Should Universities Smooth Faculty Hiring? Theory and Evidence

Zhizhong Pu, *Harvard University*[*]

December 2025

Latest full version: https://zhizhongpu.github.io/research/

Abstract

U.S. universities often adjust operations sharply in response to financial shocks. Yet this practice may be costly when shocks affect peer institutions: contraction by peers reduces competition for key inputs, creating favorable conditions to expand rather than contract. I study whether universities could improve outcomes by smoothing faculty hiring over time. I first show that hiring responds to endowment revenue shocks. I then build models of faculty hiring and characterize the identified set of gains from adopting counterfactual, shock-independent hiring schemes. Applying my models to two decades of hiring data at top sociology departments, I find little evidence that alternative hiring schemes improve hire quality.

JEL Codes: G30, I22, I23

Keywords: Faculty Hiring, Higher Education Management, University Endowment

1 Introduction

U.S. universities often adjust operations sharply in response to financial shocks. For example, the Harvard Faculty of Arts and Sciences recently slashed Ph.D. admissions slots by at least 50% for the next two years (Mao and Paulus, 2025) in response to financial uncertainty; Stanford too paused staff hiring, following reduced revenue (Levin and Martinez, 2025).[1] Some adjustments are even reversed soon after. For example, in October 2025, University

[*]I thank Luis Armona, Christopher Avery, Christopher Campos, Catharine Hill, Tim Hua, Christopher Liao, Yujia Wang, Joshua Weidmann, Ruby Zhang, and participants at the HBS Researcher Seminar for helpful comments and suggestions. I thank Zhiyi Li for research assistance. I am especially indebted to Jesse M. Shapiro for guidance and mentorship. E-mail: zpu@hbs.edu.

[1]In a different direction, after receiving a $40 million gift, the University of South Florida aims to boost its new AI college's student enrollment from 3,000 to 5,000 in 3 years (Rundle, 2025).

of Pennsylvania decided to increase Ph.D. admissions just eight months after slashing it by one-third (Chitirala and Ryan, 2025).

While an individual institution may justifiably adjust to financial shocks[2], it can be suboptimal when shocks affect peer institutions simultaneously: contractions by peers reduce competition for key inputs, creating opportunities to expand rather than contract. For example, financially unconstrained airlines capitalize on the low aircraft prices during market downturns to expand their fleets (Pulvino, 1998).[3] Similar contrarian opportunities may exist for universities: lab equipment may be cheaper to purchase when peers cut funding, and skilled technicians easier to recruit when peers freeze hiring.

Why do universities not adopt these contrarian strategies, and instead still react sharply to shocks? One hypothesis is that the gains from contrarian strategies are small.[4] But quantifying such gains is empirically challenging because many operational adjustments—staff cuts, deferred maintenance, program freezes—are hard to observe or consistently measure. Even when visible, such adjustments rarely map cleanly onto the university's core outputs, making it hard to evaluate an intertemporal allocation of resources. Partly due to these challenges, no paper to my knowledge has studied this topic.

One domain of university operations that overcomes these challenges is faculty hiring: hiring decisions are public, timestamped, and tied to measurable research outcomes central to universities' function. This feature makes it possible to evaluate the performance of a *hiring scheme*, that is, an allocation of open positions across years, using measures of hire quality at the university-year level. Moreover, studying faculty hiring is meaningful because faculty searches can be delayed or advanced with limited disruption, unlike critical functions like student health services or campus security.

In this paper, I study whether universities can hire better faculty if they decouple hiring from financial shocks. I begin by showing that faculty hiring reacts to a type of common and industry-wide shock: market value fluctuations of university endowments. I combine post-2000 endowment data from the National Association of College and University Business Officers (NACUBO) with university financials data from the Integrated Postsecondary

[2]For example, Dixit and Pindyck (1994) highlight the value of waiting for uncertainty to resolve before making irreversible investments.

[3]By contrast, financially constrained airlines do not exhibit this pattern. Relatedly, a large body of the finance literature studies similar topics, including the superior performance of contrarian stock investment strategies (Lakonishok et al., 1994)—see Shleifer and Vishny (2011) for a review.

[4]Another potential reason is credit constraint. While this explanation has not been systematically evaluated, top universities have large endowment funds and face low cost of borrowing. For example, The Economist (2025) writes, "they borrowed at rates many firms, governments and less illustrious rivals must envy. Even the priciest recent bond issuance by Harvard, Princeton and Yale trades at 0.34 percentage points above the Treasury curve—about the same as debt issued by Johnson & Johnson, one of only two American companies with a triple A credit rating."

Education Data System (IPEDS) after 2005. I construct an endowment shock by interacting a university's baseline reliance on endowment with its fluctuating endowment revenue. I estimate an event-study-style linear model on IPEDS faculty hiring and headcount data for private institutions among the U.S. News top 50 national universities.

I show that endowment shocks affect hiring: an increase in annual endowment returns of 5 percentage points leads to an increase in the number of new hires in that year by 2.97 log points, for a university with 20% of its revenue historically from endowment. My findings motivate understanding whether (and by how much) a university could gain from alternative hiring schemes that make the amount of hiring unrelated to endowment shocks.

To quantify such gains, I build a structural model of faculty hiring that treats hiring as solving a stable matching problem for a set of *departments* and job market *candidates*. In this model, departments are ranked and candidates each have an *ability*; each department receives a *payoff* equal to the sum of its hires' ability across years, and thus seeks candidates with high expected ability; candidates prefer higher-ranked departments. I consider the matching problem under three information regimes: first, departments observe, at the time of hiring, only those signals of ability that are observable to the econometrician; second, departments additionally observe private ability signals unobservable to the econometrician; third, departments have perfect knowledge of candidates' ability. For each information regime, I characterize the identified set of payoff gains from adopting counterfactual, shock-independent hiring schemes under the stable matching equilibrium.

My model implies that a payoff-maximizing department subject to a fixed hiring capacity (across years) could gain from a smoothed hiring scheme, and that the size of such gains depends on patterns of hiring by peer departments, as well as the distribution of candidate ability in each year. For instance, gains from smoothing are high for a department whose hiring pattern coincides with peers, and in settings where candidate abilities are varied rather than similar. Intuitively, a department benefits the most from smoothing when the marginal opportunities available in low-hiring years are meaningfully better than those available in high-hiring years, which in turn requires that feasible hires are different in predictable ways, and that peers forgo the hiring opportunities.

I complement the structural hiring model with two reduced-form models that depend on different assumptions. In the first model, the sum of ability in a department-year cohort is a linear function of the number of hires, hiring competition, and their interaction.[5] The second

[5] I model hiring competition in three ways. First, as the sum of hires at higher-ranked departments, which is justified by a setting in which department-propose serial dictatorship yields a stable matching. Second, as the sum of hires at all other departments, which is justified by a uniform prior. Third, as the choice probability in a random utility model of a candidate choosing an open position offered at that department-year, where the representative utility is a convex decreasing function in the department's ranking. This

model treats the hiring variables as endogenous, instrumenting with endowment shocks (for private universities) and state fiscal appropriation shocks (for publics, following Deming and Walters 2017). I again formulate payoff gains, but this time using estimated OLS and 2SLS model parameters.

I then apply the models to a researcher-level dataset of all assistant professor hires at top 21 sociology departments in the U.S. from 1991 to 2017 (Warren, 2019). While my models are applicable in various settings, the sociology setting has two advantages over other disciplines, especially in the natural sciences. First, the sociology postdoc market is significantly smaller (Huynh and Shauman, 2021), so we can expect the factual hires to approximate well the candidate pool from which departments could counterfactually hire—an implicit requirement of my structural model. Second, sociology suffers less from a potential concern that my key measure of researcher-level ability—publication records within six years after being hired—depends on research funding, which could then depend on the budget fluctuations that drive hiring.[6] Additionally, the Warren (2019) data is uniquely well-positioned for my analyses: it is complete in that it captures the universe of hires in a well-defined market of meaningful size; and it contains precise individual-level annual affiliation data critical for my counterfactual analyses, whereas imputed yearly affiliation measures, often used in other papers on academic production (e.g. Borjas and Doran, 2012; Lerner et al., 2024), would not suffice due to noise.

I supplement the hiring data with sociology program ranking data from U.S. News (1998; 2017) and researcher-level publication records from OpenAlex (Priem et al., 2022). I measure each hire's *ability* using their social science publication counts within six years after hiring.

For each department in the data, I consider two feasible counterfactual hiring schemes. The first smooths the department's hiring across years, and the second (henceforth *countercyclical*) permutes its allocation so that it hires more when peers hire less. Both schemes hold constant the departments' total (across years) number of hires. Under each of the factual and counterfactual hiring schemes, I apply my models to compute each department's payoff.

I find little evidence that counterfactual hiring schemes could significantly improve payoff. For example, smoothing hiring improves an average department's payoff by 1.6% while countercyclical hiring reduces payoff by 5.3%, according to my structural model under the information regime that departments observe only candidates' pre-hire publication counts. By contrast, the per-hire payoff at an average department in the bottom tercile of the sample is 18.5% lower than that in the upper tercile. Allowing departments to observe private ability

measure is justified in a candidate-propose deferred acceptance process because a weakly higher choice probability implies a weakly larger set of candidates who offer to the department in the first round, and thus a weakly higher lower bound of ability in the department's feasible choice set.

[6] Granted, the sociology setting comes with limitations, which I discuss in Section 4.4.

signals unobserved by the econometrician does not qualitatively change the results.

The lack of significant gains is not explained by that candidates' pre-hire publication counts are not informative of their ability. These variables are positively correlated so departments can use pre-hire publication counts to identify high-ability candidates. Indeed, allowing departments to perfectly observe candidates' ability barely expands the positive effect of smoothing (to 1.8%) while significantly worsening the negative effect of countercyclical hiring (to 11.2%), although this model has poor fit.

Instead, two factors help explain the lack of meaningful gains. The first is that the candidates do not substantially differ from each other in ability, except among the top candidates in each year. Consequently, alternative hiring schemes marginally change the payoff for mid- and lower-tier departments in the sample; for the top-tier departments, alternative hiring schemes matter more: smoothing leads to modest gains for some top departments (and modest losses for others), whereas counterfactual hiring often leads to losses. This pattern in the top-tier departments relates to the second factor: the distribution of ability across candidates hired in low-hiring years differs from that in high-hiring years. Fewer high-ability candidates are hired in low-hiring than in high hiring years, which implies that shifting hiring from high- to low-hiring years actually *undermines* the top departments' prospect of hiring top candidates.

The principal contribution of this paper is that it is the first to study the consequences of universities' intertemporal adjustment of operations. Such adjustments often receive significant public attention and debates (Moody, 2025). In particular, systematic shock-dependent adjustments stemming from tying operational budget to endowment market value was criticized as early as in Hansmann (1990).[7] However, these spending rules are still in use and, as I show, affect real operations.[8] My structural model illustrates that such rules could be suboptimal in the domain of faculty hiring. I also contribute both empirical evidence from sociology hiring that justifies the current hiring scheme, and a theoretical framework that can be applied in other disciplines for further investigation.

This paper also relates to two other strands of literature. The first is the growing literature on the faculty labor market. While previous works focus largely on the individual- and institutional-level attributes that determine labor market outcomes of candidates (Fernandes et al., 2020; Wapman et al., 2022), I contribute additional evidence to the literature on

[7]Hansmann (1990) highlights that the endowment's growing size (and thus importance) undermines its role of smoothing universities' operations against shocks, under the (then and now) prevalent endowment spending rules that distribute a fixed fraction of the endowment's market value. Such a rule was "inconsistent with ... using the endowment as a financial buffer", in that it "commits an institution to using its operating budget as a buffer to absorb shocks to the market value of its endowment, rather than vice versa."

[8]For more on the effect of endowment, see e.g. Brown et al. (2014), Bulman (2022), and Avery et al. (2024).

financial determinants of market demand (Oyer, 2006; Turner, 2013). In particular, I extend the analysis of endowment shocks' effect by Brown et al. (2014) on faculty headcount to hiring, thus distinguishing hiring from firing.

The second is the small literature on contrarian hiring strategies (Greer et al., 2001), motivated by the fire sales literature (Shleifer and Vishny, 1992). These hiring strategies are not widely observed despite potential benefits in theory. A recent exception is Kim et al. (2025), which finds that small tech firms increase hiring after mass layoffs at larger firms, with suggestive evidence on benefits to firm-level innovation. I contribute a theoretical framework that shows that gains of these strategies can depend on the ability distribution in the candidate pool, and I contribute empirical evidence in a higher education setting.

The remainder of the paper is organized as follows. Section 2 describes institutional background and presents results on the effect of endowment fluctuations on hiring. Section 3 constructs my models with a numerical example illustrating the potential gains from alternative hiring schemes, and describes methods of quantifying such gains from data. Section 4 applies the models on sociology hiring data and discusses empirical results. Section 5 concludes.

2 Endowment Shocks and Faculty Hiring

Universities face various revenue shocks from private donations (Rundle, 2025) to government funding (Levin and Martinez, 2025). In this paper, I focus on a more prevalent and common source of shocks: market value fluctuations of universities' endowment.

2.1 Background

University endowment and spending rules

Over the past decades, top private universities increasingly depend on endowment as a source of funding (Brown et al., 2014; Bulman, 2022). For example, during the 2024-2025 fiscal year, income from endowment constitutes 37% of Harvard University's operating revenue. This number is even higher for departments less dependent on external research grants and student tuition, like the Faculty of Arts and Sciences (51%, Harvard University Financial Administration, 2025).

Since the early days of university endowments, scholars (and pundits) have been discussing the role of endowment and the optimal ways to spend it.[9] Universities formalize

[9] These discussions date back to at least Tobin (1974). For a review, see Brown et al. (2014).

endowment spending rules with the goal of balancing short-term priorities and long-term financial sustainability (Brown et al., 2014). The spending rules take on various forms, but they most generally apply a fixed percentage to a multi-year moving average of the real value of endowments balances (Avery et al., 2024).[10] While these spending rules may have started out with the goal of using endowment to buffer shocks from other revenue sources like tuition and government grants (Tobin, 1974; Black, 1976), by the 1990s endowment funds have grown so large that Hansmann (1990) remarks that the fractional-of–real-value spending rules are "inconsistent with ... using the endowment as a financial buffer", in that "[such] a rule commits an institution to using its operating budget as a buffer to absorb shocks to the market value of its endowment, rather than vice versa."

Budget process

Endowment spending fits into the broader picture of university budgeting. For private universities, the budget process follows an annual cycle beginning 12-18 months before the fiscal year start, which is usually July 1st, with planning occurring in fall/winter, development in winter/spring, and final approvals in late spring. The process typically involves multiple organizational levels, starting with individual departments submitting requests to their colleges or schools, which then forward consolidated proposals to university-wide administrators (provosts, chief financial officers, presidents), who ultimately present final budgets to boards of trustees for approval. Under this timeline, the endowment market-value shocks driven by capital market fluctuations tend to enter university budgets with a lag, generally one to two years after the shock occurs (Yale University, 2022; Avery et al., 2024).

Faculty hiring

The typical hiring process at top private research universities begins at the department level, where faculty identify staffing needs, and the department chair submits a request to the dean for authorization to hire a new faculty member. Deans and provosts make the final decision on whether to authorize and fund the position based on budget. Once authorized, the department forms a search committee that manages the recruitment process. The search committee typically posts job advertisements in fall, evaluates applications and conducts interviews. At the end of the evaluation process, the department ranks candidates and makes offers in spring, subject to final approval by the dean.

[10] The exact form differs across institutions. For example. Yale University (2023) states that "[the] payout ... is equal to 80% of the prior year's spending plus 20% of the long-term spending rate applied to the previous year's beginning endowment market value, with the sum adjusted for inflation."

The rest of this section studies the effect of endowment market value shocks on faculty hiring.

2.2 Data

I obtain data on end-of-fiscal-year[11] market value of endowment funds from annual surveys conducted by National Association of College and University Business Officers (NACUBO, 2023), the most common dataset in studies of endowment (e.g. Brown et al. 2014; Goetzmann and Oster 2015; Bulman 2022). The data cover almost all university endowment funds. I collect data from 2000 to 2023 (where 2023 refers to the 2022–2023 fiscal year).

I supplement this data with Integrated Postsecondary Education Data System (IPEDS, National Center for Education Statistics, 2005–2023) institution-year panel on financials. The key variables are the breakdown of revenue by sources like tuition and investment. IPEDS data also cover endowment market value for some years and some institutions, which I use to impute missing values from NACUBO data.[12]

For faculty hiring and headcount, I use IPEDS annual survey data, which cover faculty headcount by tenure status from 2005 to 2023 (where 2023 refers to the 2023-2024 academic year). IPEDS also contains data on the number of faculty hired in the fall (on payroll by November 1st, for years before 2017) or during the academic year (for years after 2018). I drop the pre-2010 data from my main analysis due to unsatisfactory data quality.[13] Appendix Figure 1 shows that IPEDS and official university announcements yield similar time trends in faculty hires for four universities.

2.3 Construction of Endowment Shocks

I construct a university-year panel of endowment shocks. I follow Deming and Walters (2017) to leverage the variation across universities in reliance on a particular funding source to identify the impact of shifts in that funding source. This endowment shock variable interacts changes in each university's investment revenue with the institution's *ex ante* share of revenue from endowment. Numerically, this shock variable is

$$z_{i,t} := \frac{R_{i,t}^{(e)}}{E_{i,t-1}} \times \frac{R_{i,pre}^{(e)}}{R_{i,pre}} \times 100 \tag{1}$$

[11] Fiscal year runs from July 1st to June 30th for most universities in my sample.

[12] For a limited number of institution-years, end-of-fiscal-year endowment market value is missing from NACUBO data, either due to omission in the raw data or imperfect crosswalk between NACUBO institution name and IPEDS unit ID. In those cases, I hand-collect endowment value from university websites.

[13] For example, number of hires are reported once every two years prior to 2010 for a significant number of private universities.

Figure 1: Time-series of revenue shocks

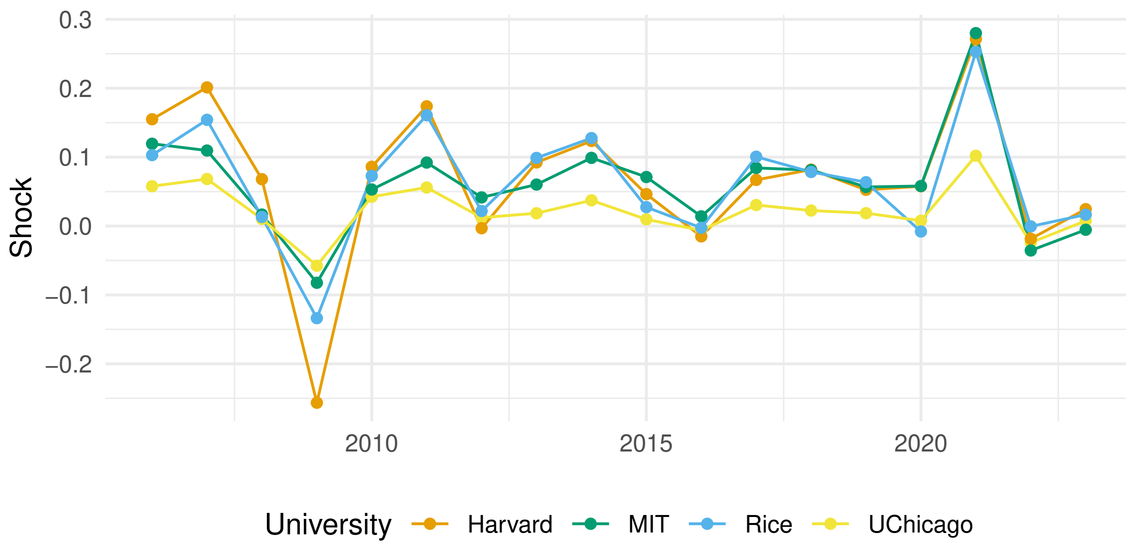

Notes: Panel (a) displays the time series of state-appropriation-induced revenue shocks, as defined in Section 2.3, for four public universities. In Panels (b)-(e), each plot displays the time series of two measures of endowment-induced shocks defined in Section 2.3 for a private university.

where i indexes university; $E_{i,t}$ is i's market value of endowment at the end of the fiscal year ending in year t; $R_{i,pre}$ is university i's revenue in an average year in the pre-period (2005 and 2006 fiscal years[14]) and $R^{(e)}_{i,pre}$ is the component of $R_{i,pre}$ from investment. Figure 1 visualizes the time series of this shock measure for four universities. Across these universities, shock time trends are correlated, and there is variation in the intertemporal variability across institutions, stemming from, for example, differential pre-period reliance on endowment.

2.4 Econometric Analysis

To identify the effect of endowment shocks, I link shocks in Equation 1 to IPEDS data on faculty hiring and headcount from 2005 to 2023. I restrict the sample institutions to the top 50 institutions according to the 2023 U.S. News Best National Universities ranking (Reiter, 2025). I specify the following event-study model (similar to Freyaldenhoven et al. 2019 Section III.B):

$$\Delta y_{i,t} = \phi_t + \sum_{k=-K}^{K} \lambda_k z_{i,t-k} + \xi_{i,t} \tag{2}$$

[14] I choose these years because they are the earliest years with available data.

where y is my outcome of interest, ϕ_t is year fixed effect, and $(-K, K)$ specify window starting and end points. The parameter of interest is the vector $\{\lambda_k\}_{-K}^{K}$ which captures the reduced-form dynamic effects of endowment revenue shocks.

This interpretation relies on three identifying assumptions. First, the shocks $\{z_{i,t-k}\}_{-K}^{K}$ must be exogenous in that it does not correlate with $\xi_{i,t}$. While it is not clear whether the "share" component of the shocks are, the "shift" component of the shock is plausibly exogenous since it originates from capital market fluctuations.[15] Second, the pre-period share of endowment $\frac{R_{i,pre}^{(e)}}{R_{i,pre}}$ approximates exposure to endowment changes in all sample periods[16]. Third, the shocks' effects need to be homogeneous across institutions and years (conditional on event time)[17] and linear in its size.[18]

Figure 2 Panel (a) shows the estimated coefficients with the outcome variable $y_{i,t}$ being the log[19] number of tenure-system faculty hires at university i in the academic year[20] starting in year t. Panel (b) replaces the outcome variable with the log number of tenure-system faculty headcount. The results show that a five percentage point of extra endowment returns during the course of the financial year ending in year 0 increases the number of hires and headcount at year 2 by 1.36 and 0.2 log points for an institution with 20% reliance on endowment. The time pattern is consistent with the timeline of university budgeting discussed in Section 2.1. Notably, the magnitude of the effect on hiring is the largest for year 0. Given the institutional knowledge about the timeline of faculty hiring described in Section 2.1, this is more likely driven by the negative effect of negative revenue shocks than the positive effects of positive shocks.

The results on the effect of endowment revenue shocks on faculty hiring motivate understanding whether a university could improve by adopting shock-independent alternative hiring schemes.

3 Models of Faculty Hiring

I move on to identify gains from alternative hiring schemes under various models.

[15] See, for example, Goldsmith-Pinkham et al. (2020) for a discussion within the shift-share instruments literature.

[16] Since my sample period span over a decade, long-term adjustments (e.g., increasing reliance on endowment) could weaken this assumption later in the sample years. If, for example, universities increasingly rely on endowment, the estimated model coefficients can overstate the effect magnitude.

[17] See, for example, Goodman-Bacon (2021) for a discussion on using this model to capture heterogeneous effects.

[18] In reality, Brown et al. (2014) show that important thresholds exist and matter for the effect of endowment revenue shocks on university operations.

[19] For my sample of top universities, the log-0 issue discussed in Chen and Roth (2023) does not apply.

[20] or by November 1st for earlier periods - see Section 2.2 on data availability

Figure 2: Dynamic effects of endowment revenue shocks

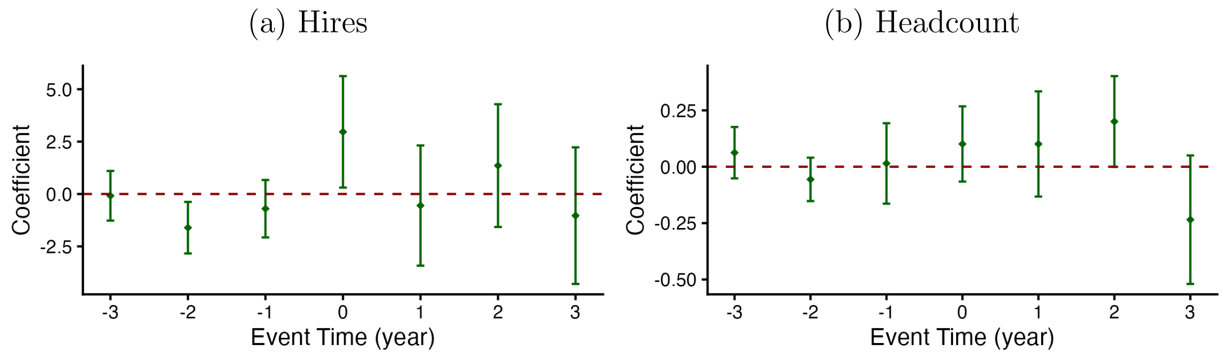

Notes: Panel (a) displays the coefficients of an event study model with dynamic policy effects as specified in Equation 2, with bars representing 95% pointwise confidence intervals based on asymptotic inference clustered by university. The outcome variable is the log number of new tenure-system faculty. This variable captures the number of newly-hired faculty on payroll in year 0 between summer and November 1st for years before 2017, and during the entire academic year for years after 2018. The construction of endowment shock variable is described in Section 2.3 using mainly data from IPEDS and NACUBO. The sample institutions are the top 50 institutions according to the 2023 U.S. News Best National Universities ranking (Reiter, 2025). Panel (b) replaces the outcome variable with the log number of tenure-system faculty headcount. Data on faculty hires and headcount are from the Integrated Postsecondary Education Data System (IPEDS).

3.1 Set-up and Notations

Unless otherwise noted, the norm notation $||\cdot||$ in this paper denotes L1 norm.

Departments and candidates

Consider a set of *departments* indexed by $i \in \{1, \ldots, I\}$. Let there be a ranking over them and denote this ranking by $r(i)$ where $r(i) \equiv 1$ denotes the best department. Department i in year $t \in \{1, \ldots, T\}$ has a budget to fill s_{it} open positions (henceforth *slots*). We can summarize slots across departments and years with a $I \times T$ slot matrix $S = [\vec{s_1}, \ldots, \vec{s_T}]$, where $\vec{s_t}$ is a $I \times 1$ vector capturing the number of slots in all departments in year t.

Departments need to hire job market *candidates* to fill the slots. There are $J_t \in \mathbb{N}_+$ candidates available for hire in year t. Let $H_{it} \subset \{1, ..., J_t\}$ denote department i's set of hires in year t. No candidate can be hired by multiple departments in the same year.[21]

Each candidate $j \in \{1, ..., J_t\}$ has three attributes: first, an ability θ_j observed by the econometrician;[22] second, a public ability signal p_j (e.g. number of publications) observable to all, and a private ability signal q_j (e.g. recommendation letters) unobservable to the econometrician.

Turning slots into payoffs

My goal of evaluating hiring decisions motivates assigning a measure of *hiring outcome* to each cohort of hires at department-year level. Let $\vec{v_t}$ be a $I \times 1$ vector capturing hiring outcomes for all departments in year t. For the remainder of this paper, I define v_{it} to be the sum of ability θ_j across hires in each department-year cohort: $v_{it} := \sum_{j \in H_{it}} \theta_j$.[23]

But H_{it} is a function of $\vec{s_t}$. To model directly how slots affect hiring outcomes, I define $f : \mathbb{R}^I \to \mathbb{R}^I$ as a mapping from slots $\vec{s_t}$ to outcomes $\vec{v_t}$. I impose the following restrictions on f:

First, f is *market-dependent* and is *not* an entry-wise function. Each department's outcome $[\vec{v_t}]_i$ depends not only on its own slots $[\vec{s_t}]_i$ but also other departments' slots (i.e. entire

[21] Thus, when the market clears, we have $J_t \equiv \sum_{i=1}^{I} |H_{it}|$.

[22] See, however, e.g., Zhang et al. (2022) on the effect of institutional environment on productivity. In reality, this could complicate the econometrician's ability to observe the candidates' ability. I ignore this complication.

[23] The identification results in this section hold under mild regularity conditions even if we broaden the definition of v_{it} to allow it to be additively separable in not just θ_j but the class $\mathcal{G}$ of increasing functions of θ_j; that is, $v_{it} := \sum_{j \in H_{it}} g(\theta_j)$ where $g(\cdot) \in \mathcal{G}$. That would accommodate log-multiplicative forms like $v_{it} := \ln\left(\prod_{j \in H_{it}} \theta_j\right)$ where $g(\cdot)$ is the natural log function. Such broader options to define v_{it} are motivated by the literature on productivity spillovers (see, for example, Azoulay et al. 2010), but they add expositional complexity and are thus not considered here.

$\vec{s_t}$). Formally, there exist $i \neq i'$, $\vec{s}$, and $\vec{s}'$ such that $s_i = s'_i$, $s_{i'} \neq s'_{i'}$, and $[f(\vec{s})]_i \neq [f(\vec{s}')]_i$.

Second, any department with 0 slots produces 0 outcome. Formally, for any t, we have $s_{it} = 0$ implies $[f(\vec{s_t})]_i = 0$.

Third, outcome depends only on contemporaneous slots, and not on other years'. This rules out, for example, productivity spill-overs across department-year cohorts, or positively- or negatively-selected candidates delaying job market entry due to market demand. Formally, this allows us to define a hiring outcome matrix $V := f(S) = \left[f(\vec{s_1}), \ldots, f(\vec{s_T})\right] =: [\vec{v_1}, \ldots, \vec{v_T}]$.

Fourth, hiring more weakly increases own outcome. Formally, $[f(\vec{s} + \vec{e_i})]_i \geq [f(\vec{s})]_i$ for any $\vec{s}$, where $\vec{e_i}$ denotes a standard basis vector where the i-th component is 1 and all others are 0.

Fifth, for a fixed year, any allocation across departments of a fixed total number of slots yields the same total output across departments. Formally, for any $\vec{s}, \vec{s}'$ satisfying $||\vec{s}|| = ||\vec{s}'||$, we have $||f(\vec{s})|| = ||f(\vec{s}')||$. In other words, $f(\cdot)$ is L1-isometric.

I define each department's *payoff* u as a linear combination of hiring outcomes across years.[24] Although not necessary, I assume for simplicity for the rest of this paper that all departments have the same weights $\vec{w} \in \mathbb{R}^T$, so that the $I \times 1$ payoff vector $\vec{u} := V\vec{w}$ summarizes all departments' payoffs.

Gains from alternative hiring schemes

Now I can formulate gains from alternative hiring schemes. For example, $f(\vec{s_t} + \vec{e_i}) - f(\vec{s_t})$ is the marginal change in all departments' outcomes as a result of department i expanding hiring by one slot in year t.[25] Generally, $[f(S') - f(S)]\vec{w}$ is the net payoff gain from an alternative slot allocation S' over S.

The general formulation clarifies that the gains can be identified by separately identifying $f(S)$ and $f(S')$. The former is straightforward when the econometrician defines the outcome to be a known function of observable factual metrics about the candidates, such as ability θ. For example, each element of $f(S)$ could be the sum ability of the each university-year cohort. By contrast, the latter term $f(S')$ is more challenging to identify, and is the focus of the rest of this section.

[24]Importantly, my framework can accommodate non-uniform weights to, e.g., counteract the secular increase in measured θ_j, or reflect the departments' time-varying preferences.

[25]This is assuming that other departments do not strategically adjust their hiring patterns in response to the expansion by i.

3.2 Structural Identification via Stable Matching

I build a model motivated by that faculty hiring can be considered a matching problem, where two sets of agents with complete and transitive preferences form matches.

3.2.1 Information and preferences

In this model, departments have an information set $\mathbf{X}_j$ about any candidate j. This information set always contains the candidate's public ability signal p_j. Depending on the exact setting such as the discipline and the time period, $\mathbf{X}_j$ can contain private signal q_j or even ability θ_j. I require, however, that all departments share the same $\mathbf{X}_j$ for the same candidate j.[26] Using $\mathbf{X}_j$, departments form an expectation of each candidate's ability $\hat{\theta}_j(\mathbf{X}_j) := \mathbb{E}[\theta_j \mid \mathbf{X}_j]$. I impose the following restriction on this expectation function.

Assumption 1. *Better public ability signals weakly increase expected ability of any candidate. Formally,* $\frac{\partial}{\partial p_j}\hat{\theta}_j(\mathbf{X}_j) \geq 0$ *for all* j.

For candidates, I assume that they prefer departments with better ranking. For departments, I assume that they maximize expected payoff $\mathbb{E}[u]$. This leads to the following result with proof in Appendix Section A.1.

Claim 1. A department given its slots and information about candidates prefers candidates with high expected ability $\hat{\theta}_j(\mathbf{X}_j)$.

3.2.2 Stability

In this matching problem, the solution concept is **stability**. A matching is stable if there is no blocking pair (i, j)—a department and candidate who both prefer each other to their current matches, with at least one preference being strict. Gale and Shapley (1962) show that with complete and transitive preferences, a stable matching always exists. The following claim follows, with proof in Appendix Section A.2.

Claim 2. A top-down serial dictatorship produces one stable matching. That is, if departments move sequentially from top to bottom by ranking $r(\cdot)$ to hire the top $\vec{s_t}$ (remaining) candidates by $\hat{\theta}(\mathbf{X})$, then the resulting matching is stable.

An additional result follows.

Claim 3. Higher-ranking departments hire candidates with higher $\hat{\theta}(\mathbf{X})$ up to ties.

[26] That is, $\mathbf{X}_{i,j} \equiv \mathbf{X}_j$ for all i.

3.2.3 Identifying gains from alternative hiring schemes

Claim 2 presents a recipe for the econometrician to produce a stable matching under any legitimate slot allocation $\vec{s_t}$, as long as the econometrician observes the department's information set. For example, if departments observe candidates' ability ($\theta \in \mathbf{X}$), then they can rank candidates by θ, and following the procedure in Claim 2 produces a stable matching. Similarly, if departments instead observe only the candidates' public ability signal ($p \in \mathbf{X}$), Assumption 1 guarantees that letting the departments rank candidates by p produces a stable matching.[27]

With a known stable matching, the econometrician knows exactly the model-predicted hires $H_{it}^{(\text{model})}$ for all (i, t) given any hiring scheme S. Using the individual-level hires data in turn leads to the model-predicted hiring outcome matrix $V^{(\text{model})} = f^{(\text{model})}(S)$. The gains from an alternative hiring scheme follow naturally from comparing $f^{(\text{model})}(S')$ to $f^{(\text{model})}(S)$.

3.2.4 Assessing model fit

In reality, hiring may differ from the model above. This motivates assessing the model fit. Since by Claim 2 we know one stable matching, we can perform an outcome test: denote the empirical (in the data) slot matrix as $S^{(\text{data})}$ and the empirical hiring outcome matrix as $V^{(\text{data})}$. Feeding $S^{(\text{data})}$ into the model produces $V^{(\text{model})} = f^{(\text{model})}\left(S^{(\text{data})}\right)$. Then we can assess the model fit at department-year level by comparing $V^{(\text{model})}$ to $V^{(\text{data})}$, and aggregate this fit measure into department- (or year-) level fit measures.

3.2.5 Bounds under private information

If departments observe candidates' private ability signals ($q \in \mathbf{X}$), the econometrician does not have data to produce a stable matching using the serial dictatorship above without imposing extra structure. However, the econometrician can bound the counterfactual outcome $f(S')$ under the assumption that the factual matching in the empirical data is stable. To do this, we need the following definition.

Definition 1. A *marginal hire* of department i in year t is one that i would have hired had it been given another slot in that year. Denote this hire by $m_{it} \notin H_{it}$.

With this definition, we have the following result.

Claim 4. Any marginal hire m_{it} at department i in year t has weakly smaller $\hat{\theta}(\mathbf{X})$ than any hire in H_{it}. That is, $\hat{\theta}_{m_{it}}(\mathbf{X}_{m_{it}}) \leq \min_{h \in H_{it}} \hat{\theta}_h(\mathbf{X}_h)$.

[27] The case where the departments observe private ability signals ($q \in \mathbf{X}$) is more intricate and I discuss it in Section 3.2.5.

Now we move to bound the gains. For the ease of exposition, below I consider a case S' where a department i—and only department i—unilaterally changes its hiring scheme relative to a baseline S. That is, S' satisfies $[S']_i \neq [S]_i$ and $[S']_\iota \equiv [S]_\iota$ for all $\iota \in \{1, ..., I\}$.

The key problem is still to calculate the counterfactual outcome $[f(S')]_i$. This is a $1 \times T$ vector. For each year t, define department i's *hiring adjustment* as the difference between its counterfactual and factual slots: $\Delta_{it} := s'_{it} - s_{it}$. For t where $\Delta_{it} \equiv 0$, since the factual matching is stable, it is one stable matching and it has the characteristic of $[f(S')]_{it} = [f(S)]_{it}$. For t where $\Delta_{it} < 0$, the set of counterfactual hires is just the factual hires H_{it} short of $-\Delta_{it}$ marginal hires. These marginal hires have ability bounded by the empirical set of hires' ability. So

$$[f(S)]_{it} - [f(S')]_{it} \in \left[\sum_{k=1}^{-\Delta_{it}} \theta_{(k)}, \sum_{k=|H_{it}|+\Delta_{it}+1}^{|H_{it}|} \theta_{(k)} \right]$$

, where $\theta_{(1)} \leq \theta_{(2)} \leq \cdots \leq \theta_{(|H_{it}|)}$ are the ordered abilities within H_{it}. I thus bound $[f(S')]_{it}$. Similarly, for t such that $\Delta_{it} < 0$, the set of counterfactual hires is just the factual hires H_{it} plus additional Δ_{it} marginal hires. From Claim 4, we know that if department i were to hire more in a year, those marginal hires would have lower $\hat{\theta}$ compared to i's factual hires. This means that the marginal hires are drawn from the next best department(s) and thus have ability bounded by the empirical set of abilities at the next best department(s). So

$$[f(S')]_{it} - [f(S)]_{it} \in \left[\sum_{k=1}^{\Delta_{it}} \theta_{(k)}, \sum_{k=|H_{i't}|-\Delta_{it}+1}^{|H_{i't}|} \theta_{(k)} \right]$$

, where i' represents the index of the next best department to i.[28] Then we can bound $[f(S')]_{it}$.

Because departments' payoff gains from alternative smoothing are monotonic (in fact linear) functions of the individual $\{[f(S')]_{it}\}_{t=1}^{T}$ that are bounded, we can bound the payoff gains from above and below by aggregating the above-mentioned bounds across years.

[28]To be exact, it is possible that $|H_{i't}| < \Delta_{it}$. Here I ignore this complexity for expositional simplicity. The precise procedure under $|H_{i't}| < \Delta_{it}$ goes as follows: one needs to iteratively move through departments ranked lower than i to exhaust finding marginal hires. In this process, the bounds on ability still applies to individual departments from which marginal hires are taken, and such bounds at department level can still be aggregated linearly across partitions of marginal hires.

3.3 Numerical Example

Primitives

Here I demonstrate the stable matching model in action in one reasonable scenario. I show that a department's gains from one feasible alternative hiring scheme—perfect smoothing—depend on department's own hiring pattern relative to its peers', as well as candidate heterogeneity in ability.

Consider $T = 2$ and each department i has a count-metric fixed hiring quota $b_i \in \mathbb{N}$ to allocate over two years $(s_{i,1}, s_{i,2}) \in \mathbb{N}^2$. Let $J_1 = J_2 = J$ and denote the j-th highest ability in set $\{\theta\}_{j=1}^{J}$ as $\theta_j = \theta(j) = -\gamma j$, where parameter $\gamma \geq 0$ controls the rate of ability decline across candidates.[29] Let the departments observe θ (so $\theta \in \mathbf{X}$).

To produce a stable matching, I let departments hire sequentially by rank $r(i)$. Each department i chooses $(s_{i,1}, s_{i,2})$ under the constraint of $s_{i,1} + s_{i,2} \leq b_i$ to maximize its payoff u_i, defined as the total (across years) ability of its hires:

$$u_i(s_{i,1}, s_{i,2}; i) = \sum_{t=1}^{2} \sum_{k=1}^{s_{i,t}} \theta \left(j_{i,t} + k - 1\right)$$

where $j_{i,t}$ is the starting position of candidates available to i in year t. Note that $j_{i,t}$ is determined by the number of hires at all other departments $\{i' : i' < i, i' \in \mathbb{N}_+\}$ ranked higher than i.

Payoffs under different hiring schemes

I compare the payoffs from two different hiring schemes: an unsmoothed baseline hiring scheme of $(s_{i,1} s_{i,2}) = (b_i, 0)$ versus a smoothed hiring where $s_{i,1} = s_{i,2} = b_i/2$, assuming for expositional simplicity that $(b_i/2) \in \mathbb{N}$. The payoff under the baseline hiring scheme is

$$u_i' = \sum_{k=1}^{s_{i,1}} \theta \left(j_{i,1} + k - 1\right) = \sum_{k=1}^{b_i} \left[-\gamma \left(j_{i,1} + k - 1\right)\right]$$

, while for the smoothed hiring scheme it is

$$\tilde{u}_i = \sum_{t=1}^{2} \sum_{k=1}^{b_i/2} \left[-\gamma \left(j_{i,t} + k - 1\right)\right]$$

, so the difference simplifies to:

[29] That is, I hold constant the distribution of θ across years with the same J.

Figure 3: Numerical Examples of Gains from Smoothing

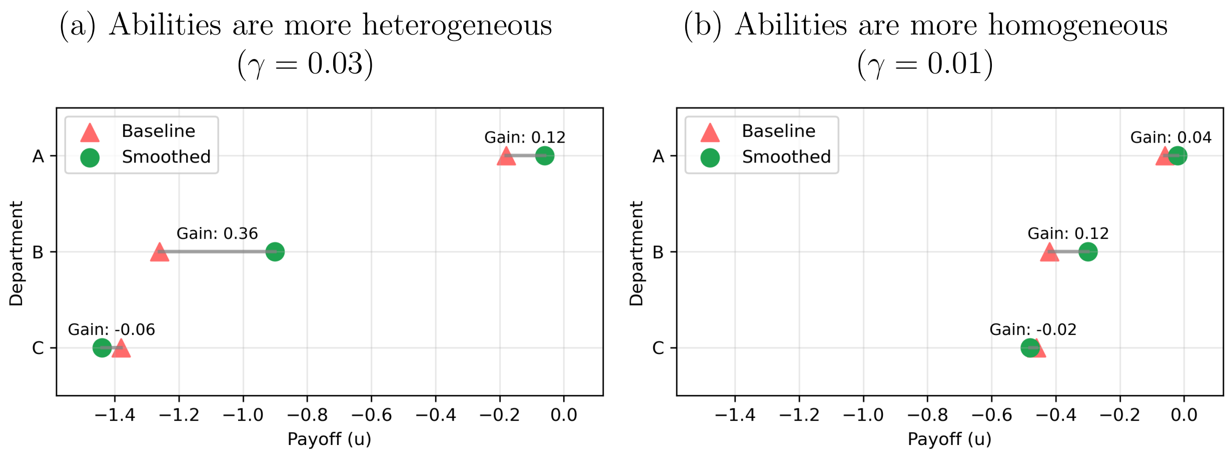

Notes: Panel (a) illustrates three example departments' payoffs under baseline vs smoothed hiring scheme. Sections 3.1-3.3 describe the set-up, primitives, notations, and procedures. Each department has a budget $b = 4$. Department A is a top-ranking department. Department B and C rank lower. Departments ranked higher than B hire 9 candidates in the first year and 5 in the second. Departments ranked higher than C hire 10 candidates in the first year and 13 in the second. The red triangle shows each department's payoff u under the baseline hiring scheme of loading the entire budget to the first year; the green dot shows u with smoothed hiring of evenly distributing the budget across years. Departments face a market supply where candidates' ability are different and declines steeply ($\gamma = 0.03$) along the ranking of candidates. Panel (b) repeats the exercise in Panel (a) but in a setting where candidates are more homogeneous in their ability ($\gamma = 0.01$).

$$\tilde{u}_i - u_i' = \gamma \frac{b_i}{2} \left(j_{i,1} - j_{i,2} + \frac{b_i}{2} \right)$$

Three observations follow. First, when candidates have homogeneous ability ($\gamma = 0$), both hiring schemes produce the same payoff. Second, gains from smoothing depends on the time-relative market demand for candidates, $(j_{i,1} - j_{i,2})$ and could be negative if smoothing shifts hiring to the year with high market demand. Third, smoothing improves payoff when candidates have heterogeneous productivity ($\gamma > 0$), holding constant other departments' demand in each year. Figure 3 offers numerical examples illustrating these observations.

3.4 Reduced-form Identification

3.4.1 Set-up

If, contrary to the assumptions in Section 3.2, candidates or departments have heterogeneous preferences[30], serial dictatorship no longer produces stable matching and we can no longer use the described point or partial identification strategies.

This concern motivates a parsimonious linear model that approximates the possibly non-linear mapping of $V = f(S)$:

$$v_{it} = \beta_1 s_{it} + \beta_2 d(\vec{s_t}; i) s_{it} + \beta_3 d(\vec{s_t}; i) + \epsilon_{it} \tag{3}$$

where $d(\vec{s_t}; i)$ summarizes hiring competition faced by department i from other departments in year t.

The choice of $d(\cdot)$ depends on assumptions about the matching process. I consider two such $d(\cdot)$ functions.[31]

3.4.2 Sum of slots at better ranked departments

Fix year t, define $d(\vec{s_t}; i) := \sum_{j:r(j)>r(i)} s_j$ as the sum of the slots in departments ranked higher than i. This is justified under the assumption that ranking $r(\cdot)$ captures the homogeneous preference of the candidates over the departments. Under this assumption, department-propose serial dictatorship following ranking $r(\cdot)$ yields a stable matching, and in this matching, department i's hiring is affected by and only by departments ranked higher than i.

3.4.3 Sum of the slots in all other departments

We may define $d(\vec{s_t}; i) := \sum_{i' \neq i} s_{i't}$ as the sum of the slots in all other departments. This is justified under assumptions that at least one side's preferences are heterogeneous. For example: when

- departments prefer idiosyncratic fit in candidates
- different departments have different information sets about candidates
- different departments aggregate rich set of information differently.

[30] or, that information set $\mathbf{X}_j$, or the way a department aggregates elements of $\mathbf{X}_j$, varies across departments

[31] In Appendix Section C.1, I consider an additional choice of $d(\cdot)$ function rooted in candidates' choice problem and a random utility model.

3.4.4 Formulating hiring outcomes and gains

I take these $d(\cdot)$ to Equation 3.4.1. To calculate hiring outcomes $V^{(\text{model})}$ stemming from a slot allocation, we can estimate this model on the data to obtain parameters $\left\{\hat{\beta}_\iota^{(\text{OLS})}\right\}_{\iota\in\{1,2,3\}}$ and residuals $\hat{\epsilon}_{it}$ for all i, t, and then sub in a (factual or counterfactual) slot matrix S. This $V^{(\text{model})}$ matrix has components:

$$\hat{v}_{it} := \hat{\beta}_1^{(\text{OLS})} s_{it} + \hat{\beta}_2^{(\text{OLS})} s_{it} d(\vec{s}_t; i) + \hat{\beta}_3^{(\text{OLS})} d(\vec{s}_t) + \hat{\epsilon}_{it}$$

where s_{it} and $\vec{s}_t$ are from S. We can then compute gains from S' relative to S via a simple difference.[32]

3.5 Identification Using Instruments

We may have concerns that slots s_{it} are endogenous, in that its factual value is not independent of the number of candidates or distribution of candidates' θ in year t. This could occur via, for example, failed faculty searches (Reed, 2014) driven by, for example, candidate pool quality. One solution is provided by exogenous budget shocks z_{it} constructed in Section 2. Define

$$\vec{\beta}^{(\text{2SLS})} := \left[\beta_1^{(\text{2SLS})}, \beta_2^{(\text{2SLS})}, \beta_3^{(\text{2SLS})}\right]$$
$$\Gamma := [s_{it}, s_{it} d(\vec{s}_t; i), d(\vec{s}_t)]$$
$$\Omega := [z_{it}, z_{it} d(\vec{z}_t; i), d(\vec{z}_t)]$$

. Then we can write the first stage as :

$$\Gamma = \Phi\Omega + \Xi$$

where Φ is a 3×3 real coefficient matrix and Ξ is a 3×1 error vector. Similarly, we an write the second stage as

$$v_{it} = \vec{\beta}^{(\text{2SLS})}\hat{\Gamma} + \varepsilon_{it}$$

where $\hat{\Gamma}$ is predicted using the first stage. For this 2SLS model to satisfy exclusion restriction, budget shocks z can affect hiring outcomes only through slots. Then the rest follows Section 3.4.

[32] As before, we can also assess model fit by comparing $V^{(\text{model})}$against V^{data}.

We now apply these identification strategies on a dataset.

4 Empirical Evidence

4.1 Data on Model Primitives

To apply the models described in Section 3, I use individual-hire-level data on the universe of assistant professor hires at 21 top sociology departments in the U.S. from 1991 to 2017 collected manually from the American Sociological Association's (ASA) *Guide to Graduate Departments of Sociology* by Warren (2019). The data contain each hire's name, year and department of hire, and the number of publications by the year of hire, which I use as the measure of public signal p_j in the model in Section 3.

At department level, I obtain ranking and the underlying scores from the U.S. News *Best Sociology Programs* ranking (U.S. News & World Report, 1998, 2017). I construct a strict ranking using the average scores, and break ties using the 1998 scores. This matches to the $r(\cdot)$ in Section 3.

At individual level, I use name and institutional affiliation to obtain each hire's publication history from OpenAlex (Priem et al., 2022). OpenAlex inherits the publication data from the discontinued Microsoft Academic Graph and is increasingly used in the economics literature (e.g. Ito et al., 2025). Figure 4 Panel (a) cross-validates OpenAlex data against the Warren (2019) data on the number of pre-hire publications. From the publication history data I compute the number of social science publications published within 6 years after the year of hire, and use this as ability measure θ_j.

4.2 Alternative Hiring Schemes

For each department i, I consider two alternative hiring schemes. The first one smoothes i's hiring across years. Formally, I minimize the variance of the slot vector $\{s_{i\tau}\}_{\tau=1}^{T}$ subject to integer and non-negativity constraints.[33] The second one, henceforth *countercyclical* hiring, is hires more when other departments hire less. Formally, for this department i, I permute elements of the slot vector $\{s_{i\tau}\}_{\tau=1}^{T}$ so that larger values appear in years when other departments (factually) together hire fewer. Importantly, both counterfactual schemes hold constant the total (across years) number of slots $||\vec{s_i}||$ as the factual number, as well as the slot allocations of all other departments $\vec{s}_{k\neq i}$.

[33] I break ties randomly.

Figure 4: Features of Sociology Hiring Data

(a) Number of pre-hire publications in OpenAlex and Warren (2019)

(b) Number of publication at hiring and during tenure clock

(c) Number of publication at hiring and during tenure clock

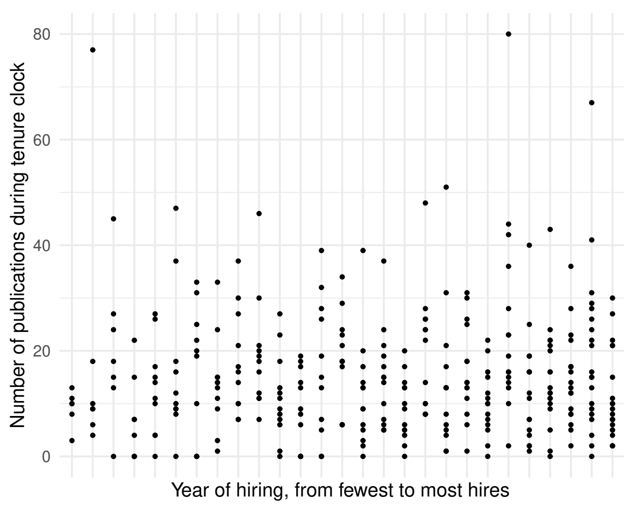

(d) Number of publications during tenure clock, by with-year rank

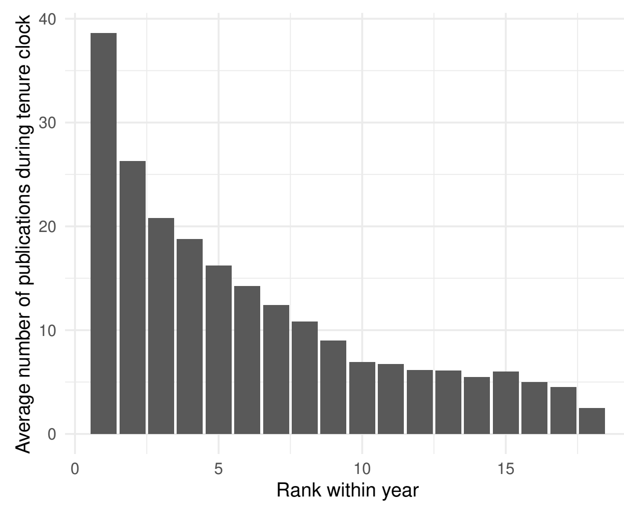

Notes: Panel (a)-(c) are scatter plots. Each scatter is an assistant professor hire in Warren (2019). Panel (a) compares the number of publications at the year of hiring according to Warren (2019) on the x-axis against OpenAlex (Priem et al. 2022) on the y-axis. Panel (b) compares the number of publications at the year of hiring (according to Warren) on the x-axis against the number of publications during the tenure clock (according to OpenAlex) on the y-axis. Panel (c) plots each hire's number of publications during tenure clock (y-axis) against their year of hire (axis), where the years are arranged from the year where departments in the sample hire the fewest (left) to the year where departments hire the most (right). Panel (d) shows, for each rank of the candidate pool by the number of publications during the tenure clock, the average number of such publications across years.

4.3 Results

I first apply to the data the stable-matching models described in Section 3.2. For each department, I calculate its payoff u as defined in Section 3.1 using uniform weights $w_i = \mathbf{1}_I \in \mathbb{R}^I$. Since larger departments that hire more may mechanically receive higher payoff, I further calculate the per-slot payoff $u_i/||\vec{s_i}||$.

Figure 5 Panel (a) shows results on per-slot payoffs in four scenarios for each department ranked from top to bottom[34]: the blue bars show empirical metrics, calculated from the factual data; the red triangles show the metrics from data simulated following the model under the departments' factual slot allocation across years (that is, using the factual S matrix); the green dots show the metrics from data simulated following the model under the departments' factual slot allocation across years, except for the department on that row where I replace its slots with a smoothed vector; the yellow triangles denote the countercyclical analog of the smoothed scheme represented by the green dots. The departments observe the candidates' number of pre-hire publications ($\mathbf{X}_j = p_j$).

[34]The bottom four departments in the Warren (2019) dataset are excluded because in some years serial dictatorship attempts to hire more candidates than available for hire when it is those departments' turn to hire.

Figure 5: Per-slot Payoff by Department under Stable Matching

(a) Observe pre-hire publications

(b) Observe ability

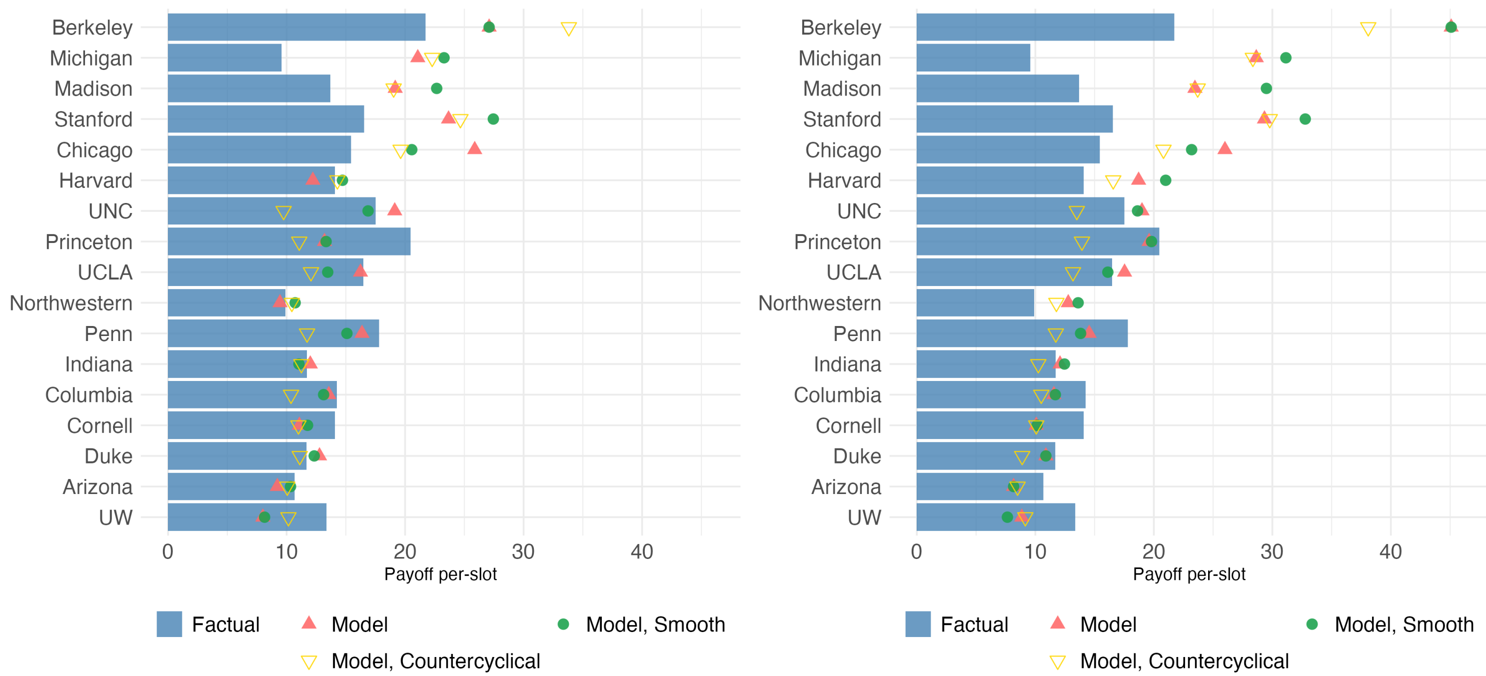

Notes: this figure shows results on per-slot payoff in four scenarios for each department ranked from top to bottom: the blue bars show empirical metrics, calculated from the factual data; the red triangles show the metrics from data simulated following the model under the departments' factual slot allocation across years (that is, the factual S matrix) - see details on this simulation in Section 3. For each department, the green dot shows that department's counterfactual per-slot payoff when I replace that department's slot vector with a smoothed vector, while holding other departments' the same; the yellow triangle is the countercyclical analog of the green dot - see Section 4.2 for details on the construction of the smoothed and countercyclical hiring schemes. A department's payoff is calculated as the sum of ability θ across all its hires (across all years), as described in Section 3.1. Ranking of departments is from the 1998 and 2017 U.S. News *Best Sociology Programs.* Ability θ is measured by the number of social science publications within 6 years after hiring according to OpenAlex (Priem et al., 2022). Public attribute p is measured by the number of publications at the time of hiring, according to Warren (2019) - see Section 4.1 for details on data and variable construction. Results for the bottom four departments in the sample are not visualized because in some years serial dictatorship attempts to hire more candidates than available for hire when it is those departments' turn to hire. The departments' information set is pre-hire public signal ($\mathbf{X}_j = p_j$) in Panel (a) and ability ($\mathbf{X}_j = \theta_j$) in Panel (b).

Several observations follow. First, the blue bars show that factually, higher-ranked departments generally hire candidates with higher ability θ and thus receive higher per-slot payoffs; meanwhile, this is also true for the scatters, reflecting the positive correlation between pre-hire ability signal p_j and ability θ_j across candidates, as illustrated in Figure 4 Panel (b). Second, the blue bars and the red triangles show that the model-predicted payoff can deviate from the factual payoff even under the same (factual) hiring scheme. The difference reflects the model fit measure discussed in 3.2.4.

Third, and importantly, the scatters show that alternative hiring schemes do not generally improve payoff. Smoothing hiring improves an average department's payoff by just 1.6%, whereas countercyclical hiring reduces payoff by 5.3%. By contrast, the per-hire payoff at an average department in the bottom tercile in the sample is 18.5% lower than that in the upper tercile.

The lack of significant gains is not explained by that the pre-hire ability signals are not informative of the hires' ability. Figure 5 Panel (b) shows that, allowing departments to perfectly observe candidates' ability barely expands the positive effect of smoothing (to 1.8%) while significantly worsening the negative effects of countercyclical hiring (to 11.2%), although this model is rejected by the data due to poor fit.

What, then, explains the absence of meaningful gains? Recall that gains from smoothing are small when candidates do not substantially differ from each other in ability, as illustrated in my numerical example in Section 3.3. This is true in the data. Figure 4 Panel (d) shows, for each within-year candidate rank, the average (across years) ability of the candidates at that rank. The illustrated pattern confirms that the ability heterogeneity is small except among the upper-tier of candidates in each year. Correspondingly, Figure 5 shows that alternative hiring schemes marginally change the payoff for mid- and lower-tier departments in the sample; for the top-tier departments, alternative hiring schemes matter more: smoothing leads to modest gains for some top departments (and modest losses for others), whereas counterfactual hiring often leads to losses.

Another contributing factor to the lack of meaningful gains could be that the ability distribution in the candidate pool varies across years. Figure 4 Panel (c) plots, for each hire in my sample, their ability against their year of hiring, arranged by the number of hires across departments in that year from the fewest (left, low-hiring years) to the most (right, high-hiring years). The illustrated pattern shows that fewer high-ability candidates are hired in low-hiring than in high-hiring years. This means that shifting hiring from high- to low-hiring years actually undermines the top departments' prospect of hiring top candidates, consistent with the loss from adopting countercyclical hiring observed in Figure 5.

Alternative models do not yield qualitatively different results. Figure 6 reports results when we allow departments to observe private ability signals unobserved by the econometrician, following Section 3.2.5. The blue bars again show the departments' factual per-slot payoff, and in each panel, the illustrated intervals bound the per-slot payoff from the corresponding counterfactual hiring scheme. We see that the bounds are often too wide to refute that the alternative hiring allocations significantly change the outcomes for most departments. In other cases, the bounds suggest that alternative hiring schemes significantly improve some departments' payoff - this is especially so when the immediately lower-ranked departments have better payoffs, which suggests that this finding could be sensitive to the choice of ranking of departments.

Figure 6: Bounds on Per-slot Payoff by Department under Stable Matching with Private Information

(a) Smoothed hiring (b) Countercyclical hiring

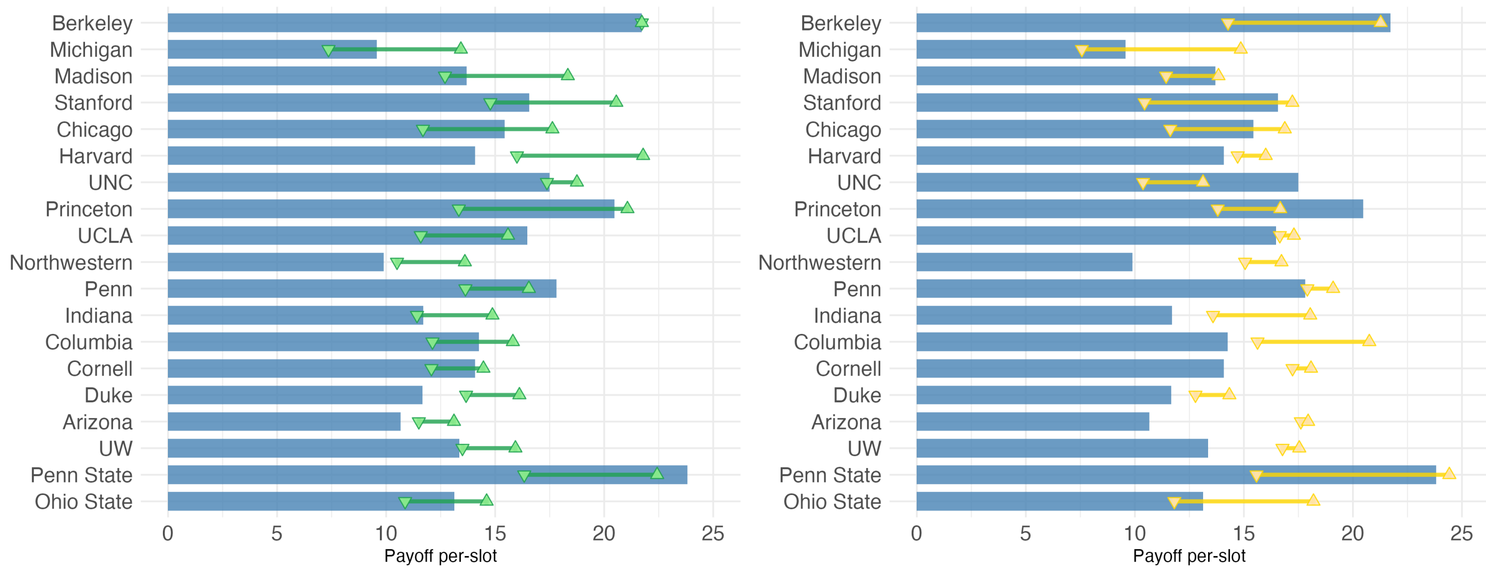

Notes: The blue bars show the per-slot payoff for each department ranked from top to bottom. The bars show bounds on the per-slot payoff as predicted by a stable matching model described in Section 3.2.5. In the data input to the matching model, all departments use the factual hiring scheme, except for the department on that row: I replace its slots with an alternative hiring scheme. The alternative hiring scheme is a smoothed vector in Panel (a) and a countercyclical permutation across years of the original factual slots - see Section 4.2 for details on the construction of the smoothed and countercyclical hiring schemes. A department's payoff is calculated as the sum of ability θ across all its hires (across all years), as described in Section 3.1. Ranking of departments is from the 1998 and 2017 U.S. News Best Sociology Programs. Ability θ is measured by the number of social science publications within 6 years after hiring according to OpenAlex (Priem et al., 2022) - see Section 4.1 for details on data and variable construction. Results for the bottom four departments in the sample are not visualized because in some years serial dictatorship attempts to hire more candidates than available for hire when it is those departments' turn to hire. The departments' information set includes private information q not observed by the econometrician.

I then apply the linear model described in Section 3.4. Figure 7 presents the results. Panel (a) shows the results when I model the hiring competition function $d(\cdot)$ as the leave-out sum of other departments' factual slots in a given year, as described in Section 3.4.2. The color schemes follow those in Figure 5. The OLS model fit is perfect because the prediction incorporates the estimated residual $\hat{\epsilon}_{it}$. Panels (b) repeats the exercises in Panel (a) but replaces the hiring competition function with the sum of the factual slots at higher-ranked departments in a given year, as described in Section 3.4.3.[35] The results show no evidence that alternative hiring schemes improve payoff.[36]

Lastly, I add the revenue shocks instruments constructed in Section 2[37] to the linear model as described in Section 3.5. Figure 8 presents the results, where each panel uses a different hiring competition function. I find no evidence that smoothing slots across years improves payoff. If hiring competition function is correctly specified as the leave-out sum as in Section 3.4.2, then the countercyclical hiring scheme significantly improves all departments' payoffs, but the effect is the implausibly large and sensitive to the specification of market competition function: the effect is the opposite for other specifications of the hiring competition function in Panel (b) and Appendix Figure 2.

4.4 Limitations

Overall, I do not find strong evidence that alternative hiring schemes significantly improves departments' payoff. Two empirical limitations could have complicated the finding.

The first is measurement error from OpenAlex. This is evident in two places in Figure 4: the existence of hires with 0 publications during their tenure clock according to OpenAlex, and the discrepancy between the number of publications according to OpenAlex and that according to Warren (2019). Future work can address this issue by improving data quality, for example, by manual correction or incorporating other datasets like the Web of Science.

The second is that the number of publications does not measure ability well, to the extent that departments value it. We can see this in Figures 5-8 where the ranking of departments from U.S. News data differs from the ranking by per-slot payoffs measured by average hire's ability θ. This could mean that alternative hiring schemes may still improve payoffs, but that gain is not captured in the results where payoffs aggregate publication counts. Future work could explore plausibly better measures of researcher ability, such as

[35] Appendix Figure 2 additionally uses a hiring competition function rooted in candidates' choice problem in a random utility model, as described in described in Appendix Section C.1.

[36] Appendix Figure 3 shows the corresponding prediction 95% confidence interval.

[37] For public universities, I construct the shocks analogously following Deming and Walters, 2017 replacing endowment shocks with state appropriation shocks data from the State Higher Education Executive Officers Association (SHEEO, 2025).

number of publication weighted by authorship order, journal prestige, and year- and topic-normalized citation counts. Future work could also explore empirical settings, such as some natural or computational sciences disciplines, where publication count is a better measure of researcher ability. The discrepancy in ranking by U.S. News and factual per-slot payoff could also reflect that the U.S. News ranking does not approximate well the ranking of departments in candidates' preferences. Future work can explore alternative rankings, e.g., by per-slot payoff, or adopting year-specific dynamic rankings.

5 Concluding Remarks

Faculty hiring responds to endowment revenue shocks. While this practice may be theoretically suboptimal, my empirical results show little evidence of significant gains from adopting plausibly feasible alternative shock-independent hiring rules. I present suggestive evidence that the finding is driven by a combination of the lack of heterogeneity in ability among the candidate pool, and by the varying distribution of candidate ability across years of different hiring activities.

Figure 7: Per-slot Payoff by Department, Reduced-Form Model

(a) $d(\cdot)$ is leave-out sum of slots

(b) $d(\cdot)$ is sum of slots across better-ranked departments

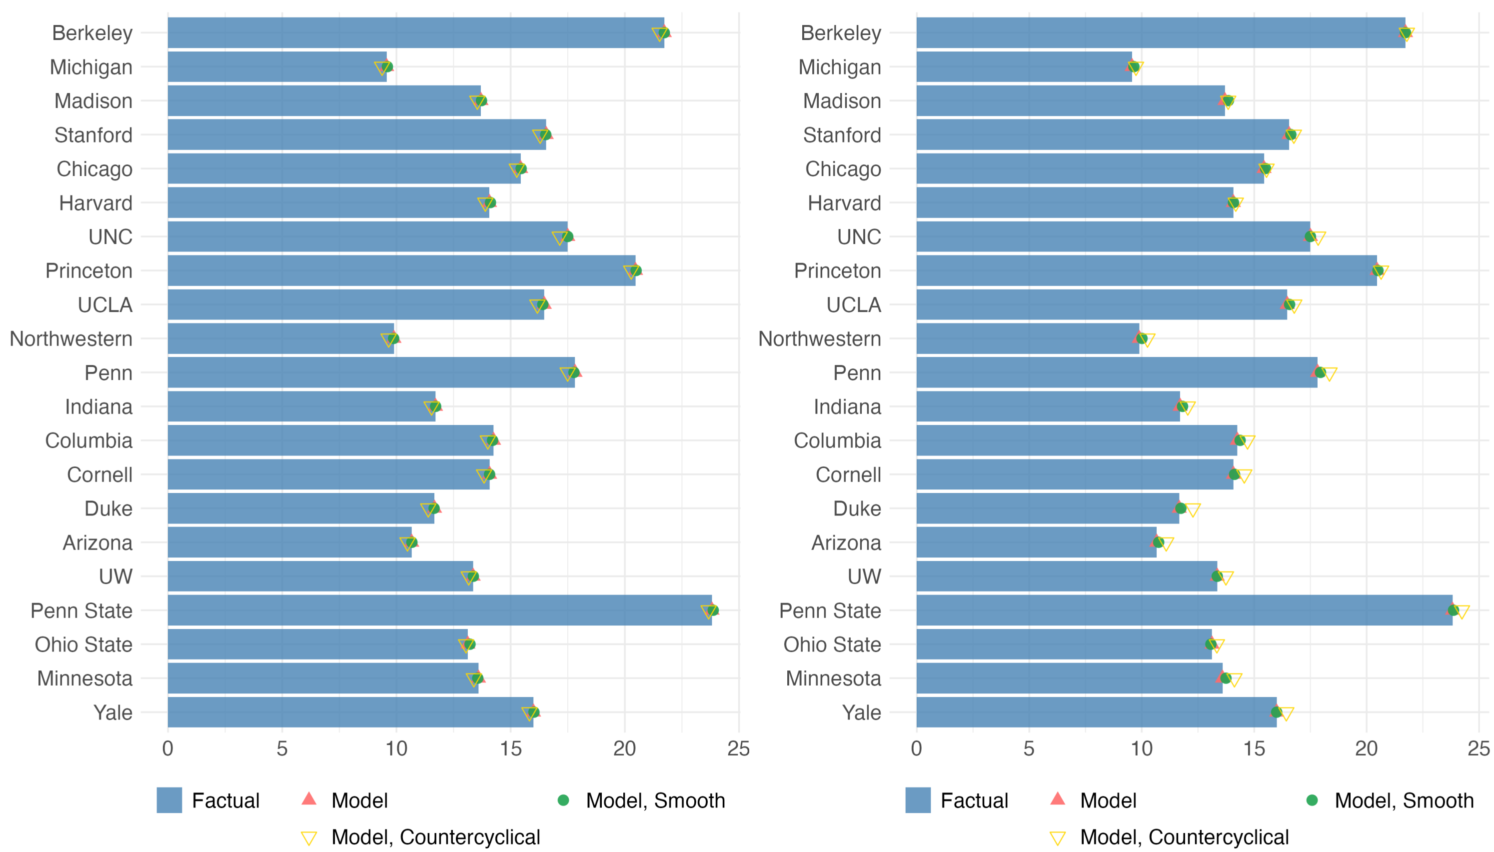

Notes: this figure shows results on per-slot payoff in four scenarios for each department ranked from top to bottom: the blue bars show empirical metrics, calculated from the factual data; the red triangles show the metrics from data simulated following the OLS model in Section 3.4 under the departments' factual slot allocation across years (that is, the factual S matrix). For each department, the green dot shows that department's counterfactual per-slot payoff when I replace that department's slot vector with a smoothed vector, while holding other departments' the same; the yellow triangle is the countercyclical analog of the green dot - see Section 4.2 for details on the construction of the smoothed and countercyclical hiring schemes. A department's payoff is calculated as the sum of ability θ across all its hires (across all years), as described in Section 3.1. Ranking of departments is from the 1998 and 2017 U.S. News *Best Sociology Programs*. Ability θ is measured by the number of social science publications within 6 years after hiring according to OpenAlex (Priem et al., 2022) - see Section 4.1 for details on data and variable construction. Panel (a) and (b) use the same linear model but different definitions of the market competition function $d(\cdot)$ as described in the subfigure title - see Section 3.4 for details on function definition.

Figure 8: 2SLS Model

(a) $d(\cdot)$ is leave-out sum of slots

(b) $d(\cdot)$ is sum of slots across better-ranked departments

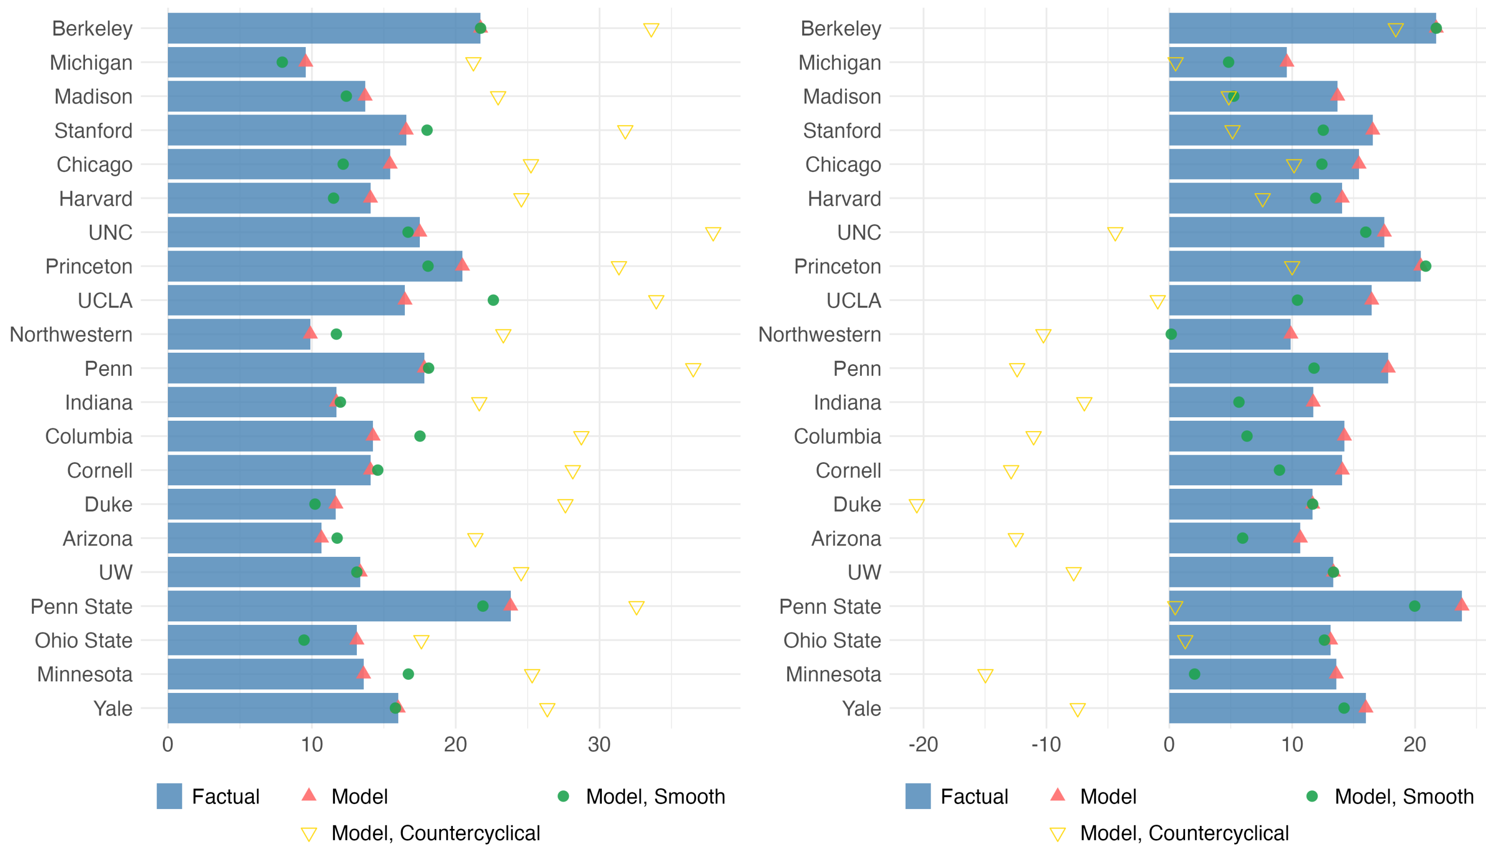

Notes: this figure shows results on per-slot payoff in four scenarios for each department ranked from top to bottom: the blue bars show empirical metrics, calculated from the factual data; the red triangles show the metrics from data simulated following the 2SLS model in Section 3.4 under the departments' factual slot allocation across years (that is, the factual S matrix). For each department, the green dot shows that department's counterfactual per-slot payoff when I replace that department's slot vector with a smoothed vector, while holding other departments' the same; the yellow triangle is the countercyclical analog of the green dot - see Section 4.2 for details on the construction of the smoothed and countercyclical hiring schemes. A department's payoff is calculated as the sum of ability θ across all its hires (across all years), as described in Section 3.1. Ranking of departments is from the 1998 and 2017 U.S. News *Best Sociology Programs*. Ability θ is measured by the number of social science publications within 6 years after hiring according to OpenAlex (Priem et al., 2022) - see Section 4.1 for details on data and variable construction. Instruments are endowment budget shocks defined in Sections 2.3 using the NACUBO and IPEDS data described in Section 2.2. Panel (a) and (b) use the same linear model but different definitions of the market competition function $d(\cdot)$ as described in the subfigure title - see Section 3.4 for details on function definition.

References

Avery, C., R. Ehrenberg, C. Hill, and D. A. Webber (2024): *Endowment Spending Rules*: University of Chicago Press, http://www.nber.org/chapters/c14978.

Azoulay, P., J. S. Graff Zivin, and J. Wang (2010): "Superstar extinction," *Quarterly Journal of Economics*, 125, 549–589, 10.1162/qjec.2010.125.2.549.

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

Borjas, G. J., and K. B. Doran (2012): "The Collapse of the Soviet Union and the Productivity of American Mathematicians*," *The Quarterly Journal of Economics*, 127, 1143–1203, 10.1093/qje/qjs015.

Brown, J. R., S. G. Dimmock, J.-K. Kang, and S. J. Weisbenner (2014): "How University Endowments Respond to Financial Market Shocks: Evidence and Implications," *American Economic Review*, 104, 931–62, 10.1257/aer.104.3.931.

Bulman, G. (2022): "The Effect of College and University Endowments on Financial Aid, Admissions, and Student Composition," Working Paper 30404, National Bureau of Economic Research, 10.3386/w30404.

Chen, J., and J. Roth (2023): "Logs with zeros? Some problems and solutions," *Quarterly Journal of Economics*, 139, 891–936, 10.1093/qje/qjad054.

Chitirala, I., and F. Ryan (2025): "SAS Ph.D. Admissions Set to Partially Rebound in Upcoming Cycle Following Fall 2025 Cuts," *The Daily Pennsylvanian*, https://www.thedp.com/article/2025/10/penn-arts-and-sciences-phd-admissions-increase.

Deming, D. J., and C. R. Walters (2017): "The Impact of Price Caps and Spending Cuts on U.S. Postsecondary Attainment," Working Paper 23736, National Bureau of Economic Research, 10.3386/w23736.

Dixit, A. K., and R. S. Pindyck (1994): *Investment Under Uncertainty*: Princeton University Press.

Fernandes, J. D., S. Sarabipour, C. T. Smith et al. (2020): "Research Culture: A survey-based analysis of the academic job market," *eLife*, 9, e54097, 10.7554/eLife.54097.

Freyaldenhoven, S., C. Hansen, and J. M. Shapiro (2019): "Pre-event trends in the panel event-study design," *American Economic Review*, 109, 3307–3338, 10.1257/aer.20180609.

Gale, D., and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, 9–15, 10.2307/2312726, Publisher: [Taylor & Francis, Ltd., Mathematical Association of America].

Goetzmann, W. N., and S. Oster (2015): "Competition Among University Endowments," in *How the Financial Crisis and Great Recession Affected Higher Education*, United States: University of Chicago Press.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624, 10.1257/aer.20181047.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277, 10.1016/j.jeconom.2021.03.014.

Greer, C. R., T. C. Ireland, and J. R. Wingender (2001): "Contrarian human resource investments and financial performance after economic downturns," *Journal of Business Research*, 52, 249–261, 10.1016/S0148-2963(99)00108-3.

Hansmann, H. (1990): "Why Do Universities Have Endowments?" *The Journal of Legal Studies*, 19, 3–42, `https://www.jstor.org/stable/724411`, Publisher: [University of Chicago Press, University of Chicago Law School].

Harvard University Financial Administration (2025): "Harvard's Endowment," `https://finance.harvard.edu/endowment`.

Huynh, J., and K. Shauman (2021): "Postdoctoral Hiring & Equity Issues in STEM: Employment Trends, Policy, and Research."

Ito, R., D. Chavarro, T. Ciarli, R. Cowan, and F. Visentin (2025): "From drains to bridges: The role of internationally mobile PhD students in linking non-mobile with foreign scientists," *Journal of Development Economics*, 177, 103577, 10.1016/j.jdeveco.2025.103577.

Kim, J. M., S. Roychowdhury, and S. Zheng (2025): "Talking to Hire: How Small Firms Pave their Way to Hire When Large Firms Lay Off," October, 10.2139/ssrn.5687022.

Lakonishok, J., A. Shleifer, and R. W. Vishny (1994): "Contrarian Investment, Extrapolation, and Risk," *The Journal of Finance*, 49, 1541–1578, 10.2307/2329262, Publisher: [American Finance Association, Wiley].

Lerner, J., H. J. Manley, C. Stein, and H. L. Williams (2024): "The Wandering Scholars: Understanding the Heterogeneity of University Commercialization," Working Paper 32069, National Bureau of Economic Research, 10.3386/w32069.

Levin, J., and J. Martinez (2025): "An Update on Staff Hiring," *Stanford Institutional News*, https://news.stanford.edu/stories/2025/02/staff-hiring.

Mao, W. C., and V. H. Paulus (2025): "Harvard FAS Cuts Ph.D. Seats By More Than Half Across Next Two Admissions Cycles," October, https://www.thecrimson.com/article/2025/10/21/fas-phd-admissions-cuts/.

Moody, J. (2025): "October Brought Deep Cuts at Multiple Campuses," *Inside Higher Ed*, https://www.insidehighered.com/news/business/cost-cutting/2025/11/07/october-brought-deep-cuts-multiple-campuses.

NACUBO (2023): "Historic Endowment Study Data," https://www.nacubo.org/Research/2024/Historic-Endowment-Study-Data.

National Center for Education Statistics (2005–2023): "Integrated Postsecondary Education Data System (IPEDS): Institutional Characteristics, Finance, and Fall Staff, 2005–2023," Washington, DC: U.S. Department of Education, Retrieved March 2025 from https://nces.ed.gov/ipeds/datacenter/DataFiles.aspx?

Oyer, P. (2006): "Initial Labor Market Conditions and Long-Term Outcomes for Economists," *Journal of Economic Perspectives*, 20, 143–160, 10.1257/jep.20.3.143.

Priem, J., H. Piwowar, and R. Orr (2022): "OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts," https://arxiv.org/abs/2205.01833.

Pulvino, T. C. (1998): "Do Asset Fire Sales Exist? An Empirical Investigation of Commercial Aircraft Transactions," *Journal of Finance*, 53, 939–978, 10.1111/0022-1082.00040, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/0022-1082.00040.

Reed, M. (2014): "When Searches Fail," https://www.insidehighered.com/blogs/confessions-community-college-dean/when-searches-fail.

Reiter, A. G. (2025): "U.S. News & World Report Historical Liberal Arts College and University Rankings," http://andyreiter.com/datasets/, Accessed: 2025-05-03.

Rundle, J. (2025): "University of South Florida Gets $40 Million to Start Cyber and AI College," *Wall Street Journal*, https://www.wsj.com/articles/university-of-south-florida-gets-40-million-to-start-cyber-and-ai-college-2a257e9a.

SHEEO (2025): "Grapevine Report," https://shefstage.sheeo.org/grapevine/.

Shleifer, A., and R. Vishny (2011): "Fire Sales in Finance and Macroeconomics," *Journal of Economic Perspectives*, 25, 29–48, 10.1257/jep.25.1.29.

Shleifer, A., and R. W. Vishny (1992): "Liquidation Values and Debt Capacity: A Market Equilibrium Approach," *The Journal of Finance*, 47, 1343–1366, 10.2307/2328943, Publisher: [American Finance Association, Wiley].

The Economist (2025): "Ivy League universities are on a debt binge," *The Economist*, https://www.economist.com/finance-and-economics/2025/08/13/ivy-league-universities-are-on-a-debt-binge.

Tobin, J. (1974): "What Is Permanent Endowment Income?" *The American Economic Review*, 64, 427–432, https://www.jstor.org/stable/1816077, Publisher: American Economic Association.

Turner, S. E. (2013): *The Impact of the Financial Crisis on Faculty Labor Markets*, 175–207: University of Chicago Press, 10.7208/chicago/9780226201979.003.0006.

U.S. News & World Report (1998): "The Best Sociology Programs in America, Ranked."

——— (2017): "The Best Sociology Programs in America, Ranked."

Wapman, K. H., S. Zhang, A. Clauset, and D. B. Larremore (2022): "Quantifying hierarchy and dynamics in US faculty hiring and retention," *Nature*, 610, 120–127, 10.1038/s41586-022-05222-x, Publisher: Nature Publishing Group.

Warren, J. R. (2019): "How much do you have to publish to get a job in a top sociology department? Or to get tenure? Trends over a generation," *Sociological Science*, 6, 172–196.

Yale University (2022): "Overview of Yale's Endowment," https://www.yale.edu/funding-yale-home/overview-yales-endowment.

——— (2023): "2202 Endowment Spending and Distributions | It's Your Yale," https://your.yale.edu/policies-procedures/2202-endowment-spending-and-distributions.

Zhang, S., K. H. Wapman, D. B. Larremore, and A. Clauset (2022): "Labor advantages drive the greater productivity of faculty at elite universities," *Science Advances*, 8, eabq7056, 10.1126/sciadv.abq7056, Publisher: American Association for the Advancement of Science.

Appendix for

"Should Universities Smooth Faculty Hiring? Theory and Evidence"

Zhizhong Pu, *Harvard University*[1]

A Additional Theoretical results

A.1 Proof of Claim 1

Proof. Fix department i. Its payoff is, by definition

$$u_i = \sum_{t=1}^{T} w_t \, v_{it} = \sum_{t=1}^{T} w_t \sum_{j \in H_{it}} \theta_j$$

. Then its expected payoff given its information set $\mathbf{X}_j$ is

$$\mathbb{E}[u_i \mid \mathbf{X}_j] = \sum_{t=1}^{T} w_t \sum_{j \in H_{it}} \mathbb{E}[\theta_j \mid \mathbf{X}_j]$$

Note that, by definition, $|H_{it}| \equiv s_{it}$ for each t. Since slots are given, maximizing $\mathbb{E}[u_i \mid \mathbf{X}_j]$ amounts to selecting the s_{it} candidates with largest $\mathbb{E}[\theta_j \mid \mathbf{X}_j]$ in each year. Equivalently, departments rank candidates by $\hat{\theta}_j(\mathbf{X}_j)$. □

A.2 Proof of Claim 2

Proof. Department i hires s_i candidates. Let $\mathcal{J}_i$ denote the set of candidates matched to department i, with $|\mathcal{J}_i| \equiv s_i$. Let i' denote the department matched to candidate $j \notin \mathcal{J}_i$.

Suppose for contradiction that (i, j) is a blocking pair. This requires that: (i) candidate j prefers i, so $r(i) < r(i')$; (ii) department i prefers j to its worst match: $\hat{\theta}_j > \min_{j' \in \mathcal{J}_i} \hat{\theta}_{j'}$

From (i), department i moved before i' in the serial dictatorship process. Since j is matched to i', j must have been available when i' moved. For j to be available when i' moved, no department moving before i' selected j. Since i moved before i', j was available when i moved.

From (ii), we have $\hat{\theta}_j > \min_{j' \in \mathcal{J}_i} \hat{\theta}_{j'}$. Since j was available and i would prefer j over at least one member of $\mathcal{J}_i$, i would have selected j. This contradicts with the fact that i did not select j. Therefore, no blocking pair exists. □

[1] E-mail: zpu@hbs.edu

A.3 Static Model of Budget-Dependent Hiring

Consider a university that maximizes its utility by choosing faculty headcount H and non-faculty expenditures Z, subject to a budget constraint. The model makes three key assumptions:

1. Complementary inputs: Faculty and non-faculty resources are complements in production
2. Fixed faculty wages: The cost per faculty member is institutionally fixed[2]
3. Revenue shocks: The university faces exogenous changes in available revenue

The university's preferences are represented by a CES utility function:

$$U(H, Z) = [\alpha H^{\rho} + (1-\alpha)Z^{\rho}]^{\frac{1}{\rho}} \tag{4}$$

where $\rho < 1$ governs the elasticity of substitution[3]. The university faces budget constraint:

$$wH + pZ \leq B$$

where w is the fixed cost per faculty member; p is the unit cost of non-faculty expenditures, and B is total available budget. The university's optimization problem thus yields first-order condition:

$$\frac{\partial U/\partial H}{\partial U/\partial Z} = \frac{w}{p}$$

Solving this gives the optimal faculty headcount:

$$H^{*} = \frac{\alpha^{\frac{1}{1-\rho}} B}{\alpha^{\frac{1}{1-\rho}} w + (1-\alpha)^{\frac{1}{1-\rho}} p}$$

The key result shows that faculty hiring increases with budget constraint:

$$\frac{\partial H^{*}}{\partial B} = \frac{\alpha^{\frac{1}{1-\rho}}}{\alpha^{\frac{1}{1-\rho}} w + (1-\alpha)^{\frac{1}{1-\rho}} p} > 0$$

Proposition 1. *For $\rho < 1$ and fixed w, an exogenous increase in budget constraint B leads to an increase in optimal faculty headcount H^{*}.*

[2]This is reasonable because it's difficulty to either raise wage for faculty to raise marginal utility or replace low-wage low-utility faculty with high-wage high-utility ones. This assumption rules out adjustment on the intensive (quality) margins.

[3]The elasticity of substitution is $\sigma = 1/(1-\rho)$, with $\rho \to -\infty$ representing perfect complements, and $\rho \to 1$ representing perfect substitutes

Figure 1: Number of Faculty Hires by Year, IPEDS vs. University Announcement

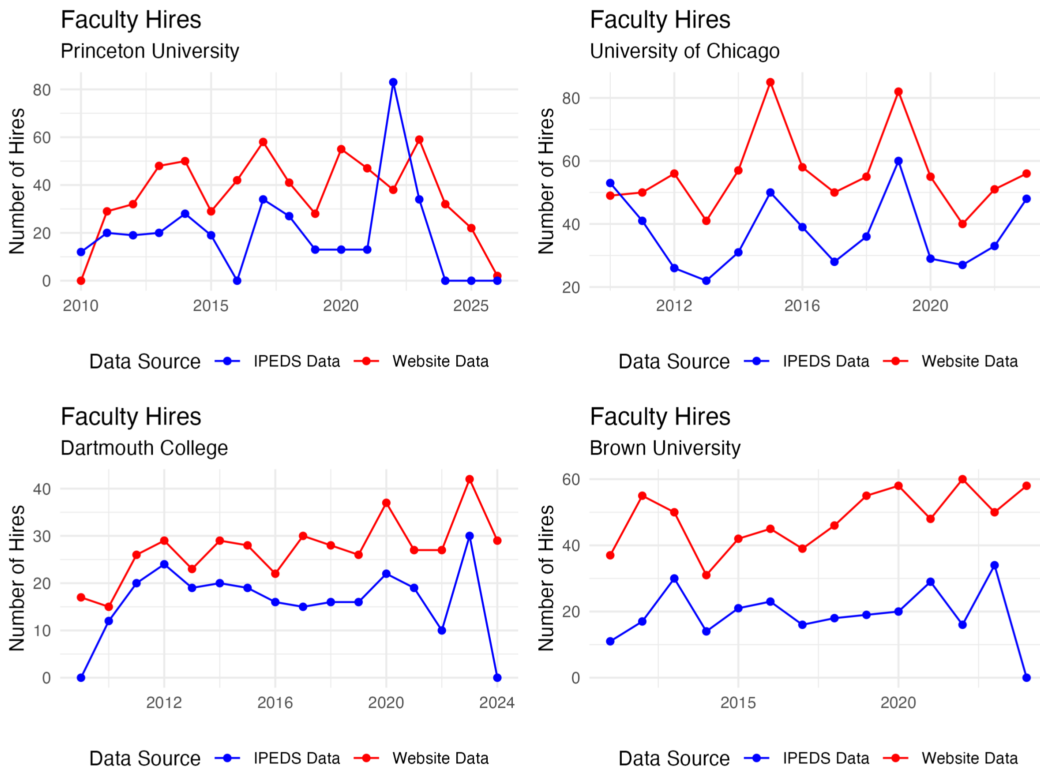

Notes: This plot shows the time series of the number of faculty hires in four private universities. The blue series are based on data from IPEDS. The red series are based on my hand-collected hiring announcement data from university websites.

Given this institutional background and the static model in Section A.3, I sketch a dynamic hiring model. For private universities, endowment market-value shocks occurring at time T generally translate into observable budget changes around $T + 1$ or $T + 2$ due to payout smoothing (Yale University, 2022; Avery et al., 2024). However, given labor market frictions, universities may adjust faculty hiring decisions proactively in anticipation of these expected budget changes after capital market conditions are realized. This could shift hiring responses forward relative to the realized expenditure reactions. For public institutions, state appropriations influence budgets contemporaneously. However, the multi-year state budgeting process allows universities to anticipate funding changes. Therefore, faculty hiring at public universities may reflect anticipatory adjustments influenced by early legislative discussions and economic forecasts rather than solely finalized budget outcomes.

B Methodological Details

C Additional Empirical Results

C.1 Additional Market Competition Function

In the reduced-form identification in the main text, I consider two parsimonious market competition functions. Here I consider another one motivated by candidates' choice problem under a random utility model set-up.

Let $d(\vec{s_t}; i) := C_{it}$ represent the market competition faced by department i in year t stemming from the number of slots at all departments, where C_{it} is a choice probability in the following set-up.

Fix year t and drop year subscripts. Consider candidate j's choice problem among a set of $||\vec{s}||$ alternatives each being a slot at a department in year t.

Assumption 2. *Candidate j's utility for department i is*

$$\mu_{ji} = \nu_i + \varepsilon_{ji}$$

where ν_i is a representative utility and ε_{ji} is a stochastic taste shock that follows i.i.d. Type-I extreme value distribution independent of j's characteristics.

Then a candidate's probability of choosing a specific slot at department i is

$$\frac{\exp(\nu_i)}{\sum_{\iota=1}^{I} s_\iota \exp(\nu_\iota)} =: C_i$$

where exp denotes the exponential function.

Two features make the choice probability C_i an appealing summary of the inter-department hiring competition. First, increasing $||\vec{s}||$ strictly decreases C_i while strictly increases competition. Second, increasing C_i weakly increases the $\hat{\theta}$ of the candidates that fill department i's slots. To see this, fix a given slot and assume a deferred acceptance process where the candidates propose. A larger C_i weakly expands the set of candidates who offer to (the given spot at) i in the first round, which translates to a weakly higher lower bound of $\hat{\theta}_j$ across i's feasible option set. Thus, i hires candidates with weakly higher $\hat{\theta}$.

This model has two additional features. First, a unit change in $s_{i'}$ changes C_i more when i' is generally more preferred (i.e. has higher ν). Second, department-candidate fit does not matter conditional on ν_i, in that the effect of a unit change in $s_{i' \neq i}$ on C_i does not depend on j's own characteristics.

Appendix Figure 2 presents the results when applying this model to the data.

C.2 Sampling Uncertainty in Reduced-form Model

Figure 2: Per-slot Payoff by Department, Reduced-Form Models with Alternative Competition Function

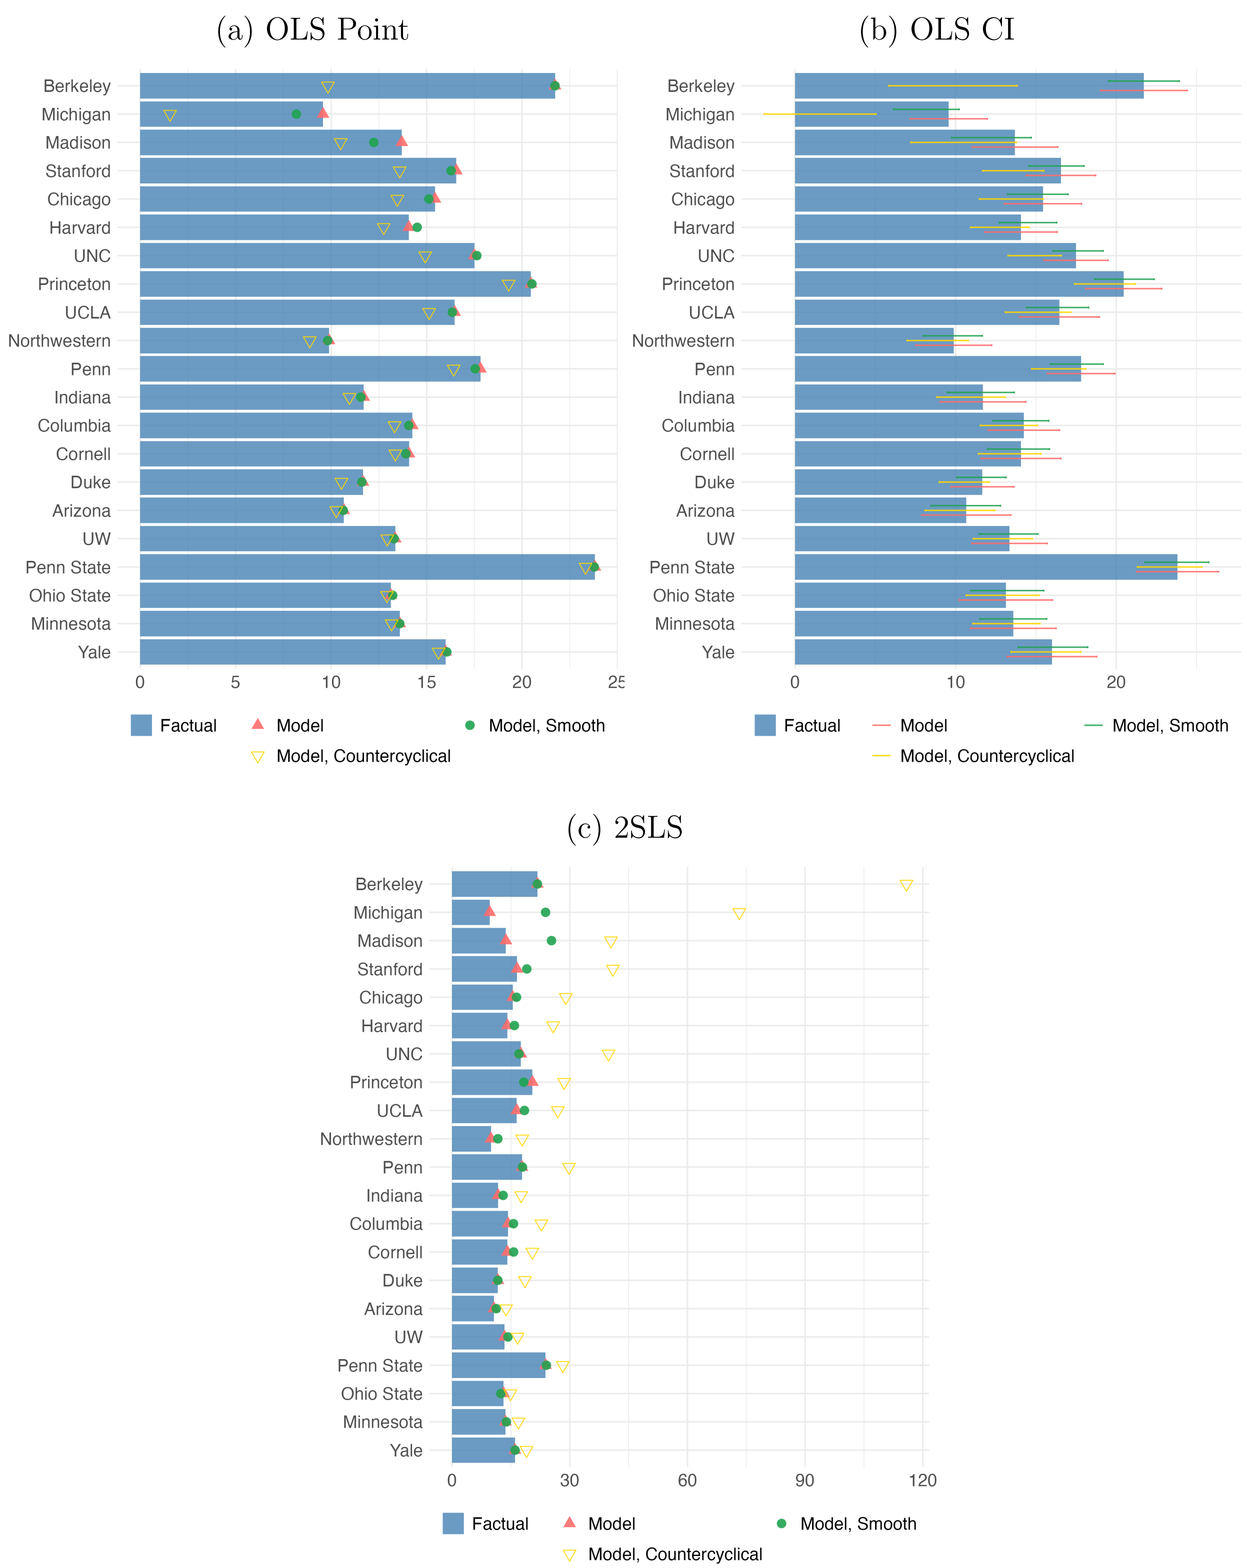

Figure 3: Per-slot Payoff by Department, Reduced-Form Model with Confidence Intervals

(a) Leave-out (CI) (b) Better (CI)