Jeffrey R. Brown Discussion of Avery, Ehrenberg, Hill, and Weber

This version: July 31, 2024

It is a privilege to have the opportunity to discuss a paper by an incredibly talented set of authors on a topic that is of both intellectual and practical interest. This paper by Avery, Ehrenberg, Hill, and Weber certainly qualifies on all counts. They study endowment spending during the academic years 2008-09 through 2021-22 using Schedule D from the informational IRS Form 990. The authors do a nice job of explaining their data source's relative advantages (e.g., possibly less measurement error due to official tax filings) and disadvantages (e.g., exclusion of public institutions) to that of data used in prior studies (e.g., NACUBO).

In the early part of the paper, the authors clearly explain how endowments and quasi-endowments work, the constraints managers face, and provide useful summary statistics of endowment growth over their sample. The authors take a slightly different approach from prior research in that the authors estimate payout functions rather than examine empirical deviations from stated payout rules (as done, for example, in Brown et al 2014). Using a simple empirical model and limited assumptions, they approximate spending rules for 80 percent of the sample. Among the important results of this paper are that endowments grew considerably over the period of study, with growth rates exceeding that of the spending of the universities they support. The authors discuss the extent to which their findings could be consistent with various models of endowments.

Rather than provide a detailed critique of their data and methods, I am going to invoke the discussant's prerogative to speak more broadly about the topic. Part of my motivation is there is still important work to be done: there are still many unanswered questions about how endowments should and do operate, and this ought to be of special interest to those of us employed by universities that partly rely on endowments for funding.

A useful place to start evaluating the behavior of university endowments is to examine their mission statements, since these presumably are highly correlated with what economists would characterize as the objective function of an endowment. There is a noticeable similarity across a range of public and private institutions. For example, the Foundation supporting the University of Illinois (where I serve as a dean) says it exists "to advance the interests and welfare of the University." Yale University's massive endowment has a mission that is stated as "to best serve the university's mission and ensure its longevity." The mission of Harvard's endowment is "to support the educational and research goals of the University." Many other mission statements read similarly.

This focus on supporting, serving, or advancing the university suggests that to understand how to strategically manage an endowment, one should focus on the objective function of the university that the endowment serves. It turns out, however, that this is not so easy. There is no consensus in the literature on how to model what it is that universities seek to optimize, with possibilities ranging from expanding educational access to maximizing the prestige of the university, among other possibilities. Of course, universities can differ from one another in their relative priorities (e.g., consider a large public land-grant research university versus a small, elite liberal arts college). Even within the same institution, leaders might simultaneously pursue multiple objectives, such as expanding knowledge through research, delivering high quality education, and creating innovative

solutions to societal problems. These objectives can be difficult to measure. For example, rankings are widely considered a highly flawed measure of prestige, and research is measured by most universities based on expenditures rather than on the social value created. Additionally, different parts of the university often pursue very different objectives: the head of a research institute in the sciences might prioritize external research funding, a dean might emphasize academic program success, a department chair in social sciences might care about publications in a small set of elite journals, and an athletic director might care about football wins and losses. The relative priority of these objectives can vary in importance over time and under different leadership regimes. And, of course, there are a myriad of constraints – funding, legal, regulatory, political, societal, and so on – that make it difficult to characterize optimal behavior of a university.

I will assert from casual observation (and without formal evidence) that there exists considerable heterogeneity in the "effective" objective functions of universities. This heterogeneity makes it difficult to make precise predictions about how university endowments "ought" to behave with respect to their two key decisions: i) how to allocate their portfolio and ii) how to pay out income from the endowment to the university over time.

So, what is a researcher to do? One approach is to observe how universities manage payouts from their endowments over time and through different economic environments to see what we can infer about what they value. This paper is in that spirit, and it supports what earlier papers have also shown: that most endowments use some form of "smoothing" mechanism to determine how much to pay out to a university each year. This paper uses a flexible empirical model that captures both moving average rules (i.e., a target that calculates payouts to some number of years of prior endowment values) and hybrid rules (i.e., rules that take into account past endowment values but also current endowment spending) and estimates the parameters for each private university in their sample.

The predominance of smoothing functions shown by this and prior papers means that payouts are less volatile than the returns on the endowment's portfolio. Such smoothing makes intuitive sense insofar as the university is making long-term commitments based on expected payouts. For example, if a sub-account of an endowment is dedicated to funding a faculty position with a steady stream of expenses, it is reasonable that the university may want to have a lower-volatility revenue stream to fund it over time.

But are these smoothing rules really optimal? In my view, the preponderance of the evidence suggests that although they may be simple and convenient, they are unlikely to be optimal. Let us take the idea that universities desire stable revenue streams to support long-term commitments. If one narrowly frames the expenditure needs and thinks only in terms of a particular endowment sub-account funding a particular expenditure, then one might expect to see more "asset-liability matching," i.e., investing in a portfolio that de-risks the income used to fund the expenditure, such as by investing primarily in fixed income instruments. This is akin to an insurance company investing primarily in fixed income instruments to back annuity liabilities. Yet this is not what we observe: most university endowments have portfolios that are heavily invested in public and private equity, and many follow the Yale example of being extremely heavily invested in alternative asset classes. These alternative investments have high expected returns, but they also carry elevated risk, especially when one defines risk as the ability to match the desired payout stream.

One could also take a broader view of investing, which recognizes that the university's endowment is only one part of its overall revenue portfolio. Universities also receive funding from tuition, federal research grants, gifts, contracts, and public funds for state institutions, among many other sources. From the perspective of a university's larger portfolio of revenue, one might want to invest the endowment in a manner that considers the covariances of endowment returns with the variability of other sources. For example, federal research funding varies with changes in Presidential and Congressional leadership. Tuition, although set by the institution, can be subject to macroeconomic conditions and political or social pressures. Alumni giving varies with economic conditions (see, for example, Brown, Dimmock and Weisbenner, 2015). Athletic ticket and sponsorship revenue varies with team performance. And for public institutions, state budget appropriations are subject to political uncertainty. If one recognizes that income from the university endowment is only one part of an overall revenue portfolio, one might expect endowments to choose investments that partially hedge some of these other risks. Yet we observe little of using endowment investments to hedge other risks.

What is going on? Mission statements aside, it may be that endowments are not actually optimizing in pursuit of the objective function of the university. The likely reason is that principal-agent problems abound. As mentioned earlier, University objective functions are often difficult to measure, university and foundation presidents may have limited time horizons, foundations have their own boards who might have interests that do not perfectly align with the university's, it may be easier to communicate to donors in terms of return performance of the endowment than to communicate in terms of how well the endowment took into account covariances with other revenue streams, and so on. University and endowment presidents have an incentive to care about the growth of the endowment, *per se*, as this is an easy metric to which to point as evidence of success. (Indeed, a simple search of news articles and press releases about presidents and chancellors stepping down often point to their success at growing the endowment as a key accomplishment.)

Indeed, this was one of the insights of the work the Brown, et al (2014) paper in which we showed that universities have an asymmetric response to positive versus negative endowment shocks. In the face of significant negative shocks, endowments often reduce their subsequent payout amount by even more than their own smoothing formula would have indicated. We went on to show that this behavior was concentrated among the subset of universities for which the endowment was within ten percent of its value when the current president took office. This is consistent with a model in which university presidents are rewarded for endowment growth and thus are willing to take a deeper hit to their operational budget to maintain the value of the endowment, a phenomenon we called "endowment hoarding." Notably, the current paper is unable to further test this hoarding hypothesis in the face of negative shocks because the only negative returns in their sample occurred in the first and last period of their data, suggesting that future work might be appropriate to support or refute this hypothesis.

Overall, I still view the behavior of university endowments to be somewhat of a puzzle, both in terms of how they allocate their portfolios and in terms of how they choose their payout rules. More research is needed on this topic, as our understanding is still limited, especially in comparison to how important these questions are to many of the institutions that employ academic economists! I

welcome this additional work on the topic, and I genuinely hope that it spurs additional research going forward by these authors and others.

References:

Brown, Jeffrey R., Stephen G. Dimmock, Jun-Koo Kang, and Scott J. Weisbenner. 2014. "How University Endowments Respond to Financial Market Shocks: Evidence and Implications." *American Economic Review*, 104 (3): 931–62.

Brown, Jeffrey R., Stephen G. Dimmock, and Scott J. Weisbenner: 2015. “The Supply of and Demand for Charitable Donations to Higher Education.” In Brown, Jeffrey R. and Hoxby, Caroline M, editors, How the Financial Crisis and Great Recession Affected Higher Education. University of Chicago Press.