Volume Title: How the Financial Crisis and Great Recession Affected Higher Education

Volume Author/Editor: Jeffrey R. Brown and Caroline M. Hoxby, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 978-0-226-20183-2 (cloth); 978-0-226-20197-9 (eISBN)

Volume URL: http://www.nber.org/books/brow12-2

Conference Date: September 27–28, 2012

Publication Date: December 2014

Chapter Title: Endowment Management Based on a Positive Model of the University

Chapter Author(s): Caroline M. Hoxby

Chapter URL: http://www.nber.org/chapters/c12856

Chapter pages in book: (p. 15 – 41)

1

Endowment Management Based on a Positive Model of the University

Caroline M. Hoxby

1.1 Introduction

In this chapter, I propose a positive model of the university in which many apparently peculiar features of universities—such as endowments and tuition subsidies—are generated by the internal logic of the model. The model proposes a specific objective function for universities and demonstrates how it is enforced. That is, the objective function leads to actions that reinforce the initial selection of the objective function. The model offers important predictions for the decisions that universities should make on many fronts if they are behaving in accordance with it. In this chapter, I focus on the implications for financial decisions, especially universities' endowment spending rules and portfolio allocations. The model is designed to explain America's great private research universities and very selective liberal arts colleges and—with modest adaptations—institutions that are partly

Caroline M. Hoxby is the Scott and Donya Bommer Professor in Economics at Stanford University, a senior fellow of the Hoover Institution, and a research associate and director of the Economics of Education Program of the National Bureau of Economic Research.

I am grateful to many people with whom I have had discussions that helped me refine ideas on this topic. I am similarly grateful to people who have written previously on this and related topics. These include William Bowen, Jeffrey Brown, Henry Hansmann, and Robert Merton. Eric Bettinger, Keith Brown, David Chambers, Stephen Dimmock, Elroy Dimson, William Goeztmann, Bridget Long, Sharon Oster, Cristian Tiu, Sarah Turner, and Scott Weisbenner gave me extremely helpful comments on an early outline of this chapter. This chapter draws very substantially on my 2011 Clarendon Lecture in Economics entitled "The Complex and Increasingly Global Market for Elite Higher Education," which is rewritten in the form of two chapters in the Clarendon Lectures monograph *The Role of Markets in Education* (Oxford University Press [OUP], under contract). I am grateful for the support of the OUP, especially Adam Swallow. I am solely responsible for the content of this chapter. For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see http://www.nber.org/chapters/c12856.ack.

controlled by the state such as America's and Britain's great public research universities. The model is designed to explain them precisely because many people find them hardest to explain, even enigmatic. An ancillary benefit of the model is that it provides a justification for the existence of all of the aforesaid institutions by assigning them a unique role in the creation of intellectual capital. A simpler version of the model could explain most other colleges and universities, and this is a topic to which I return briefly.

A very brief intuition into the model, laid out in section 1.3, is as follows. Each university maximizes its contribution to society's intellectual capital by making two types of investments: (a) investments in advanced human capital embodied in people and (b) investments in new knowledge embodied in research. These investments have peculiar properties that make them unlikely to be financed by conventional means. In making these investments, universities play a role that closely parallels that of a venture capitalist. They invest not only money, but also expertise and infrastructure. The university gets what is essentially an equity stake in its intellectual capital investments. Collecting these (equity-type) returns is a key difficulty for universities, and it is in overcoming this difficulty that universities (a) generate an endowment, not merely gifts and grants for current use; and (b) reinforce the initial choice of their objective function so that the model is, as a logical matter, closed.

This chapter draws heavily upon Hoxby (forthcoming) in that the model of the university is the same. However, those chapters present empirical evidence that justifies how the model characterizes universities. They also explain how universities evolved over decades into the institutions that they are modeled as being. In this chapter, I present the model without much empirical justification and without history. Hoxby also examines the implications of the model for universities' policies on tuition, admissions, and a variety of other outcomes. In this chapter, I touch on these policies briefly and focus on the implications for endowment management.

The papers to which this study is most closely related are Hansmann (1990) and Merton (1993). The Hansmann paper is a masterful review of many explanations of why universities have endowments and what each justification implies for endowment management. He carefully analyzes and rejects both the intergenerational equity model of Tobin (1974) and extensions of the intergenerational equity model that presume that there will be an increasing cost of education for succeeding generations. The latter type of model produces spending rules along the lines of the Stanford Rule (Merhling, Goldstein, and Sedlacek 2005). Hansmann also analyzes several partial explanations for endowments and their management rules—for instance, programming models—such as Grinold, Hopkins, and Massy (1978)—which do not maximize an objective function that values education and research but merely assess whether a rule can attain a metric such as maintaining a certain ratio of expenses to endowment. Hansmann even

carefully assesses casual explanations for endowments such as the character and incentives of people who serve as university trustees.

Following his review—which I do not imitate in this chapter, partly because his review is so masterful and partly for reasons of length—Hansmann concludes:

> The argument that has been offered most frequently in recent years to explain the accumulation of endowments—that they are a means to intergenerational equity—is unpersuasive. More compelling reasons to accumulate endowments are that they serve as a financial buffer against periods of financial adversity, that they help to insure the long-run survival of the institution's reputational capital, that they protect the institution's intellectual freedom, and that they assist in passing on values prized by the present generation. It is not clear, however, that these are today the reasons why endowments are accumulated. Nor is it clear that the sizes of existing endowments, and the ways in which they are managed, are well chosen to serve these goals. In particular, prevailing endowment spending rules seem inconsistent with most of these objectives. We cannot say, from the arguments and evidence surveyed here, that the endowments of the major private universities today are either too big or too small. It does appear, however, that surprisingly little thought has been devoted to the purposes for which endowments are maintained and that, as a consequence, their rate of accumulation and the pattern of spending from their income have been managed without much attention to the ultimate objectives of the institutions that hold them. (1990, 39–40)

In other words, Hansmann concludes that there is not a satisfying, positive model of universities in which endowments play a logical and necessary role. He does not himself propose one.

In many ways, Merton takes off where Hansmann ends. Indeed, his paper begins with part of the above quotation from Hansmann. Merton's key contributions are (a) to focus attention on the nonendowment sources of cash flow for universities (tuition, gifts, and so on), and (b) to solve the problem of how to manage an endowment given these nonendowment sources of cash flow and flows of expenses. Merton treats the cash flow and expenses as exogenous because "that abstraction does not significantly alter the optimal portfolio allocations." Merton's article is a great contribution that is not only useful to the broad scholarly community, but that is specifically helpful to this study because it allows me to focus on questions that he leaves unresolved.[1]

I focus on what the university's objective is; why that objective generates sources of cash flow like gifts, tuition, and grants; why endowments exist;

1. I believe that this is a fair characterization of Merton's paper because he writes, "The course taken here is to address this question in the middle range: it does not attempt to specify . . . the objective function for the university, but it does derive optimal investment and expenditure for [the] endowment in a context which takes account of overall university objectives and the availability of other sources of revenue besides endowment" (1993, 212).

and how the endowment and cash flows are interdependent. Essentially, my model sets up the structural version of the university's investment problem. For each problem there exists a reduced-form version that can be written strictly in terms of cash flows. This reduced-form problem is solved in Merton's paper—except for some issues that I take up later.

In addition to the closely related papers by Hansmann and Merton, there is a rich but less related literature on the management of endowment portfolios (the types of securities in which they should invest) and on spending rules (the amount of the annual return on the endowment that should be spent rather than saved in the endowment).[2] Many of these papers provide important estimates of elasticities or other parameters—for instance, the sacrifice in average returns and decrease in risk that universities experience when they invest more of their portfolio in bonds. Or, to take another example, Brown, Dimmock, and Weisbenner (chapter 5, this volume) estimate the response of donations to the rate of return on the endowment. However, as noted by both Hansmann and Merton, this rich array of papers tend to take no stand on why endowments exist or the purpose they serve and instead rely—explicitly or implicitly—on a rule of thumb such as needing to maintain a perpetual level flow of expected real income. Thus, the body of papers referenced in this paragraph, although terribly useful given an objective function and reduced-form rules based on it, are logically posterior to this chapter.

This chapter contrasts with the small existing literature that does propose objective functions for universities—most notably prestige.[3] These papers strike me as unsatisfactory in that prestige is a fundamentally incomplete explanation. The authors never demonstrate why institutions should compete on the particular metrics of prestige—endowments and admissions thresholds—that the authors choose. They never explain why endowments or admissions thresholds exist in the first place so that they *could* be grounds for prestige competitions. Indeed, saying that universities compete for prestige seems merely to be conceding that there is no positive model of universities, with the result that universities must necessarily be said to pursue essentially superficial goals that happen to have predictive power.

The structure of the remainder of the chapter is as follows. In section 1.2, I review the venture capital problem because it clarifies important features of the university's investment problem. In section 1.3, I lay out a positive model of the university. Section 1.4 presents a less abstract, more down-

2. See, for example, the Advisory Committee on Endowment Management (1969); Black (1976); Donaldson, Lufkin, and Jenrette (1969); Ennis and Williamson (1976); Eisner (1974); Grinold, Hopkins, and Massy (1978); Litvack, Malkiel, and Quandt (1974); Malkiel and Firstenberg (1976); Tobin (1974); and Williamson (1975). More recent papers include Blume (2009); Brown et al. (2010); Brown, Dimmock, and Weisbenner (chapter 5, this volume); Brown, Garlappi, and Tiu (2010); Dimmock (2012); Lerner, Schoar, and Wang (2008); and Campbell (2012).

3. This literature is best exemplified by Ehrenberg (2000). A related but different and incomplete explanation is the ever-increasing demand for elite higher education (see Clotfelter 1996).

to-earth portrait of a university than is described by the model. In section 1.5, I describe the implications of the model for endowment spending and portfolio allocation. This is where Merton (1993) plays an important role. Finally, in section 1.6, I conclude by explaining what the model has to say about some questions that commonly arise.

1.2 The Venture Capital Problem

It may appear to be a digression to review the venture capital problem and how it is solved before moving to universities. However, it is efficient to conduct this review because it is easier to see the structure of the problem when it is presented in its more abstract form, without all of the details that often confound analysis of universities. Although economists who work on the venture capital problem usually lay it out in a fairly different way from the presentation here, the fundamental problem and aspects of the solution I present draw heavily upon this literature.[4]

1.2.1 Problematic Investment Projects

There exists a set of investment projects with the following characteristics:

- They are brought forward by a person (or persons) whose unique capacities are necessary for the success of the project. For instance, this person could be someone who has the idea for an invention. Hereafter, I call this person the "agent."[5]
- To be successful, the project requires assets—usually described as expertise and infrastructure—whose building or adjustment costs are convex in the per-period increase in their amount or diversity. This is just to say that it is expensive to build the relevant assets very quickly, especially if one does not already have a foundation of closely related assets. For instance, the project may require expertise in finance, marketing, the law, certain engineering, or some other body of knowledge. The infrastructure required may be the plant and equipment necessary to construct and test a prototype of the invention. What is important is that the project requires assets that have costs that are convex in the way described. Hereafter, I use the phrase "convex adjustment costs" to refer to such costs.
- Unless an assessor has expertise in areas related to the agent's unique capacity, his estimates of the potential profitability of the project will be extremely imprecise.

4. I rely particularly on Repullo and Suarez's (2004) elegant exposition. Other helpful papers include: Bergemann and Hege (1998), Berglof (1994), Casamatta (2003), Gompers (1995), Gorman and Sahlman (1989), Kaplan and Strömberg (2001, 2003), Sahlman (1990), and Trester (1998).

5. He is typically called the entrepreneur in the venture capital literature, but this conveys a false sense that the project is routinely a start-up business.

- The project is risky. This almost goes without saying since (a) it depends on the existence of the agent who could, say, be killed in an accident; and (b) something about the project must be advanced if only an expert can assess its potential profitability.

These conditions generate a problem. The agent cannot get financing for the project from a bank or raise money in some other conventional way because (a) the bank or other conventional financier has insufficient expertise to assess the project and cannot build this expertise without incurring prohibitive costs, and (b) the agent's costs of building the expertise and infrastructure himself are prohibitive. Even if the agent were able to get financing from a bank, he might substantially underfinance the project because he is risk averse and his investment would entail a great deal of undiversified risk.

1.2.2 The Potential for a Venture Capital Solution

Venture capitalists are a potential solution to this problem. They are modeled as people who have a pool of financial capital to invest and who have already built expertise and infrastructure in areas relevant to certain types of projects. A venture capitalist invests not only money but also expertise and infrastructure. Using his expertise, he selects projects whose expected rate of return exceeds a threshold based on the opportunity costs for his financial capital, expertise, and infrastructure. In return for making an investment, the venture capitalist gets what is essentially an equity position in the project (more on this below).

The venture capital problem exists precisely because the costs of building expertise and infrastructure are convex as described. This has a few implications. First, a venture capitalist's portfolio is necessarily limited in its breadth and size by the cost of building diverse expertise and infrastructure. His portfolio may contain many projects and be much more diversified than any one of these projects, but it will be far less diverse than, say, a market index fund. Second, the venture capitalist will want to finance a stream of projects over a number of years—not finance a single generation of projects and then get out of the business. This is because the costs of building expertise and infrastructure are convex in the speed of doing it. Thus, the venture capitalist's investments in expertise and infrastructure are more likely to pay off if he builds them gradually and employs them for several generations of projects. Third, the venture capitalist has strong incentives to maintain his expertise and infrastructure, as opposed to letting them depreciate. He will fear situations in which his returns from projects are insufficient to support maintenance-type investments in his expertise and infrastructure. For instance, a venture capitalist who was short on current income might lose his experts to other employers. If much of this were to occur, his ability to continue his venture capital business would be in doubt. Thus, venture capitalists may wish to have a smoother stream of investments than an inves-

tor in financial markets. Fourth, the cost structure implies that projects will crowd one another out at the margin. That is, two projects may have similar potential profitability and might each pass the venture capitalist's threshold if each could use his existing stock of expertise and infrastructure. However, neither project might justify the venture capitalist's quickly building enough new expertise and infrastructure to accommodate it. The potential for crowding out means that venture capitalists will acquire a good deal of information about an agent before making an offer to him.

The venture capital solution to the investment problem contains a two-sided moral hazard: the agent must make an effort and contribute his unique capacity but the venture capitalist must also contribute his expertise and infrastructure. This moral hazard has two important implications. First, the problems caused by moral hazard decrease as the project's potential profitability rises: both parties will be glad to make an effort when they realize that they are in a situation with vast potential profits. Thus, the optimal contract concentrates the returns to the venture capitalist in the best states of the world. For instance, a contract might specify that the venture capitalist receives no share of the returns if the profits are below some threshold, a certain share if the profits are in some intermediate range, and a higher share if profits are in a high range. Second, owing to the moral hazard, the optimal contract will make the agent fund some share of the project if he can. This is because the agent's moral hazard shrinks as he himself invests more: he will lose his own stake if he makes no effort. Thus, within a relevant range, the venture capitalist is willing to invest more when the agent invests more, and the entire project can be more optimally financed and managed.

A final thing worth mentioning is that venture capitalists can bring "outside money" to projects to improve their solution. That is, the venture capitalist may find himself in situations where he has sufficient expertise and infrastructure for a project but insufficient funds. In such cases, he may be able to get external financing for the project that the agent could never get on his own because (a) the venture capitalist is known to have the expertise to assess the project, and (b) the credibility of the venture capitalist's announced assessment is increased by his investing his own expertise, infrastructure, and money.

Summing up, an investment project of the type described is more likely to be optimally funded—that is, solved by venture capital—when:

- the expertise and infrastructure required are such that a venture capitalist can pay for the cost of building and maintaining them by earning returns on a sufficiently large and diverse portfolio of (somewhat related) projects;
- the venture capitalist does not so discount the future that he is unwilling to use his expertise and infrastructure over multiple generations of projects;

- the expected returns to the project are sufficient to clear the threshold set by a venture capitalist's opportunity costs of financial capital, expertise, and infrastructure; and
- reasonably optimal contracts are possible. These will entail the venture capitalist taking a sufficiently small equity share that the agent's moral hazard problem is not crippling. The agent will typically make some contribution (even if it is only his time and effort), and the venture capitalist will typically hold a nonlinear equity position in the project.

1.3 The University's Problem

In this section, I create a convenient fiction, the "founder," to embody the people whose assets and actions have shaped what universities have become. While actual universities have often developed what they do and how they finance it through a combination of vision, trial, muddle, competition, and imitation, the founder is blessed with extremely clear vision, rational expectations, and well-defined objectives. Obviously, the founder is not meant to be taken literally, but he makes it easier to lay out a positive model of the university because he is the prime mover and has an abstract understanding of the problem he is trying to solve. It is by no means necessary for actual university leaders to have been so concerted or to articulate similarly abstract models of their universities—just as it is not necessary for actual corporate leaders to manage or speak about their firms in highly abstract ways.[6]

The founder's goal is to use his initial capital, acquired through some exogenous means, to make the maximum contribution to society's intellectual capital, each unit of which he values at its social (private plus public) return. It is important that the founder wants to add intellectual capital, not merely fund intellectual capital that society would readily create in his absence. For instance, if a person's optimal education is such that he could finance it on his own using conventional means, that person is not of interest to the founder. In practice, this puts aside a good share of human capital investments that earn private returns, are not terribly risky, do not require

6. Although the founder is probably best thought of as an abstract person or persons, there are people who have founded universities recently enough for us to observe their motives and actions: Leland and Jane Stanford (founders of Stanford University); Andrew Carnegie and Richard and Andrew Mellon (founders of the two parts of what is now Carnegie-Mellon University); Peter and Hannah Widener (founders of Widener University); John D. Rockefeller (founder of Rockefeller University); and so on. All of these people took funds that they had earned in enterprises that were privately profitable—railroads, steel, chemicals, shipping, oil—and dedicated them to society's good by founding a university. Their first moves were hiring faculty from other universities, accumulating libraries, building laboratories, and otherwise creating expertise and infrastructure. Of course, great private universities like Harvard and Yale also had founders, but they are sufficiently "historical" for their motives to be less easily characterized. This is not to say that these universities do not fit the model, but that explaining them requires a more historical approach.

great expertise for a loan officer to assess, and do not require highly expert instructors or infrastructure with convex adjustment costs.

Rather, the founder takes as his problem some part of the market for intellectual capital that will be missing without his intervention. This is parallel to the venture capitalist taking on some set of investments that will not occur if they are constrained to use conventional financing. Advanced intellectual capital has properties that make investments in it likely to be missing or undersized.

1.3.1 The Peculiar Properties of Advanced Intellectual Capital

I define intellectual capital as knowledge that generates welfare for society in any way—private earnings, private nonmonetary benefits, public benefits, and so on. Intellectual capital takes one of two forms: human capital and research, which I differentiate for clarity, as follows. Human capital is knowledge that is stored in a person's brain and that cannot be used without his making complementary effort. Human capital need not be new knowledge: a engineer who knows how to solve a certain well-understood type of problem has human capital. Research is a contribution to the stock of world knowledge that can be embodied in some form—a publication, for instance—that it is independent of any particular person. This is not to say that research generates returns without some person's effort. A publication that contains new knowledge does not generate returns unless someone interprets and employs it.

Creating advanced intellectual capital generally requires three factors: human capital, research, and infrastructure. It is hard to get to the frontier of knowledge without benefitting from previous knowledge in the form of research and/or the human capital of instructors, who help people learn how to interpret research for themselves and who transfer the knowledge in their own brains. It is hard to conduct research without infrastructure such as laboratories in which experiments are conducted, libraries in which research is stored, or classrooms in which people with human capital instruct others. Moreover, these three factors—human capital, research, and infrastructure—interact in complex ways so that, for instance, there is little point in building a laboratory without the right complement of researchers to go in it. The result of all this is that it is hard to build the capacity to create advanced intellectual capital "from scratch." Building this capacity has costs that are convex in the per-period increase in amount and in diversity.

Even if they have access to the same human capital, research, and infrastructure, people vary in their capacity to store human capital, generate research, and earn returns on knowledge to which they have access. To store, create, or use advanced intellectual capital, a person may need to be intelligent, have original ideas, or have the will to exercise human capital or develop research into useful products. For convenience, I call this array of capacities "aptitude" and note that there is ample evidence that the distribu-

tion of it has a long right-hand tail.[7] Thus, a person who has 99th percentile aptitude may generate returns that are an order of magnitude greater for a given investment in intellectual capital than a person with 90th percentile aptitude.

Estimating where a person is in the long right-hand tail of aptitude is a costly exercise. It generally requires the person's own time or effort, the time and effort of an assessor who himself has high aptitude, and the use of infrastructure. Furthermore, even a good estimate of aptitude is only an estimate: we expect aptitude to be revealed over much of a person's adult life. For instance, whether a person's contributions to knowledge are worthy of a Nobel Prize or only a solid academic post is something usually revealed after midlife. Similarly, whether a person will use his knowledge to create a Fortune 500 company or to conduct a very successful but largely unremarkable career is usually revealed slowly.

Because aptitude varies and is revealed slowly, an individual should be risk averse when making human capital investments in himself. If the person knows that his aptitude is at least fairly high, then the risk of such investments is heightened by the long right-hand tail of aptitude, the fact that his optimal investment may be very large relative to average income, and the fact that completely unrelated risks (a debilitating accident, for instance) may destroy his ability to exercise the advanced intellectual capital in which he has invested.

The prohibition of slavery or, more prosaically, forced labor has the consequence that people cannot easily take equity positions in one another's advanced intellectual capital and thereby share risk. Similarly, even if a bank wished to invest in a portfolio of advanced intellectual capital (presumably by financing many persons' education and research), its inability to take equity positions would cause it to underinvest greatly: it could not earn more on persons whose returns happened to be very high than it could earn on persons whose returns were only moderate. A bank or person who attempted to hold a diverse, advanced intellectual capital portfolio would also find it challenging to acquire the expertise to assess each person's aptitude well.

Finally, it is hard to keep intellectual capital—especially research—fully private. It tends to get revealed in the process of being used. Thus, a good share of the returns to intellectual capital accrue to society, not to the people who produce them. The returns to human capital are typically far more private but still not fully private because colleagues of a person with advanced human capital typically acquire some of his knowledge through spillovers.[8]

7. "Aptitude" is *not* the same as IQ or any other narrow definition of ability. It includes forms of motivation, management and leadership skills, interpersonal skills, linguistic and rhetorical abilities, political skills, luck, and so on.

8. Of course, some individuals may have human capital that they use mainly for private returns and may also produce research that earns mainly public returns. A physician who does both clinical and research work would be an obvious example.

1.3.2 The Founder's Problem

If he is to maximize his contribution to society's intellectual capital, the founder must invest in such a way that he deals with the peculiar properties of advanced intellectual capital, as described above. In particular, he cannot add much to society's intellectual capital all by himself. He needs people who can store and exercise human capital. He also needs people who can create new knowledge. He needs the expertise to assess which people have the highest aptitude. He needs to combine the people with expertise and infrastructure. His expertise and infrastructure assets must be sufficient in breadth and scale to support a portfolio of intellectual capital investments with some diversity in risks and opportunities. Since the costs of building expertise and infrastructure are convex in the per-period increase in amount and diversity, he will want to keep his investments going for some time, not make a one-off investment in intellectual capital and then get out of the activity.

In all these ways, the founder is in a position parallel to the venture capitalist. We can give a name to his two types of projects: (a) students, people who have the capacity to store and use human capital as alumni; and (b) researchers, people who have the capacity to invent new knowledge. (To be clear, the founder of a university would have hundreds or thousands of projects, not two projects called students and researchers.) High-aptitude students are parallel to agents with projects because, like them, they have a unique capacity (aptitude) to acquire and exercise human capital that the founder cannot do without. Even the most able student requires instruction and infrastructure, and his human capital pays out over a long career in which his aptitude to generate social returns with it is slowly revealed. Even if they are able to capture all of the returns to their human capital (none are public), high-aptitude students should be risk averse about making very large investments in their own human capital and will find it hard to get external financing for such investments without being vetted by an expert. Research projects are also parallel to agents except that their problem is aggravated by the more public nature of research. A researcher has a unique capacity to invent knowledge. He typically requires an array of infrastructure and other experts to do so. The social benefits of the knowledge he creates are likely to be revealed slowly and he will probably find it hard to capture more than small share of them. Thus, he will underinvest and be risk averse if asked to make large investments by himself. Getting external financing will be stymied by others' lack of ability to assess his ideas and by their (often justified) conviction that most of the returns will be public, not private.

1.3.3 The Founder's Solution

Since the problem is parallel to the venture capitalist's, the founder's solution should logically also be parallel. In particular, he would like to be in a situation where:

- He can build a foundation of experts and infrastructure useful for making investments in intellectual capital. This is a university.
- He can use the experts to assess which students and researchers have the greatest expected social returns to investments in their intellectual capital. His threshold rate of (social) return should be based on the opportunity costs associated with his funds and the infrastructure and expertise he has built. His threshold is embodied in a university's admissions and hiring standards.
- He can invest money, expertise, and infrastructure in researchers and students who meet his threshold. His investment is the "tuition subsidy" or difference between the cost of a student's education and what he actually pays for it—as much as 80 percent on average in the case of very high aptitude students (Hoxby 2009).
- He can demand that students who benefit from his investment contribute whatever they are able to invest in their own human capital to reduce the moral hazard problem. This contribution is tuition, which is often need-based for high-aptitude students.[9]
- He can ensure that his experts and infrastructure will be employed over sufficient cohorts of students and researchers to justify their costs.

It is on the final point that the founder runs into difficulties that the venture capitalist does not have. The venture capitalist can expect to earn enough on one generation of projects to finance the next, thereby keeping his stock of expertise and infrastructure employed long enough to make it worthwhile to have built them. The founder cannot write explicit contracts with his students in which he takes equity stakes in their earnings since such contracts would entail illegal forms of forced labor. Also, unlike the venture capitalist who is concerned with purely private returns, the founder cannot write contracts in which he gets an equity stake in public returns to research or human capital.

This is where the founder's solution is, depending on how one thinks about it, clever or fortuitous. It is also where endowments become necessary.

Students

Consider what the founder would like to do with his students. He would like to convince them that it is their social and moral obligation to give the university the equivalent of what the venture capitalist would require from them. This is the amount that will maximize the intellectual returns generated by his initial capital. Because some of the students' returns will presumably be social, the obligation on them can be multifaceted. Those who earn private returns on their human capital can be asked for gifts in

9. Because much research earns only public returns, there is no obvious parallel for a researcher's contribution to reduce moral hazard—unless his research is likely to generate private returns from, say, a patent.

the form of money. Others could be asked for gifts in the form of expertise. Others could be asked to use their political influence to ensure that the university is treated well by the government. And so on. The founder will not couch their obligation as an obligation to provide him with returns so as to maximize the impact of his initial capital. He will instead couch it as an obligation to provide for the next generation of students as they were provided for themselves.

There are a variety of "soft" ways in which the founder could hope to convince students that this is their obligation. The university could use exhortation and inculcation. It could develop loyalty by creating athletic teams, clubs, traditions, songs, events, and special apparel. It is my belief that many of these soft methods work because the founder himself and the people whom he recruits to run the university *truly* believe in the obligation and take a *true* interest and pleasure in providing for the next generation.

Nevertheless, it will be easier to convince a student of his obligation if, when he donates, he knows that the university has tied its own hands in such a way that the gift will be worth much less if it is not actually used to invest in the human capital of future generations. This is why endowments make sense. Even if it has no restrictions other than that the principal must be preserved in perpetuity and spent for university activities (broadly construed), much of the value of a gift to the endowment is in the future. Therefore, a university that receives a gift as endowment has a greater incentive to exist, as a university, in the future. It is less likely to so exist if it uses its funds in some manner that does not generate solid returns to intellectual capital.

To see this, consider a university that, starting with some cohort of students, begins to instruct students in vacuous material that will not be valued by employers or society. Even if these students and their equally vacuously educated successors feel an obligation to donate a share of their earnings to the university, their donations will be small because they do not have valuable human capital. The university will become increasingly unable to maintain its infrastructure and expertise—losing experts to competing institutions, for instance. Eventually, the university will find it hard to attract students because its infrastructure and expertise are so depreciated, and at this point it will fold. The larger the share of its (prevacuity) gifts in the form of endowment, the faster the university will fold and more value it must sacrifice when it folds.[10] In short, by moving much of the value of a gift into the future, an endowment gives a university's leaders an immediate incentive to engage in investments that can sustain the university. These are

10. There are numerous colleges in the United States that have arguably suffered precisely the fate described. There are examples in which the curriculum truly became vacuous. However, a more common problem is curriculum not having kept pace with changes in what employers and the rest of society valued.

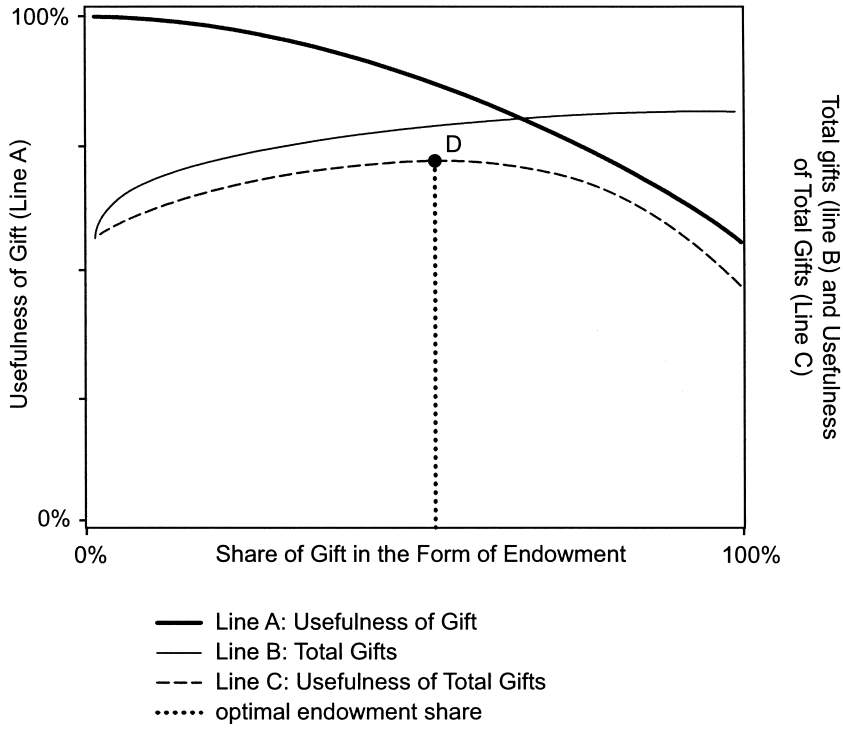

Fig. 1.1 Heuristic diagram of university's preferred endowment share

necessarily investments in intellectual capital that are *valued*, socially if not privately. (Below, I consider how a university might sustain investment in intellectual capital that earns only public returns.)

A simple figure (figure 1.1) illustrates the point that the university will accept gifts in the form of endowment up a point. Line A shows how the usefulness—for the purpose of maximizing the contribution to intellectual capital—of a dollar of gift falls the larger the share of it that arrives in the form of endowment. This line necessarily slopes down simply because the timing of the use of the endowment is relatively inflexible. It will slope more steeply if there are additional restrictions on how the endowment is used. Line B shows how total gifts increase as the share of gift dollars going into the endowment rise. The logic of the shape of this line is described in the previous paragraph, but it is not clear whether it merely flattens or actually falls as the endowment share gets close to 100 percent. This is not important for line C, which is the product of lines A and B: the usefulness of total gifts for contributing to intellectual capital. Line C peaks at a point like D, and this point determines the share of gifts that the university accepts in the form of endowments.

Researchers

Now consider researchers, the returns on whose intellectual capital are more likely to be public, not private. Of course, the university can explicitly contract with researchers who are likely to generate private returns using university-derived expertise or infrastructure. For instance, the university can share in patents. In such cases, the university should solve the problem no differently than a venture capitalist would. In addition, the university can oblige its researchers to devote part of their effort to instruction so that they produce human capital in students, like any other input at the university. In this case, we are back to the problem of students, with that problem's solution of obligations and implicit contracts. Indeed, the optimal ratio of researchers' instructional-to-research time depends on the next part of the founder's solution.

This is the interesting problem of maximizing the university's ability to generate research that earns only public returns. The founder could simply sink his investment into a few generations of research, but these investments would quickly peter out because the purely public returns to this research would not generate university funds for new ones. The solution is a semiimplicit/semiexplicit contract—this time with external funders—that is made more credible by the endowment.

The founder is not the only entity interested in knowledge that is a public good. There are other philanthropists who wish similarly to maximize their contribution to intellectual capital but who lack the wherewithal (or, perhaps more realistically, the timing) to build the expertise and infrastructure that would allow them to invest efficiently in intellectual capital. There are also philanthropists who wish only to contribute to intellectual capital in a narrow area and who therefore could not support expertise and infrastructure assets on an efficient scale. Finally, there is the government, which—to the extent that it maximizes social welfare—should internalize many of the public returns to research as though they were private.

To maximize his contribution to research that generates only public returns, the founder can use his initial funds to build expertise and infrastructure assets and to sponsor some first generation of research as a demonstration. He can then seek outside philanthropic and government funds to support research that can use the university's existing infrastructure and expertise or that only requires the university to build infrastructure and expertise in a cost-efficient way—that is, gradually and in a manner related to existing infrastructure and expertise. This is exactly parallel to the venture capitalist's problem. Recall, however, that the venture capitalist could credibly vouch for a project because he would have his own contract with the agent and would also stand to lose if his equity share were worthless. It is not so easy for a university to credibly vouch for research that is projected to earn purely public returns. Of course, the provisions of a grant can ensure

that, as a formal matter, every dollar is spent as the funder intends, but an outside funder should have more confidence in the research if the founder, with his expertise, has also put his own investment at stake.

This is where it will help to encourage some share of gifts for research to be in the form of endowment. It will be easier to convince outside funders that their money will be used for the intended research if the university has tied its own hands by making some research gifts less useful if the university does not, in fact, build and maintain infrastructure and expertise in the area in question over a long period. That is, the benefit of having some gifts for research go to the endowment—as opposed to current use—is that the university thereby increases the total outside funds (including government funds) for research. This benefit must be traded-off against the costs of the restrictions imposed by the endowment rules.[11]

In short, maximizing his contribution to intellectual capital through research should cause the founder to solicit an optimal share of research-related gifts in the form of endowment. The implied figure would parallel to that for student gifts in form, although—of course—the parameters of the relevant curves would differ depending on the inferences outside funders draw from the university's restricting itself via the endowment.

1.4 A Portrait of the University

The university's problem, as presented in the previous section, is so abstract that it can be hard to see how it matches the universities that we know from our own experience. What would this university, which maximizes its contribution to society's intellectual capital by acting as a social venture capitalist, actually look like?

The instigating financial capital of the universities' founders—who undoubtedly would have been less prescient than the fictional founder—would be unimportant, having long ago been dwarfed or replaced by the gifts of former students complying with their implicit contracts. Some alumni gifts would go to endowment and some to current use. The ratio would be a (surely not explicit) function of the degree to which the endowment nature of gifts increased the amount of gifts. Alumni who earned little (perhaps because they worked in public service) would not give or would give in some nonmonetary form: time, expertise, and so on. Above some earnings threshold, alumni would tend to give a share of their lifetime returns that was either steady or increasing in the scale of those returns. Thus, if the distribution of lifetime returns among alumni had a long right-hand tail, so would the distribution of donations within a cohort. (Note the word *returns*, as opposed

11. Interestingly, this is an argument for tenure of researchers that is quite independent of other arguments for and against. If the university wants to increase the credibility of its commitment to research in some area, tenuring researchers in that area is an effective way to tie its own hands to supporting that area for some time.

to earnings or wealth. An alumnus who was born rich, for instance, might have modest returns on the investment that the university made in him, even if he had high earnings.)

Second, research would be supported by a few means. The university would have venture capital-like contracts with its researchers who use its assets to produce patentable or otherwise private return-generating research. Faculty time would often be split between instruction (human capital production) and research since the former activity would provide reliable returns through students. Finally, research that generated popular public goods would be supported by a combination of the university's assets (infrastructure and experts) and outside funders' money. Some gifts for research would be in the form of endowment, with the ratio again being dictated by the degree to which the endowment nature of funds increased the amount of outside funds.

The infrastructure and expertise of the university would continually, gradually drift from the areas that earned high social returns in the past to areas that were likely to earn them in the future. This would not necessarily be intentional. Rather, it would be an organic response to the project (student and research) opportunities that arose. For instance, the university might routinely accept grants that required it to add incrementally to infrastructure and experts, and the gradual piling up of these increments would slowly change the nature of these key university assets.

The higher the aptitude of the students at a university, the greater would be the investment in their education. Students would be asked to pay for part of the investment, according to their ability to do so, to reduce their moral hazard. Unless the highest-aptitude students were systematically much more able to pay than their slightly lower-aptitude peers (and so on down the aptitude spectrum), the universities that enrolled the highest-aptitude students would also subsidize the current cost of their education the most. These same universities would thus be most reliant on inculcating an sense of obligation to the next generation of students and on supporting the credibility of their implicit contracts through thoughtful use of the endowment.

A few caveats are in order here. First, it seems unlikely that a university could create a sense of loyalty or obligation if it invested substantially different amounts in different students. Thus, universities are likely to be fundamentally constrained in how heterogeneous their admissions can be.[12] Second, an alumnus's sense of obligation may be reduced if the next generation of students looks too dissimilar from his own cohort. Thus, the university may be constrained to alter the character of its student population gradually. Third, it may be easier to instill a sense of obligation in a

12. That is, a university might find two students who have the same rate of return on human capital investments but whose optimal size of investments differed substantially. If the university were constrained to make similar investments in all of its students, it could not admit both students even if both students could, in fact, benefit from the same university assets.

student if he is surrounded by a critical share of students who already have it. Thus, a university may find it difficult to move wholesale into new campuses at which none of the students already have the tradition of obligation. Fourth, nothing that has been said constrains a university to contribute intellectual capital through *both* human capital and research. Indeed, very selective liberal arts colleges fit the university model except that the research components do not apply.[13]

It is important to observe that the university modeled here might make investments in intellectual capital that simply do not get made in countries in which no one has solved the founder's problem in the founder's way. That is, there is no reason to think that every country will develop very selective private universities that make investments in human capital on the same per-person scale as America's top private universities do. If some mechanism that mimics the founder's solution does not arise, some intellectual capital investments will probably just not be made because conventional financing would not work. Of course, in theory, it would be possible for a government to mimic the founder's solution. However, governments have political constraints that private universities do not.

1.5 Implications for Endowment Management

It will be observed that the model implies that university endowments came into being for fairly simple reasons. If the object of the university is to maximize its contributions to intellectual capital, its social venture capital activities are its reason for existence. The endowment arises simply because returns must enter an investment pool before being reinvested and the university is able to make greater investments in intellectual capital if it accepts some returns (and some gifts to support research) in the form of endowment. Thus, the main thing that is true of the endowment is the obvious thing: it represents a series of commitments that the university has made to invest funds in a set of intellectual capital projects. Importantly, the endowment has no independent purpose or objective of its own. This is why it should *not* be possible to define an objective for the endowment that is independent of the other side of the university's portfolio—the social venture capital side.

1.5.1 The University's Investment Problem, the Basic Version

Consider a university that is deciding what to do with its available investment pool. This pool consists of all current-use funds, all endowment income, and other funds that the university could legally use, including various forms of decapitalization, and even borrowing. (In this subsection,

13. A handful of all-research universities appear to be sustainable (Rockefeller University, for instance), but they are narrowly focused. The fact that many universities combine human capital creation and research suggests that the activities enjoy substantial production synergies.

I ignore specific restrictions on the use of funds but will return to this problem.) The university's problem is fairly straightforward to describe, though not straightforward to solve in practice.

The university knows the current set of intellectual capital projects that it could pursue—that is, all of the students and research in which it could potentially invest in the current period. Let us suppose it has rational expectations about the intellectual capital projects that will arise in each future period. For each project, current or future, the university can assign probabilities to each possible scenario of expenses and returns. That is, the university can reduce each project to the market value and investment risk characteristics it would have if it were a traded asset: it has a certain expected return, variation around that expected return, and correlations between its return and those of other projects owing to exogenous events. For instance, some projects might be sensitive to the price of oil. Other projects might be sensitive to which party controls Congress.

Suppose, for a moment, that projects did not need to use the assets specific to the university—its infrastructure and expertise—that have the convex adjustment costs we have emphasized. That is, suppose that projects only required inputs (in addition to the student or researcher himself) that could be bought at a world market price. Then, the projects could be evaluated fairly independently. The university would compute the properties of each potential project as an asset and find the optimal path of current and future portfolios, knowing that funds would be available in each future period for projects in one of three ways: (a) there would be returns generated by the endowment principal, (b) there would be funds if some money currently in the investment pool (such as current endowment returns) were not invested in current intellectual capital projects but "saved" for future projects by being invested in financial markets, and (c) there would be returns from past intellectual capital projects or cash flow from outside funders who were committed to research projects. Then, the procedure for finding the optimal *financial market* portfolio would be what Merton outlines—except that his solution is framed quite differently because he does not think of projects/assets with expected investments flowing out and returns flowing in. He thinks of cash flows, including expenses, which he models as an exogenous, dynamic, stochastic process. However, for any given path of current and future intellectual capital projects, there is a reduced-form that can be expressed purely as cash flows. Thus, Merton's fundamental solution method carries over, although of course the solution would look different with cash flows that are more realistic than the processes he actually uses. Moreover, the insights he derives carry over as well, namely:

- The spending rate on the endowment falls out of the optimization process and depends on the social venture capital side of the portfolio. In fact, no rule that is independent of the intellectual capital side of the

portfolio makes sense. This excludes simple rules such as "4.5 percent of the last three years' rolling average of total endowment market value."
- The optimal portfolio allocation for the endowment will be such that the university's financial assets have risks and returns that offset those of the university's intellectual capital portfolio. For instance, if the university's intellectual capital is disproportionately in projects that depend on demand for computer software, its financial assets ought to diversify its total portfolio away from that industry.
- An important function of endowment investments is to hedge against unanticipated changes in the costs of university activities. For instance, a university might invest some of its financial portfolio in real estate near campus that could be used for faculty housing.[14]

1.5.2 Incorporating Restrictions on the Uses of Funds

In fact, a university's pool of available investment funds is not unrestricted as to uses. Sizeable parts of it are restricted, and returns and cash flows associated with current intellectual capital projects are often restricted as well. Flows from outside funders are often highly restricted, but alumni sometimes place restrictions on their gifts as well.

Thus, instead of solving a single investment problem as described above, in the model the university instead solves a series of investment problems, sequentially, as follows:

- Begin with the most restricted investment pools and the potential projects that qualify for them. Solve for the optimal current and future portfolio of these projects. Remove the projects that get funded from consideration. Remove the funds that are invested in these projects from the pool that could be invested in financial markets instead of the intellectual capital side of the portfolio.
- Proceed to the next most restricted investment pools and consider all of the remaining potential projects that qualify for them. Again, solve and then remove the relevant projects and funds from further consideration.
- Continue sequentially until the unrestricted investment pool is reached, at which time all remaining potential projects are considered. Solve the problem as described in the last section, noting that all investment pool funds that were not committed (that is, those from all of the restricted pools as well as those from the unrestricted pool, if any) should be invested in the financial market side of the portfolio to attain goals like diversification and cost hedging.

Although the sequential process may appear at first to be onerous, in fact it should reduce the scale of the problem at the final stages greatly because

14. This is an example drawn from Merton (1993).

the university will need to consider a smaller number of potential projects for unrestricted funds.

1.5.3 Adding Adjustment Costs to the University's Problem

Unfortunately, if we make the problem more realistic by allowing for convex costs of adjusting the university's expertise and infrastructure, the problem could go from being tractable (though certainly demanding) to intractable. I first describe the problem broadly, at which point it should appear intractable, and then discuss why it might be fairly tractable in practice.

Suppose that the university fully understands the costs of adjusting its expertise and infrastructure.[15] Then, when it considers every potential intellectual capital project, it must consider the costs of adjusting its expertise and infrastructure to serve the project. Those adjustment costs will, in turn, depend on which other projects are being implemented concurrently and which projects were done in the past. If we consider all of the ways in which projects' use of expertise and infrastructure could interact, the number of portfolios that we should consider would grow explosively large. That is, the state space would explode.

To make this concrete, suppose that the same faculty member can either engage in important instruction or lead an important research project. If the university considers the research project first, it might allocate her to that project and invest in it. It might then find that the instruction was a poor investment because it would require the university to hire someone else just like her, and the cost of making this hire might be very expensive. If the university had considered the same projects in the other order (instruction first, research second), it might have come to a different conclusion—for instance, investing in the instruction but not the research. Moreover, although the difference in these two decisions might seem small, they would not be because whichever project gets done would not only occupy the faculty member's time but would also absorb some other infrastructure and expertise—which would have domino effects on other projects. These other affected projects would affect yet more projects, and so on.

For another example, consider a university that foresees that some area of research that is not productive now will possibly be highly productive in twenty years—if the university has certain expertise and infrastructure by then. Then, the university must consider all of the possible paths:

- starting projects now that are unproductive but that build some infrastructure and expertise in the relevant areas. This will make related future projects more attractive, leading to more infrastructure and expertise being built;

15. Of course, it might not understand its adjustment costs fully, but allowing for this would merely add a layer of difficulty without adding clarity.

- starting projects at some intermediate time (ten years, say) and incurring greater adjustment costs then, in the hope that the ten-year-out projects are sufficiently productive to justify the costs, especially when the twenty-year-out projects' dependence on that infrastructure and expertise is considered; and
- waiting for twenty years to start the relevant projects and incurring massive adjustment costs then.

Each of these paths will have domino effects on other potential projects that overlap in time and use of the university's infrastructure and expertise.

Thus, with adjustment costs, the investment portfolios that should be considered would appear to be endless in number. I would argue, however, that the situation is not so dire because: (a) many of tomorrow's projects are extremely similar to today's projects so that the number of future contingencies is much smaller than it appears, and (b) although universities' per-period adjustment costs at the extensive margin are large and highly convex, its adjustment costs at intensive margins are small.

The first of these points is easy to see: next year's cohort of potential students is so much like this year's cohort of potential students that they are nearly interchangeable. Thus, it is *possible* that a biology course developed for this year's cohort of students (because it was a good investment in them) might be a bad investment for next year's cohort of students. However, possible is not likely. Similarly, a good share of the "new" research ideas that arise each year are, in fact, very closely related to the "old" research ideas that arose in the previous year.

The second of these points—extensive versus intensive margin adjustment costs—means that it is costly for universities to quickly acquire expertise and infrastructure that forces them to grow (in size or area of expertise), but not costly for universities to transform existing expertise and infrastructure among uses that are already in their project portfolios. For instance, if a university decided to double the size of its entering class overnight, this would entail tremendous adjustment costs owing to the need to build student housing, increase faculty, and so on very quickly. In contrast, if a university decided to stop admitting classes that were one-third humanities majors, one-third sciences, and one-third social sciences and start admitting classes that were two-thirds humanities, one-sixth sciences, and one-sixth social sciences, this would generate costs but they would not be nearly so dramatic. Some libraries, classrooms, housing, administrators, student support staff, and so on would flow fairly easily from, say, social science to humanities use. Thus, the university should treat as similar various project portfolios that have the same implications for the extensive margins of adjustment, even if they have different implications for the intensive margins of adjustment.

In short, including adjustment costs in the university's investment problem does make the problem more difficult, but perhaps not absolutely

intractable. A practical person solving the university's investment problem might do well to focus only on the adjustment cost implications of (a) new projects that are very unlike existing projects, and (b) projects that involve building substantial expertise or infrastructure at the extensive margin.

1.6 Conclusions

The proposed, positive model of the university has implications for numerous questions, and many of the more abstract questions have already been addressed. For instance, section 1.5 described how the spending rule—which probably should be called a "function" rather than a "rule"—would arise in this model. It also described how optimal portfolio allocations should be determined by a university following this model—on both the intellectual capital and the financial capital side. But these answers, which include the insights first derived by Merton, are very abstract. It may therefore be useful to answer a few, less abstract but related questions.

In the context of the model, what would a university do if, for exogenous reasons, its financial market portfolio were temporarily earning a very high rate of return?

This is the circumstance in which a university thinks about keeping funds in the financial market portfolio that might, in an ordinary period, be invested in the intellectual capital side of the portfolio. The logic is that the university could make so much money on financial assets that reducing its current intellectual capital investments would make sense: it would be enabled, thereby, to increase future intellectual capital investments by so much that the university's total contribution to world intellectual capital would increase.

This logic is, indeed, what falls out of the model: current intellectual capital projects face a higher hurdle rate. *However*, the university's managers would take account of the fact that, if they restricted or reduced intellectual capital projects now, they would face even higher adjustment costs of implementing intellectual capital projects in the future, when the period of remarkable financial returns was over and they needed to start making intellectual capital investments at a higher rate than before.

A related question that often occurs to potential donors is whether they should donate now or later—because they believe that their *personal* rate of return will be high in the near future, relative to the university's. Of course, potential donors are not always correct that their personal rate of return will be higher than the university's. But, even if they are, potential donors probably do not take account of how they affect adjustment costs when they give in the future instead of now. In the context of the model, this is problematic. The model implies that universities would try to make potential donors internalize the adjustment costs they impose.

In the context of the model, what would a university do if, for exogenous reasons, considerable parts of its financial market portfolio were underwater and could not be liquidated except at very high cost?

This is another circumstance in which a university thinks about keeping funds in the financial market portfolio that might, in an ordinary period, be invested in the intellectual capital side of the portfolio. The logic is that it is very expensive, in terms of foregone future intellectual capital, to invest in intellectual capital because the cost of funds is the rate at which the institution can borrow.

This logic is, indeed, what falls out of the model: current intellectual capital projects would face a higher hurdle rate. Most obviously, the university might wish to shed expenses that are not associated with assets that have convex adjustment costs—certain staff who have fairly generic skills, for instance. *However*, the university's managers would again take account of the fact that, if they shed infrastructure or expertise, they might face high costs of reacquiring it in the future. In short, the model implies that universities will gain from developing a clear understanding of their adjustment costs, especially those that occur after downward adjustments in infrastructure and expertise (since these are not common).

In the context of the model, what would a university do if its financial market portfolio seemed always to be earning a substantially higher rate of return than its intellectual capital portfolio?

Because the universities do not actually attempt to compute the value of their intellectual capital projects as assets, it is quite possible that this question only appears to be common. Thus, one response to this question is that it should not be answered until the university has made the calculations to show that it is, in fact, a question.

However, assuming that it is a legitimate question, the model quite possibly implies "do nothing." In the model, the funds given to universities are there for the purpose of investing in intellectual capital, not running a financial business on the side. Thus, the only legitimate, alternative use of current-investment funds would be later investments in intellectual capital. It is possible that the intellectual capital side of the portfolio will always earn lower returns.

More specifically, the model implies that the university should assess its intellectual capital investments in a slightly different way, depending on whether they do or do not have related financial assets. Consider first the intellectual capital projects that should be related to financial market assets. For instance, if some firms are making tremendous profits, why are some of the university's alumni not leaders of or investors in these firms? If some inventors/entrepreneurs are making great profits, why does the university not have a stake in them, either through a patent or through having educated the inventors? Perhaps the university trains future managers and investors poorly. Perhaps the university's research in areas that generate private profits

is not conducted well. Perhaps the university is not pursuing patents in which it has a legitimate stake. Perhaps the university is not succeeding in convincing its students of their obligation to the next generation.

Now consider the intellectual capital investments that have scarcely related or no financial assets. The lack of related financial assets is prima facie evidence that the returns are largely public. However, even these investments can be subdivided. For some, external funders are in as good a position as university experts to evaluate the social return to the project—curing a disease that affects a population of poor people, say. For other projects, external funders may lack the expertise to evaluate the social return to the project—for instance, moving the frontier in an abstruse area of math or deciphering a papyrus that contains a lost Greek tragedy. If external funders are not funding a public returns-producing project that they can evaluate, the model implies that the university ought to improve its communications with them or reconsider the public value of the project. However, if the only people expert enough to understand the project's public value are experts who dwell in universities, then the model suggests that the university can do no better than consulting its own understanding and that of experts at peer institutions.[16]

In the context of the model, what would a university do if its financial market portfolio seemed always to be growing as a share of its total (financial plus intellectual capital) portfolio?

Again, because universities do not actually attempt to compute the value of their intellectual capital projects as assets, it is quite possible that this question is wrongheaded. This is another argument for making the calculations, even if they are imperfect.

Assuming that the question is not wrongheaded, however, the model implies that the financial side would only grow persistently as a share of the total portfolio if (a) the returns on future intellectual capital projects were *substantially* higher than those of today's intellectual capital projects, and (b) adjustment costs were such that when those future days arrive, the cost of suddenly needing to provide infrastructure and expertise to them would not be exorbitant.

Since circumstances (a) and (b) probably do not often arise in conjunction, the model implies that a university should examine itself if its financial market portfolio's share of its total portfolio rises very persistently. It seems most likely that this would occur if the university were not solving its investment problem correctly according to the model.

1.6.1 Key Implications of the Model

A few key implications of the model should now be evident. First, it implies that an optimizing university would actively attempt to compute

16. In the long run, there may be returns to universities from educating more people to understand the public value of projects of the type in question.

the values in its intellectual capital portfolios. This may be difficult, but the best available estimates would be better than none for the purposes of university and endowment management. Second, it implies that an optimizing university would attempt to calculate its adjustment costs. Knowing these costs is important for making every investment decision, whether on the intellectual or financial portfolio side. Third, the model implies that an optimizing university's leaders (the president, provost, and so on) would be in close communication with its endowment managers, with the specific goal of exchanging information on the value of existing intellectual capital investments, on returns to prospective investments in intellectual capital, and on adjustment costs. These exchanges should generate spending rates on the endowment that depend on and routinely adjust to circumstances on the social venture capital side of the portfolio. Overall, the model implies that highly endowed universities can make greater contributions to world intellectual capital—arguably, contributions that no other entities can make—if they manage both sides of their portfolios in a rigorous and coordinated manner.

References

Advisory Committee on Endowment Management. 1969. *Managing Educational Endowments: Report to the Ford Foundation*. New York: Ford Foundation.

Bergemann, D., and U. Hege. 1998. "Venture Capital Financing, Moral Hazard and Learning." *Journal of Banking and Finance* 22:703–35.

Berglof, E. 1994. "A Control Theory of Venture Capital Finance." *Journal of Law, Economics, and Organization* 10:247–67.

Black, Fischer. 1976. "The Investment Policy Spectrum: Individuals, Endowment Funds and Pension Funds." *Financial Analysts Journal* 32 (1): 23–31.

Blume, Marshall E. 2009. "Institutional Spending Rules and Asset Allocation." Wharton Working Paper no. 01-09, University of Pennsylvania.

Brown, Jeffrey, Stephen Dimmock, Jun-Koo Kang, and Scott Weisbenner. 2010. "How University Endowments Respond to Financial Market Shocks: Evidence and Implications." NBER Working Paper no. 15861, Cambridge, MA.

Brown, Keith, Lorenzo Garlappi, and Cristian Tiu. 2010. "Asset Allocation and Portfolio Performance: Evidence from University Endowment Funds." *Journal of Financial Markets* 13:268–94.

Campbell, John. 2012. "Investing and Spending: The Twin Challenges of University Endowment Management." In *Forum Futures*, edited by Maureen Devlin. Washington, DC: American Council on Education.

Casamatta, C. 2003. "Financing and Advising: Optimal Financial Contracts with Venture Capitalists." *Journal of Finance* 58:2059–86.

Clotfelter, Charles. 1996. *Buying the Best*. Princeton, NJ: Princeton University Press.

Dimmock, Stephen. 2012. "Background Risk and University Endowment Funds." *Review of Economics and Statistics* 94 (3): 789–99.

Donaldson, Lufkin, and Jenrette, Inc. 1969. *Managing Endowment Capital*. New York: Donaldson, Lufkin, and Jenrette.

Ehrenberg, Ronald G. 2000. *Tuition Rising: Why College Costs So Much*. Cambridge, MA: Harvard University Press.
Eisner, Robert. 1974. "Endowment Income, Capital Gains and Inflation Accounting: Discussion." *American Economic Review* 64:438–41.
Ennis, Richard, and J. Peter Williamson. 1976. *Spending Policy for Educational Endowments*. Wilton, CT: The Common Fund.
Gompers, Paul A. 1995. "Optimal Investment, Monitoring, and the Staging of Venture Capital." *Journal of Finance* 50:1461–89.
Gorman, M., and W. A. Sahlman. 1989. "What do Venture Capitalists Do?" *Journal of Business Venturing* 5:77–90.
Grinold, Richard, David Hopkins, and William E. Massy. 1978. "A Model for Long-Range University Budget Planning under Uncertainty." *Bell Journal of Economics* 9:396.
Hansmann, Henry. 1990. "Why do Universities Have Endowments?" *Journal of Legal Studies* 19 (1): 3–42.
Hoxby, Caroline M. 2009. "The Changing Selectivity of American Colleges." *Journal of Economic Perspectives* 23 (4): 95–118.
———. Forthcoming. *The Role of Markets in Education*. The 2011 Clarendon Lectures in Economics, Oxford University Press.
Kaplan, S. N., and P. Strömberg. 2001. "Venture Capitalists as Principals: Contracting, Screening, and Monitoring." *American Economic Review Papers and Proceedings* 91:426–30.
———. 2003. "Financial Contracting Theory Meets the Real World: An Empirical Analysis of Venture Capital Contracts." *Review of Economic Studies* 70:281–315.
Lerner, Joshua, Antoinette Schoar, and Jialan Wang. 2008. "Secrets of the Academy: The Drivers of University Endowment Success." *Journal of Economic Perspectives* 22 (3): 207–22.
Litvack, James M., Burton G. Malkiel, and Richard E. Quandt. 1974. "A Plan for the Definition of Endowment Income." *American Economic Review* 64:433.
Malkiel, Burton G., and Paul Firstenberg. 1976. *Managing Risk in an Uncertain Era: An Analysis for Endowed Institutions.* Princeton, NJ: Princeton University.
Merhling, Perry, Paul Goldstein, and Verne Sedlacek. 2005. *Endowment Spending, Goals, Rates, and Rules*. Forum for the Future of Higher Education.
Merton, Robert C. 1993. "Optimal Investment Strategies for University Endowment Funds." In *Studies of Supply and Demand in Higher Education*, edited by Charles T. Clotfelter and Michael Rothschild. Chicago: University of Chicago Press.
Repullo, Rafael, and Javier Suarez. 2004. "Venture Capital Finance: A Security Design Approach." *Review of Finance* 8:75–108.
Sahlman, W. A. 1990. "The Structure and Governance of Venture-Capital Organizations." *Journal of Financial Economics* 27:473–521.
Tobin, James. 1974. "What Is Permanent Endowment Income?" *American Economic Review* 64:427.
Trester, J. J. 1998. "Venture Capital Contracting under Asymmetric Information." *Journal of Banking and Finance* 22:675–99.
Williamson, J. Peter. 1975. "Background Paper." In *Funds for the Future: Report of the Twentieth Century Fund Task Force on College and University Endowment Policy*, edited by Peter J. Williamson. New York: McGraw-Hill.