Journal of Financial Markets 2 (1999) 49–68

Journal of
FINANCIAL
MARKETS

The alpha factor asset pricing model: A parable

Wayne E. Ferson*, Sergei Sarkissian, Timothy Simin

Department of Finance and Business Economics, University of Washington, Box 353200, Seattle, WA 98195-3200, USA

Received 11 August 1998

Abstract

Recent empirical studies use the returns of attribute-sorted portfolios of common stocks as if they represent risk factors in an asset pricing model. If the attributes are chosen following an empirically observed relation to the cross-section of stock returns, such portfolios will appear to be useful risk factors, even when the attributes are completely unrelated to risk. We illustrate this result using a parable and argue that the moral of the story is important in practice. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: C5; G12

Keywords: Asset pricing; Factor models; Cross-sectional regressions; Arbitrage portfolios; Anomalies

1. Introduction

"Had the author's intentions met with a more candid interpretation from some whom out of respect he forbears to name, he might have been encouraged to an examination of books written by some of those authors above described, whose errors, … he thinks he could have detected and exposed in such a manner that the persons who are most conceived to be affected by them, would soon lay them aside and be ashamed. But he has now given over those thoughts; since the weightiest men in the weightiest stations are pleased to think it a more dangerous point to laugh at those corruptions in Religion, which they themselves must disapprove, than to endeavor pulling up those very foundations wherein all Christians have agreed". – Jonathan Swift, from *A Tale of a Tub*, 1704.

* Corresponding author. Tel.: (206) 543-1843; fax: (206) 685-9392; e-mail via: ss@u.washington.edu

PII: S1386-4181(98)00005-6

Recent empirical studies have used portfolios formed by sorting individual stocks on the attributes of firms, such as market capitalization, lagged returns and ratios of earnings or accounting value to the price per share. Ferson (1996) argues theoretically, that such an approach can be problematic. Spread portfolios, long the high-return attribute stocks and short the low-return attribute stocks, can appear to serve as risk factors in a beta pricing framework, even when the attributes bear no relation to the cross-section of risks.

In this paper we use a parable to illustrate serious pitfalls in the use of attribute-sorted portfolios. A numerical example provides insights about cross-sectional regression coefficients, as in Fama and MacBeth (FM, 1973), and the spread portfolios used by Fama and French (FF, 1993, 1996). We emphasize the distinction between the conditional and unconditional variances of the portfolios. We explain why the market beta is needed in the FF 'three-factor model', even though the beta has no cross-sectional explanatory power for expected returns. The model illustrates several relations among spurious factors, sampling error and large-market arbitrage.

In our story, a hypothetical researcher discovers the alphabet effect; that is, a pattern in stock returns related to the first letter of a firm's name. The alphabet effect is an anomaly, completely unrelated systematic risk, and otherwise calibrated to resemble the book-to-market effect in stock returns. The researcher uses alphabet-sorted portfolios to explore the alphabet effect, first running monthly cross-sectional regressions of the returns on their position in the alphabet. The time series of these monthly cross-sectional regression coefficients are similar to spread portfolios that are long early-in-the-alphabet firms and short late-in-the-alphabet firms. The analyst uses an alphabet-spread portfolio as the *alpha risk factor*. He estimates time-series regression betas of assets on his factors and empirically examines the cross-sectional explanatory power of the betas for stock portfolio returns. He finds that they work very well, which gives rise to the '*alpha-factor asset pricing model*'.

Since the attributes are completely unrelated to risk, our parable illustrates that a researcher who uses attribute-sorted portfolios as if they were risk factors may have merely transported the anomaly from one measure to another. We argue that this raises serious concerns about some recent research. Recent findings (e.g. Fama and French, 1996) that size and book-to-market sorted portfolios appear to 'explain' the effects of price-to-book ratios, firm size, cash-flow-to-price, or earnings-to-price ratios are exactly what one would expect to find if these are true anomalies, unrelated to risk.[1]

---

[1] Of course, our story does not imply that attributes such as price-to-book are *in fact* unrelated to risk. However, there has been some debate about this, with studies on both sides of the issue. See, for example, Fama and French (1993, 1995, 1996), Lakonishok et al. (1994) and Daniel and Titman (1997). While the similarities to the work of Fama and French (1993, 1996) will be readily apparent, we do not imply that their three factor model' is spurious.

Our point is to raise a caution flag about the pitfalls of using attribute-sorted portfolios. In order to discover whether regularities such as the book-to-market effect are related to risk, it is necessary to model how the attributes are cross-sectionally related to explicit risk factors. Fama (1997) also emphasizes the necessity of identifying the risk factors in asset pricing tests. In the case of size and book-to-market, a number of recent studies have made some progress in the direction of measuring economic risks. Although our parable suggests that the profession should hold to the highest standards of evidence in this particular case, the jury is still out.[2]

Section 1 lays out the logic of the problem with the use of spread portfolios as risk factors, when stocks are sorted based on an anomaly. Section 2 provides the details of the data generating mechanism. The parable itself is in Section 3, written as a self-contained story. Section 4 explores more deeply the moral of the story. Section 5 concludes. Appendices provide supporting proofs and discuss the sensitivity of the results.

1.1. The logic of the analysis

We start with a three step procedure using a time-series of monthly, Fama–MacBeth style cross-sectional regression coefficients. Following Ferson (1996), we then argue that spread portfolio returns may be viewed as similar to the FM coefficients. Our numerical example shows that the critique is valid for both FM coefficients and spread portfolios.

In the first step, stock returns are sorted into portfolios on the basis of some attribute of the stocks that is related to the cross-section of returns. To make the analysis as stark as possible, we assume that the attribute is an anomaly, completely unrelated to the risk of the stocks. The anomaly may represent some irrational mispricing, or alternatively, it may be the result of 'data snooping', as modelled by Lo and MacKinlay (1990). (Section 4 explores these alternatives in more detail.) The anomalous attribute of the sorted portfolios is δ, an N-vector,

[2] Fama and French (1995) find size and book-to-market factors in accounting earnings and sales, which is suggestive of a risk interpretation. Jagannathan et al. (1997) and Kubota and Takehara (1995) provide direct evidence for Japanese stocks that book-to-market ratios are related to risk. On the other hand, Lakonishok et al. (1994), present evidence for the U.S. which they view as inconsistent with the idea book-to-market is a proxy for risk. Daniel and Titman (1997) also consider an 'attribute based' model for returns, and conclude that the book-to-market effect is not explained by loadings on the Fama-French 'factors'. However, if these loadings are not measures of economic risk exposure such evidence does not resolve the issue. Berk (1997) criticizes the power of the sorting procedures used by Daniel and Titman (1997), and Davis et al. (1998) find book-to-market loadings to be significant in a longer sample of U.S. data. Cheung et al. (1995), Ferson and Harvey (1998) and Brennan et al. (1998) provide models of the relation of book-to-market and other attributes, to specific risk factors.

where N is the number of assets. Suppose that the elements of the vector are ordered, starting with the high-return values of the attribute, then moving to the low-return values.

In the second step, regressions of the cross-sectional vector of portfolio returns at time t, R_t, are run on the attribute vector. For simplicity, assume that δ is constant over time and centered at zero (a more general analysis follows below). The Fama–MacBeth slope coefficient for month t is a portfolio, R_{pt}, with weights that are proportional to the attribute vector: $R_{pt} \propto \delta' R_t$.

In the third step, the time-series of R_{pt} is used as a 'risk factor'. In particular, time-series regressions of the N returns in R_t on R_{pt} deliver estimates of the factor betas. The N-vector of the factor betas is proportional to $\mathrm{Cov}(R_{pt}, R_t)$ and therefore to $\delta'\mathrm{Var}(R_t)$. The problem arises when a cross-sectional analysis of returns on these factor betas is conducted. A regression of returns on the factor betas will essentially reproduce the original anomaly when the factor betas are proportional to δ, or $\delta'\mathrm{Var}(R_t) \propto \delta'$. This is likely to occur when the original pattern in the returns is an anomaly, unrelated to risk. For example, suppose that returns are independent and identically distributed, with $\mathrm{Var}(R_t) = \sigma^2 I$. In this case there is no systematic risk and there should be no risk premiums, but $\delta'\mathrm{Var}(R_t) = \sigma^2\delta'$, so the spurious factor betas will appear to exactly price the assets!

Similar results should be obtained when spread portfolios replace the FM coefficients. The argument recalls Fama's (1976) interpretation of cross-sectional regression coefficients, and is similar to the one made by Fama and French (1993). A cross-sectional regression coefficient for stock returns on some attribute is a linear combination of the returns, with weights that sum to zero. The portfolio has a positive (say, unit) value of the attribute and thus is long in high-attribute stocks and short in low-attribute stocks. If a multiple regression is used, it has zero exposure to the other regressors. Subject to these conditions, it has minimum variance. A spread portfolio is also long in the high-attribute stocks and short in the low-attribute stocks, with weights that sum to zero. If multiple independent sorts are used, as in Fama and French (1996), it is orthogonal to the other attributes. While a spread portfolio does not explicitly minimize variance subject to these conditions, it avoids estimation error. Thus, it should not be surprising to find that the empirical results for spread portfolios are similar to those for FM coefficients.

2. Data generation

The data include 377 monthly observations from August of 1963 through December of 1994. For each month t each NYSE and AMEX firm in the CRSP database is allocated to one of 100 equally weighted portfolios based on the first two letters of the firm's name, with numbers or special symbols (e.g., '&')

preceding letters and special symbols preceding numbers. When more than an even multiple of 100 firms are reported for a month, the extra firms at the end of the alphabet are removed from the sample that month. The 100 monthly portfolio returns are measured net of the return of a one-month Treasury bill, from Ibbotson Associates data.

We generate excess returns, r_t, for each month t by the following formula:

$$r_t = \mu 1 + \delta_0 + \delta'_1 z_{t-1} + \varepsilon_t, \quad t = 1, \ldots, T, \tag{1}$$

where μ is a scalar, 1 is an N-vector of ones, δ_0 is an N-vector, δ_1 is a $L \times N$ matrix, z_{t-1} is an L-vector of lagged economy-wide instruments, with $E(z_{t-1}) = 0$, and ε_t is an N-vector of identically distributed disturbances with covariance matrix, $V(\varepsilon)$.

The scalar μ is chosen to equal the grand mean of the monthly excess returns of the $N = 100$ portfolios. The N-vector δ_0 is centered around zero with 99 equally spaced increments. The absolute value of the $\max(\delta_{0i})$ is equal to one-half of the difference in the average returns between the high return, high book-to-market portfolios (1.09%) and the low return, low book-to-market portfolios (0.31%) reported by Fama and French, 1996, Table 1, panel A). The unconditional means of z and ε are zero, so the vector δ_0 determines the average, or unconditional alphabet effect. Since Eq. (1) is meant to represent an anomaly, the attributes (δ_0, δ_1) which describe the cross-section of returns are completely unrelated to systematic risk. There is no statement about equilibrium. Section 4 discusses how the example is related to equilibrium models.

In most stories of 'mispricing', the deviations of prices from fundamental values are time-varying and transitory. The distribution of pricing errors are likely to be time-varying; thus, inducing predictable time-variation in the anomalous returns. Such a component is captured by the predetermined term $\delta'_1 z_{t-1}$ in Eq. (1). While $\delta'_1 z_{t-1}$ contributes to the unconditional variance of returns, it does not contribute to the conditional covariance matrix, and thus is not systematic risk as defined in asset pricing models. We include both a predetermined, time-varying component and a fixed component in the anomaly, to highlight the distinction between conditional and unconditional variance.[3] In our example, $L = 2$ and the predetermined instrumental variables, z_{t-1}, are lagged monthly observations of the 3-month Treasury bill yield and dividend yield on S&P 500. Both instruments are demeaned.

The 2×100 matrix δ_1 is formed as follows, to resemble the book-to-market effect. First, we obtain two regression slope coefficients from a time-series regression of r_t^{hml} on a constant and the lagged instrument vector, where r_t^{hml} is

[3] In the 'characteristics based' model of Daniel and Titman (1997) the non-risk attributes are designed to be unrelated to the unconditional covariance matrix of returns, and there is no explicit distinction between conditional and unconditional risk measures.

the difference, for month t, between the returns on high book-to-market and low book-to-market portfolios.[4] These slope coefficients are used to define the vector of elements of the δ_1 matrix, each with 99 equally spaced increments centered around zero, where the difference in the coefficients for the early-in-the-alphabet portfolio and the late-in-the-alphabet portfolio is equal to the coefficients from the r_t^{hml} regression. Then, the columns of δ_1 are rescaled, multiplying them component-wise by a vector with 100 elements, formed by taking a sequence of equally spaced numbers between 1 and 55 inclusive, each raised to the power 0.5. This is done, in combination with the structure of δ_0, to insure that our generated data will roughly mimic the asymmetry of book-to-market sorted portfolio returns (e.g. Fama and French, 1996, Table 1 panel A).

Finally, we generate the $T \times 100$ matrix of unexpected returns, ε, where T is the number of months in the sample ($T = 377$). To accomplish this, we first regress our actual excess return data on the instrument vector and obtain the residuals, e. We then regress e on the residuals from a time-series regression of S&P500 excess returns on the lagged instruments. The slope coefficients determine the vector, b. We model the conditional covariance of ϵ according to the 'single index' model, as $V(\varepsilon) = bb' + \sigma^2 I$, where $\sigma^2 = 0.0013$ is the median eigenvalue of the difference between the actual sample covariance matrix of e and bb'. Since the betas on the S&P 500 bear no relation to the position in the alphabet (see Table 1), choosing $V(\varepsilon) = bb' + \sigma^2 I$ guarantees that the conditional covariance matrix of the generated returns is unrelated to the alphabet effect. Choosing σ^2 as an eigenvalue minimizes the distance between $V(\varepsilon)$ and the sample covariance matrix of the residuals.[5] We transform the residuals e from the actual data into our generated unexpected returns, ε, by multiplying them by the inverse of the square root of their sample covariance matrix and then by the square root of $V(\varepsilon)$. The transformed residuals have sample covariance matrix, $V(\varepsilon)$.

2.1. Properties of the spurious factors

In our parable a researcher uses the artificial data to construct γ_{1t}, the FM regression coefficients, and a spread portfolio, which is long 'A' firms and short

[4] The book-to-market sorted portfolios are formed similar to Fama and French (1992) and are kindly provided by Campbell Harvey. Stocks are ranked on the total market value of equity in June of the year $t+1$ and sorted into three portfolios: Small (30%), Medium (40%) and Big (30%). Stocks are also ranked separately on NYSE BE/ME (Low(30%), Medium(40%), or High(30%)), where BE is book equity for June of year t and ME is market equity at the end of December of year t. Monthly portfolio returns are computed from July of year $t+1$ through June of $t+2$. All portfolio returns are value-weighted. HML is $(S/H + B/H)/2 - (S/L + B/L)/2$.

[5] In order to minimize $|\mathrm{Cov}(e) - V(\varepsilon)| = |\mathrm{Cov}(e) - bb' - \sigma^2 I|$, choosing σ^2 as any eigenvalue of $\mathrm{Cov}(e) - bb'$, gives $|Cov(e) - bb' - \sigma^2 I| = 0$. Since the solution is not unique and all the eigenvalues are positive, we use the median of the eigenvalues as σ^2.

Table 1
Descriptive statistics for 100 portfolios formed on name ranking from artificial data: August, 1963–December, 1994

Each of 100 portfolio returns is generated by the process $r_t = \mu 1 + \delta_0 + \delta_1' Z_{t-1} + \varepsilon_t$, where μ is the grand mean of actual excess returns, 1 is the size 100 vector of ones, δ_0 is the size 100 vector of equally spaced numbers centered around zero, δ_1 is a 2×100 matrix of coefficients, and Z_{t-1} is the vector of demeaned instruments which include lagged monthly observations of the 3-month Treasury bill yield and dividend yield on S&P 500. There are 377 months in the sample. Mean monthly returns (Mean) and standard deviation of monthly returns (Std. Dev.) are shown for each portfolio. The coefficients from the regression of monthly returns for a portfolio on the monthly excess return of the equally weighted market portfolio are shown as the Market Beta. The final four rows present average slope coefficients, per decile or quintile, resulting from the time-series regressions of monthly returns on the alpha risk factor (the monthly time-series Fama–MacBeth coefficients corresponding to the alphabet effect) or the spread portfolio returns A minus Z (AMZ) along with the excess equally weighted market return. Firms with different names enter a given portfolio in different months. The headers indicate the maximum range of letters within a portfolio.

	Alphabet quintiles						
	1a	1b	2	3	4	5a	5b
Name range	AA-BE	AT-CO	CH-HA	GE-NA	MA-SI	RE-TR	ST-ZW
Mean	0.0108	0.0104	0.0098	0.0090	0.0082	0.0076	0.0072
Std. Dev.	0.0555	0.0574	0.0576	0.0574	0.0551	0.0578	0.0593
Market Beta	0.9497	1.0214	1.0212	1.0262	0.9463	1.0005	1.0626
Alpha Beta	0.4235	0.4016	0.2856	0.0626	− 0.2214	− 0.4546	− 0.6239
Market Beta	0.9634	1.0226	1.0234	1.0167	0.9413	1.0099	1.0413
AMZ Beta	0.4656	0.5109	0.4713	0.0808	− 0.3886	− 0.5793	− 0.7243
Market Beta	0.9500	1.0093	1.0130	1.0146	0.9495	1.0250	1.0614

'Z' firms. These portfolios are spurious factors, constructed to have zero conditional market betas and thus, no systematic risk. Since they have positive expected returns they would represent an arbitrage opportunity in a large, perfect market. Given a finite sample of stocks, there is an estimation error in the cross-sectional regression coefficients, which implies volatility in the portfolios. The appendix provides a general analysis of the variance of FM coefficients.

In the artificial sample, the correlation of the constructed γ_{1t} with the 'market' (S&P 500) is 0.07, while that of the spread portfolio, AMZ, is 0.16. The unconditional standard deviations are $\sigma(\gamma_{1t}) = 0.015$ and $\sigma(\text{AMZ}) = 0.0113$, which may be understood as the result of two effects. The first is the sampling variation, which approaches zero as N becomes large. From Eq. (1), the variance due to sampling variation is approximately equal to $w'(\sigma^2 I)w$, where w is the portfolio weight (since $w'b \approx 0$). With $N = 100$, the standard deviation due to sampling error is approximately $\sigma/10 = 0.0036$, where $\sigma^2 = 0.0013$ is the residual variance of the generated returns, as described above.

The second source of variance in the spurious factors is the common time-variation in the anomaly, contributing to the unconditional variance an amount equal to $w'\mathrm{Var}(\delta'_1 z_{t-1})w$. As this part contributes nothing to the conditional variance, it is not risk in the conventional sense. However, if unconditional betas on the spurious factors are used, this common variation will be picked up. Fama and French (1996) use unconditional betas.

In our example the two components of unconditional volatility in the spurious risk factors contribute similar amounts. A time-series regression of γ_{1t} on the lagged instruments delivers the predictable part of the variance, due to the common time-variation in the anomaly. The regression produces $R^2 = 0.52$. Based on the equal-weight approximation, the R-square should be about $1 - (0.0036/0.015)^2 = 0.51$, which shows that these approximations provide an adequate description.

In forming the spread portfolio, AMZ, we take an equally weighted average of the alphabet portfolios numbered 1 through 40, and subtract from it the return of an equally weighted average of the portfolios numbered 61 through 100. The portfolio weights of γ_{1t}, estimated from the cross-sectional regression, follow a more complicated pattern. As noted by Jagannathan et al. (1997), there is not an exact linear relation between returns and the betas on these portfolios. Therefore, with an infinite sample of stocks or a sufficiently powerful test it would be possible to reject the 'alpha factor asset pricing model'. Our point, however, is to illustrate what can occur in a realistic finite sample.

Finally, our data generating process incorporates a grand mean, μ, in the returns. Since the grand mean is not zero, it will not be possible to fit the returns with a factor model based on the spread portfolios alone. This is because the average loading on the spread portfolio is zero, while the average return is not zero. Including a market factor allows the model to capture the grand mean, since the average market beta is 1.0. Thus, a market factor is useful even if the cross-section of the market betas has no explanatory power.

3. Let the story begin

Once upon a time, not too long ago, a financial economist was sorting stocks on the tapes from Chicago according to the attributes of the firms whose equities were represented there. This sort of sorting had proved profitable in his field of endeavor, and indeed had lead to famous effects on equity returns such as size, yield and price. One day our relentless investigator focused on the letters of the alphabet, and caused to discover an amazing pattern in the returns of stocks so sorted. This result, as depicted in Table 1, he named the 'alphabet anomaly'.

The figures in Table 1 certainly suggested an anomalous situation, by the following logic. If returns are seen to be related to the position of the firm in the alphabet, but said position bears no relation to either the total risk (standard deviation) or to the market risk (beta on the Standard and Poors index), then it would seem unlikely that some intermediate definition of 'factor' risk should be able to explain this puzzling pattern. As the table so clearly illustrates, our scientist discovered an apparent departure from the laws of asset pricing, which require that expected returns should be related to risk.

Now some observers would likely suspect our daring detective of having merely dredged these patterns from the data. But, in fact, it was the logic of economics that stimulated the inquiry; for it is widely felt that firms with early-in-the alphabet appellations derive various advantages, such as a predominant position in lists employed when searching for a vendor, as are likely to be prepared by potential customers of their goods and services. Indeed, equity market lore had long embraced the value that can be attached to a good corporate name, and a recent spate of mergers had generated new entities, their titles strategically coined to occupy a dominant, early-in-the-alphabet position.

Lest the simple patterns in the table prove deceptive, our empiricist turned to a more rigorous analysis. For each of the 377 months in the sample, a cross-section of the 100 alphabet-sorted portfolio returns is duly regressed on the position in the alphabet, the latter instrument taking on the values from 0.6 to − 0.59 in equal increments. The time-series means of the cross-sectional regression coefficients accordingly, are displayed under regression Eq. (2), and the Fama-MacBeth t-ratio in parenthesis.[6]

$$\begin{array}{l} r_{it} = \underset{0.009}{\gamma_{0t}} + \underset{\underset{(4.24)}{0.0033}}{\gamma_{1t}}A_i + \varepsilon_{it}, \quad i = 1, \dots, 100. \end{array} \tag{2}$$

The average value of the premium associated with early-in-the-alphabet firms, as compared with late in the alphabet firms (the alphabet premium), was calculated to be about one third of a percentage point per month, a magnitude of certain economic significance to investors. Its t-ratio was of such a magnitude as to rule extremely improbable, the event that the result had occurred simply by chance. The striking statistical and economic significance of the alphabet effect so confirmed by the regression, our analyst was fervent indeed. Subsequent regressions (not reported here, but certainly, available by request) established that neither the market beta nor the total risk, as measured by the standard

[6] The Fama–MacBeth t-ratios are constructed as the time–series mean of the cross-sectional regression coefficients, divided by the standard error of the mean.

deviation of the return, bore a marginal relation of any significance to returns, in the presence of the alphabet effect.

If our story ended here, several predictable conclusions would likely emerge. An obligatory call would echo throughout the legions of financial economists, for future research would sorely be needed to provide an explanation for the alphabet effect. Some would take up the search for the underlying, rational risk factors. Some would see the effect, given their previous convictions and natural inclinations, as further convincing evidence of irrationality and market inefficiency. Still others, among these scribes sympathetic to each of the above mentioned views, would wish to exploit the alphabet anomaly by forming investment funds and money management services. It would be argued, whether mispricing or systematic risk underlay the effect, that wise investors should, in either event, wish to structure their portfolios in regards to it. Fortunately, however, there is more to this tale.

3.1. The alpha-risk factor

Our researcher was unable to so lightly dismiss the possibility that the alphabet effect merely proxied for some unidentified risk factor. Assuming such a factor to exist, he reasoned, the trick must be to obtain an empirical proxy for the factor. Recalling a classic interpretation of regressions like (2) (Fama, 1976), the slope coefficients γ_{1t} are zero net investment portfolio excess returns, constructed for maximum correlation with the unobserved risk that must be presumed to underlie the alphabet effect. It became evident to our financial econometrician that the time series of the coefficients γ_{1t} can themselves represent the excess returns on factor-mimicking portfolios. So encouraged, he considered the employment of the coefficients to this service, coining the appellation *alpha risk factor* to describe their newly discovered role.

Whilst cross-sectional coefficients may properly be called portfolios, it is in the details of their construction that some would find cause for consternation. For in following the investment strategy so implied, one would be forced to extreme positions, both long and short, and thereby incur in the actual practice, such high costs of transacting that the classical asset pricing paradigm did not contemplate. A simpler approach, in the spirit of Fama and MacBeth (1973), would simply subtract from the return on an early-in-the-alphabet portfolio, the return on a similarly constructed late-in-the-alphabet portfolio. Our researcher decided to call this version of the alpha risk factor A minus Z (AMZ). Henceforth, subsequent tests would be conducted using, alternatively, the variable γ_{1t} or AMZ, for in comparing the results of the two, some measure of the robustness of the findings is provided.

A careful analysis was clearly in order, of the extent to which the alpha risk factor could, indeed, explain expected returns. Accordingly, as a first step in the analysis, our empiricist set forth time-series regressions to estimate risk exposures:

$$r_{it} = \alpha_i + \beta_{im} r_{mt} + \beta_{i\alpha} \gamma_{1t} + u_{it}, \quad t = 1, \ldots, T. \tag{3}$$

In this equation r_{mt} is the excess return of the 'market'and γ_{1t} represents the alpha risk factor premiums or, in their stead, the spread portfolio, AMZ. The coefficients to be estimated include β_{im}, which is the market beta for portfolio i, and $\beta_{i\alpha}$, the slope coefficient on the alpha risk factor, now more succinctly referred to as the *alpha beta*. An intercept, α_i, would be included to absorb any such failure of the means of the variables to coincide as may occur.

The alpha betas estimated by Eq. (3) were found to be strongly significant, with large t-ratios. Curiously, their magnitudes were arrayed in close accord with those very same average returns that initially exposed the alphabet anomaly, as is so clearly in evidence in the summary figures reported in Table 1. The results using AMZ to measure the alpha risk factor were similar. Thus, the regressions seemed to strongly suggest a risk-based explanation for the alphabet anomaly.

At this point our scholar made another interesting observation as to the interpretation of the regressions of Eq. (3). As excess returns over a Treasury bill are used as the returns measures, r_{it}, then according to the paradigm of beta pricing, it follows that the intercepts, α_i, should be zero if the market and alpha betas provide a complete and accurate description of the expected returns. Across the 100 portfolios, the intercepts ranged from -0.0019 at the low extreme, to $+0.0028$ at the high extreme, with the maximum absolute t-ratio being only 1.45. Using AMZ, the results were again, similar. It seemed that the time series regressions provided most encouraging support for a risk-based model of expected returns, using the market and the alpha risk factors in a two-factor model. So encouraged, our investigator deemed this new model should be known as the *alpha factor asset pricing model*.

3.2. The acid test

A more aggressive and convincing attack on the central issue seemed desireable. Could a risk-based explanation, using the market and alpha betas, fully explain the alphabet anomaly? To this crucial question our investigator next turned. The natural test, it seemed, was whether the betas could capture the cross-sectional information in returns represented by the position of the firm in the alphabet.

Our empiricist then conducted another cross-sectional regression analysis, using the alphabet-sorted portfolios, and similar to regression (2), the Fama–MacBeth method, excepting, in this case, also including the alpha betas as a cross-sectional predictor:

$$r_{it} = \gamma_{0t} + \gamma_{mt} \beta_{im} + \gamma_{1t} \beta_{i\alpha} + \gamma_{2t} A_i + v_{it}, \quad i = 1, \ldots, 100. \tag{4}$$

Table 2
Cross-sectional regressions for the acid test

Cross-sectional regressions are estimated in two passes. In the first pass, portfolio returns are regressed on a constant, market beta, and alpha or AMZ betas. Then in the second pass, the residuals plus the intercept, taken from the first pass, are regressed on a constant and the position in the alphabet, A_i. The coefficients on the market beta and alpha or AMZ betas come from the first pass regression; the coefficients for A_i come from the second pass regression. The fit of the overall procedure is summarized by the maximum R-square (maxR^2) and minimum R-square (minR^2). For each month, the R-square is one minus the ratio of the cross-sectional variance of residuals from the second pass regression over the cross-sectional variance of the returns used in the first-pass regression.

$$r_{it} = \gamma_{0t} + \gamma_{1t}\beta_i^M + \gamma_{2t}\hat{\beta}_i^A + \gamma_{3t}A_i + \varepsilon_{it}, \quad i = 1, \ldots, 100$$

$$r_{it} = \gamma_{0t} + \gamma_{1t}\beta_i^M + \gamma_{2t}\hat{\beta}_i^{\mathrm{AMZ}} + \gamma_{3t}A_i + \xi_{it}, \quad i = 1, \ldots, 100$$

Intercept	β_i^M	$\hat{\beta}_i^A$	$\hat{\beta}_i^{\mathrm{AMZ}}$	A_i	minR^2	maxR^2
− 0.0000	− 0.0003	0.0032		0.0001	0.0	0.36
(− 1.0651)	(− 0.0788)	(4.0446)		(1.0651)		
− 0.0000	− 0.0012		0.0024	0.0002	0.0	0.33
(− 1.5820)	(− 0.3635)		(3.9943)	(1.5820)		

The evidence of regressions (4), as summarized in Table 2, appeared striking indeed, so effectively did the alpha betas subsume the alphabet effect; for, in their presence, the position in the alphabet was rendered insignificant, both economically, by virtue of its small coefficient, and statistically, by virtue of its small t-ratio.[7] The average values of the intercept and of the premium associated with alpha risk were virtually identical to those obtained in the initial investigation of the alphabet effect. Using AMZ to measure the alpha risk factor, similar findings were again obtained, as may be seen in Table 2. The evidence in favor of the alpha factor asset pricing model was impressive, indeed.

The results must be seen as especially supportive of the model, taken in view of the errors-in-variables problem, for which it was a commonly held view that errors in measuring regression betas should, by all accounts, be larger than errors in measuring the firms' position in the alphabet. Such errors would tend to bias regression (4) against the alpha betas.

[7] Indeed, the alpha betas work so well that the $x'x$ matrix of the multiple regression is ill-conditioned when both A_i and $\beta_{i\alpha}$ are present, their sample correlation being 0.984. Accordingly, Eq. (3) is estimated in a step-wise fashion. First, the returns are regressed on the alpha betas and the market betas, then the intercept + residuals from that regression are run on the A_i, to examine the null hypothesis that the A_i have no marginal explanatory power, given the alpha betas and market betas.

3.3. Implications

The potential importance of our scholar's findings to the everyday affairs of commerce seemed enormous, for since the early years of Sharpe (1964) the equity market beta had served as a theoretical cornerstone of many financial endeavors. But as of late the beta had fallen into disrepute, felled largely by the aforementioned effects of size, yield and price. Some said that beta lay languid, and if so, then by its imminent death had left many such applications in the lurch, for no easy successor lay waiting in the wings.

Would the alpha factor asset pricing model provide the foundation on which investors could build their portfolios in the future, confident in its ability to summarize the risks for which the market would bestow its expected rewards? Could firms now determine their costs of capital using the alpha betas, in such a way as to direct their investment resources toward the highest valued use as signaled by the equity markets? And what of some deeper theory, leading to a more fundamental understanding of *why* the higher risks and rewards should fall to those firms whose names placed them in the early portions of telephone books and, as had become more relevant in recent times, in many searchable databases? The answers to these questions are another story, as yet untold.

4. The moral of the story

Our parable is a numerical example in which artificial data are constructed to resemble actual monthly stock returns, and the alphabet effect resembles the book-to-market effect. However, we use a covariance matrix in our generating process that guarantees that the alphabet attribute is unrelated to systematic risk. We show that when a risk factor is constructed in a fashion similar to that of Fama and French (1993, 1996) it will appear to explain the cross-section of average returns, when combined with market index betas in a 'two-factor' model. Given this caveat, the crucial question is: How relevant is the example to the actual practice of asset pricing research? We argue that it is potentially highly relevant.

The likelihood that a result like our parable occurs in the real world can be considered as the sum of two probabilities. The first is the probability that the result occurs, given the existence of a large-market arbitrage opportunity, multiplied by the probability that such an anomaly exists. The second is the probability that the result occurs given no arbitrage, multiplied by the probability that there is no arbitrage. We do not presume to influence the reader's prior on the existence of arbitrage opportunities, so we take each case in turn.

First, if the data generating process in our example were the 'true' process, there would be a large-market arbitrage opportunity based on the position in the alphabet. Our example shows that an arbitrary attribute, bearing such an anomalous relation to returns, can be repackaged as a spurious risk factor. However, recent studies do not use arbitrary attributes. Some of the most empirically powerful characteristics for the cross-sectional prediction of returns are ratios, with market price per share in the denominator. Berk (1995) emphasizes that the price of any stock is the value of its future cash flows discounted by future returns, so an anomalous pattern in the cross-section of returns would produce a corresponding pattern in book-to-market ratios or other proxies of cash-flow-to-price. A cross-sectional regression of returns on these ratios will pick out the anomalous patterns. Thus, the use of valuation ratios such as book-to-market as a sorting criterion increases the risk of creating a spurious risk factor when there is an arbitrage opportunity in returns.

Second, even with no arbitrage the results we illustrate here can occur in the presence of data mining. If the ‘anomaly’ is discovered through data mining, then we interpret our generating process (1) as representing the in-sample returns. The structure of the δ coefficients across the attribute-sorted portfolios is an artifact of the particular sample. Future returns, outside of the sample period, would not follow the same process. Our result should obtain if the analyst uses the same data to form spurious factors and test the model, as he uses to find the original pattern in returns. For example, in our parable monthly stock returns for the U.S. from 1964 to 1994 are used both to discover a pattern, and to generate the spurious factors, which appear to work in the same sample. In this case, the factors should not work with fresh data.[8]

Fama and French sort stocks based on attributes with an observed empirical relation to the cross-section of returns. On the one hand, if there is no arbitrage or data mining, so that the relation reflects expected return differences driven by exposure to risk factors, their spread portfolio approach should deliver proxies for the unobserved factors. On the other hand, we generate data with no relation between systematic risk and the attributes. We show that when there is an anomaly, their spread portfolio approach is likely to transform the anomaly into spurious factors. In the real world, attributes may be correlated with systematic risk and also with anomalous patterns in return. The net result of the two effects, risk versus anomaly, is complicated and model specific.

[8] While the evidence of Fama and French (1996) may be subject to this criticism, Fama and French (1998) document book-to-market effects in international equities, which reduces the risk of a data-mined anomaly. Note that if an anomaly was ‘discovered’ on a given data sample, then when splitting that sample in two there is no reason to think that the anomaly would be confined to a subsample or be independent across the subsamples.

5. Conclusions

Recent empirical studies use the returns of attribute-sorted portfolios of common stocks to represent risk factors in an asset pricing model, where the attributes are chosen following an empirically observed relation to the cross-section of stock returns. This story has illustrated, using artificial data, that such attribute-sorted portfolios will appear to be useful risk factors, even when the attributes are *completely* unrelated to risk exposure. We hope that our parable serves to raise the level of skepticism in the profession about such approaches.

Our analysis focuses on one of a number of issues that may arise when portfolios are formed by sorting on the basis of securities' attributes. Equity market databases are inherently unbalanced panels, with more stocks than quarters or months. As new data on equity attributes becomes widely accessible, we expect to see more studies that sort equities according to their attributes. We have seen that sorting procedures are subtle and easily abused. More work is needed to improve our understanding of the properties of such approaches.

Acknowledgements

The seeds of this paper were sown some years ago by Robert Litzenberger. The paper has also benefited from comments by Jonathan Berk, Nai-fu Chen, Eugene Fama, Kenneth French, Ravi Jagannathan, Raymond Kan, Jaemin Kim, Owen Lamont, Bruce Lehmann, Ed Rice, Avanidhar Subrahmanyam, Sheridan Titman and Guofu Zhou. It was presented at the 1998 CIBER-UCLA Doctoral Consortium in Finance, at the Spring, 1998 Conference of the Institute for Quantitative Research in Finance and at the 1998 Western Finance Association Meetings. Ferson acknowledges support from the Pigott-Paccar professorship at the University of Washington and the National Bureau of Economic Research.

Appendix A. Supporting analysis[9]

Consider a more general *θ-anomaly*, as represented in Eq. (A.1):

$$E_t(r_{t+1}) = \mu + \theta C_t, \tag{A.1}$$

where r_{t+1} is the n-vector of asset returns at time $t+1$, measured in excess of a reference asset, such as a Treasury bill. θ is an $N \times K$ matrix of K attributes of the N assets, normalized to have a cross-sectional mean of zero. For simplicity

[9] This part of the appendix is drawn from Ferson (1996).

we assume that θ is fixed.[10] To accommodate a time-varying expected return anomaly, the $K \times 1$ coefficient vector C_t is observed at time t and is allowed to vary over time.

The unexpected returns at date $t+1$ are given by the vector ε_{t+1}, defined by Eq. (A.2);

$$r_{t+1} = E_t(r_{t+1}) + \varepsilon_{t+1}. \tag{A.2}$$

Assume that the covariance matrix of the unexpected returns, $V(\varepsilon)$, may be written according to a general factor structure as in Eq. (A.3):

$$V(\varepsilon) = bb' + D, \tag{A.3}$$

where b is a matrix of factor loadings and D is an $N \times N$ matrix of the residual risks. Well-known asset pricing theories, such as the Arbitrage Pricing Model (APT, Ross, 1976) use Eq. (A.3) to represent asset risks. The idea is that the common, or systematic risk of the unexpected returns is captured by the loadings, b, while the idiosyncratic and diversifiable risks are captured by the residual covariance matrix, D.[11] In a risk-based asset pricing model, the cross-section of the expected returns are determined by the systematic risks, b. In contrast, our anomalous expected returns are constructed to be unrelated to the systematic risks. The following definition is used:

Definition: the attributes θ are said to be unrelated to risk iff $\theta' b = 0$.

This says that the cross-section of the attributes is orthogonal to the matrix of the systematic risks, or factor loadings (θ has cross-sectional mean zero). We can account for specific economic risk factors, without loss of generality, by simply replacing the term bb' by the term $bV(f)b'$, where $V(f)$ is the covariance matrix of the systematic risk factors.[12] A special case is the *standard factor analysis model*, used in the parable. Here, the residual covariance matrix of the asset returns, $D = \sigma^2 I$, where σ^2 is a scalar and I is the identity matrix.

[10] A fixed attribute matrix involves only a small loss of generality when the returns r_{t+1} are measured for portfolios sorted according to quantiles of the attributes for individual stocks, as is the common practice in empirical studies. This is because using the quantile measure in a cross-sectional regression effectively abstracts from time-series variation in the distribution of the attributes. Cross-sectional regression coefficients of returns on the attributes are unaffected by subtracting any cross-sectional constant from the measure of the attribute, even if that constant varies over time.

[11] For example, if D has bounded eigenvalues as n grows large, then the residual risk is negligible in a large portfolio. See Chamberlain and Rothschild (1983).

[12] This case would be relevant in the context of a model like Merton's (1973) intertemporal asset pricing model. In our example, $V(f)$ is the variance of the Standard and Poors 500 excess return and $K = 1$.

Consider a cross-sectional regression of the returns r_{t+1} on the attributes, θ. For simplicity, assume that an OLS regression is used. The coefficient estimator is given by:

$$\hat{C}_t = (\theta\theta)^{-1}\theta' r_{t+1}, \tag{A.4}$$

and the estimation error is:

$$\hat{C}_t - C_t = (\theta\theta)^{-1}\theta'\varepsilon_{t+1}. \tag{A.5}$$

The estimator $\hat{C}_t$ is like the alpha risk factor in our parable. Eq. (A.4) shows that $\hat{C}_t$ is a portfolio of the returns, r_{t+1}. Eq. (A.5) implies that the unconditional time-series variance of this portfolio is $\mathrm{Var}(C_t) + \mathrm{Var}[(\theta'\theta)^{-1}\theta'\varepsilon_{t+1}]$. The first component is due to the common time-variation in the anomaly, and the second component is due to sampling variation. Note that as N becomes infinite, the second component converges to zero when the cross-sectional regression coefficients converge in mean square to the true parameters. We discuss this decomposition of variance in the context of our numerical example, in Section 2.1.

Eq. (A.5) implies that FM regression coefficients are large-market arbitrage portfolios. As N becomes infinite, the estimates converge, so the portfolios have zero residual risk in a large market. According to the Arbitrage Pricing Theory (APT), such portfolios should have vanishingly small returns in a large market. Thus, a lack of arbitrage implies $C_t = 0$. In the presence of an arbitrage opportunity C_t is not zero, and the arbitrage portfolios have nonzero expected returns. The portfolios thus capture the anomaly that reveals the arbitrage.[13] An arbitrage opportunity correlated with an observable attribute is likely to present an opportunity to sort portfolios, similar to the ones in our parable. However, using 100 portfolios, our example suggests that it will be difficult to discern the arbitrage opportunity in the face of the sampling error.

Consider a time-series regression of the excess returns, r_{t+1}, on the series of the $\hat{C}_t$, and let the $K \times N$ matrix β denote the regression slopes for a fixed value of N. These are the 'loadings' on the attribute portfolios, when the portfolios are treated as if they were asset-pricing factors (in the parable, the β are the alpha betas or AMZ betas). Under standard assumptions, the time series slopes converge in probability as T becomes infinite:

$$\beta \to \mathrm{Cov}(\hat{C}_t)^{-1}\mathrm{Cov}(\hat{C}_t, r_{t+1}). \tag{A.6}$$

[13] Huberman's (1982) proof of the APT makes use of similar portfolios, and shows that, when there is an arbitrage opportunity, the ratio of expected return to risk is unbounded as the number of assets grows. MacKinlay (1995) similarly defines 'optimal orthogonal portfolios', and uses them to study the power of asset pricing tests to detect an anomaly. His analysis suggests that a risk-based explanation of the reward-to-risk ratios of book-to-market sorted portfolios is unlikely.

Substituting from Eqs. (A.1), (A.2) and (A.4) it is easy to show that this expression is equivalent to:

$$\beta \to [(\theta'\theta)^{-1}\theta' V(\varepsilon)\theta(\theta'\theta)^{-1}]^{-1}[\mathrm{Var}(C_t)\theta' + (\theta'\theta)^{-1}\theta' V(\varepsilon)]. \tag{A.7}$$

Eq. (A.7) determines when the cross-section of loadings on the attribute-sorted portfolios will have explanatory power for asset returns. Consider a cross-sectional regression of the returns r_{t+1} on the loadings, β. The term $[(\theta'\theta)^{-1}\theta' V(\varepsilon)\theta(\theta'\theta)^{-1}]^{-1}$ is a $K \times K$ constant in a cross-section of N returns, so the β are proportional to $[\mathrm{Var}(C_t)\theta' + (\theta'\theta)^{-1}\theta' V(\varepsilon)]$. Thus, there are two terms that determine the cross-sectional explanatory power of the loadings. The first term, $\mathrm{Var}(C_t)\theta'$, is proportional to θ'. Therefore, by Eq. (A.1), the θ-anomaly says that if the first term was the only term, and provided $\mathrm{Var}(C_t) > 0$, the cross-sectional relation between the loadings and expected returns would be an exact one. This term relies on the common time-variation in unconditional returns to generate explanatory power for the loadings. However, the second term can also give rise to a cross-sectional relation when $\mathrm{Var}(C_t) = 0$.

If returns follow the standard factor analysis model, then using the fact that $\theta' b = 0$ when the attributes are unrelated to risk, the second term of Eq. (A.7) becomes:

$$(\theta'\theta)^{-1}\theta' V(\varepsilon) = (\theta'\theta)^{-1}\theta'[bb' + \sigma^2 I] = \sigma^2(\theta'\theta)^{-1}\theta',$$

which is again proportional to θ'. The cross-sectional relation of expected returns to the loadings is in this case an exact one, even when $\mathrm{Var}(C_t) = 0$.

Appendix B. Sensitivity analysis

In this appendix we address the sensitivity of the results to variations in the way we generate the data. There are three steps taken by the analyst in the parable. First, a set of regressions attempt to explain returns with the position of a firm's name in the alphabet. While particular slope coefficients and t-statistics in these regressions are affected by a specific calibration scheme, qualitatively the results are not sensitive to a wide range of variation in the definitions of δ_0 and δ_1, the key parameters in the artificial data generating process.

The second step is the time-series regressions of returns on the market index and the alpha factor or the spread portfolio, AMZ. The patterns observed in the AMZ betas and alpha betas, across the 100 portfolios, are sensitive to variations in δ_0 and δ_1. This allows us to pick these coefficients to roughly match the patterns in book-to-market portfolios. The second step regressions are also sensitive to the inclusion of the grand mean return, μ, in the data generating process, and to the presence of the market beta. Without these features, the intercepts become statistically significant.

The third step uses the betas in cross-sectional multiple regressions to explain returns. These results are not very sensitive to variations in the δ_0 and δ_1 parameters when the alpha-betas are used. There is some sensitivity when the AMZ betas are used. This may be understood by noting that, unlike the FM coefficients, the spread portfolio AMZ bears no relation to the 20% of the portfolios that are not used in its construction, except through the time-varying part of the anomaly. Thus, it is more sensitive to the specification of the parameter, δ_1.

The final aspect of our data generating process is the specification of the unexpected returns. We find that achieving independence between the alphabet anomaly and the systematic risk part of the conditional variance matrix is crucial. When we generate unexpected returns using a covariance matrix with dependence, the alpha-factor or AMZ betas are still 'priced', but we no longer observe that they can completely capture the alphabet anomaly. This may be understood from the right-most term of Eq. (A.7). In our example, we generate data with no relation between expected returns and systematic risk. Dependence between systematic risk and θ distorts the cross-sectional relation away from that of the underlying anomaly.

References

Berk, J.B., 1997. Sorting out sorts. Working paper, University of Washington.

Berk, J.B., 1995. A critique of size-related anomalies. Review of Financial Studies 8, 275–286.

Brennan, M., Chordia, T., Subrahmanyam, A., 1998. Alternative factor specifications, security characteristics and the cross-section of expected stock returns. Journal of Financial Economics 49(3), 345–373.

Chamberlain, G., Rothschild, M., 1983. Arbitrage, factor structure and mean variance analysis on large asset markets. Econometrica 51, 1281–1304.

Cheung, Yin-Wong, He, J., Ng, L., 1995. What are the global sources of rational variation in international equity returns? Working paper, City University of Hong Kong.

Daniel, K., Titman, S., 1997. Evidence on the characteristics of cross-sectional variation in stock returns. Journal of Finance 52, 1–33.

Davis, J., Fama, E., French, K., 1998. Characteristics, covariances and average returns: 1929–1997. Working paper, University of Chicago.

Fama, E., 1976. Foundations of Finance, New York, Basic Books.

Feema, E., 1997. Determining the number of priced state variables in the ICAPM. Working paper, University of Chicago.

Fama, E., French, K., 1992. The cross-section of expected returns. Journal of Finance 47, 427–465.

Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Fama, E., French, K., 1995. Size and book-to-market factors in earnings and returns. Journal of Finance 50, 131–155.

Fama, E., French, K., 1996. Multifactor explanations of asset pricing anomalies. Journal of Finance 51, 55–87.

Fama, E., French, K., 1998. Value versus growth: The international evidence. Journal of Finance, forthcoming.

Fama, E., MacBeth, J., 1973. Risk, return and equilibrium: Empirical tests. Journal of Political Economy 81, 607–636.
Ferson, W., 1996. Warning: Attribute-sorted portfolios can be hazardous to your research. In: Saitou, S., Sawaki, K., Kubota, K. (Eds.), Modern Portfolio Theory and its Applications. Center for Academic Societies, Osaka Japan, pp. 21–32.
Ferson, W., Harvey, C., 1998. The fundamental determinants of international equity returns: A perspective on international asset pricing. Journal of Banking and Finance 21, 1625–1665.
Huberman, G., 1982. A simple approach to arbitrage pricing. Journal of Economic Theory 28, 183–191.
Jagannathan, R., Kubota, K., Takehara, H., 1997. The relation between labor income, risk and average return: Empirical evidence from the japanese stock market. Journal of Business 71(3), 319–347.
Kubota, K., Takehara, H., 1995. Common risk factors of Tokyo stock exchange firms: In finding mimicking portfolios. Discussion paper no. 20, Faculty of Economics, Musashi University.
Lakonishok, J., Shliefer, A., Vishny, R., 1994. Contrarian investment, extrapolation and risk. Journal of Finance 49, 1541–1578.
Lo, A., MacKinlay, A.C., 1990. Data snooping biases in tests of financial asset pricing models. Review of Financial Studies 3, 431–468.
MacKinlay, A.C., 1995. Multifactor models do not explain deviations from the CAPM. Journal of Financial Economics 38, 3–28.
Merton, R., 1973. An intertemporal capital asset pricing model. Econometrica 41, 867–887.
Ross, S., 1976. The arbitrage theory of capital market asset pricing. Journal of Economic Theory 13, 341–360.
Sharpe, W., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance 19, 425–442.
Swift, J., 1704. A Tale of a Tub, in: Ross, A., Woolley, D., eds., The Oxford Authors: Jonathan Swift, 1984 (Oxford University Press, New York).