Evolution of firm size

Lukas Gonon*and L. C. G. Rogers†

April 15, 2014

Abstract

In this paper we develop the idea that firm sizes evolve as log Brownian motions $dS_t = S_t(\sigma dW_t + \mu dt)$ where the constants μ, σ are characteristics of the firm, chosen from some distribution, and that the firms are wound up at some random time. At any given time, we see a firm of a given size. What can we say about its characteristics given its size? How would we invest in such a market? What do these assumptions imply about the distribution of sizes? By making simple and well-chosen modelling assumptions, we are able to develop quite concrete forms of the dependence of firm characteristics on size, from which we are able to deduce optimal investment weights as a function of size alone. As in the approach of Fernholz [Fernholz, 2002], this avoids the need to estimate growth rates of stocks in order to decide on investment strategy.

1 Introduction.

Problems of dynamic optimal investment have been studied for many years now, going back at least to the landmark paper of Merton [Merton, 1971]; but the theory does not always tie up well with practice. One of the main conclusions of [Merton, 1971] is that for a constant relative risk aversion agent trying to optimize his utility of terminal wealth in a log-Brownian market, the optimal strategy is to invest his wealth in fixed proportions in the available assets. This is memorable and simple, but applying it is problematic. It is problematic because the optimal proportions depend on the growth rates of the assets and their volatilities, and these will not be known. The growth rates in particular cannot even be estimated with any degree of confidence (see, for example, Section 4.2 of [Rogers, 2013].) Nevertheless, the mean-variance type of strategy which Merton's result implies remains a commonly-used approach to portfolio management, with point estimates substituted for the unknown true values, and little account taken of the estimation error.

*Department of Mathematics, ETH Zürich.
†Statistical Laboratory, University of Cambridge.

In contrast, the stochastic portfolio approach of Fernholz [Fernholz, 2002] aims to base portfolio choice on market observables, particularly, the relative sizes of firms in the whole market. While this approach does not entirely escape from the need to know something about the parameters of the asset dynamics, it advances a very different way of thinking of the problem, and raises a number of interesting questions: Why should the size of a firm be an indicator for its performance? How could we explain how size and performance may be related? How would we exploit any information that may be found in the firm size? It is generally believed that the returns of small firms are superior to those of large firms; one possible explanation is as a liquidity premium [Fernholz and Karatzas, 2006], but there may be other features at work here, such as survivorship bias.

What we propose in this paper is a very simple mechanism to explain why size could be informative about the characteristics of a firm. We shall suppose that independent firms of unit size are created at the times of a Poisson process, and at the time of their creation they receive a randomly-chosen drift and volatility. They then evolve until a random time at which they disappear from the market[1]. In such a model, the size of a live firm carries information about the random growth and volatility parameters, and if the distributions are set up suitably, it turns out to be possible to extract this information in relatively simple closed forms. This is the business of Section 2, where we show how to derive the conditional distribution of the growth rate and volatility of a stock, given its size. We then weaken the independence assumption in Section 3, and show how to use the firm size in the construction of portfolios. The first approach is to assume that we use only the firm sizes, and deduce the conditional distribution of the growth rate and the volatility; and the second would be to assume that the volatility could be estimated with sufficient precision that it might be assumed known, and then use the size and the volatility to deduce the conditional distribution of the growth rate.

The model studied here has implications for the distribution of firm sizes[2], and can be used to explore the distribution of the sizes of firms of different ages[3]. This is investigated in Section 4, and conclusions follow in Section 5.

There is already an extensive literature on the distribution of firm sizes, much of which refers back to a principle called Gibrat's Law, which begins from the statement[4] that the log of a firm's size should evolve as a Lévy process, and concludes that the firm size distribution should be log-normal, as would be expected from the Central Limit Theorem[5].

[1]This could be interpreted as bankruptcy, or merger, or any other way that a firm could vanish. Interestingly, Hopenhayn [Hopenhayn, 1992] cites evidence that in the US, forty percent of manufacturing firms disappear and are replaced by new ones over a five-year period.

[2]See the plots in [Fernholz, 2002] which display the remarkable stability of plots of log size against log rank for US equities.

[3]The paper of Cabral & Mata [Cabral and Mata, 2003] presents evidence based on a study of Portuguese firms that the distribution of the sizes of older firms tends to have higher mean and less skew.

[4]We paraphrase here, to give the hypothesis in contemporary form.

[5]In some quarters, Gibrat's Law is stated as saying that the tail of the firm-size distribution should behave as a power law; see, for example, [Simon and Bonini, 1958]. In practice, it is hard to distinguish between

It would be impossible to conduct a full literature survey, but we mention here a few particularly interesting or influential contributions. Hopenhayn [Hopenhayn, 1992] presents a model where individual firms are subject to Markovian shocks and then optimally choose whether or not to continue. If they do, their outputs for the next period are optimally chosen. New firms enter until their marginal profitability falls to zero, and Hopenhayn derives a steady-state equilibrium solution. The steady-state nature of the solution seems to be the only feature shared with the present study. Ijiri & Simon [Ijiri and Simon, 1967] propose a model for firms in a sector where the log sizes evolve as AR(1) processes, together with a common sector-wide deterministic growth effect. In a similar vein, Hashemi [Hashemi, 2000] proposes to model the evolution of the log size as an Ornstein-Uhlenbeck process, and then compares the theoretical log-gaussian steady-state distribution of firm size with empirical data. In an empirical study of data on Portuguese firms, Cabral & Mata [Cabral and Mata, 2003] show that the density of firm size varies with age, generally moving to the right and flattening as the firms get older. What differentiates our study from most of the existing literature is that we aim to use firm size as an indicator for investment; the distribution of firm sizes as such is for us merely a staging point on our route.

2 Modelling assumptions.

Suppose that firms are born at the times of a Poisson process of intensity ρ. At the time of birth, a firm is of unit size, and is given a drift μ and a volatility σ drawn from some density. The size S_t of a firm evolves as

$$dS_t = S_t(\sigma dW_t + \mu dt) \tag{2.1}$$

where W is a Brownian motion independent of the values of (μ, σ) of the firm, and independent of all Brownian motions of all other firms. The firm is killed at rate $\varepsilon > 0$. This rate will be allowed to depend on σ as

$$\varepsilon(\sigma) = \tfrac{1}{2}\omega_0^2 \sigma^{2\psi} \tag{2.2}$$

for some $\omega_0 > 0$ and $\psi \in [0, 1)$. Suppose that the times of the Poisson process before 0 are given by $0 > -\tau_1 > -\tau_2 > \ldots$. Letting $X := \log S$ denote the log size of a firm, we have that X is a Brownian motion with volatility σ and drift $c := \mu - \frac{1}{2}\sigma^2$. It turns out to be convenient to reparametrize the characteristics of the firm as $(\theta, \sigma) := (c/\sigma^2, \sigma)$. In these terms, we shall let m denote the density of the pair (θ, σ) selected independently for each firm.

At time 0, there will be a random number of firms alive[6]. By considering the possible

a log-normal distribution and a Pareto distribution. Kaizoji *et al.* [Kaizoji et al., 2006] find that while US firm sizes have an approximately log-normal distribution, a power-law tail better explains firms size data in Japan.

[6]If $\psi = 0$ and therefore ε is constant, then the distribution of firms still alive will be Poisson with mean ρ/ε, since this is the limiting distribution of the length of an $M/M/\infty$ queue with arrival rate ρ and service rate ε.

survival of each of the firms born, it is not hard to see that the time-0 density of firms of log-size x and characteristics (θ, σ) will be

$$
\begin{aligned}
q(x,\theta,\sigma) &:= \sum_{n\geq 0}\int_0^\infty \frac{(\rho t)^n e^{-\rho t}}{n!}\,\rho\exp(-\varepsilon t-(x-ct)^2/2\sigma^2 t)\,\frac{dt}{\sqrt{2\pi\sigma^2 t}}\;\; m(\theta,\sigma)\\
&= \int_0^\infty \rho\exp(-\varepsilon t-(x-ct)^2/2\sigma^2 t)\,\frac{dt}{\sqrt{2\pi\sigma^2 t}}\;\; m(\theta,\sigma)\\
&= \frac{\rho\exp(x\theta-|x|\sqrt{\theta^2+\delta^2})}{\sigma^2\sqrt{\theta^2+\delta^2}}\;\; m(\theta,\sigma), \qquad (2.3)
\end{aligned}
$$

where

$$\delta := \sqrt{2\varepsilon}/\sigma = \omega_0\,\sigma^{\psi-1}. \qquad (2.4)$$

What choices of m could be made to ensure that the conditional density of (θ, σ) given x would be tractable? We shall consider two particularly amenable alternatives[7]

$$m(\theta,\sigma) = B_\lambda \exp(\xi\theta - A\sqrt{\theta^2+\delta^2})\;\{\theta^2+\delta^2\}^{\lambda/2}\;\sigma^{-\nu} \qquad (2.5)$$

for $\lambda = 0,\ 1$, for some $A > 0$, $\nu > 1$, $|\xi| < A$, and normalization constant B_λ. The integral with respect to θ of (2.5) can be expressed in terms of Bessel functions, using the identities (A.7), (A.7). Recalling that $\delta = \omega_0\,\sigma^{\psi-1}$, we see from (A.10) that the integral over σ will be finite only if $\nu > 1$. Lengthy but routine calculations lead us to explicit expressions for the normalizing constants $B_0,\ B_1$; we defer these calculations to Appendix A.

The point of these choices is that in either case, *conditional on x and σ, θ has a generalized hyperbolic distribution*; specifically, the conditional distribution of θ given x and σ is $GH(\cdot \mid \lambda, \alpha, \beta, \delta, 0)$ for $\lambda = 0,\ 1$, where $\alpha = A + |x|$, $\beta = x + \xi$, $\delta = \omega_0\,\sigma^{\psi-1}$. The mean is given by the formula

$$E[\,\theta \mid x,\sigma\,] = \frac{\delta\beta K_{\lambda+1}(\gamma\delta)}{\gamma K_\lambda(\gamma\delta)}, \qquad (2.6)$$

where $\gamma := \sqrt{\alpha^2-\beta^2}$.

We can also obtain the marginal distribution of x by integrating out the variables θ and σ. Writing

$$\alpha = A+|x|, \quad \beta = x+\xi, \quad \gamma = \sqrt{\alpha^2-\beta^2}, \qquad (2.7)$$

[7]Integer λ all lead to tractable forms, because $e^x\,K_{\lambda-\frac{1}{2}}(x)\,x^{\lambda+\frac{1}{2}}$ is a polynomial of degree λ for any non-negative integer λ. Going beyond $\lambda = 1$ does not seem to add much.

we have

$$
\begin{aligned}
q(x) &= \iint \rho B_\lambda \exp\{\,\beta\theta - \alpha\sqrt{\theta^2+\delta^2}\,\}\,(\theta^2+\delta^2)^{(\lambda-1)/2}\sigma^{-\nu-2}\,d\theta d\sigma && (2.8)\\
&= \iint \rho B_\lambda\, GH(\theta\,|\,\lambda,\alpha,\beta,\delta,0)\; 2\alpha^\lambda\, K_\lambda(\gamma\delta)(\delta/\gamma)^\lambda \sigma^{-\nu-2}\,d\theta d\sigma && (2.9)\\
&= \int_0^\infty \rho B_\lambda\; 2\alpha^\lambda\, K_\lambda(\gamma\delta)\;(\delta/\gamma)^\lambda\,\sigma^{-\nu-2}\,d\sigma\\
&= 2\rho B_\lambda(\alpha/\gamma)^\lambda \int_0^\infty \sigma^{-\nu-\psi}(\omega_0\,\sigma^{\psi-1})^\lambda K_\lambda(\gamma\delta)\;\sigma^{\psi-2}d\sigma\\
&= 2\rho B_\lambda(\alpha/\gamma)^\lambda \int_0^\infty \left(\frac{v}{\gamma\omega_0}\right)^{(\nu+\psi)/(1-\psi)}\left(\frac{v}{\gamma}\right)^\lambda K_\lambda(v)\;\frac{dv}{\gamma\omega_0(1-\psi)}\\
&= \frac{2\rho B_\lambda\alpha^\lambda}{1-\psi}\,\omega_0^{-(1+\nu)/(1-\psi)}\,\gamma^{-2\lambda-1-(\nu+\psi)/(1-\psi)}\;\int_0^\infty v^{\mu'-1}K_\lambda(v)\;dv\\
&= \frac{\rho B_\lambda\alpha^\lambda}{1-\psi}\,\omega_0^{-(1+\nu)/(1-\psi)}\,\gamma^{-2\lambda-1-(\nu+\psi)/(1-\psi)}\;2^{\mu'-1}\,\Gamma\left(\frac{\mu'-\lambda}{2}\right)\Gamma\left(\frac{\mu'+\lambda}{2}\right), && (2.10)
\end{aligned}
$$

where $\mu' = \lambda + (\nu+1)/(1-\psi)$. This is an unusual density; it has polynomial tails and is not differentiable at zero, but is otherwise not too bad.

3 Investing based on firm size.

At time 0, we see a (Poisson) number N of firms of different sizes, and from the sizes we can deduce the conditional distribution of their characteristics $(\theta,\sigma) \equiv (\mu\sigma^{-2} - \frac{1}{2}, \sigma)$. How would an agent choose to invest in the available assets?

To answer this, let us suppose that the size at time t of the i^{th} firm which was available[8] at time zero is denoted by S^i_t. Our modelling hypothesis tells us that

$$
dS^i_t = S^i_t(\sigma_i dW^i_t + \mu_i dt), \qquad (i = 1,\dots,N) \tag{3.1}
$$

where the W^i are independent Brownian motions, and the σ_i and μ_i are unknown parameters having the conditional distributions implied by the joint density (2.3), and the marginals (2.5). Suppose that an agent starts with wealth 1, and may invest in these N risky assets, as well as a bank account delivering interest at constant continuously-compounded rate r; his objective is to achieve

$$
\sup E\log(w_T) \tag{3.2}
$$

where w_t is his wealth at time t. We shall also suppose that the agent only considers fixed-mix rules, where his wealth evolves as[9]

$$
dw_t = w_t\big[\; rdt + \pi\cdot(\sigma dW_t + (\mu - r\mathbf{1})\,dt)\;\big] \tag{3.3}
$$

[8] We shall suppose that t is small enough that we may neglect the deaths and births of firms during $(0,t)$.

[9] We use $\mathbf{1}$ to denote the column vector of 1's.

where $\pi \in \mathbb{R}^N$ is a fixed vector, and we (ab)use the symbol σ to denote the diagonal matrix $\mathrm{diag}(\sigma_i)$. The reason to restrict to such rules is that if we were to consider the classical Merton [Merton, 1971] investment problem (see [Rogers, 2013] for a recent account) where the agent knew the constant growth rates and constant volatility of the assets, and if his utility was CRRA, then he would invest fixed proportions of his wealth in the risky assets. We are here in a situation where the growth rates and volatilities are assumed constant but not known, the agent's utility is CRRA, so it is a reasonable approximation to restrict to such strategies.

Accepting this, if the agent invests as in (3.3) his wealth at time T will be

$$\log w_T = \pi \cdot (\sigma W_T + (\mu - r\mathbf{1})T) + (r - \tfrac{1}{2}|\sigma^T \pi|^2)T. \tag{3.4}$$

Therefore, his objective (3.2) is simply

$$E \log w_T = \pi \cdot (E\mu - r\mathbf{1})T + (r - \tfrac{1}{2}\pi \cdot E(\sigma\sigma^T)\pi)T. \tag{3.5}$$

Optimizing this over π leads to the optimal choice

$$\pi^* = E(\sigma\sigma^T)^{-1}(E\mu - r\mathbf{1}), \tag{3.6}$$

where the expectations are expectations conditional on what the agent is assumed to know. The message therefore is that all we need is to calculate the conditional expectations of μ_i and σ_i^2 given the log-size x_i, and then deduce π^* from (3.6).

We shall consider two situations, the first where the agent observes only the log-sizes x_i; and the second where the agent observes the log-sizes x_i and also the covariance structure of the assets. In principle, by observing the evolution of the S^i, the agent is able to refine his knowledge of σ_i, μ_i as time passes, but this would be quite intractable if treated completely; we therefore resort to the simplifying assumption that there is no change in the agent's beliefs about the model parameters over the investment window.

3.1 Case 1: Observation of log sizes only.

Here we shall suppose that the *only* information which the agent exploits is the initial log-sizes $x_i \equiv \log S_0^i$, $i = 1, \ldots, N$, and he calculates the conditional means and variances given those observations. This modelling story is not wrong, but it is limited, because our assumptions make all of the assets *independent*. Empirical evidence firmly rejects this, and suggests instead some generally positive correlation between different stocks. Introducing dependence into the model has to be done with some care. The obvious first attempt is just to make the evolutions of the different log firm sizes to be correlated Brownian motions, but this is completely intractable. It is completely intractable because conditional on the sizes of the surviving firms at time 0 we have information about the joint moves of the underlying Brownian motions; and so the mean growth rates and volatilities of each of the individual stocks will be conditionally dependent on the observed values of *all* of the stocks. The very simple forms of the joint distributions that we carefully constructed in Section 2 are destroyed.

So what we shall suppose is that we observe log firm sizes X_t^i which evolve as

$$Y_t^i = X_t^i + Z_t \tag{3.7}$$

where the processes X^i are independent, and evolve as presented in Section 2, and Z is a stationary Ornstein-Uhlenbeck process independent of the $\tilde{X}^i$, evolving as

$$dZ_t = \bar{\sigma}\, d\bar{W}_t - bZ_t\, dt \tag{3.8}$$

for some constants $\bar{\sigma}$ and $b > 0$. We could interpret Z as some overall business cycle effect if we so desired; this would give us some notion of the magnitude of b, since business cycles are generally held to last on the order of three to five years. Whatever the interpretation of Z, including it in the modelling story has two effects. The first is that values $E\,\mu$ and $E(\sigma\sigma^T)$ appearing in the expression (3.6) for π^* have to be replaced by their conditional expectations given the observed value of Y_0^i. The effect of this is to replace the conditional expectation $E[\cdot|x]$ by the convolution integral

$$E[\,\cdot \mid y\,] = \int E[\,\cdot \mid y - z\,]\, p_M(z)\, dz \tag{3.9}$$

where p_M is the invariant $N(0, \bar{\sigma}^2/2b)$ density of the stationary process Z. The second effect of this modelling assumption is to change the covariance structure[10] of the available assets to be positively correlated.

Changing the story in this way alters the wealth dynamics (3.3) to[11]

$$dw_t = w_t\big[\, rdt + \pi \cdot (\sigma dW_t + \bar{\sigma} d\bar{W}_t + (\mu - r\mathbf{1})dt)\,\big]. \tag{3.10}$$

The analysis that led to the expression (3.6) for π^* goes through with minor modification to give us

$$\pi^* = \bar{V}^{-1}(E\mu - r\mathbf{1}), \tag{3.11}$$

where

$$\bar{V} \equiv V + \bar{\sigma}^2 \mathbf{1}\mathbf{1}^T \tag{3.12}$$

and V is the diagonal matrix $E[\mathrm{diag}(\sigma_i^2)]$. The form of $\bar{V}$ is simple enough to invert explicitly, and we find after a few calculations that

$$\pi^* = V^{-1}(E\mu - r\mathbf{1}) - \bar{\sigma}^2 V^{-1}\mathbf{1}\, \frac{\mathbf{1}\cdot V^{-1}(E\mu - r\mathbf{1})}{1 + \bar{\sigma}^2\, \mathbf{1}\cdot V^{-1}\mathbf{1}}. \tag{3.13}$$

[10]We shall assume that the time scale T of the investor is sufficiently small relative to the business cycle time scale that we may neglect the effect on the assets' drifts due to the presence of Z.

[11]For simplicity, we shall ignore the mean-reversion effect in the drift of Z. We have insisted that the investment horizon is small enough that the hypothesis that there is no change in the population of firms is a reasonable one, and such time scales would be small compared to the time scales of a business cycle.

We therefore see that in order to calculate the optimal investment proportions, we need to be able to calculate the conditional expectations $E[\mu|x]$ and $E[\sigma^2|x]$.

Calculating an expression for $E[\sigma^2|x]$ requires minor modification of an earlier calculation. At (2.8) we integrated out θ and σ to obtain the marginal density of x, and to obtain the conditional variance we simply need to insert a factor of σ^2 into the integrand. The effect of this is to change ν to $\nu-2$, and the final expression (2.10) with this substitution gives $q(x)E[\sigma^2|x]$:

$$q(x)E[\sigma^2|x] = \frac{\rho B_\lambda \alpha^\lambda}{1-\psi}\, \omega_0^{-(\nu-1)/(1-\psi)}\, \gamma^{-2\lambda-(\nu-1)/(1-\psi)}\, 2^{\mu_*-1}\, \Gamma\Big(\frac{\nu-1}{2(1-\psi)}\Big)\, \Gamma\Big(\frac{\nu-1}{2(1-\psi)}+\lambda\Big), \tag{3.14}$$

where $\mu_* = \lambda + (\nu-1)/(1-\psi)$.

Calculating $E[\mu|x] = E[\sigma^2(\theta+\frac{1}{2})|\, x\,]$ is a bit more involved, but we can use the known distribution of θ conditional on x and σ to express $E[\sigma^2\theta\,|x] = E[\sigma^2 E\{\theta|x,\sigma\}|x]$. We have

$$E[\,\theta\,|\,x,\sigma\,] = \frac{\delta\beta K_{\lambda+1}(\delta\gamma)}{\gamma K_\lambda(\delta\gamma)}, \tag{3.15}$$

so we can use the notations (2.7) and modify the expression at (2.9) to give

$$\begin{aligned} q(x)E[\sigma^2\theta|x] &= \int \rho B_\lambda(x+\xi)\, 2\alpha^\lambda\, K_{\lambda+1}(\gamma\delta)(\delta/\gamma)^{\lambda+1}\sigma^{-\nu}\, d\sigma \\ &= \frac{\rho B_\lambda\, \beta\omega_0^{-(\nu-1)/(1-\psi)}}{1-\psi}\, \frac{2^{\mu_*-1}\,\alpha^\lambda}{\gamma^{\mu_*+\lambda+1}}\, \Gamma\Big(\frac{\mu_*-\lambda-1}{2}\Big)\, \Gamma\Big(\frac{\mu_*+\lambda+1}{2}\Big). \end{aligned} \tag{3.16}$$

where $\mu_* = \lambda + (\nu-\psi)/(1-\psi)$.

3.2 Case 2: Observation of log sizes and covariance structure.

In this subsection, we shall suppose that the agent knows for each of the available firms not just the size, but also the volatility. This is a reasonable assumption to make, since it is typically much easier to make a reliable estimate of volatility than growth rate; see, for example, [Rogers, 2013], Chapter 4.2. However, it would be a mistake to assume that the assets evolve independently, and it is therefore necessary to say what is assumed about their co-movement. What we shall suppose is that the log-size x^i of firm i evolves as

$$dx_t^i = \sigma_0 \exp(-Rx_t^i)\, dW_t^i + \bar\sigma\, d\bar W_t + \text{FV term}, \tag{3.17}$$

where $R > 0$ is a constant, and the Brownian motions W^i, $\bar W$ are independent. This modelling story reflects the idea that all firms are affected by market-wide movements, but smaller firms are more influenced by their own idiosyncratic effects. The assumed form (3.17) of the evolution of the x^i leads in turn to the form

$$a \equiv \sigma\sigma^T = \bar\sigma^2\, \mathbf{1}\mathbf{1}^T + \sigma_0^2\, \mathrm{diag}(\, \exp(-2Rx^i)\,) \tag{3.18}$$

for the instantaneous covariance for this example. The optimal investment proportions are again determined by (3.6), this time taking the form

$$\pi^* = a^{-1}(E[\,\mu\,|\,x,\sigma\,] - r\mathbf{1}) \tag{3.19}$$

where the conditional expectation $E[\,\mu\,|\,x,\sigma\,] = \sigma^2\{\,E[\theta|x,\sigma] + \frac{1}{2}\,\}$ is determined using (2.6).

4 Numerical studies.

Here we present some results of a numerical exploration of the solutions derived in the previous Sections of the paper. We supposed that N firms of a variety of sizes were available for investment, ranging from -4 to 12, corresponding to a range of actual sizes from 0.018 to 162754. If we imagine that the starting size of a firm was USD 1M, then this would cover firms worth anything from about USD 20,000 up to USD 160bn, which is a realistic range.

In all the plots we kept fixed the values $\rho = 1$, $r = 0.05$, $\lambda = 0$, $\omega_0 = 0.75$, and $\psi = 0.5$.

Case 1. The first plot Figure 1 shows portfolio chosen in Case 1, and the expected growth rates and expected volatilities[12] as they vary with firm size in this example; we see that the expected growth rate and volatility both have a unique minumum near to zero, and are monotone away from that minimum.

The portfolio weights are those which would be chosen by a log agent who was presented with the opportunity to invest in $N = 1000$ firms in the market, where the firm sizes follow the distribution given by the density q from (2.10). As would be expected, the agent invests more in the assets with higher growth rates, but this is not the whole story. The small firms are also present in the portfolio with positive weight, while mid-cap firms are shorted.

Case 2. We present two plots for Case 2, the first, Figure 2, with $\sigma_0 = 0.25$, and the second, Figure 3, with $\sigma_0 = 0.65$. In both examples we took $\bar{\sigma} = 0.3$ as with the Case 1 examples.

The plots show the portfolio weights for the two examples, and these are quite strikingly different. In the first, the agent should be short small stocks, and long the larger stocks. For the second example, the agent should be long the small stocks and short the large ones, exactly the other way round! Under our modelling assumptions, smaller firms have higher volatility, but they also give increased diversification benefit, and the first two examples presented illustrate how these two effects work against each other: for smaller idiosyncratic volatility, we see that the agent will put more weight on the larger firms, but for the higher idiosyncratic volatility the diversification benefits prevail, and the agent chooses *less* of the larger firms, even going negative in the biggest firms. We should be careful not to read too much into these results before we have some convincing econometric justification for the assumed form (3.17) of the asset volatilities, but the point being made here is that given some reasonable estimates of the asset covariance we can calculate the optimal positions to

[12] In fact, the square root of $E[\sigma^2]$

Figure 1: Expected volatilities and growth rates for Case 1. The first panel shows the portfolio size (3.13) against log size, the second shows the expected volatility conditional on log size, and the third show the expected growth rate conditional on log size.

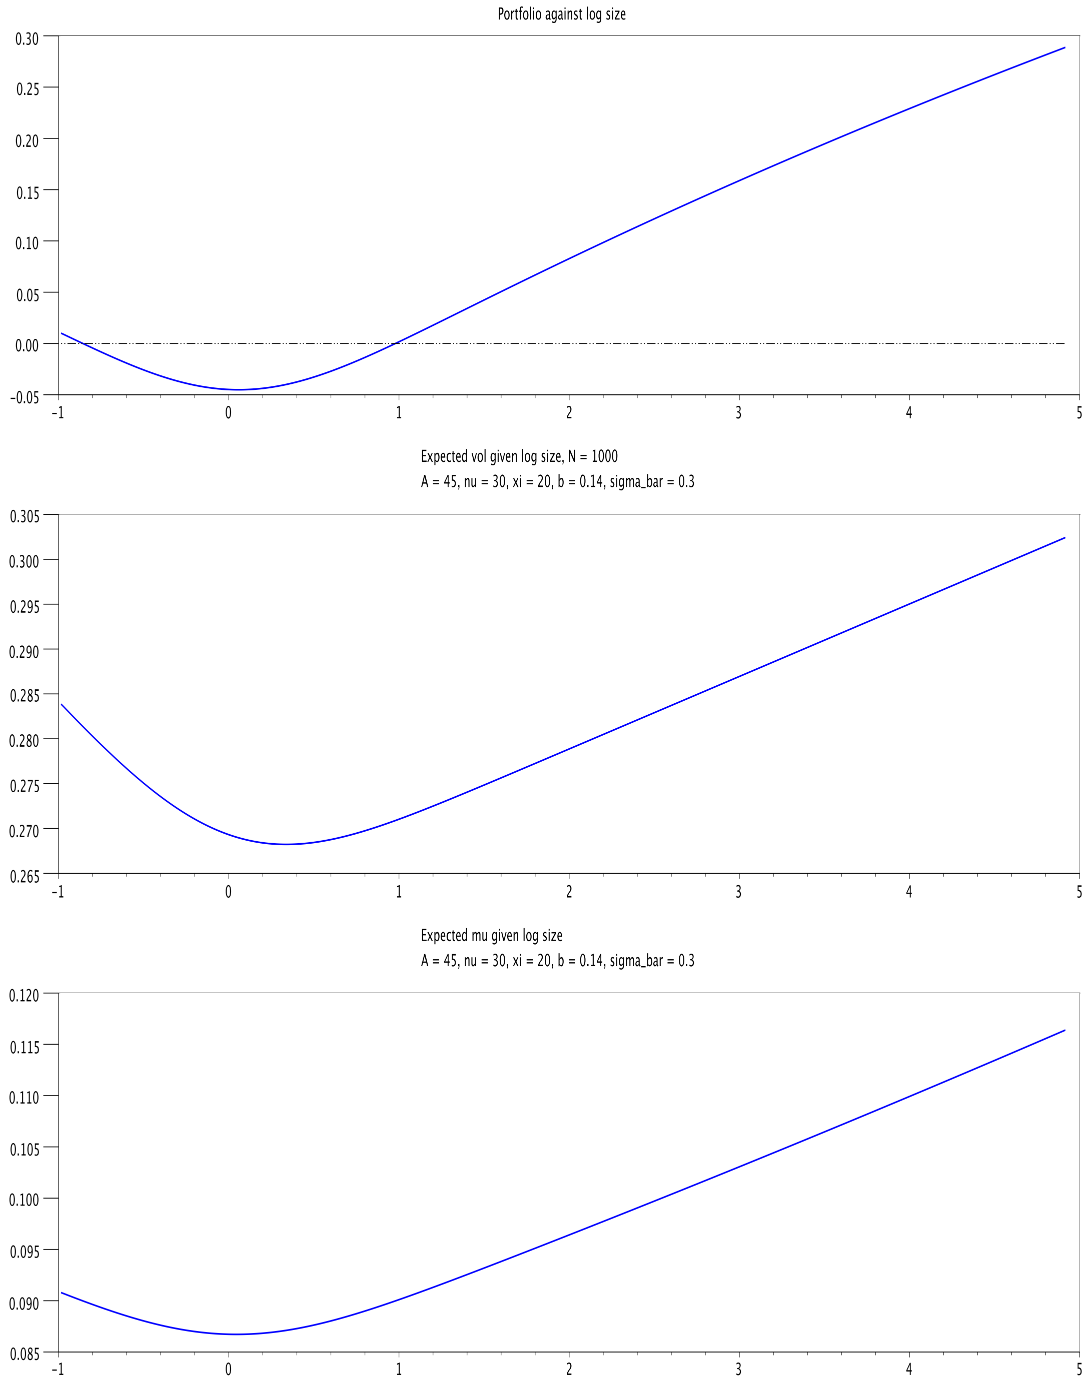

Figure 2: Case 2 (known variance), plot of portfolio against log size, taking idiosyncratic volatility $\sigma_0 = 0.25$, common volatility $\bar{\sigma} = 0.3$.

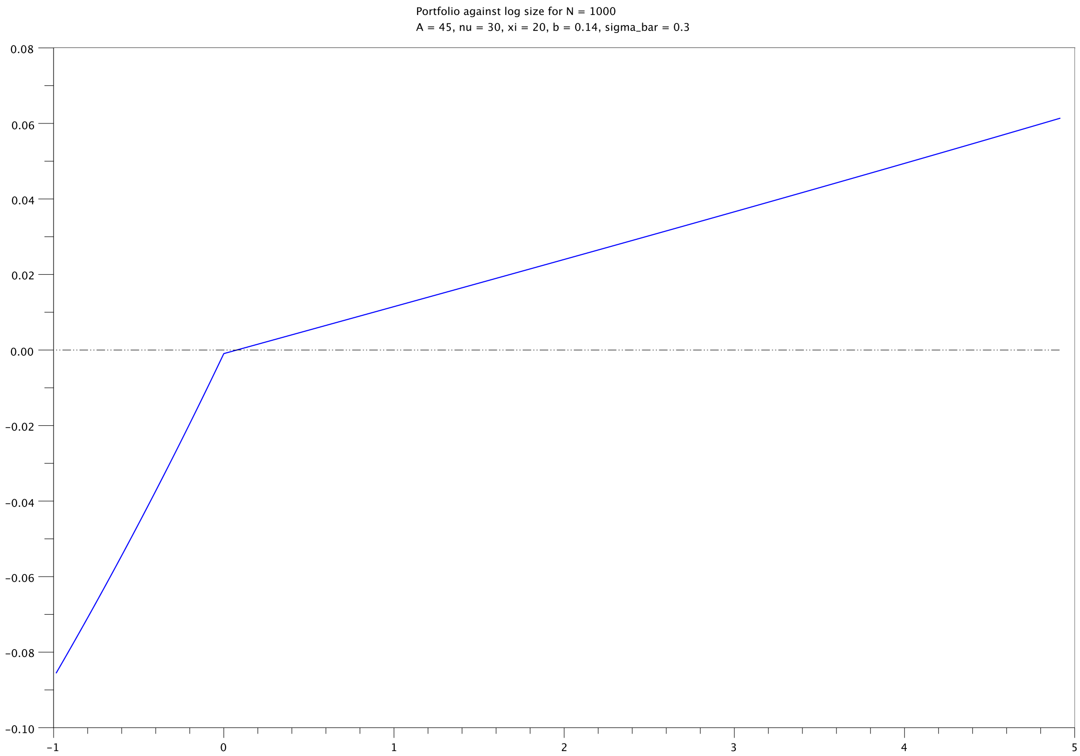

be held *without* having to make any estimates of growth rate, which is problematic if we ever try to apply the Merton solution in practice. It is intriguing that the portfolio weights we see can favour the small firms, or the large firms, depending on the values of some of the model parameters. Thus these two very different investment styles can be accommodated in our modelling assumptions.

The final plot Figure 4 is inspired by the remarkable plot on page 95 of Fernholz [Fernholz, 2002] of the log of market capitalization against the log of the rank of the firm. The plot in [Fernholz, 2002] shows how this looks for the US equity market at different points in time, separated by decades, and the striking fact is that the plots are very similar. All the plots fall away linearly for most of the range of log rank, and then curve down more steeply as the ranks get larger. Our Figure 4 is a plot of the theoretical relationship of log size against log

Figure 3: Case 2 (known variance), plot of portfolio against log size, taking idiosyncratic volatility $\sigma_0 = 0.65$, common volatility $\bar{\sigma} = 0.3$.

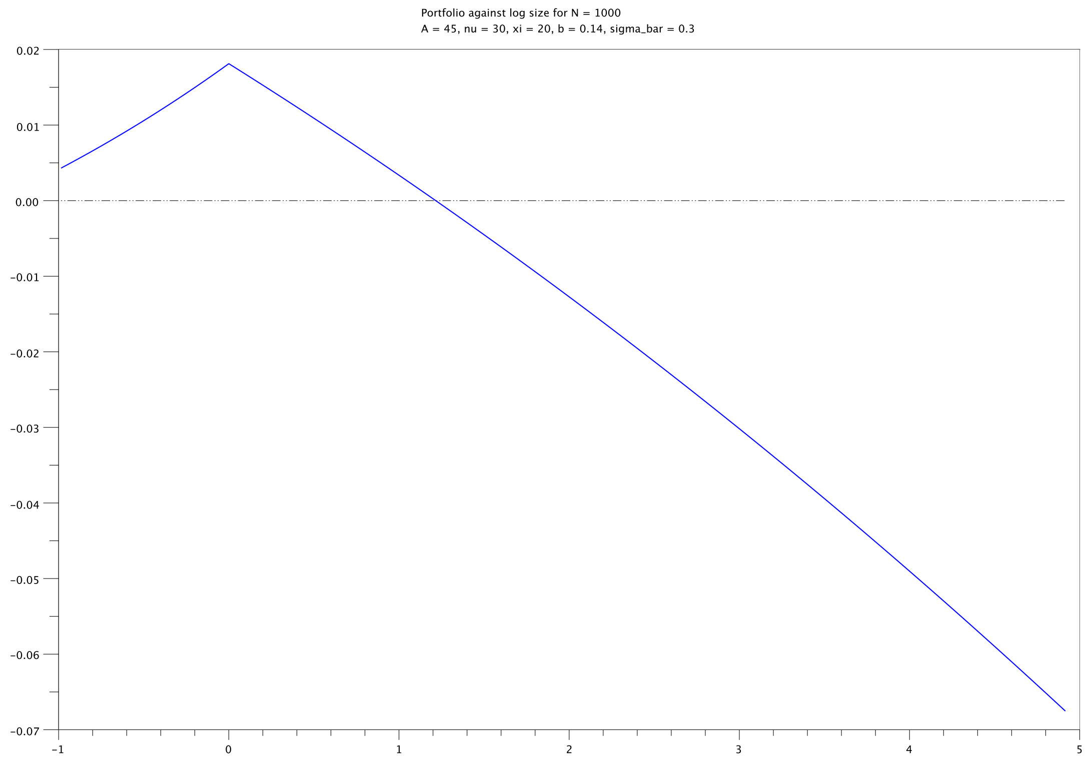

Figure 4: Log size against log rank

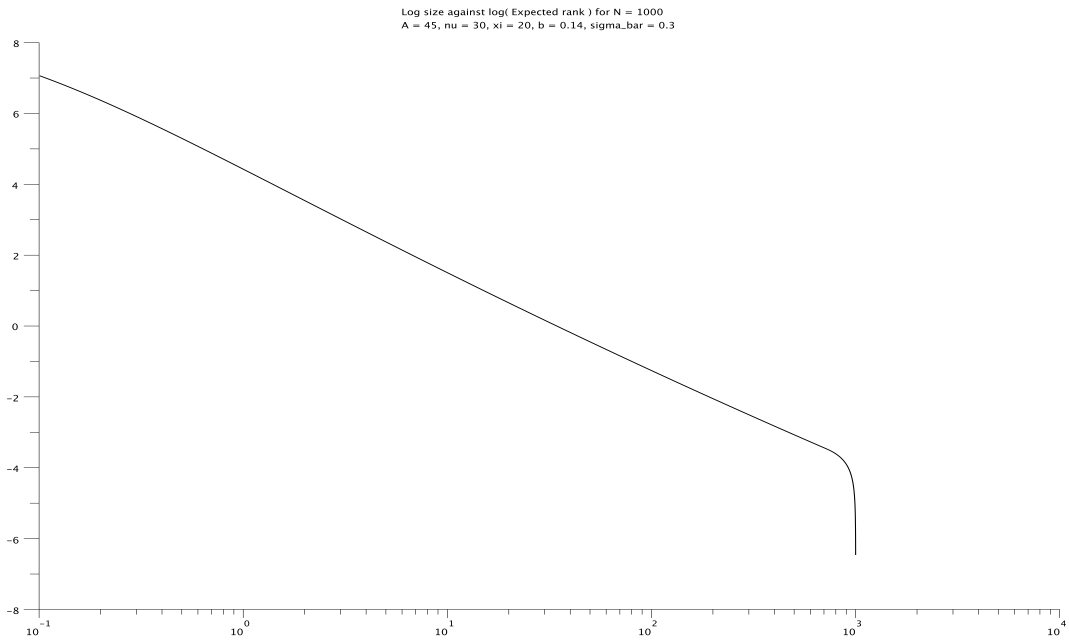

rank implied by the modelling assumptions of this paper, and shows a very striking similarity to the empirical plots from Fernholz.

5 Conclusions

We have in this paper presented and analyzed a simple yet tractable model which tries to explain why it might be that firm size is a relevant indicator for investment. The investment implications and econometric evidence remain to be evaluated. One important advantage of this approach is that we are freed from any need or temptation to estimate growth rates. If we allow ourselves only to use firm sizes (Case 1), we have a recipe for choosing portfolio proportions which requires the input of only a few parameters (recall that our numerical study was investing in 1000 assets), though it does make some bold simplifying assumptions about the dependencies between assets. If we allow ourselves to input some information about asset covariances (Case 2), then we can also find investment proportions, again without need to estimate growth rates. The qualitative forms of the solutions found suggest that the model will adapt well to different situations. Finally, when we examine the relationship implied between log size and log rank, we find it closely matches the empirical relationship, which has been shown to be stable over long periods of time. These observations lead us to conclude that the model presented here has a number of advantages, and deserves to be explored further.

A Generalized hyperbolic distributions.

The generalized hyperbolic distribution has density

$$\begin{aligned} y &\mapsto GH(y \,|\, \lambda, \alpha, \beta, \delta, \mu) \\ &\equiv \frac{(\gamma/\delta)^\lambda}{\sqrt{2\pi} K_\lambda(\delta\gamma)} \frac{\exp(\beta(y-\mu))\, K_{\lambda-\frac{1}{2}}(\alpha\sqrt{\delta^2+(y-\mu)^2})}{(\sqrt{\delta^2+(y-\mu)^2}/\alpha)^{\frac{1}{2}-\lambda}} \end{aligned} \tag{A.1}$$

where $\gamma = \sqrt{\alpha^2 - \beta^2}$, and all the parameters $\lambda, \alpha, \beta, \delta, \mu$ are real, with the condition that $|\beta| \leq |\alpha|$. The MGF is of the form

$$z \mapsto \frac{\gamma^\lambda e^{\mu z}}{\{\alpha^2 - (\beta+z)^2\}^{\lambda/2}} \frac{K_\lambda(\delta\sqrt{\alpha^2-(\beta+z)^2})}{K_\lambda(\delta\gamma)}. \tag{A.2}$$

In this note, we are concerned only with the case where the offset μ is zero, and with $\lambda = 0,\, 1$. This is because

$$K_{\frac{1}{2}}(y) = K_{-\frac{1}{2}}(y) = e^{-y}\sqrt{\frac{\pi}{2y}}, \tag{A.3}$$

which fits nicely with the form of the density of the size. We find that for $\lambda = 0, 1$ the generalized hyperbolic density is

$$GH(y \,|\, \lambda, \alpha, \beta, \delta, 0) = \frac{(\gamma/\delta)^\lambda \exp(\beta y - \alpha\sqrt{y^2+\delta^2})}{2\alpha\, K_\lambda(\gamma\delta)} \left(\frac{\sqrt{y^2+\delta^2}}{\alpha}\right)^{\lambda-1}. \tag{A.4}$$

Using the fact that the density integrates to 1, we see that for $\alpha > 0$, and for $|\beta| < \alpha$, writing $\gamma = \sqrt{\alpha^2 - \beta^2}$,

$$\int \exp(\beta y - \alpha\sqrt{y^2+\delta^2})\, dy = (2\alpha\delta/\gamma)\, K_1(\delta\gamma). \tag{A.5}$$

Differentiating this identity with respect to α and using the Bessel function identity

$$2K_\nu'(z) = -K_{\nu+1}(z) - K_{\nu-1}(z). \tag{A.6}$$

gives us

$$\begin{aligned} \int \exp(\beta y - \alpha\sqrt{y^2+\delta^2})\, \sqrt{y^2+\delta^2}\, dy &= (2\delta\beta^2/\gamma^3)\, K_1(\delta\gamma) + (\delta\alpha/\gamma)^2 \left[\, K_0(\delta\gamma) + K_2(\delta\gamma) \,\right] \\ &= \delta^2 \left(K_0(\delta\gamma) + K_2(\delta\gamma)(1 + \frac{2\beta^2}{\gamma^2}) \right), \end{aligned} \tag{A.7}$$

where the step to the final form (A.7) utilizes the identity

$$K_\nu(z) = K_{\nu+2}(z) - \frac{2(\nu+1)}{z} K_{\nu+1}(z). \tag{A.8}$$

In what follows, we will often use that for $C > 0$, $\mu > \lambda \geq 0$, we have (with $\psi \in [0, 1)$)

$$\int_0^\infty K_\lambda(Cs^{\psi-1})s^{-\mu-1}ds = (1-\psi)^{-1}C^{\frac{-\mu}{1-\psi}}\, 2^{\frac{\mu}{1-\psi}-2}\,\Gamma\left(\frac{\mu-\lambda+\lambda\psi}{2(1-\psi)}\right)\Gamma\left(\frac{\mu+\lambda-\lambda\psi}{2(1-\psi)}\right) \quad \text{(A.9)}$$

which by a simple substitution follows from the identity valid for all $\mu > \nu \geq 0$:

$$\int_0^\infty K_\lambda(t)t^{\mu-1}dt = 2^{\mu-2}\,\Gamma\left(\frac{\mu-\lambda}{2}\right)\Gamma\left(\frac{\mu+\lambda}{2}\right); \quad \text{(A.10)}$$

see equation (8) on page 388 of [Watson, 1995].

We are now in a position to evaluate the normalization constants B_0, B_1 from (2.5), making the obvious notational substitutions, A for α, ξ for β, $\omega_0\sigma^{\psi-1}$ for δ (see (2.4)), and writing $\gamma = \sqrt{A^2-\xi^2}$.

THE CASE $\lambda = 0$. Using (A.5) and (A.9), we obtain

$$\begin{aligned}
1 &= B_0\int_0^\infty \sigma^{-\nu}\int_{-\infty}^\infty \exp\left(\xi\theta - A\sqrt{\theta^2+2\varepsilon/\sigma^2}\right)\,d\theta d\sigma \\
&= 2AB_0\omega_0\gamma^{-1}\int_0^\infty \sigma^{-(\nu-\psi)-1}K_1(\gamma\omega_0\sigma^{\psi-1})\,d\sigma \\
&= \frac{AB_0}{(1-\psi)\gamma^2}\left(\frac{\gamma\omega_0}{2}\right)^{-(\nu-1)/(1-\psi)}\Gamma\left(\frac{\nu-2\psi+1}{2(1-\psi)}\right)\Gamma\left(\frac{\nu-1}{2(1-\psi)}\right).
\end{aligned}$$

THE CASE $\lambda = 1$. We set $\bar{\mu} = (\nu-2\psi+1)/(1-\psi)$ and now use (A.7) and (A.9):

$$\begin{aligned}
1 &= B_1\int_0^\infty \sigma^{-\nu}\int_{-\infty}^\infty \exp\left(\xi\theta - A\sqrt{\theta^2+2\varepsilon/\sigma^2}\right)\sqrt{\theta^2+2\varepsilon/\sigma^2}\,d\theta d\sigma \\
&= B_1\omega_0^2\int_0^\infty \sigma^{-(\nu-2\psi+1)-1}\left[K_0(\gamma\omega_0\sigma^{\psi-1}) + K_2(\gamma\omega_0\sigma^{\psi-1})(1+\frac{2\xi^2}{\gamma^2})\right]\,d\sigma \\
&= \frac{B_1\omega_0^2}{1-\psi}(\gamma w_0)^{-\bar{\mu}}2^{\bar{\mu}-2}\left[\Gamma\left(\frac{\bar{\mu}}{2}\right)^2 + (1+\frac{2\xi^2}{\gamma^2})\Gamma\left(\frac{\bar{\mu}-2}{2}\right)\Gamma\left(\frac{\bar{\mu}+2}{2}\right)\right] \\
&= \frac{B_1}{(1-\psi)\gamma^2}\left(\frac{\gamma\omega_0}{2}\right)^{-(\nu-1)/(1-\psi)}\Gamma\left(\frac{\bar{\mu}}{2}\right)\Gamma\left(\frac{\bar{\mu}-2}{2}\right)\left[\,\bar{\mu}-1+\bar{\mu}\,\frac{\xi^2}{\gamma^2}\,\right].
\end{aligned}$$

The normalization constants can now be read off.

References

[Cabral and Mata, 2003] Cabral, L. and Mata, J. (2003). On the evolution of the firm size distribution: Facts and theory. *The American Economic Review*, 93:1075–1090.

[Fernholz, 2002] Fernholz, E. R. (2002). *Stochastic Portfolio Theory*. Springer.

[Fernholz and Karatzas, 2006] Fernholz, R. and Karatzas, I. (2006). The implied liquidity premium for equities. *Annals of Finance*, 2(1):87–99.

[Hashemi, 2000] Hashemi, F. (2000). An evolutionary model of the size distribution of firms. *Journal of Evolutionary Economics*, 10(5):507–521.

[Hopenhayn, 1992] Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60:1127–1150.

[Ijiri and Simon, 1967] Ijiri, Y. and Simon, H. A. (1967). A model of business firm growth. *Econometrica*, pages 348–355.

[Kaizoji et al., 2006] Kaizoji, T., Iyetomi, H., and Ikeda, Y. (2006). Re-examination of the size distribution of firms. *Evolutionary and Institutional Economics Review*, 2(2):183–198.

[Merton, 1971] Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3:373–413.

[Rogers, 2013] Rogers, L. C. G. (2013). *Optimal Investment*. Springer.

[Simon and Bonini, 1958] Simon, H. A. and Bonini, C. P. (1958). The size distribution of business firms. *The American Economic Review*, 48(4):607–617.

[Watson, 1995] Watson, G. N. (1995). *A treatise on the theory of Bessel functions*. Cambridge University Press.