the level of that anomaly, and theories that explain autocorrelations with exogenous variables that are naturally time varying may hold more promise. An example is the investor recognition hypothesis' possible attribution of autocorrelation to information costs, which have likely decreased for a variety of reasons over the last few decades. This explanation is consistent with Hou and Moskowitz's (2003) finding that firms with low measures of investor recognition have the most delayed response to new information.

One possible explanation of the persistence observed in some anomalies is that securities are mispriced, but that transactions costs prevent a small number of informed traders from immediately arbitraging anomalies away. Only as knowledge of the inefficiency gradually disseminates across large numbers of investors, who trade on the anomaly in the process of implementing more conventional investment strategies, is the anomaly eroded. Another alternative is that it is transactions costs themselves that have declined over time. As investing in the anomaly becomes cheaper, it naturally disappears. While this could be true for some cases, other anomalies, such as value-weighted index autocorrelations, appear to shrink too quickly for this to be a reasonable possibility.

Understanding why some anomalies disappear quickly (e.g., weekly autocorrelation of the value-weighted index), while others last longer (e.g., the January effect) or do not seem to decline at all (e.g., momentum, see Jegadeesh and Titman, 2001) seems an important matter that this study leaves untouched. Such an investigation might improve our understanding of market efficiency and would likely have implications for the current debate over rational and behavioral asset pricing models. We plan to pursue these issues in future work.

A Student-t / stochastic volatility sampling algorithm

In this appendix we adapt the methods of Geweke (1993) and Jacquier, Polson, and Rossi (1994, hereafter JPR) to estimate the stochastic volatility model

$$R_t - \mu = \rho_0 \delta^{(t-\tau)^+} (R_{t-1} - \mu) + \epsilon_t \quad (26)$$

$$\ln h_t = a + b \ln h_{t-1} + c R_{t-1} + \eta_t, \quad (27)$$

where $\epsilon_t \sim t(0, h_t, \nu)$, $\eta_t \sim N(0, \sigma_\eta^2)$, and $t \in \{1, 2, \dots, T\}$.

Following Geweke (1993), assuming that $\epsilon_t \sim t(0, h_t, \nu)$ is equivalent to assuming that $\epsilon_t = \sqrt{h_t} \sqrt{\omega_t} \epsilon_t^*$, where ϵ_t^* is a standard normal and $\nu/\omega_t \sim \chi^2(\nu)$. We adopt this latter representation here.

Introducing stochastic volatility to the framework outlined in Section 2 requires adding an additional component to the Gibbs sampling algorithm. First, conditioning on the time series of h_t and of ω_t , we draw the parameters σ_η , a , b , c , and ν . Second, conditional on all the parameters as well as the asset returns, we draw values of h_t and ω_t .

Conditional on the h_t , drawing the parameters σ_η , a , b , and c uses standard regression results such as those found in Zellner (1971). In particular, the distribution of σ_η is an inverted gamma, and the vector $[a, b, c]$ given σ_η is multivariate normal.

Conditional on the ω_t , drawing ν is also fairly straightforward. Following Geweke (1993), the conditional density of ν_t is

$$(\nu/2)^{T\nu/2} \Gamma(\nu/2) \exp(-\xi\nu), \quad \text{where } \xi = \frac{1}{2} \sum_{t=1}^T [\ln(\omega_t) + \omega_t^{-1}] + \lambda. \quad (28)$$

We sample from this univariate density using the griddy Gibbs sampler.

To draw the latent variable ω_t given the all the parameter values and the time series of h_t and R_t , we use the result from Geweke that

$$(\epsilon_t^2/h_t + \nu) / \omega_t \sim \chi^2(\nu + 1). \quad (29)$$

The last step is to draw the h_t conditional on the model parameters and the ω_t . As in JPR, we draw the entire time series of h_t by cycling through each element one at a time. In effect,

this step actually consists of T separate draws from the densities

$$p(h_t|h_1, h_2, \dots, h_{t-1}, h_{t+1}, \dots, h_T, R_1, R_2, \dots, R_T, \omega_1, \omega_2, \dots, \omega_T, \theta), \quad (30)$$

where θ represents the vector of all model parameters. Similar to JPR, the Markovian nature of the h_t process and Bayes rule together imply that this density is proportional to

$$p(h_t|h_{t-1}, \theta)p(h_{t+1}|h_t)p(R_t|h_t, \omega_t). \quad (31)$$

Furthermore, this product of densities is proportional to

$$f(h_t) \equiv \frac{1}{h_t} \exp\left(-\frac{(\ln h_t - m_t)^2}{2s_t^2}\right) \frac{1}{\sqrt{h_t}} \exp\left(-\frac{\epsilon_t^2}{2\omega_t h_t}\right), \quad (32)$$

where

$$\epsilon_t = (R_t - \mu) - \rho_0 \delta^{(t-\tau)^+} (R_{t-1} - \mu) \quad (33)$$

$$m_t = [a(1-b) + b(\ln h_{t-1} + \ln h_{t+1}) + cR_{t-1} - bcR_t] / (1 + b^2) \quad (34)$$

$$s^2 = \sigma_\eta^2 / (1 + b^2) \quad (35)$$

When $c = 0$ and $\omega_t = 1$, these match the formulas found in JPR.

Similarly to JPR, we draw from this density using the Metropolis Hastings algorithm with the candidate generating density

$$q(h_t) \propto h_t^{-(\phi+1)} \exp(-\psi_t/h_t), \quad (36)$$

where $\phi = -1.5 + (1 - 2 \exp(s^2)) / (1 - \exp(s^2))$ and $\psi_t = .5\epsilon^2/\omega_t + (\phi + 1) \exp(m_t + .5s^2)$. This produces a inverse gamma candidate generating density that approximates the target density in (32). JPR show that this candidate generator has relatively thick tails and demonstrates good convergence properties.

A candidate draw h_t^* from this inverse gamma is then accepted, replacing the current draw h_t , with probability

$$\min \left\{ \frac{f(h_t^*)/q(h_t^*)}{f(h_t)/q(h_t)}, 1 \right\} \quad (37)$$

If the draw is rejected, the current draw h_t is kept.

By drawing each h_t in turn, from $t = 1$ to $t = T$, a new draw of the time series of h_t is obtained. While the algorithm must be modified for $t = 1$ and $t = T$, this is straightforward following the procedure of Jacquier, Polson, and Rossi (2002).

References

- Ahn, D., J. Boudoukh, M. Richardson, and R.F. Whitelaw, 2002, "Partial Adjustment or Stale Prices? Implications from Stock Index and Futures Return Autocorrelations," *Review of Financial Studies* 15, 655-689.
- Andrews, D.K., 1993, "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica* 61, 821-856.
- Avramov, D., 2002, "Stock Return Predictability and Model Uncertainty," *Journal of Financial Economics* 64, 423-458.
- Badrinath, S.G., J.R. Kale, T.H. Noe, 1995, "Of Shepards, Sheep, and Cross-Autocorrelations in Equity Returns," *Review of Financial Studies* 8, 401-430.
- Banz, R.W., 1981, "The Relationship Between Return and Market Value of Common Stocks," *Journal of Financial Economics* 9, 3-18.
- Blume, M.E. and R.F. Stambaugh, 1983, "Biases in Computed Returns: An Application to the Size Effect," *Journal of Financial Economics* 12, 387-404.
- Booth, D.G. and D.B. Keim, 2000, "Is There Still a January Effect?" in: D.B. Keim and W. T. Ziemba (eds.), *Security Market Imperfections in Worldwide Equity Markets* Cambridge, University Press, UK, 169-178.
- Brennan, M.J. and Y. Xia, 2001, "Assessing Asset Pricing Anomalies," *Review of Financial Studies* 14, 905-942.
- Casella, G. and E.I. George, 1992, "Explaining the Gibbs Sampler," *American Statistician* 46, 167-174.
- Cochrane, J., 1999, "New Facts in Finance," *Economic Perspectives* 23, Federal Reserve Bank of Chicago, 36-58.
- Conrad, J. and G. Kaul, 1989, "Mean Reversion in Short-Horizon Expected Returns," *Review of Financial Studies* 2, 225-240.
- Cremers, M., 2002, "Stock Return Predictability: A Bayesian Model Selection Perspective," *Review of Financial Studies* 15, 1223-1249.
- Fama, E., 1965, "The Behavior of Stock Prices," *Journal of Business* 38, 34-105.

- Fama, E., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25, 383-417.
- Fama, E. and K. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics* 33, 3-56.
- Fisher, L., 1966, "Some New Stock Market Indexes," *Journal of Business* 39, 191-225.
- French, K., 1980, "Stock Returns and the Weekend Effect," *Journal of Financial Economics* 8, 55-69.
- Gallant, A.R., D. Hsieh, and G. Tauchen, 1997, "Estimation of Stochastic Volatility Models with Diagnostics," *Journal of Econometrics* 81, 159-192.
- Geman, S. and D. Geman, 1984, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721-741.
- Geweke, J., 1993, "Bayesian Treatment of the Independent Student-t Linear Model," *Journal of Applied Econometrics* 8, S19-S40.
- Grinblatt, M. and T.J. Moskowitz, 2003, "Predicting Stock Price Movements from the Pattern of Stock Returns: The Role and Consistency of Tax Loss Selling," *Journal of Financial Economics*, forthcoming.
- Haugen, R.A. and P. Jorion, 1996, "The January Effect: Still There After All These Years," *Financial Analysts Journal*, January-February, 27-31.
- Haugen, R.A. and J. Lakonishok, 1988, *The Incredible January Effect*, Dow Jones-Irwin, Homewood, IL.
- Hawawini, G., 1980, "Intertemporal Cross Dependence in Securities Daily Returns and the Short-Term Intervailing Effect on Systematic Risk," *Journal of Financial and Quantitative Analysis* 15, 139-149.
- Hawawini, G. and D.B. Keim, 1995, "On the Predictability of Common Stock Returns: World-Wide Evidence," in: R.A. Jarrow, V. Maksimovic, and W.T. Ziemba (eds.), *Handbooks in Operations Research and Management Science* 9, North Holland, Amsterdam, 497-544.

- Hong, H. and J.C. Stein, 1999, "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance* 54, 2143-2184.
- Hou, K. and T.J. Moskowitz, 2003, "Market Frictions, Price Delay, and the Cross-Section of Expected Returns," working paper, Ohio State University (Hou) and University of Chicago (Moskowitz).
- Jacquier, E., N.G. Polson, and P.E. Rossi, 1994, "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics* 12, 371-389.
- Jacquier, E., N.G. Polson, and P.E. Rossi, 2002, "Bayesian Analysis of Stochastic Volatility Models with Fat Tails and Correlated Errors," *Journal of Econometrics*, forthcoming. Chicago (Polson and Rossi).
- Jegadeesh, N. and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance* 48, 65-91.
- Jegadeesh N. and S. Titman, 2001, "Profitability of Momentum Strategies: An Evaluation of Alternative Explanations," *Journal of Finance* 56, 699-720.
- Kandel, S. and R.F. Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset Allocation Perspective," *Journal of Finance* 51, 385-424.
- Keim, D.B., 1983, "Size-Related Anomalies and Stock Return Seasonality: Further Empirical Evidence," *Journal of Financial Economics* 12, 13-32.
- Llorente, G., R. Michaely, G. Saar, and J. Wang, 2002, "Dynamic Volume-Return Relation of Individual Stocks," *Review of Financial Studies* 15, 1005-1047.
- Lo, A.W. and A.C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies* 1, 41-66.
- Lo, A.W. and A.C. MacKinlay, 1990, "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies* 3, 175-205.
- Lo, A.W. and A.C. MacKinlay, 1999, *A Non-Random Walk Down Wall Street*, Princeton University Press, Princeton, NJ.
- McQueen, G. and S. Thorley, 1997, "Do Investors Learn? Evidence from a Gold Market Anomaly," *Financial Review* 32, 501-525.

- Mech, T.S, 1993, "Portfolio Return Autocorrelation," *Journal of Financial Economics* 34, 307-344.
- Merton, R.C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance* 42, 483-510.
- Mittoo, U. and R. Thompson, 1990, "Do the Capital Markets Learn from Financial Economists?," working paper, University of Manitoba (Mittoo) and Southern Methodist University (Thompson).
- Newey, W.K. and K.D. West, 1987, "A Simple Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55, 703-705.
- Pastor, L., 2000, "Portfolio Selection and Asset Pricing Models," *Journal of Finance* 40, 179-223.
- Pastor, L. and R.F. Stambaugh, 1999, "Cost of Equity Capital and Model Mispricing," *Journal of Finance* 54, 67-121.
- Pastor, L. and R.F. Stambaugh, 2000, "Comparing Asset Pricing Models: an Investment Perspective," *Journal of Financial Economics* 56, 335-381.
- Pastor, L. and R.F. Stambaugh, 2001, "The Equity Premium and Structural Breaks," *Journal of Finance* 56, 1207-1239.
- Reinganum, M.R., 1983, "The Anomalous Stock Market Behavior of Small Firms in January: Empirical Tests for Tax-Loss Selling Effects," *Journal of Financial Economics* 12, 89-104.
- Reinganum, M.R., 1990, "Market Microstructure and Asset Pricing: An Empirical Investigation of NYSE and NASDAQ Securities," *Journal of Financial Economics* 28, 127-147.
- Ritter, C. and M.A. Tanner, 1992, "The Gibbs Stopper and the Griddy Gibbs Sampler," *Journal of the American Statistical Association* 87, 861-868.
- Rozeff, M. and W. Kinney, 1976, "Capital Market Seasonality: The Case of Stock Returns," *Journal of Financial Economics* 3, 379-402.
- Schwert, G.W., 2001, "Anomalies and Market Efficiency," in: G. Constantinides, M. Harris, and R.M. Stulz (eds.), *Handbook of the Economics and Finance*, North-Holland, Amsterdam.
- Shapiro, A., 2002, "The Investor Recognition Hypothesis in a Dynamic General Equilibrium: Theory and Evidence," *Review of Financial Studies* 15, 97-141.

- Sias, R.W. and L.T. Starks, 1997, "Return Autocorrelation and Institutional Investors," *Journal of Financial Economics* 46, 103-131.
- Watts, R., 1978, "Systematic 'Abnormal' Returns after Quarterly Earnings Announcements," *Journal of Financial Economics* 6, 127-150.
- Zellner, A., 1971, *An Introduction to Bayesian Inference in Econometrics*, Wiley, New York.
- Zhang, L., 2003, "The Value Premium," working paper, University of Rochester.

Table 1
Frequentist analysis of equity index autocorrelations

	Equally-Weighted Index	Value-Weighted Index	Decile 1-8 Index
Sample autocorrelations			
OLS (in parentheses) and Newey-West (in brackets) standard errors			
1962 - 2001	0.257 (0.021) [0.102]	0.038 (0.022) [0.029]	0.234 (0.021) [0.090]
1962 - 1987	0.280 (0.026) [0.138]	0.089 (0.027) [0.041]	0.257 (0.027) [0.121]
1988 - 2001	0.186 (0.036) [0.112]	-0.064 (0.037) [0.043]	0.155 (0.036) [0.089]
Structural break test <i>t</i> -statistics			
Using OLS (top number) and Newey-West (bottom number) standard errors			
	1.808 0.420	3.147 2.569	1.935 0.537

Table 2A
Posterior summary of disappearing autocorrelation, constant volatility
Medians and 90% confidence intervals (in parentheses), 1962-2001 sample

	Diffuse prior	LM (1988) prior	No-decline prior	No-anomaly prior
<i>Equally-Weighted Index</i>				
$\mu \times 52$	0.182 (0.132, 0.232)	0.182 (0.132, 0.233)	0.184 (0.133, 0.234)	0.184 (0.144, 0.222)
ρ_0	0.284 (0.236, 0.352)	0.269 (0.232, 0.306)	0.258 (0.222, 0.293)	0 (n/a)
δ	0.999 (0.950, 1.000)	0.999 (0.998, 1.000)	1 (n/a)	
$\sigma \times \sqrt{52}$	0.146 (0.142, 0.149)	0.146 (0.142, 0.150)	0.146 (0.142, 0.150)	0.151 (0.147, 0.155)
<i>Value-Weighted Index</i>				
$\mu \times 52$	0.120 (0.081, 0.160)	0.120 (0.081, 0.159)	0.118 (0.079, 0.158)	0.119 (0.080, 0.156)
ρ_0	0.123 (0.036, 0.271)	0.083 (0.038, 0.126)	0.039 (0.002, 0.074)	0 (n/a)
δ	0.582 (0.059, 0.998)	0.462 (0.044, 0.987)	1 (n/a)	
$\sigma \times \sqrt{52}$	0.143 (0.140, 0.147)	0.143 (0.140, 0.147)	0.143 (0.140, 0.147)	0.143 (0.140, 0.147)

Table 2B**Posterior summary of disappearing autocorrelation, stochastic volatility**

Medians and 90% confidence intervals (in parentheses), 1962-2001 sample

	Diffuse prior	LM (1988) prior	No-decline prior	No-anomaly prior
<i>Equally-Weighted Index</i>				
$\mu \times 52$	0.210 (0.172, 0.247)	0.207 (0.170, 0.244)	0.210 (0.172, 0.248)	0.224 (0.193, 0.253)
ρ_0	0.362 (0.292, 0.444)	0.292 (0.248, 0.337)	0.255 (0.216, 0.293)	0 (n/a)
δ	0.999 (0.998, 1.000)	0.998 (0.567, 0.999)	1 (n/a)	
a	-0.408 (-0.577, -0.264)	-0.364 (-0.518, -0.235)	-0.370 (-0.545, -0.238)	-0.323 (-0.478, -0.191)
b	0.949 (0.928, 0.966)	0.954 (0.936, 0.970)	0.953 (0.932, 0.970)	0.959 (0.940, 0.975)
c	-4.668 (-5.846, -3.498)	-4.589 (-5.77, -3.461)	-4.482 (-5.738, -3.307)	-3.829 (-5.044, -2.691)
σ_η	0.243 (0.192, 0.296)	0.220 (0.181, 0.270)	0.229 (0.176, 0.277)	0.226 (0.175, 0.275)
ν	27 (11, 46)	29 (14, 47)	29 (15, 46)	24 (13, 47)
<i>Value-Weighted Index</i>				
$\mu \times 52$	0.129 (0.097, 0.161)	0.128 (0.095, 0.161)	0.130 (0.097, 0.162)	0.133 (0.103, 0.162)
ρ_0	0.197 (0.107, 0.282)	0.110 (0.063, 0.155)	0.041 (0.002, 0.079)	0 (n/a)
δ	0.562 (0.059, 0.996)	0.448 (0.044, 0.949)	1 (n/a)	
a	-0.214 (-0.368, -0.127)	-0.213 (-0.358, -0.130)	-0.206 (-0.325, -0.121)	-0.212 (-0.311, -0.135)
b	0.973 (0.953, 0.983)	0.973 (0.955, 0.983)	0.973 (0.959, 0.984)	0.973 (0.960, 0.982)
c	-4.931 (-6.643, -3.780)	-4.946 (-6.347, -3.626)	-4.639 (-6.065, -3.362)	-4.810 (-6.108, -3.620)
σ_η	0.144 (0.110, 0.199)	0.144 (0.118, 0.190)	0.141 (0.111, 0.179)	0.143 (0.116, 0.169)
ν	32 (18, 47)	32 (19, 48)	36 (20, 49)	30 (16, 47)

Table 3
Allocations to the market index as of December 31, 2001

<i>Prior:</i>	Diffuse	LM (1988)	No decline	No anomaly	Diffuse	LM (1988)	No decline	No anomaly
Constant Volatility				Stochastic Volatility				
Following a -4% return	-3.93	-4.13	-7.63	1.91	-0.88	-0.49	-6.00	2.32
Following a zero return	1.54	1.53	1.27	1.91	2.15	2.18	1.60	2.71
Following a +4% return	7.00	7.18	10.17	1.91	6.44	5.89	12.42	3.17
<i>Value-Weighted Index</i>								
Constant Volatility				Stochastic Volatility				
Following a -4% return	0.99	1.01	-0.46	1.00	0.94	0.93	-0.28	1.00
Following a zero return	1.03	1.03	0.92	1.00	1.16	1.14	1.08	1.22
Following a +4% return	1.07	1.05	2.30	1.00	1.43	1.40	3.01	1.48

Table 4
Average certainly equivalent losses as of December 31, 2001

<i>Prior used to find weights:</i>	Diffuse	LM (1988)	No decline	No anomaly	Diffuse	LM (1988)	No decline	No anomaly
	<i>Equally-Weighted Index</i>				<i>Value-Weighted Index</i>			
<i>Prior of investor</i>	Constant Volatility				Stochastic Volatility			
Diffuse	0	0.0004	0.1046	0.2457	0	0.0022	0.2592	0.0896
LM (1988)	0.0004	0	0.0918	0.2628	0.0023	0	0.3125	0.0654
No decline	0.1012	0.0897	0	0.6511	0.2573	0.3064	0	0.6461
No anomaly	0.2559	0.2753	0.6940	0	0.0890	0.0643	0.6477	0
Diffuse	0	0.0000	0.0128	0.0000	0	0.0000	0.0141	0.0001
LM (1988)	0.0000	0	0.0133	0.0000	0.0000	0	0.0144	0.0002
No decline	0.0129	0.0134	0	0.0137	0.0141	0.0145	0	0.0145
No anomaly	0.0000	0.0000	0.0137	0	0.0001	0.0002	0.0144	0

Table 5
Maximum likelihood estimates of the January effect
Asymptotic standard errors in parentheses

$R_{\text{spr},t} = \alpha_0 + \alpha_1 I_{J,t} + \beta R_{\text{em},t} + \epsilon_t$			
	1962-1981	1982-2001	
α_0	-0.0014 (0.0036)	-0.0085 (0.0034)	
α_1	0.0964 (0.0093)	0.0633 (0.0073)	
β	0.4029 (0.0534)	-0.1490 (0.0745)	
σ_ϵ	0.0493 (0.0018)	0.0434 (0.0016)	
$R_{\text{spr},t} = \alpha_0 + \alpha_1 \delta^{(t-\tau)^+} I_{J,t} + \beta R_{\text{em},t} + \epsilon_t$			
	1962-2001		
α_0	0.0007 (0.0028)	-0.0057 (0.0026)	-0.0056 (0.0024)
α_1	0 (n/a)	0.0798 (0.0057)	0.1022 (0.0114)
δ		1 (n/a)	0.9678 (0.0138)
τ			1976.0 (5.0898)
β	0.1669 (0.0409)	0.1278 (0.0388)	0.1254 (0.0400)
σ_ϵ	0.0532 (0.0011)	0.0485 (0.0012)	0.0481 (0.0013)
AIC	-2299.471	-2385.583	-2387.464
BIC	-2278.665	-2360.616	-2354.175
BIC*	-2291.153	-2375.601	-2374.156
LRT (<i>p</i> -value)	93.9933 (0.0000)	5.8815 (0.0528)	

Table 6

Posterior summary of the January effect under alternative priors

Medians and 90% confidence intervals (in parentheses), 1962-2001 sample, for the model

	$R_{\text{spr},t} = \alpha_0 + \alpha_1 \delta^{(t-\tau)^+} I_{j,t} + \beta R_{\text{rem},t} + \epsilon_t$				$R_{\text{rem},t} \sim N(\mu_m, \sigma_m^2)$	
	Diffuse prior	2% CAPM prior	1% CAPM prior	No-decline prior	No-anomaly prior	
α_0	-0.0056 (-0.0094, -0.0016)	-0.0046 (-0.0084, -0.0006)	-0.0028 (-0.0065, 0.0011)	-0.0058 (-0.0095, -0.0019)	0.0008 (-0.0033, 0.0048)	
α_1	0.0976 (0.0749, 0.1208)	0.0711 (0.0569, 0.0864)	0.0469 (0.0358, 0.0581)	0.0802 (0.0660, 0.0931)	0 (n/a)	
δ	0.9666 (0.8758, 0.9913)	0.9710 (0.5740, 0.9975)	0.9505 (0.3662, 0.9982)	1 (n/a)		
τ	1976 (1967, 1994)	1991 (1971, 2000)	1995 (1975, 2000)			
β	0.1267 (0.0417, 0.2094)	0.1335 (0.0478, 0.2207)	0.1436 (0.0568, 0.2325)	0.1315 (0.0454, 0.2111)	0.1678 (0.0704, 0.2615)	
σ_ϵ	0.0485 (0.0459, 0.0512)	0.0485 (0.0460, 0.0513)	0.0495 (0.0468, 0.0524)	0.0485 (0.0461, 0.0512)	0.0533 (0.0507, 0.0564)	
$E[R_{\text{spr},2002}]$	0.0396 (0.0103, 0.0679)	0.0472 (0.0024, 0.0674)	0.0336 (0.0005, 0.0490)	0.0747 (0.0611, 0.0879)	0.0017 (-0.0026, 0.0058)	
μ_m			0.0053 (0.0022, 0.0085)			
σ_m			0.0433 (0.0411, 0.0456)			

Table 7
Investing in the January portfolio as of December 31, 2001

<i>Prior used to find weights:</i>	Diffuse	2% CAPM	1% CAPM	No decline	No anomaly
Portfolio Holdings					
Allocation to the spread portfolio	5.158	5.663	3.979	10.784	0.092
Allocation to the market	0.344	0.239	0.421	-0.428	0.970
Certainty Equivalent Losses					
<i>Prior of investor</i>					
Diffuse	0	0.001	0.005	0.119	0.097
2% CAPM	0.001	0	0.011	0.100	0.118
1% CAPM	0.005	0.011	0	0.174	0.057
No decline	0.109	0.090	0.159	0	0.392
No anomaly	0.104	0.126	0.061	0.463	0

Figure 1: Posterior distributions of τ for disappearing autocorrelations under diffuse priors

1962-2001 sample

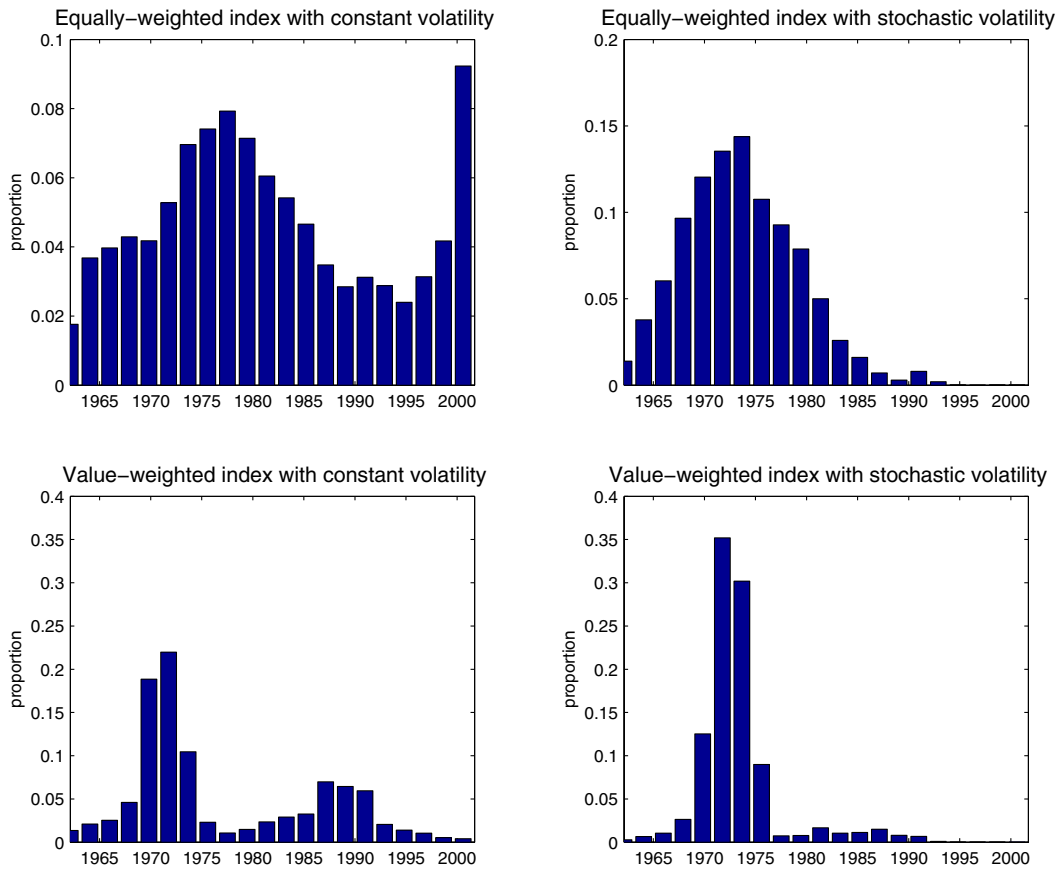


Figure 2: Posterior time series of autocorrelations under diffuse priors
 posterior medians and 90% confidence intervals, 1962-2001 sample

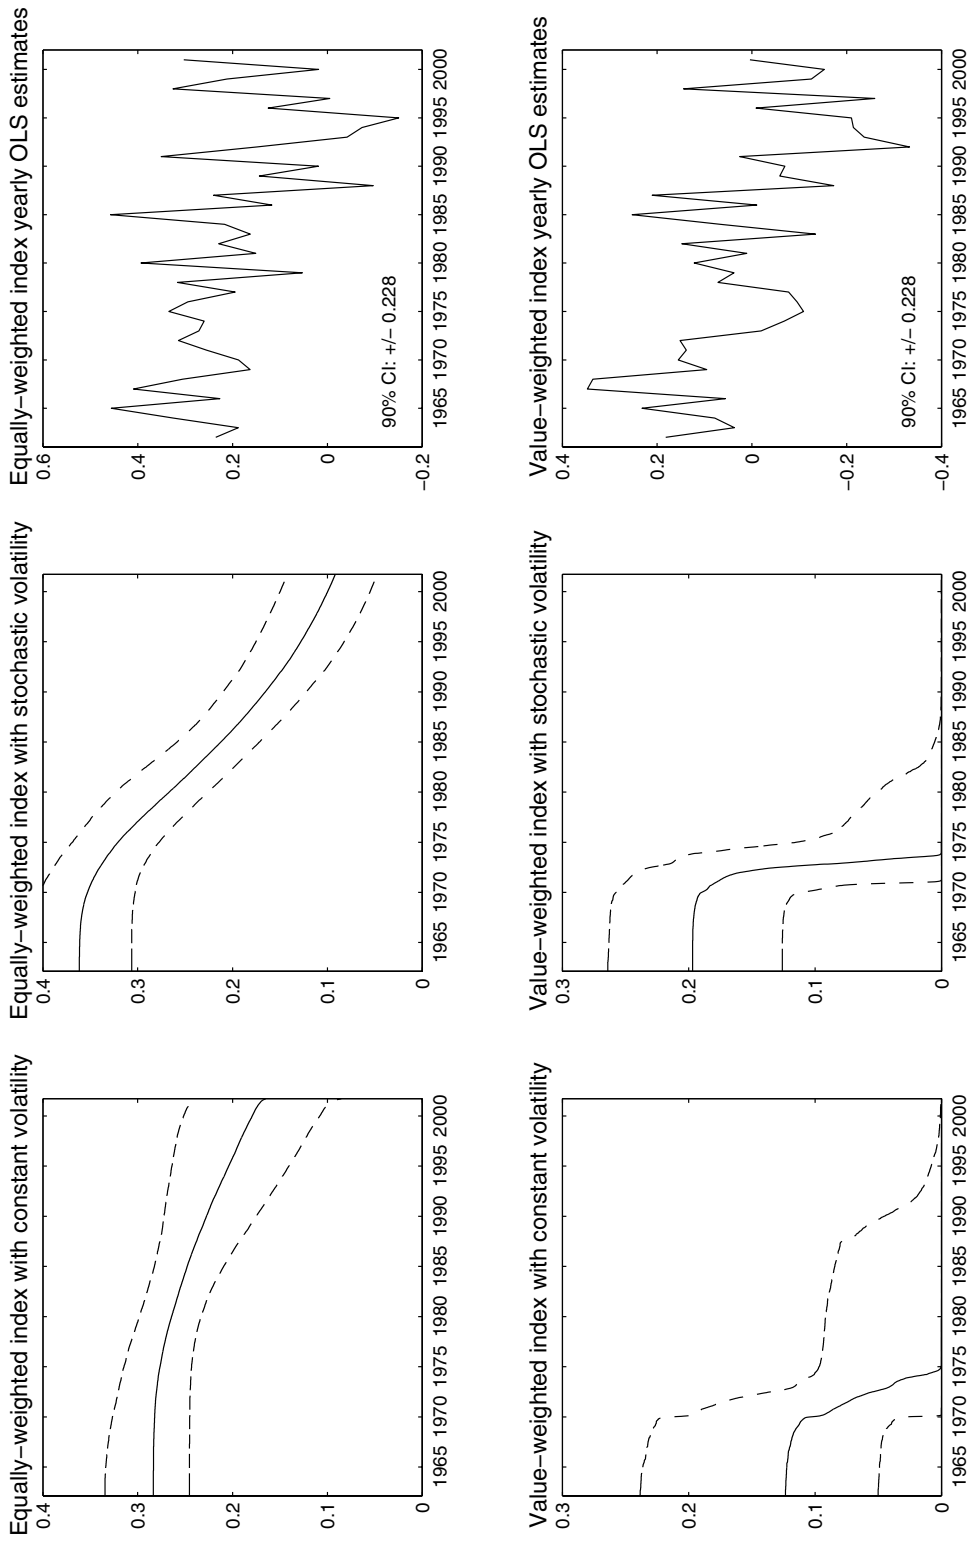


Figure 3: Rolling sample estimates of return autocorrelations

Samples starting in 1962 and ending between 1970 and 2001

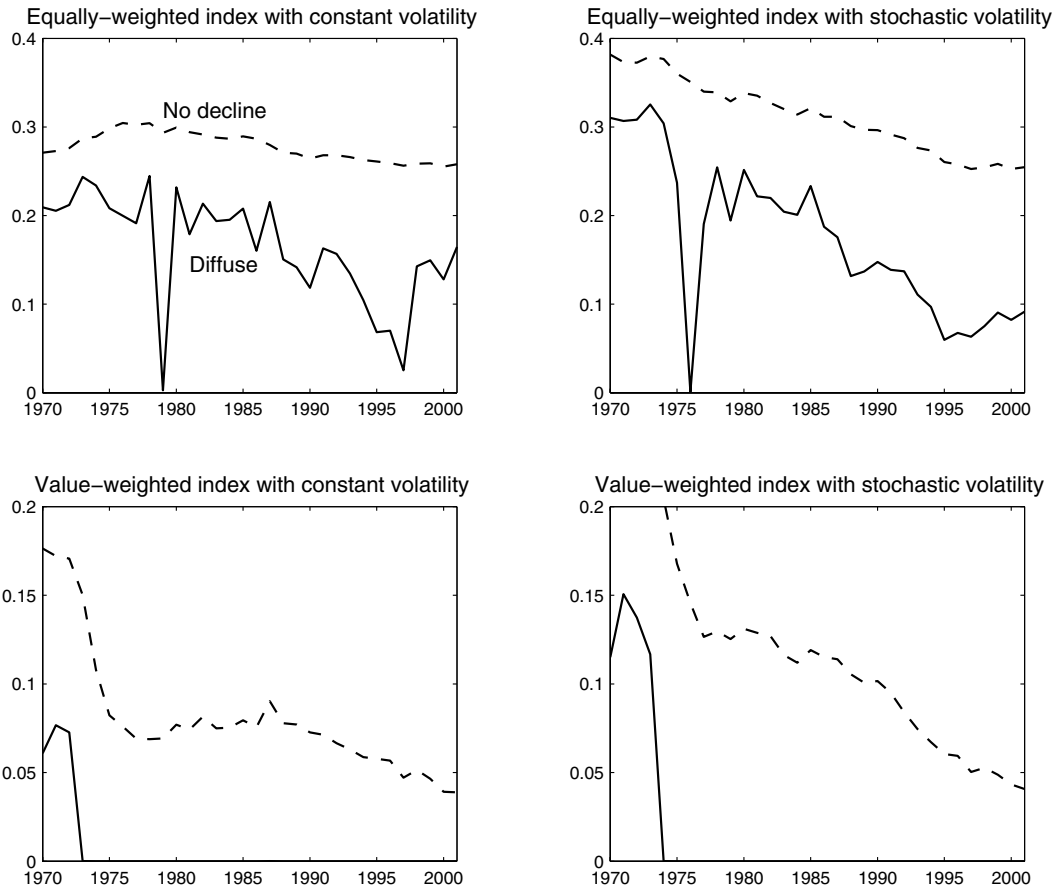


Figure 4: Rolling allocations to the equally-weighted market portfolio

Samples starting in 1962 and ending between 1970 and 2001

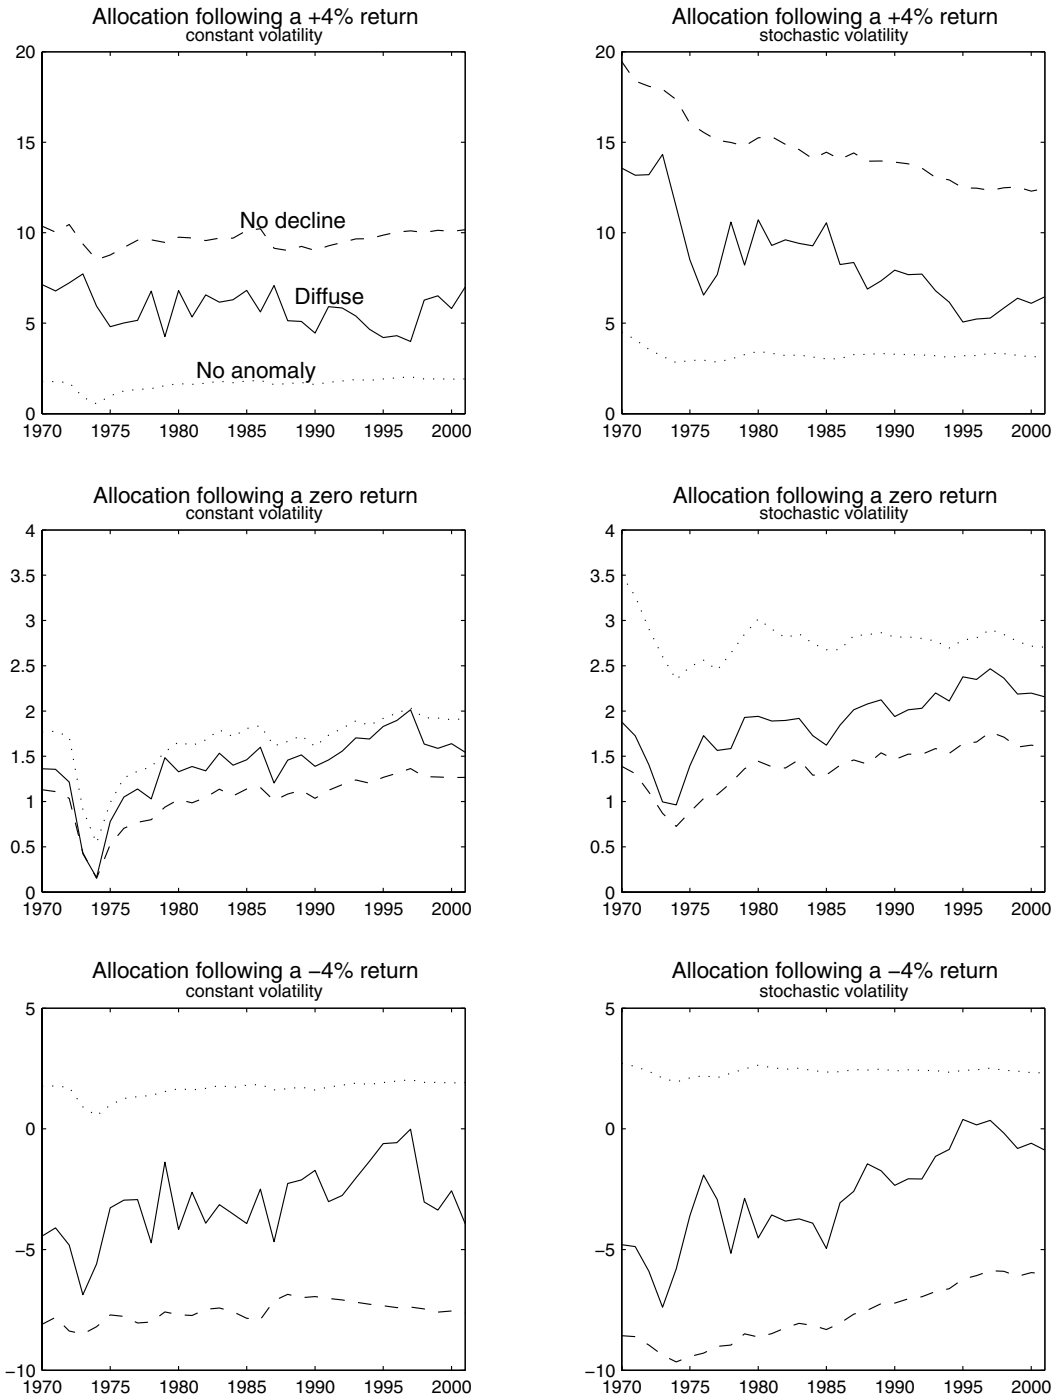


Figure 5: Rolling allocations to the value-weighted market portfolio

Samples starting in 1962 and ending between 1970 and 2001

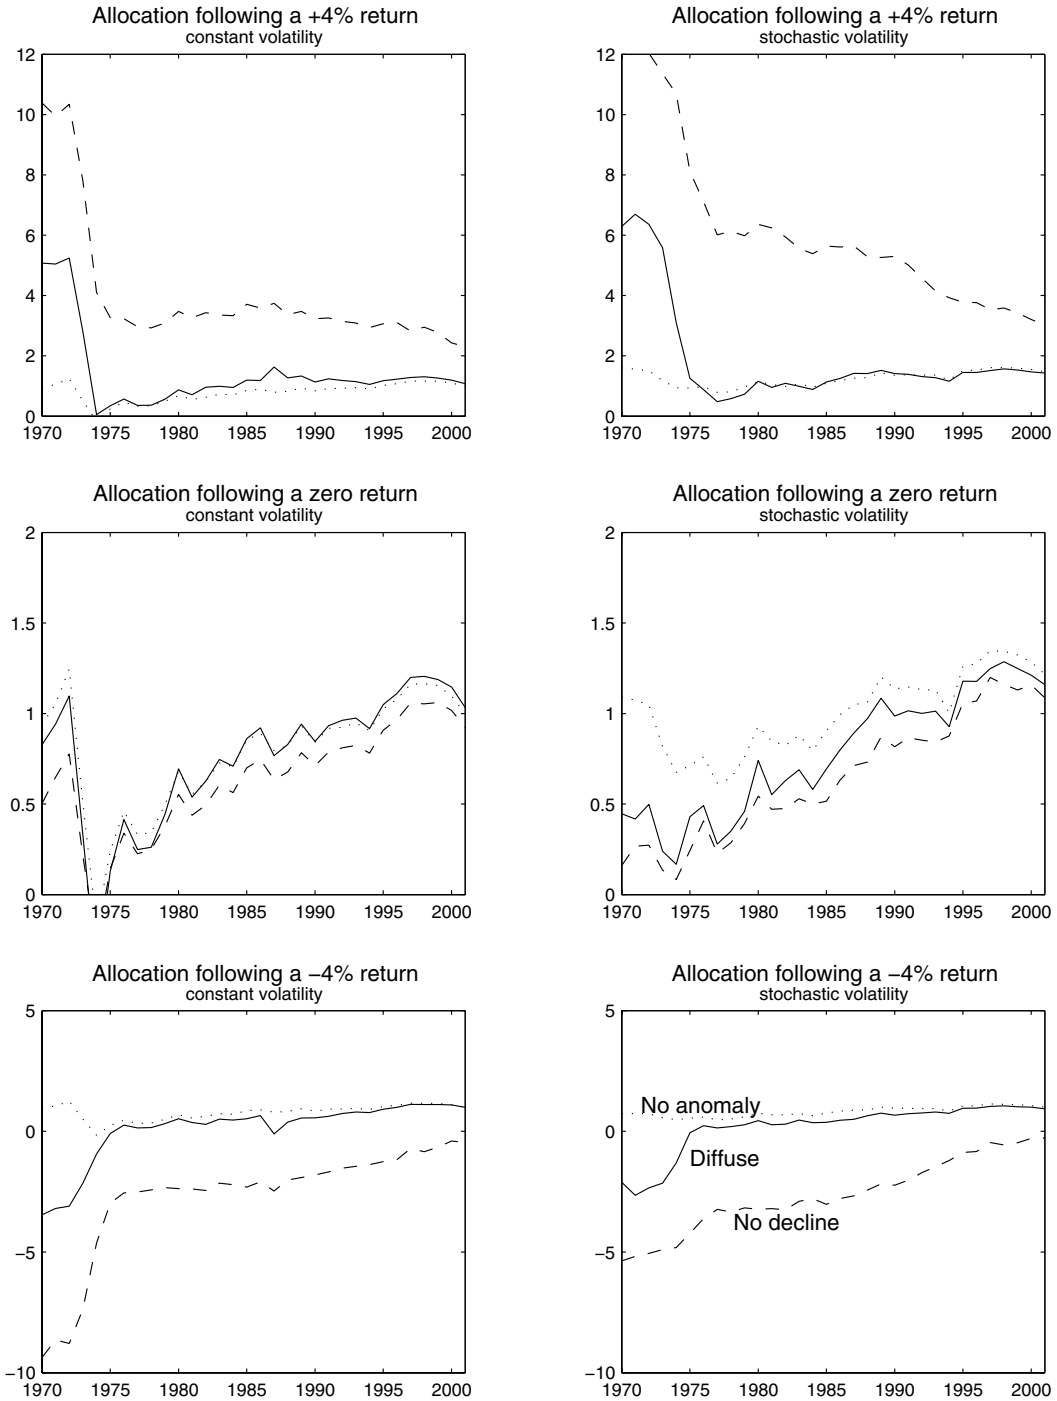


Figure 6: Posterior distributions for the January effect
1962-2001 sample





Figure 7: Rolling allocations to the January spread portfolio

Samples starting in 1962 and ending between 1972 and 2001

