

THE SELF-FULFILLING PROPHECY OF POPULAR ASSET PRICING MODELS

Bradford Cornell^a and Jason Hsu^b

The assumption that asset prices are determined by the efforts of end investors to maximize intertemporal utility supports a pricing theory that is both elegant and intuitive. Unfortunately, the assumption is counterfactual. End investors, with few exceptions, lack the capacity to behave in a fashion consistent with the theory. More to the point, they don't try. Instead, they delegate investment decision-making. Thus, it is important to understand the investment management ecosystem. Is it a simple pass-through mechanism? We do not believe so and argue, instead, that the lack of alignment implies the cross-section of asset returns is significantly influenced by active money managers and deviates from the predictions of the consumption-based model. Using a simple thought experiment, we demonstrate that the widely adopted discounted cash flow model is likely both to drive prices and to determine the cross-section of average returns. This leads to a self-fulfilling feedback loop in which once an asset pricing model is adopted by active managers as a means of estimating the discount rate, it becomes a determinant of expected returns.



1 Introduction

A question that lies at the core of academic finance and is of intense interest to practitioners is why different assets have different expected returns. The simple answer is differences in risk. That, then, raises the central question in asset pricing—what is “risk” and how is it “priced”? Early answers to this question, such as CAPM, focused

on the second moments of asset returns. As finance evolved to become an applied branch of economics, financial economists began to assert that individuals mostly care about utility from consumption rather than the nominal balance of their portfolio. Risk, therefore, should be defined in terms of shocks to the investor’s consumption profile caused by fluctuations in asset prices. Following this logic, modern asset pricing models focus on the relation between consumption and asset returns. In this context, the goal of the end investor is to maximize the utility of intertemporal consumption.

^aCalifornia Institute of Technology, Pasadena, CA 91125, USA. E-mail: bcornell@hss.caltech.edu

^bRayliant Global Advisors, Research Affiliates, LLC, and Anderson Graduate School of Management, UCLA, Newport Beach, CA 92660, USA. E-mail: hsu@rallc.com.

For consumption-based models to explain the observed cross-section of returns, however, there must be some mechanism through which the end investors' utility functions impact asset prices. Here, we argue that there is no evidence such a mechanism exists. This should not be interpreted to mean that asset prices can become totally divorced from end-investor utility. Our point is that the bounds on prices imposed by end-investor utility maximization are weak, and end-investors' participation in the market is, generally, not the mechanism which determines the cross-section of expected returns.

Instead, we argue that asset prices, and thereby expected returns, are set by active fundamental managers, who are delegated agents of the end investors.¹ As a result, the cross-section of expected returns will be heavily influenced by the discount rate models used by these active fundamental managers. Insofar as these managers do not adopt a discount rate model based on the end-investors' utility functions, the latter become disconnected from asset prices. While our comments are general to the utility of the end investors, we narrow our exposition to focus on the standard consumption model, which we review in the next section.² This narrowing of focus puts us in the context of the traditional asset pricing literature and makes our critique more concrete, without loss of generality.

2 The consumption-based asset pricing model

The standard consumption-based model begins with the assumption that prices are set by end investors, who maximize the expected utility of stochastic intertemporal consumption. More specifically, the investor's problem is choosing assets and consumption so as to maximize

$$E_0 \sum_{t=0}^{\infty} B^t u(c_t) \quad (1)$$

where $u(c)$ is the representative investor's utility from consumption, B is the representative investor's subjective rate of discount, and the expectation is defined over states of the world at each date. Solving the intertemporal optimization problem, asset prices are then given by

$$p_0 = E_0 \sum_{t=1}^{\infty} B^t \frac{u'(c_t)}{u'(c_0)} x_t \quad (2)$$

where p_0 is the price of the asset, $u'(c)$ denotes the first derivative, and x_t is the future random cash flow produced by the asset in period t .³ Defining

$$m = B^t \frac{u'(c_t)}{u'(c_0)}$$

to be the stochastic discount factor, Equation (2) can be written more compactly as

$$p_0 = E_0 \sum_{t=1}^{\infty} m_t x_t. \quad (3)$$

The relation between the consumption-based model and other asset pricing models, such as CAPM and other linear factor models, can be seen more directly by rewriting Equation (3) in terms of returns (and dropping the time subscripts). Following Cochrane (2005), the return version of the consumption-based model is given by:

$$E(R^i) - R^f = -R^f \text{cov}(m, R^i). \quad (4)$$

Equation (4) makes it clear that the ultimate source of risk is the covariance of asset returns with the marginal utility of consumption as given by the stochastic discount factor.⁴ A risky security is one that pays out more in good times (when the marginal utility of consumption is low) and less in bad times.⁵ To compensate for this risk such securities have to offer higher expected returns.

Despite its elegance, the consumption-based model has one glaring deficiency. The standard model and its more conventional variants have failed at explaining the cross-section of

returns; even tortured versions of the standard model have struggled to match data. This is perhaps unsurprising. For Equations (3) and (4) to “work” there must be a mechanism by which they are enforced. Somehow investors must compare market prices with the prices generated from Equation (3), buying those securities that are underpriced and selling those that are overpriced. Cochrane (2005) describes how that mechanism works in the context of the consumption-based model. To paraphrase Cochrane,

If the consumption-based estimate of value is higher than the market value, and if the investor can buy more of the asset, she will. As she buys more, the market price of the asset will rise and her current and future consumption will increase. The increase in consumption drives down the marginal utility of consumption. This process continues until Equation (4) is satisfied.

The problem is that the mechanism described by Cochrane is counterfactual. Instead, due to cognitive and information constraints, end investors delegate investment decision-making to an asset management ecosystem which services them. In that ecosystem, end investors are not the price setting marginal investors. Instead, as French (2008) documents, investment decisions are largely delegated to professional money managers who are not incentivized to make sure that Equations (1) to (4) hold, and would not be so incentivized even it were possible to do that. The relevant question, therefore, is what determines the cross-section of expected returns given actual investment practice.⁶ Addressing that question requires a more detailed description of the institutional structure of investment management.

3 End investors and the institutional structure of investment decision-making

Surveys of financial literacy (see Lusardi and Mitchell, 2013; Securities and Exchange Commission, 2012) find that typical investors have

nowhere near the information, training, or skills to operationalize Equations (1) to (4) even if they are not subject to psychological biases. For instance, Lusardi and Mitchell (2013) report that only about half the individual investors surveyed understood that diversification reduces risk, and only 21% understood the inverse relation between bond prices and interest rates. Similarly, the SEC (2012) reports that most investors fail to understand how compound interest is calculated or the distinction between real and nominal returns.⁷

The survey results are not surprising. Investment knowledge is complex and costly to acquire. For decades, asset pricing was poorly understood even by finance scholars and professional money managers. The customary response to such complexity is not to incur the enormous financial and mental cost to become an expert in order to operationalize a sophisticated model. Instead, the majority of investors delegate investment decision-making by buying professionally managed funds or other investment products. Even investors who hold equity shares directly often delegate their investment decision-making. The top 1% of households, which hold approximately 50% of the total household direct equity holdings, generally hire private bankers or financial advisors to manage securities on a fully discretionary basis in their accounts.⁸

However, while professional money managers may possess the skill to evaluate securities based on Equations (3) and (4), there is little, if any, evidence that they actually have a specific mandate from investors to do so. In our experience, correlation with consumption does not enter the evaluation of manager performance.

Our approach should also be distinguished from that taken in the behavioral finance literature. The standard behavioral models assume that asset prices are determined by irrational individual agents and thus deviate from what are predicted

by the standard consumption-based model.⁹ By contrast, we do not link deviation in prices from the standard consumption-based model to end-investor trading mistakes. Instead, we argue that individual investors are aware of their own lack of knowledge and discipline and thus delegate investment decision-making to professional money managers. It is the lack of alignment between the manager and the end investor, along with the practical difficulty of actually operationalizing the consumption model for a large number of heterogeneous end investors, which drives a wedge between observed cross-sectional returns and Equation (4).¹⁰ Ignoring the asset management ecosystem is akin to analyzing political actions without considering politicians as self-interested entities—it overlooks the actual behavior of the dominant actors and ignores the effect of the mechanism for delegation on eventual outcomes.

Of course, it is possible that end investors delegate investment decision-making with the understanding that professional managers will attempt to solve the intertemporal optimization problem on their behalf—that is, managers are incentivized to care deeply about the covariance between stock returns and consumption growth fluctuations. Unfortunately, as noted above, we find little evidence that the incentives of the money managers are related to solving the end investor's utility maximization problem. Instead, nearly all managers are incentivized to outperform a stated benchmark rather than attempting to operationalize Equation (4) for their end investors.^{11,12} In this regard, it is useful to remember that professional money management services are sold through multiple intermediaries to end investors. To be marketed successfully, investment strategies need to be understood and accepted by end investors, who, as we have established already, often have little financial knowledge and, at worst, can have a very flawed understanding

of the capital market. What they understand and desire, in turn, is influenced by popular financial market media, which have neither the aptitude nor the appetite for Equations (1) to (4). To this end, final investors' aversion to handling the complexity of investing leads to the delegation of investment decision-making, and their limited investment knowledge, in conjunction with information costs, leads to imperfect contracting.

To further complicate matters, the delegation that occurs in financial markets is not a straightforward principal-agent relationship. For example, corporate pension fund investments are typically overseen by a chief investment officer (CIO) appointed by a board of trustees which generally includes a few corporate treasury officers and labor union representatives. The treasury officers represent the corporate pension sponsor, whose interest is to minimize pension expense; their interests are not aligned directly with the pension beneficiary. For public pension funds, the places of such corporate officers are taken by government officials, whose incentives are potentially further removed from the plan beneficiaries. The pension organizations then further delegate investment decisions to investment consultants, who then delegate to asset managers. Each layer of delegation gives rise to new conflict of interest.¹³ Additionally, pension benefits to the end beneficiary may be implicitly or explicitly guaranteed by governments regardless of the investment result; this means that the end investor's consumption planning is further delinked from the actions of the delegated agents in the investment ecosystem.

Given the environment described, in which money managers are incentivized to deliver out-performance relative to peers and/or benchmarks, such as the S&P500 Index, rather than to operationalize Equations (1) to (4), the question is what

then determines the cross-section of expected returns?

4 Money managers and their influence on market prices

To address the question, we begin by arguing that active fundamental managers play the dominant role in setting market prices. This follows almost by definition. Obviously, passive managers are price takers, who by and large maintain a buy-and-hold portfolio; at best, they impact prices, transiently, through managing inflows and outflows. Among active money managers, some invest based on relative valuation (e.g., using various valuation metrics relative to industry peers) while others invest based on fundamental value analysis. Even though relative valuation traders do not attempt to compute a “fair value” for stocks, Summers (1985) argues they nonetheless help discipline prices by anchoring on the price levels influenced by the fundamental value analysts.¹⁴ Ultimately, the prices of fundamental assets must be set by investors who compare prices with estimates of fundamental value and trade on the discrepancy. In this sense, the mechanism is similar to that described by Cochrane, but the actors have changed. Instead of end investors using personal utility to make the evaluation, we argue the key agents are active fundamental investors.¹⁵ It is critical to understand how these active fundamental managers estimate fair value, because the process they use in their assessments is the key to modeling the cross-section of expected returns.¹⁶

It is generally accepted that the ultimate source of value in any equity investment depends highly on the cash flows the company is expected to produce. Consequently, it is not controversial that fundamental valuations are often derived through discounted cash flow (DCF) analyses of some type. This general proposition is consistent

with the evidence. We find that discounted cash flow analysis is the method of choice for estimating fair value in the investment management industry. Brotherson *et al.* (2013, 2014) report that leading investment banks use DCF to estimate the enterprise value of a company. Furthermore, DCF analysis is the valuation approach advocated in standard MBA finance texts including Brealey *et al.* (2011) and Berk and Demarzo (2013). The DCF method is also the valuation approach featured in books such as by Pratt (2008), Damodaran (2012), Koller *et al.* (2010), and Rosenbaum and Pearl (2013), which are written for investment professionals. The DCF analysis is also the method taught in the CFA Program curriculum as well as tested in its certifying exams. In addition, Asquith *et al.* (2005) report that security analysts rely on DCF models for performing fundamental valuations.

Assuming that fundamental investors use DCF models to assess fundamental value, it is helpful to think in terms of a representative investor to simplify the analysis while noting that the conclusions are not dependent on this simplification. In this context, the conclusion is straightforward; the cross-section of expected returns is determined by the discount rates used by the representative investor in valuing individual common stocks. This follows from the mathematical fact that if the market price of a stock equals the discounted present value of expected future cash flows, the expected return on the stock, over the life of the investment, equals the discount rate. Consequently, to explain the cross-section of expected returns, it is necessary to learn how fundamental investors choose the discount rates used in their valuation analyses. This is the empirical question to which we turn to in the next section. For now it is surprising to observe that virtually no research has been done on the subject. This may be in part because of confidentiality limitations with regard to the valuation models employed by

active fundamental managers, but to our knowledge virtually no effort has been made to address the question. In light of our work, we see this as a significant shortcoming.

5 Self-fulfilling prophecy in asset pricing

The importance of understanding how fundamental managers choose the discount rates to use in their valuation analyses is underscored by what we call the self-fulfilling prophecy of asset pricing. The prophecy is best explained by a simple example.

Suppose, hypothetically, that based on academic research fundamental managers have decided to use the Fama–French three-factor model to compute discount rates. It follows that the cross-section of expected returns will reflect the size and value effects predicted by Fama and French (1993). This would occur even if Fama and French had hit upon the size and value factors by data mining. Once the model is established and accepted, it will continue to hold because it is established and accepted—making it even more likely to be accepted in the future. This is the self-fulfilling prophecy of asset pricing.

To pursue the thought experiment further define time, T , to be the date on which the Fama–French model is discovered, and, to avoid unnecessary complications, assume that once it is discovered the model is accepted without a lag. Suppose that prior to its discovery, the Fama–French model was one of several as yet unknown approximate asset pricing models and *not* the most accurate one. Then the Fama–French model is discovered and accepted. At that point in time, it becomes the best model. Its discovery alters the process for estimating discount rates. Thus the model enters the real world of security valuations and portfolio decisions. Going forward, the model derived from observed historical behavior *is* the *ex-ante* correct model of expected returns.

Because such path-dependent historical developments are possible, it would be helpful not only to know how active fundamental investors estimate discount rates at the present time, but also how they did so in the past. If a three-factor model is used now, what was used prior to its discovery? For that matter, what was used prior to the discovery of CAPM? Unfortunately, the data necessary to answer such questions with precision have probably been lost, but it may be possible to make some partial headway.

6 How active managers estimate their discount rates

To summarize, our analysis implies that the cross-section of expected returns is not only described by theories of asset pricing but also *determined* by them. In the standard model with rational end investors and no delegation, Equations (1) to (4) hold whether they have been discovered by financial economists or not. Put another way, the development of the consumption-based theory has no impact on the cross-section of expected returns; it simply describes and explains investor behaviors. The same is not true of asset pricing theories in the context of our thought experiment, where end investors delegate decision-making to active fundamental managers who then play the dominant role in setting prices. Models that money managers apply to estimate discount rates *become* the determinants of expected returns.

Note this prediction is very different from what is predicted by the behavioral finance literature. In the behavioral literature, an anomaly driven by investors' mistakes would be arbitrated by rational money managers upon being discovered and becoming popular. In our thought experiment, an anomaly could actually be sustained by the operating mechanism of the modern asset management ecosystem. Whether a return factor ends up being arbitrated away or further baked into the cross-section of stock returns would depend

significantly on whether money managers have been taught to adopt a factor as a component of the discount rate used in DCF valuation or as an anomaly to be exploited.

Given the importance of the process by which fundamental managers estimate discount rates to use in valuation, we conducted a preliminary empirical investigation. We emphasize that this investigation is highly tentative. Our hope is to motivate future research on what we feel is one of the most important and least examined areas in empirical asset pricing. Whereas there is massive literature on testing asset pricing theories, there has been virtually no empirical work on how fundamental investors, in practice, estimate the discount rates used in their valuation models.

To explore how discount rates might actually be estimated by fundamental managers, we looked at a variety of sources including analyst reports; professional services such as Morningstar (2013a, 2013b) and Duff and Phelps (2013) that provide estimates of the cost of equity capital; and confidential valuation reports prepared for M&A transactions by investment banks to which we had access. Although none of these sources provides direct evidence on how major fundamental investors, such as The Capital Group, estimate discount rates, the evidence is suggestive as to what approaches are likely to be commonly employed.

In brief, our review of the documents revealed the following. First, with respect to estimation of the discount rate, the CAPM, and Fama–French three-factor model were the overwhelmingly most popular choices. However, CAPM was rarely implemented directly, but instead was adjusted. The two most common adjustments were the inclusion of a size premium and the addition of an adjustment for “firm-specific risk.” The size factor is clear. With respect to the

firm-specific risk adjustments, general conclusions are more difficult to draw because different valuation reports made different adjustments. However, the adjustments typically appeared to be related to “distress,” which sometimes can be correlated with value characteristics in the cross-section. Accepting that characterization, the adjusted CAPM approaches a version of the Fama–French model with the distress adjustment replacing the HML factor. This does not, however, prove the existence of a self-fulfilling prophecy. It may mean simply that the three-factor model accurately captures the true risk factors independent of the process we describe. We address this possibility further in the next section.

Second, we find no evidence of attempts to operationalize the consumption-based asset pricing model, nor any recommendation that it be adopted. If our self-fulfilling prophecy is correct, the features of these models that extend the simpler CAPM and Fama–French models will not be reflected in asset prices. This is consistent with our assumption that the asset management ecosystem is such that the delegation of investment management decision-making does not result in managerial behaviors that can be modeled as helping end investors solve the problem of maximizing the present value of the expected utility of consumption.

7 Challenges to the self-fulfilling prophecy hypothesis

The self-fulfilling prophecy hypothesis is not without challenges. To begin, the hypothesis does not offer a full explanation how actual pricing and pricing theories interact and change over time. The theory predicts that once effects are baked into practice, as the size and value effects are included in the discount rate, those effects should persist indefinitely. In fact, for a period of time they should become more pronounced, because, as they become more widely accepted,

their impact on expected returns increases, making the data ever more consistent with the theory and, thereby, spreading acceptance of the theory. But some priced effects appear to disappear over time. For instance, there is evidence that, although the value effect persists, the size effect has disappeared in the last 20 years (Shumway and Warther, 1999). Given the volatility of stock prices, hypotheses such as the disappearance of the size effect are very difficult to confirm. Nonetheless, if this finding is true, the self-fulfilling prophecy theory offers no explanation.

Testing the validity of the self-fulfilling hypothesis is also a challenge. To take an example, assume that empirical studies find that a factor, say HML in the Fama–French model, is correlated with returns across securities. How is one to determine whether the relation exists because the factor is a true risk factor, whether it has been incorporated into the discount rate through the self-fulfilling mechanism, or both? At this juncture, we do not have an answer, but the question highlights the importance of investigating the methods by which major fundamental investors choose discount rates and, if possible, how those methods have evolved over time.

Finally, the theory is also devoid of normative content. It predicts that the cross-section of expected returns will reflect the risk factors fundamental investors incorporate into the discount rates, but it is silent on what risk factors *should* be incorporated.

8 Implications of the analysis and conclusions

Perhaps the most interesting and troubling implication of our analysis is that statistical artifacts unrelated to fundamental risk factors based on end-investor utility can be baked into the cross-section of expected returns. Take momentum as a possible example. Jegadeesh and Titman (1993)

reported finding evidence of short-term momentum in stock prices. This led Carhart (1997) to propose a four-factor model that added a momentum factor to the Fama–French model. To date this model remains controversial, and our preliminary empirical analysis did not find evidence that it is generally used as a tool for estimating discount rates. Nonetheless, if it were to become adopted widely, our theory predicts that momentum would become baked into the cross-section of expected returns even though there is no theoretical reason it should be there and even if the original finding were a spurious artifact of the estimation period.¹⁷

In this respect, the theory implies that it is possible for the market to get stuck in suboptimal path-dependent equilibria. If there were no cognitive constraints, information processing costs, or contracting costs, presumably market prices would reflect end-investor preferences as given by the consumption-based asset pricing model. Given the constraints, however, the observed ecosystem arises and the self-fulfilling prophecy comes into play. As a result, rather than reflecting end-investor preferences, the cross-section of expected returns mirrors fundamental investors' discount rates. The process by which those discount rates are chosen depends on history—for example, whether or not a momentum effect has come to be accepted. This means it is possible for discount rates to diverge from those that maximize end-investor expected utility. Of course, there is a limit to how far expected returns can diverge from the predictions of the consumption model. Large enough discrepancies would create an incentive to develop specialized funds. But, given the volatility of asset prices and the cognitive, information processing, and contracting constraints that exist, significant discrepancies may well be possible.

In conclusion, we argue that it is dangerous to rely on theoretical models that have counterfactual

aspects and which ignore the ecosystem of modern financial management.

Taking account of that ecosystem, we argue that the cross-section of expected returns is determined not by the margin utility of end investors, but by the discount rates that active fundamental managers use in the valuation analysis. This leads to the self-fulfilling prophecy of asset pricing described here. Because the possible existence of this self-fulfilling prophecy has important implications for asset pricing, it is critical to understand precisely how active fundamental investors estimate the discount rates used in their valuation models. It is a question that remains largely unexamined.

Notes

- ¹ As discussed further later, by active fundamental managers we mean managers who estimate the fundamental value of stock based on the present value of estimated future cash flows and attempt to profit from discrepancies between the estimated fundamental value and the market price.
- ² The investor's utility function can certainly include more than consumption of goods and services.
- ³ See Cochrane (2005) for the details of the derivation.
- ⁴ Other asset pricing models can be derived from the general consumption-based model by placing restrictions on the form of the stochastic discount factor.
- ⁵ The minus sign appears in the equation because consumption and the marginal utility of consumption are inversely related.
- ⁶ In a paper similar in spirit to ours, Adrian *et al.* (2014) develop a model in which the marginal utility of wealth of financial intermediaries replaces the marginal utility of consumption of end investors on the grounds that the intermediaries are the marginal investors.
- ⁷ Surveys of individual investors' expected returns also call into question the consumption-based model. Greenwood and Shleifer (2014) and Amromin and Sharpe (2012) find that investors' expectation of future returns are procyclical rather than countercyclical as predicted by the standard consumption-based models.
- ⁸ Many investment managers offer both fund products and segregated management service based on the same

investment strategies. Ultra-high-net-worth individuals often use segregated account management offered by tax-aware managers to perform tax loss harvesting to improve after-tax performance.

- ⁹ Shleifer (2000) provides an excellent summary of the early literature on behavioral finance.
- ¹⁰ See Chevalier and Ellison (1997).
- ¹¹ Phil Edwards, the European Director of Strategic Research at Mercer, reminded us that some managers do co-invest in their own funds. As such, they may be more aligned with the end investors and more interested in operationalizing Equations (1) to (4). However, even managers who co-invest are still likely far more motivated by outperforming their benchmark index.
- ¹² A very sophisticated investor could assemble a team of managers benchmarked to different indices such that the resulting portfolio optimizes his utility. However, we simply do not find this believable for the average, if not for most of the, end investors including the highly affluent ones. Even if some investors were sophisticated enough to operationalize Equations (1)–(4) by selecting the right mixture of managers instead of stocks, it is still possible that managers' own utility maximization leads to prices which are not set by the standard consumption model.
- ¹³ Erik Knutzen, former CEO of New England Pension Consulting, confirmed that "anyone who has attended a number of [pension] plan sponsor Investment Committee or Board meetings is painfully aware of [the lack of alignment]."
- ¹⁴ Most of the assets under management in equities are in traditional long-only active fundamental equity strategies. Long-short hedge funds based on quantitative factors or statistical arbitrages represent a small minority of the managed strategies as measured by assets.
- ¹⁵ Mike Sebastian, a managing director at Aon Hewitt Investment Consulting, pointed out to us that, of the 1,117 active institutional equity products listed in the eVestment Alliance database, 730 are self-identified as fundamental-based.
- ¹⁶ Certainly there are many types of institutional investors, outside of our scope, who compete in the equity market and move prices. For example, HFT (high frequency trading) hedge funds, while controlling very small assets, generate very high trading volume; they could have a significant price impact. Quant active managers using statistical signals are often neither fundamental nor relative valuation based. Inclusion of these managers certainly makes the analysis more complex. However, the qualitative intuition remains unaffected.

¹⁷ Given the vast amount of academic research many “effects” that have been discovered are almost certain to be artifacts and many more artifacts will likely be found in the future.

Acknowledgments

We would like to thank John Cochrane, Phil Edwards, Kenneth French, Robin Greenwood, David Hirshleifer, Rajiv Gokhale, Tim Hodgson, Erik Knutsen, Philip Lawton, Neil Rue, Richard Roll, Mike Sebastian, Andrei Shleifer, and John St. Hill for their helpful comments on an earlier version of this paper entitled “The Machinery of Asset Pricing”.

References

- Adrian, T., Etula, E., and Muir, T. (2014). “Financial Intermediaries and the Cross-section of Asset Returns,” *Journal of Finance* **69**, 2557–2596.
- Amromin, G. and Sharpe, S. A. (2012). “From the Horse’s Mouth: How Do Investor Expectations of Risk and Return Vary with Economic Conditions?” Working Paper 2012-08, Federal Reserve Bank of Chicago.
- Asquith, P., Mikhail, M., and Au, A. S. (2005). “The Information Content of Equity Analyst Reports,” *Journal of Financial Economics* **75**, 245–282.
- Berk, J. and Demarzo, P. (2013). *Corporate Finance*, 3rd edition. New York, NY: Prentice-Hall.
- Brealey, R. A., Myers, S. C., and Allen, F. (2011). *Principles of Corporate Finance*, 10th edition. New York, NY: McGraw-Hill, Irwin.
- Brotherson, W. T., Eades, K. M., Harris, R. S., and Higgins, R. C. (2013). “Best Practices in Estimating the Cost of Capital: An Update,” *Journal of Applied Finance* **23**, 15–33.
- Brotherson, W. T., Eades, K. M., Harris, R. S., and Higgins, R. C. (2014). “Company Valuation in Mergers and Acquisitions: How is Discounted Cash Flow Applied by Leading Practitioners?,” *Journal of Applied Finance* **24**, 43–51.
- Carhart, M. (1997). “On Persistence in Mutual Fund Performance,” *Journal of Finance* **57**, 52–82.
- Chevalier, J. and Ellison, G. (1997). “Risk Taking by Mutual Funds as a Response to Incentives,” *Journal of Political Economy* **105**, 1167–1200.
- Cochrane, J. H. (2005). *Asset Pricing*, 2nd edition. New York, NY: Princeton University Press.
- Damodaran, A. (2012). *Investment Valuation*. Hoboken, NJ: John Wiley & Sons.
- Duff and Phelps (2013). *Risk Premium Report 2013*. New York, NY: Duff and Phelps.
- Fama, E. F. and French, K. R. (1993). “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics* **33**, 3–56.
- French, K. (2008). “Presidential Address: The Cost of Active Investing,” *Journal of Finance* **63**, 1537–1573.
- Greenwood, R. and Shleifer, A. (2014). “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, forthcoming.
- Jegadeesh, N. and Titman, S. (1993). “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *Journal of Finance* **48**, 65–91.
- Koller, T., Goedhart, M., and Wessels, D. (2010). *Valuation*, 5th edition. New York, NY: Wiley Finance.
- Lusardi, A. and Mitchell, O. S. (2013). “The Economic Importance of Financial Literacy,” Working Paper, Wharton School.
- Morningstar (2013a). *Cost of Capital Yearbook*. Chicago, IL: Morningstar.
- Morningstar (2013b). *Ibbotson 2013 Yearbook: Stocks, Bonds, Bills, and Inflation*. Chicago, IL: Morningstar.
- Pratt, S. P. (2008). *Valuing a Business*. New York, NY: McGraw-Hill.
- Rosenbaum, J. and Pearl, J. (2013). *Investment Banking*, 2nd edition. New York, NY: Wiley Finance.
- Securities and Exchange Commission (2012). *Study Regarding Financial Literacy Among Investors*. Office of Investor Education and Advocacy of the U.S. Securities and Exchange Commission.
- Shleifer, A. (2000). *Inefficient Markets: An Introduction to Behavioral Finance*. New York, NY: Oxford University Press.
- Shumway, T. and Warther, V. A. (1999). “The Delisting Bias in CRSP’s Nasdaq Data and Its Implications for the Size Effect,” *Journal of Finance* **54**, 2361–2379.
- Summers, L. H. (1985). “On Economics and Finance,” *Journal of Finance* **40**, 633–635.

Keywords: Behavioral finance; delegated agent model; asset management institutional features