Dynamic Spatial Competition in Early Education: An Equilibrium Analysis of the Preschool Market in Pennsylvania*

Pierre Bodéré†

September 4, 2026

Abstract

High-quality preschool generates large social returns, yet U.S. preschool markets feature unequal access, high prices, and low quality. Using Pennsylvania data, I develop and estimate a dynamic model of entry, exit, quality investment, and spatial price competition. I evaluate broader demand-subsidy eligibility and quality-tiered reimbursements for serving subsidized children at high quality. Dynamic provider responses shape policy incidence and cost-effectiveness over time: entry and investment expand supply, amplifying enrollment gains and extending policy benefits beyond subsidized families. Broader eligibility without quality-tiered reimbursements is most cost-effective in consumer-surplus terms, but quality targeting is crucial for directing low-income children toward high-quality centers.

*I am grateful to Chris Conlon, Mike Dickstein, Guillaume Fréchette and Daniel Waldinger for their guidance and support. I thank Brian McManus for a helpful discussion. This project benefited from conversations with Claudia Allende, Milena Almagro, Tslil Aloni, Angela Crema, Elizabeth Davis, Chris Flinn, Alejandro Ganimian, Mike Gilraine, Karam Kang, Adam Kapor, Elena Manresa, Marguerite Obolensky, Paul Scott, Fiona Scott Morton, Katja Seim, Aaron Sojourner, Teresa Steininger, Kevin Williams, Seth Zimmerman, and from audiences at various seminars and conferences. I am especially indebted to Jonathan Elliott for his help with gathering the data. All errors are my own.

†Yale School of Management, pierre.bodere@yale.edu

1 Introduction

High-quality early childhood education provides benefits that persist well into adulthood, in particular for disadvantaged children, offering the prospect of a cost-effective approach to reducing educational inequality (Elango et al., 2015; Hendren and Sprung-Keyser, 2020). Though much of the evidence on large-scale preschool expansion comes from universal pre-K, in the U.S. more than 80% of centers are private[1]. Scaling high-quality early education in decentralized markets has proved challenging: providers face high costs, and parents face low quality and high prices, resulting in low enrollment.[2] In response, most U.S. states have deployed a combination of means-tested subsidies for parents and quality incentives for private providers, with growing calls from policymakers and researchers to expand these policies to further improve access, affordability, and the quality of the early education supply.[3]

Designing a cost-effective expansion of these policies requires accounting for providers' equilibrium responses, which shape the distribution of high-quality enrollment across socioeconomic groups and the resulting educational benefits. Demand subsidies can expand the formal market, and the extent to which preschools adjust their prices in response varies with parents' price sensitivity and nearby competition. Whether and how quickly a larger subsidized market attracts new providers and induces quality improvements in turn depends on providers' entry and investment costs. Even when supply responds, high-quality provision need not expand in disadvantaged neighborhoods if local willingness to pay is low relative to the cost of quality, unless the policy raises the returns to investing there.

To evaluate the incidence and cost-effectiveness of early education policies, I develop and estimate a dynamic spatial competition model that endogenizes preschool prices, enrollment, and market structure. Families differ in how they trade off out-of-pocket prices, travel distance, and quality. On the supply side, each year, potential entrants weigh the up-front costs of licensing, staffing, and opening a center at their respective locations against the stream of profits they expect to earn after entry. Once active, providers set prices to maximize current profits, taking into account nearby demand and competition in a Nash–Bertrand

[1] See Cascio and Schanzenbach (2013) and Gray-Lobe et al. (2023) for evaluations of U.S.-based universal pre-K programs. NSECE Project Team (2014) discusses the public-private composition of U.S. formal early education providers.

[2] Grunewald and Davis (2011) review a series of facts on early education providers' profits and turnover. See Section 2.1 for a discussion on private provision and enrollment rates. According to the National Center for Education Statistics, less than a third of center-based arrangements are rated high quality (U.S. Department of Education, 2014).

[3] See Hotz and Wiswall (2019) for a general review of early education policies, and Cannon et al. (2017) for an analysis of quality rating and improvement systems (QRIS) in particular. Department of the Treasury (2021) and Davis and Sojourner (2021) provide policy proposals to expand existing policies.

pricing game. They also decide whether expected future profits justify remaining open and, if so, at what quality to operate the following year. Moving to a higher rating requires meeting progressively more demanding standards for staff qualifications and professional development, classroom practices, and leadership and management, so providers trade off the current cost of improving quality against potentially higher demand and greater subsidy payments tied to ratings in subsequent years. The returns to entry and quality investment also depend on nearby rivals, whose entry, exit, and quality choices alter a center's future demand and profits. I capture these interdependent choices in a dynamic discrete game in which providers act simultaneously and quality is a durable state that evolves through costly investment.

To bring this model to the data, I assemble a panel covering all licensed centers in Pennsylvania from 2010 to 2018, as well as detailed information on the prevailing subsidy regime in each year. The state rates centers through Keystone STARS, a four-level scale ranging from STAR 1 (lowest) to STAR 4 (highest). Pennsylvania also combines means-tested demand subsidies with quality-tiered reimbursements, known as add-ons, paid to higher-rated providers serving subsidized children. I document four facts that highlight the role of providers' equilibrium responses across neighborhoods in shaping the distributional consequences of preschool policies. First, access to early education, including to high-quality centers, is unequally distributed. Neighborhoods with more college-educated women have more centers within a 10-minute drive, particularly high-quality ones. Second, providers' dynamic decisions reinforce these disparities: centers exit more and invest less in lower-income neighborhoods. Third, upgrades raise enrollment and prices, with larger price increases in higher-income neighborhoods, consistent with higher-income families valuing highly rated care more and providers responding to local willingness to pay. Lastly, providers' investment choices respond to quality-tiered reimbursements: an additional $100 in expected daily subsidy payments raises a center's probability of upgrading from STAR 3 to STAR 4 by approximately 3 p.p.

I estimate the model in two steps. I first estimate families' preferences and providers' marginal costs within each year, holding the set and quality of centers fixed. Changes in Pennsylvania's subsidy schedules generate variation in families' out-of-pocket prices across income groups, while changes in quality-tiered reimbursements shift providers' revenues across ratings. I combine this policy variation with cost shifters such as local teacher wages, instruments based on changes in nearby competitors' characteristics, and survey data on preschool choices. The estimates imply that higher-income and college-educated families are willing to pay more and travel longer for highly rated centers. Daily marginal costs are $3.4 higher at the highest than at the lowest rating, reflecting both treatment and selection. Market power is modest at the median center, with a markup of around 15% of price, but varies substantially across centers and neighborhoods. Consistent with spatial competition, within-center profits fall as nearby rivals' quality rises but increase with local high-income demand and when the center operates at the highest rating, implying that returns to quality vary across locations.

Second, I estimate the parameters governing providers' incentives to enter, exit, and invest in quality, thereby endogenizing the preschool market structure. Nesting the rich stage-game model of preschool pricing and enrollment into a dynamic game creates a curse of dimensionality because the state space must capture numerous heterogeneous competitors and detailed local demographics. To maintain computational tractability, I train neural networks to approximate stage-game profits, learn a parsimonious representation of the state, and then apply value function approximation. I identify entry, fixed, and upgrade costs by combining observed provider transitions with the stage-game mapping from local market states into flow profits. The resulting parameters and equilibrium strategies reproduce key transition patterns in the data: exit rates decline with quality, and upgrades are rare and generally occur one rating at a time, consistent with providers building up quality over time. Combined estimates of fixed and marginal costs imply annual early education provision costs of approximately $8,000 per child, close to industry estimates for Pennsylvania (Friedman-Krauss et al., 2018). Estimated entry costs are moderate, around $450,000, in line with typical state grants supporting center openings. Investment costs of reaching the next quality tier range from 50% to 80% of entry costs, while skipping tiers is considerably more expensive. These adjustment costs imply that more generous subsidies can change prices and enrollment immediately, but expand and improve the preschool supply only gradually.

I then use the estimated model to isolate the economic forces underlying providers' equilibrium responses by comparing three regimes: no subsidies, means-tested demand subsidies without quality-tiered reimbursements (quality-neutral), and means-tested demand subsidies combined with them (quality-targeted). Holding market structure fixed, demand subsidies have two countervailing effects on provider market power. By shielding eligible families from posted-price variation, they reduce the price sensitivity of demand faced by some centers and raise markups. At the same time, they draw price-sensitive low-income families into formal care, making the marginal consumers faced by other centers more price sensitive and lowering markups. These forces generate substantial heterogeneity: moving from no subsidies to the Pennsylvania environment causes around one fifth of centers to increase markups, and more than 10% to cut them by more than 23%. As market structure adjusts, demand subsidies induce entry, while quality-tiered reimbursements steer entry and quality improvements toward high-quality care in lower-income neighborhoods, lowering access barriers. These supply responses feed back into pricing: additional nearby entry lowers incumbent markups. Targeted policies thus spill over to other families through prices on impact and through access, quality, and competition over time.

Informed by these equilibrium responses, I evaluate plausible counterfactual preschool policies that vary along two dimensions: the eligibility threshold for demand subsidies and the generosity of tiered reimbursements that make these subsidies quality-targeted. Relative to the Pennsylvania policy environment, removing subsidies reduces enrollment by approximately 37%, driven by the collapse of low-income participation, particularly at highly rated centers. Quality-tiered reimbursements are crucial for directing low-income children toward high-quality centers: without these incentives, enrollment shifts from higher- to lower-rated centers. Broadening eligibility to include middle-income families while retaining quality-tiered reimbursements raises overall enrollment by 11% and high-quality enrollment by 23% on impact. Providers' dynamic responses amplify these effects. Center closures deepen the enrollment shortfall in the no-subsidy scenario, whereas entry and quality investment compound the gains from eligibility expansion, raising high-quality enrollment by 39% after seven years alongside a 24% expansion in the number of high-quality centers. The dynamic responses also have distributional consequences. Holding eligibility fixed, more generous quality-tiered reimbursements induce providers to invest in higher quality, extending the benefits of the policy beyond directly targeted families. After seven years, the resulting expansion in access to high-quality options generates consumer surplus gains for unsubsidized high-income families that can be almost half as large as those accruing to subsidized families.

I evaluate these policies according to two measures of their marginal value of public funds (MVPF) that provide complementary perspectives on their benefits. I consider consumer surplus, which places the emphasis on parents' preferences, and calibrated educational benefits, which motivate policymakers' interest in high-quality preschool expansions. Across all policies, the measures highlight the importance of considering providers' dynamic responses when evaluating preschool policies. For instance, the consumer-surplus MVPF of Pennsylvania's policies increases from 0.32 on impact to 0.40 after seven years, a 26% improvement resulting from provider entry, quality investment, and increased price competition in response to the subsidies. From the perspective of consumer surplus, a quality-neutral expansion of demand subsidies generates the highest MVPF, as this policy expands the benefits of more affordable preschool to families with a strong demand for early education. When high-quality preschools generate larger educational gains, however, educational-benefit MVPF measures favor quality-targeted demand subsidies, as those policies successfully match children who benefit the most to high-quality centers.

Related Literature. This paper builds on an expanding literature evaluating the effects of early childhood education programs on child development in North America. Studies of targeted programs, including Perry Preschool, the Infant Health and Development Program, and Head Start, find large short- and long-run benefits for participants (Heckman et al., 2013; Duncan and Sojourner, 2013; Elango et al., 2015; Kline and Walters, 2016; Ludwig and Miller, 2007). Evidence from large-scale universal programs is more mixed, with positive effects in Georgia, Oklahoma, and Boston but detrimental effects in Tennessee and Québec (Cascio and Schanzenbach, 2013; Gray-Lobe et al., 2023; Durkin et al., 2022; Baker et al., 2008, 2019). Taken together, this evidence suggests that preschool's benefits depend on both program quality and the quality of children's available alternatives (Cascio, 2015). Motivated by these ex post evaluations, I develop an equilibrium model in which large-scale policies determine both the quality of preschool supplied and which children attend, thereby endogenizing two central determinants of their cost-effectiveness.

This paper also contributes to the growing literature on the supply side of early childhood education. Ex post studies provide empirical support for the mechanisms embedded in my equilibrium model: policy changes can crowd out private centers, alter capacity and quality, induce exit, and stimulate entry (Bassok et al., 2014; Brown, 2018; Hotz and Xiao, 2011; Lee et al., 2025). I bring these adjustment margins together in a dynamic model to evaluate how providers' entry, exit, quality investment, and pricing responses shape the effects of early education policies. A growing strand of papers featuring general-equilibrium models emphasizes parental labor-supply responses (Borowsky et al., 2022; Berlinski et al., 2020); I focus on provider-side equilibrium and the resulting spatial distribution of access and quality.

This paper extends the literature studying education markets in equilibrium by incorporating providers' dynamic decisions. Neilson (2013), Allende (2019), and Armona and Cao (2022) study schools' and colleges' responses to policies using a differentiated product model of education providers. More closely related to this paper, existing school-choice models incorporate provider entry and exit through static games (Singleton, 2019; Dinerstein et al., 2020). In contrast, I model the dynamic considerations of early education providers when deciding whether to enter, exit, or invest in quality. The combination of this dynamic model with a differentiated-product stage game makes this framework well suited to evaluating the range of policies considered in early education.

More broadly, this paper contributes to the study of subsidy design in markets with imperfect competition in which policymakers may prioritize one group of consumers. In early education, as in other markets ranging from K-12 (Neilson, 2013; Allende, 2019) to healthcare (Polyakova and Ryan, 2023; Tebaldi, 2025), targeted subsidies may change the residual demand faced by providers with market power, leading them to change prices or product characteristics, thereby benefiting or harming unsubsidized consumers. Beyond these short-run responses, I allow providers to enter, exit, and invest in quality selectively across space, so policy incidence and cost-effectiveness evolve as the supply side adjusts. Because subsidized and unsubsidized families share local providers, the resulting changes in access, quality, and competition also generate spillovers akin to the "preference externality" often at play in differentiated-product markets with spatially heterogeneous consumers (Waldfogel, 2003).

Finally, this paper contributes to the empirical literature on dynamic games, where large state spaces often motivate two-step estimators that recover value functions from first-stage conditional choice probabilities (Ryan, 2012; Collard-Wexler, 2013; Caoui et al., 2026). These papers rely on a reduced-form model of profits, but this approach is less attractive here because early education policies alter both providers' dynamic incentives and the pricing and enrollment equilibrium. I therefore follow Sweeting (2013) and rely on structural profits from the stage-game equilibrium combined with value-function approximation in a full-solution approach. I use neural networks to approximate equilibrium profits and generate data-driven bases for providers' value functions, partly circumventing the need for the researcher to hand-craft basis functions capturing complex economic interactions. This combination makes full-solution estimation and counterfactual analysis feasible in a rich spatial game with endogenous prices, entry, exit, and quality investment.

2 Institutional Setting and Data

I assemble a new administrative panel of licensed center-based preschool providers in Pennsylvania. The data combine information on providers' prices, enrollment, and measured quality with detailed records on public subsidy policies and payments. I link these records to geographically detailed Census data on preschool-aged children and their household characteristics and supplement them with survey evidence on families' early education arrangements. This novel combination makes it possible to jointly study families' preschool choices and providers' pricing, entry, exit, and quality-investment decisions. I first describe the institutional setting and then present the data.

Pennsylvania Early Education Policies. The federal and state governments have implemented several policies to bolster preschool markets and improve access, affordability, and quality of early education. The key federal funding program relevant to this paper is the Child Care and Development Fund (CCDF). The CCDF provides grants to states to subsidize early education for low-income parents and distribute grants to providers. States overwhelmingly operationalize these funding schemes through quality rating and improvement systems (QRIS). As of 2017, all U.S. states except Mississippi have a QRIS either planned or implemented (Cannon et al., 2017). Pennsylvania runs a well-established QRIS called Keystone STARS that launched in 2002. The main demand-subsidy branch is an income-based subsidy to families called Child Care Works (CCW), which serves eligible children from birth to age 13. This program represents the vast majority of subsidized preschool enrollment over my sample period.[4] This policy reduces the consumer price faced by families below 200% of the Federal Poverty Line (FPL). [5] Subsidized rates are governed by a copayment schedule calibrated to represent approximately 7% of a family's income. Providers serving subsidized children thus receive parents' copayments and a subsidy payment. These subsidy payments are capped by a Maximum Child Care Allowance (MCCA), such that, if a provider charges a high private-pay rate, the sum of the copayment and the subsidy payment may not fully cover the provider's private-pay rate. In those instances, the Pennsylvania code allows providers to charge a top-up. In the model, I assume that providers systematically do so.

The system rates centers on a four-level scale, where STAR 1 is the lowest possible rating and STAR 4 the highest.[6] The quality ratings encompass four dimensions of improvement for providers: staff qualifications, early childhood education program, partnership with families, and management practices. The two highest ratings require a fraction of employees to have completed or be enrolled in a child development program. Centers are periodically assessed on the quality of their teaching using measures such as the environment rating scale, a standard metric found to be positively associated with child development (Keys et al., 2013). Programs can apply for a higher rating at any time. Upgrading quality incurs various sunk costs, from meeting certification requirements to raising staff qualifications. Higher quality also involves higher operating costs arising from increased wages and more stringent standards. A center can downgrade to a lower rating if the requirements to maintain its current rating are not met, even though downgrading is rare in the data.

Keystone STARS also provides incentives for providers to operate at high quality when serv-

[4] The state also provides preschool via PA-PreK Counts, but its enrollment is low over my sample period, and I therefore focus on CCW (Friedman-Krauss et al., 2018).

[5] I ignore additional work and study eligibility requirements as applicant data are unavailable.

[6] Before 2016, when STAR 1 became the floor, centers could be "No STAR". For consistency I bundle "No STAR" and STAR 1.

Figure 1. Example of Consumer and Provider Prices

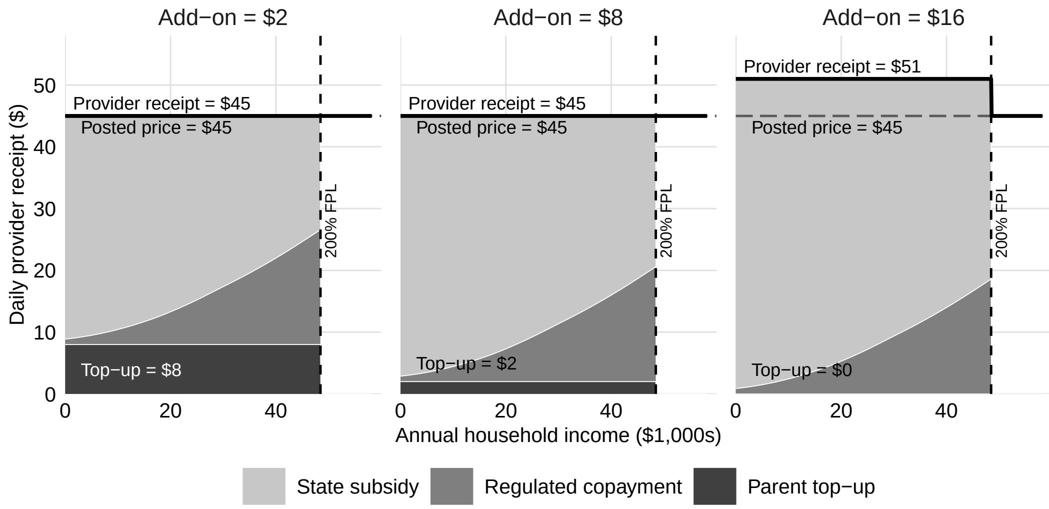

Notes: The figure illustrates the subsidy accounting formulas for a provider with a daily private-pay price of $45, in a region with an MCCA of $35, and, across panels, add-on rates of $2, $8 and $16. The vertical dashed line marks the 200% FPL eligibility threshold in the 2018 copayment schedule.

ing low-income children. This incentive scheme takes the form of quality-tiered reimbursements called add-on rates. When the subsidy payment falls short of a provider's private-pay rate, add-ons reduce the top-up that subsidized families can be charged. Once the top-up has been eliminated, I assume that further increases in add-ons raise the provider's receipts from serving subsidized children.[7] Thus, add-ons operate first as a quality-targeted demand subsidy and, once they exhaust the top-up, as a supply-side subsidy for providers serving subsidized children at high quality. Figure 1 illustrates this mechanism, holding the posted price and copayment schedule fixed while progressively increasing the add-on rate; the middle panel approximates the STAR 4 add-on rates observed near the end of the sample period. The medium- and dark-gray regions represent parents' copayments and top-ups, respectively, while the light-gray region represents the state subsidy and the black line reports total provider receipts.

Data Sources. I assemble a unique dataset covering the universe of Pennsylvania providers from 2010 to 2018, including enrollment, prices, capacity, quality ratings, staffing, and relevant policy schedules.[8] Enrollment is measured in full-time-equivalent units at the classroom level, where I also observe the number of teachers and aides. Policy variables include the income-based copayment schedule for subsidized families, the maximum child care allowance,

[7] Pennsylvania QRIS documents note that add-ons can push subsidy payments above MCCA, and there is no rule that specifies they cannot push provider payments above private-pay rates.

[8] These data were obtained through multiple Freedom of Information Act requests.

and quality-tiered add-on rates. I restrict the sample to centers, excluding home- and family-based providers, to focus on the formal early education sector and providers facing similar cost structures and incentives. I complement these data with the American Community Survey (ACS) to characterize preschool demand, defined as households with 3- or 4-year-old children. To model demand at a disaggregated spatial level, I use five-year ACS block-group estimates of population by age, gender, income, and educational attainment, supplemented with tract-level variables when necessary.[9] Finally, I use the National Survey of Early Care and Education (NSECE) to construct additional demographic moments on parents' preschool choices, including distance traveled and costs.

Descriptive Statistics. Table 1 provides descriptive statistics for centers operating in 2018 by quality rating. There are 3,929 centers serving preschool-aged children in the last year of the sample. Around half of centers operate at the lowest quality level. Provider characteristics differ substantially across quality ratings: daily prices are highest for STAR 4 centers, at $40.25 in 2015 dollars, compared with $35.00 for STAR 1 providers. Highly rated centers also tend to be larger and to enroll more children: the average licensed capacity of STAR 4 centers is 123.47 children, compared with 71.28 for STAR 1 centers, and STAR 4 centers serve 44.10 children on average, compared with 24.00 for STAR 2 centers.[10] Private-pay enrollment represents roughly one-half of preschool enrollment, consistent with industry reports stating that parent tuition provides the majority of center revenue (NSECE Project Team, 2014). Finally, the data suggest that the vast majority of centers are not close to binding licensed capacity constraints. Even when using classroom-size requirements to construct a more stringent measure of capacity, only 15 percent of providers operate at 90 percent or more of their capacity, with little difference across quality tiers. Similarly, using staff-to-child ratios to construct capacity implies that fewer than 30 percent of centers operate at 90 percent or more of capacity.[11] Appendix Figure B2 shows the distribution of enrollment relative to this classroom-capacity measure across years.

[9] See Appendix C for the imputation of education-specific income distributions at the block-group level.

[10] Enrollment in STAR 1 is not always reported. In estimation, I impute it as described in Appendix B.

[11] Staff-to-child ratios are 10:1 for preschool in Pennsylvania. Unlike classroom space or building size, however, staffing can be adjusted by hiring additional workers.

Table 1. Descriptive Statistics of Pennsylvania Centers

Centers	STAR 1	STAR 2	STAR 3	STAR 4
A. Center Characteristics				
Number of Centers	2,025	687	461	756
Preschool Daily Price	35	33.91	37.36	40.25
Licensed Capacity	71.28	85.37	112.26	123.47
Accreditation Status	0.03	0.02	0.01	0.35
B. Enrollment Characteristics				
Full Time Eq. Preschool Enrollment		24	36.23	44.1
Share Private Pay Enrollment		0.48	0.48	0.5
Number Preschool Teachers		3.58	5.17	5.97
Number of Preschool Classrooms		1.95	2.81	3.23
Share 90% of Classroom Capacity		0.16	0.15	0.16

Notes: Panel A reports center-level characteristics for all centers in each rating group. Daily prices are full-time preschool prices expressed in 2015 dollars. Panel B reports enrollment characteristics for STAR 2–4 centers with observed preschool enrollment information. "Share 90% of Max Capacity" is the share of centers with preschool enrollment above 90 percent of the capacity implied by the number of classrooms and group sizes for preschool-aged children.

3 Descriptive Evidence

This section documents four patterns that illustrate the role of providers' equilibrium responses in shaping preschool access, prices, and quality across neighborhoods. First, access to preschool, especially to highly rated centers, is unevenly distributed across space and is greater in more educated neighborhoods. Second, providers' dynamic decisions reinforce these disparities: low-rated centers exit more frequently and upgrade less frequently in lower-income neighborhoods. Third, price and enrollment responses to rating upgrades indicate that parents value Pennsylvania's STAR ratings, with stronger price responses in more advantaged neighborhoods. Fourth, providers' rating choices respond to the financial incentives embedded in quality-tiered reimbursements.

Access to Preschool is Heterogeneous Across Space. Physical distance to centers plays a key role in shaping parents' choices of provider. Surveys indicate that 55 percent of early education arrangements for children ages 3 through 5 are located within 3 miles of the child's home (OPRE, 2016). Figure 2 illustrates spatial inequality in access in Pittsburgh in 2018. Panels (a) and (b) focus on the city's East End. College-educated women are concentrated around Downtown and the Oakland university corridor,[12] where high-quality centers are also more abundant. In contrast, the southern and eastern peripheries have

[12] Oakland is home to the University of Pittsburgh and borders Carnegie Mellon University.

Figure 2. Heterogeneous Access to Early Education

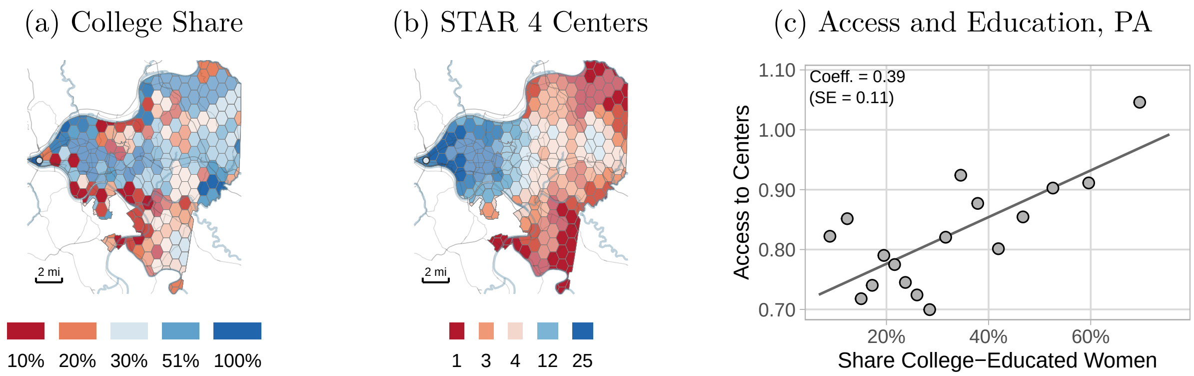

Notes: Panels (a) and (b) map the Pittsburgh East market in 2018. Panel (a) reports the share of college-educated women. Panel (b) reports the number of STAR 4 centers within a 10-minute drive of each grid-cell centroid. Panel (c) is a statewide binscatter of preschool access against the share of college-educated women at the census-tract level in 2018. Access is defined using the Enhanced Two-Step Floating Catchment Area method. Bins are selected following Cattaneo et al. (2024).

few high-quality options nearby. Panel (c) measures preschool access relative to the local population of 3- and 4-year-olds (Davis et al., 2019).[13] Access systematically increases with the share of college-educated women, yet remains low overall: the majority of Pennsylvania tracts have fewer than one seat per preschool-aged child.[14] Similar patterns hold for access to high-quality preschools. Figure B4 shows that this heterogeneity in access is also reflected in children's attendance rates. While only around 50% of children attend preschool in census tracts where less than a quarter of women are college graduates, this proportion climbs to almost 80% in tracts where more than half of women are college graduates. The model presented in Section 4 aims to discern the extent to which these differences in access and enrollment can be attributed to parents' preferences or providers' costs.

Centers' exit and quality-upgrading decisions also differ systematically across neighborhoods. Figure 3 displays the 2018 status of centers that were rated STAR 1 in 2010. By 2018, almost 60% of STAR 1 centers in the poorest neighborhoods had exited, compared with less than 30% in the richest tracts. Among surviving centers, fewer than 4% in low-income neighborhoods upgraded to STAR 4, compared with almost 15% in the richest tracts.

Parents Value Preschool Quality. To investigate whether parents value quality as measured by Pennsylvania's Keystone STARS ratings, I exploit within-center changes in ratings

[13] Appendix B provides details on the access measure construction, which consists of a demand-adjusted measure of supply expressed in seats per child.

[14] Appendix Figure B3 shows that overall access also increases with median household income after conditioning on tract population density.

Figure 3. Exit and Upgrade Patterns Differ across Neighborhoods

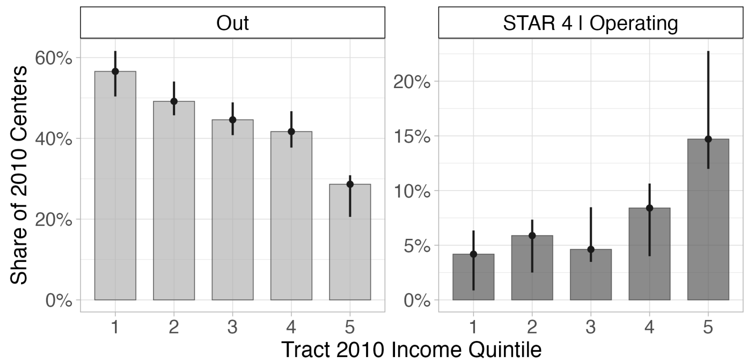

Notes: The sample consists of centers rated STAR 1 in 2010. In 2018, these centers can either be out of the market or in one of the four ratings. The figure shows, for each census tract income quintile, the proportion of these centers having exited in the left panel, and, conditional on still being in operation, the proportion rated STAR 4 in the right panel.

and examine their effects on prices and enrollment. I estimate:

$$p_{jt} = \sum_s \delta_s \mathbf{1}\{r_{jt} = s\} + \phi_j + \mu_t + \epsilon_{jt} \tag{1}$$

$$\log \frac{s_{jt}}{s_{0t}} = \alpha p_{jt} + \sum_s \beta_s \mathbf{1}\{r_{jt} = s\} + \psi_j + \nu_t + \epsilon_{jt} \tag{2}$$

where p_{jt} is the price of preschool j in year t, r_{jt} is the center's discrete STAR rating, and s_{jt} and s_{0t} are the market shares of preschool j and the outside option in year t. The two regressions include center and year fixed effects. In (2), I instrument for price following Berry (1994) to address endogeneity from time-varying unobserved quality, yielding a simple logit demand model without demographic preference heterogeneity.

Table 2 reports estimates across several specifications, all with center fixed effects controlling for time-invariant characteristics such as capacity and location. Columns (1)–(4) suggest that rating upgrades raise prices: moving from the lowest to the highest rating increases the daily price by \$1.6. Price increases are larger in more educated and higher-income neighborhoods, consistent with providers responding to greater willingness to pay for quality. Higher prices alone, however, are not sufficient to establish that parents value quality, since they may be accompanied by lower enrollment. Equation (2) addresses this concern by estimating the impact of quality upgrades on market shares, conditional on prices. Rating coefficients increase monotonically in columns (5) and (6). Once instrumented, price enters negatively, consistent with prices being positively correlated with unobserved center quality. Column (6) implies a willingness to pay of about $0.187/0.027 \approx \$6.9$ per day for an upgrade from STAR 1 to STAR 4.

Table 2. Price and Market Share After Quality Upgrade

Dependent variable	Daily Price				log(share/outside)	
	(1)	(2)	(3)	(4)	(5)	(6)
STAR 2	-0.1134				0.0294**	0.0263**
	(0.1020)				(0.0122)	(0.0124)
STAR 3	0.4875**				0.1219***	0.1352***
	(0.2126)				(0.0184)	(0.0203)
STAR 4	1.588***				0.1440***	0.1872***
	(0.2350)				(0.0205)	(0.0305)
Linear STAR rating		0.4436***	0.1919	0.1798*		
		(0.0742)	(0.1191)	(0.1014)		
Linear STAR × Above-median education			0.4107**			
			(0.1595)			
Linear STAR × Above-median income				0.4229***		
				(0.1526)		
Daily Price (2015 $)					0.0002	-0.0270**
					(0.0007)	(0.0128)
Observations	26,995	26,995	26,945	26,879	26,995	26,995
R^2	0.90015	0.89981	0.90003	0.89999	0.91450	0.90800
First-stage F-stat, Daily Price (2015 $)						177.89
Preschool fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports center-year regressions for 2010–2018. Columns 1–4 use daily price in 2015 dollars as the dependent variable. Columns 5–6 use $\log(s_j) - \log(s_0)$, where s_j is center j's market share and s_0 is the outside-option share. All specifications include preschool and year fixed effects. Columns 3 and 4 interact a linear STAR rating measure with indicators for whether the preschool is located in an above-median 2010 neighborhood, defined using the block-group share of college-educated women and median household income, respectively. Column 6 instruments price using the first principal component of the instruments whose construction is detailed in Appendix C. Standard errors are clustered at the preschool level. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Providers Respond to Financial Incentives to Operate at High Quality. As discussed in Section 2, Pennsylvania uses add-on reimbursement rates to encourage providers to serve disadvantaged children at higher quality by first reducing family top-ups, effectively making those centers more affordable, and, once those top-ups are eliminated, directly increasing provider receipts. Figure A1 shows that the schedule of add-on rates varied substantially from 2010 through 2018. Add-ons are zero at the lowest quality tier, while the rate for the highest tier rose to almost 25% of the average daily price by 2018. For example, a STAR 4 center received a daily add-on of $5 per subsidized child in 2015, rising to $7.5 in 2016 and $9.2 in 2017.

Changes in the generosity of add-on rates affect the share of revenue that providers receive from the state. In addition, centers serving more subsidized students benefit more from upgrading when add-on rates are generous. I use this variation to test whether centers respond to financial incentives by estimating a conditional logit model over the four quality ratings. Let x_{jrt} denote the expected daily add-on payment associated with choosing rating

Table 3. Upgrade Choice and Quality Incentives

Dependent variable	Rating Choice		
	(1)	(2)	(3)
A. Parameter Estimates			
Additional Subsidy Revenue (USD 100)	0.334	0.26	0.304
	(0.045)	(0.040)	(0.139)
STAR Rating fixed effects	Yes	Yes	Yes
Year fixed effects	No	Yes	Yes
Mean subsidized enrollment control	No	No	Yes
B. Absolute Marginal Effects			
Upgrade from 1 to 2	0.022	0.017	0.02
Upgrade from 2 to 3	0.035	0.027	0.032
Upgrade from 3 to 4	0.033	0.026	0.03
Observations	20,792	20,792	20,792

Notes: Estimates from a multinomial logit model of next-year STAR rating choice on expected add-on payment from subsidized preschool enrollment. All specifications use center-years from 2012 onward. Exits are omitted from the four-rating choice set. Expected add-on payments are the daily add-on rate for each possible next STAR rating multiplied by the center's average subsidized preschool enrollment in the panel, divided by 100. Column 3 controls for the center's average subsidized preschool enrollment.

r in year $t+1$, equal to the corresponding add-on rate multiplied by center j's average subsidized preschool enrollment, in hundreds of dollars. The probability that a center j rated s chooses rating r at $t+1$ is

$$\mathbb{P}(r|C_{jst}) = \frac{\exp\left(\beta x_{jrt} + \phi_{rt} + \mu_{rs}\right)}{\sum_{r'} \exp\left(\beta x_{jr't} + \phi_{r't} + \mu_{r's}\right)}. \tag{3}$$

Here, ϕ_{rt} are choice-specific year fixed effects and μ_{rs} are current-rating-by-choice fixed effects, which absorb differences in the desirability and cost of choosing each rating unrelated to the add-on schedule. The sample begins in 2012, when the schedule generates meaningful variation across ratings, and excludes center-years in which the provider exits. Identification comes from variation in expected add-on payment generated by changes in the schedule and differences in subsidized enrollment across centers.

Table 3 presents estimates and average marginal effects from three specifications. The resulting estimates show a consistent, positive effect of additional state revenue on quality-rating choices. The average marginal effect of an additional \$100 in daily add-on payments, which would, for instance, result from a \$5 add-on rate applied to a center serving 20 subsidized children, on the probability of upgrading from STAR 3 to STAR 4 ranges between 2.6 and 3.3 percentage points depending on the specification. The average upgrade rate of STAR

3 centers over the sample is 12.6%, which implies that the effects are substantial. Policies working through centers' revenues can therefore affect the structure of the early education market.

4 Model

Motivated by the evidence in Section 3, I develop a dynamic model in which entry, exit, and quality decisions evolve over time, while each period nests a stage game in which families choose preschools and incumbent centers set prices.[15] This structure is important for evaluating early-education policies whose effects unfold both within and across periods. For example, changes in parent subsidies affect the profitability of serving different neighborhoods in the short run and, through entry, exit, and quality investment, the spatial distribution of centers in the medium run. A dynamic model with reduced-form flow profits would not capture these channels. I first describe the stage game governing families' choices and centers' pricing decisions, and then turn to providers' dynamic decisions.

4.1 Stage Game: Demand and Supply of Preschool Education

Market environment. Following recent structural work on education markets, including Neilson (2013); Allende (2019), I model families' preschool choices using a spatially differentiated demand framework. I partition Pennsylvania into geographically distinct areas m, each containing a fixed set of residential nodes $\mathcal{L}_m$. A market is an area–school-year pair (m, t).[16] Families residing at node $\ell \in \mathcal{L}_m$ can choose any preschool in market m. Each family i has one preschool-aged child and two observed characteristics: maternal education $e_{it} \in \{H, L\}$, indicating whether the mother has a college degree, and income $y_{it} \sim F^y_{e\ell t}(\cdot)$. Thus, each node ℓ in year t is summarized by its location, its population $N_{\ell t}$ of 3- and 4-year-olds, the share with college-educated mothers, and the two conditional income distributions $(F^y_{L\ell t}, F^y_{H\ell t})$. This spatial aggregation captures the role of distance without requiring micro-level data on families' locations and preschool choices.

[15] Neilson (2013) and Allende (2019) model quality responses statically, while Singleton (2019) and Dinerstein et al. (2020) study school entry and exit in static games. The closest modeling approach is that of Sweeting (2013), who studies radio stations. Unlike preschools, radio stations do not compete on prices and are not spatially differentiated.

[16] To limit artificial boundaries in families' choice sets, the empirical application uses geographically large areas, constructed by aggregating school districts. In large urban areas, natural geographic features are used to define market boundaries, such as rivers in Pittsburgh; see Appendix C.

Product space and preferences. A set of preschools $\mathcal{J}_{mt}$ operates in each market mt. A preschool j at time t is characterized by a location, a price charged for full-day education services p_{jt}, and a STAR quality rating r_{jt}. In addition, preschool j in year t is characterized by a fixed unobserved quality component ξ_j, a time-specific demand shock ξ_t, and a transitory demand shock $\Delta\xi_{jt}$. These terms capture characteristics of a center that are not observed by the econometrician but are known by parents when choosing a preschool. The fixed component ξ_j in particular captures fixed features of the center, such as its location or capacity. Families have heterogeneous preferences for providers' characteristics.

Demand Subsidies. Section 2 outlined the set of incentive schemes implemented in Pennsylvania as part of the state's quality rating system. Families with income below an eligibility threshold benefit from subsidies to help pay for early education. These demand-side subsidies are governed by four policy variables depicted in Figure A2: an income threshold determining subsidy eligibility $\bar{y}_t$, a copayment schedule $copay_t(y)$, quality-tiered add-on rates $addon_t(r_{jt})$, and a local Maximum Child Care Allowance $MCCA_{mt}$. Copayment schedules are increasing with income and determine a baseline subsidized rate for a family with income y_i. These policies determine the price that parents with income y_{it} pay to send their child to preschool j, which has a market price of p_{jt} and a STAR rating r_{jt}:

$$\begin{aligned} p^c_{ijt}(p_{jt}, y_{it}) = &\underbrace{\mathbf{1}\{y_{it} \leq \bar{y}_t\}\Big(copay(y_{it}) + \max\{0, p_{jt} - MCCA_{m(j)t} - addon_t(r_{jt})\}\Big)}_{\text{subsidized}} \\ &+ \underbrace{(1 - \mathbf{1}\{y_{it} \leq \bar{y}_t\})p_{jt}}_{\text{unsubsidized}} \end{aligned} \tag{4}$$

The utility that household i, living in neighborhood ℓ, derives from sending their child to preschool j, located a drive time $d_{j\ell(i)}$ from ℓ, is given by

$$u_{ijt} = \beta^I_i + \beta^r_i r_{jt} + \alpha_i p^c_{ijt}(p_{jt}, y_{it}) + \lambda_i d_{\ell(i)j} + \xi_t + \xi_j + \Delta\xi_{jt} + \varepsilon_{ijt} = U_{ijt} + \varepsilon_{ijt}.$$

Preference shocks ε are assumed to be drawn from a Type I extreme-value distribution, which gives rise to market shares

$$s_{jt}(\boldsymbol{p}_t, \boldsymbol{r}_t, \boldsymbol{\xi}_t) = \sum_i w_i s_{ijt}(\boldsymbol{p}_t, \boldsymbol{r}_t, \boldsymbol{\xi}_t) = \sum_i w_i \frac{\exp U_{ijt}(\boldsymbol{p}_t, \boldsymbol{r}_t, \boldsymbol{\xi}_t)}{1 + \sum_k \exp U_{ikt}(\boldsymbol{p}_t, \boldsymbol{r}_t, \boldsymbol{\xi}_t)}, \tag{5}$$

where w_i are weights representing the proportion of parents with type i. More details are provided in Appendix C.

The model bundles alternatives to center-based preschool, including parental, relative, and family-based care, into the outside option. Preference heterogeneity allows the attractiveness of these alternatives to vary across socioeconomic groups. Centers compete through prices, and differ in their observed and unobserved quality, and in their location. Given the descriptive fact that capacity constraints rarely bind as discussed in Section 2, I abstract from explicit capacity constraints, with persistent differences in capacity absorbed by center fixed effects.[17] Finally, I take parents' labor-market decisions as exogenous. Changes in the opportunity cost of parental time could affect demand for center-based care—for example, through β_i^I—but I abstract from this margin to focus on competition in the early education market.

Marginal Costs of Early Childhood Education. I use the demand system to recover marginal costs of early education providers from their observed pricing behavior. In the stage game, quality ratings are taken as given, and centers compete on prices. I assume that preschools are profit-maximizing firms engaged in Bertrand competition and ignore the role of multi-establishment chains.[18] The incentives in place in Pennsylvania imply that provider j, with rating r_{jt}, serving child i with family income y_{it}, receives

$$p^p_{ijt}(p_{jt}, y_{it}) = \underbrace{\mathbf{1}\{y_{it} \leq \bar{y}_t\} \left(\max\{p_{jt}, \min\{p_{jt}, MCCA_j\} + addon_t(r_{jt})\}\right)}_{\text{subsidized}} + \underbrace{(1 - \mathbf{1}\{y_{it} \leq \bar{y}_t\})p_{jt}}_{\text{unsubsidized}}. \tag{6}$$

Given the market-level vector of prices $\boldsymbol{p}_t$, provider' profits are

$$\pi(\boldsymbol{p}_t) = N_{mt} \sum_i w_i s_{ijt}(\boldsymbol{p}_t)(p^p_{ijt}(p_{jt}, y_{it}) - mc_{jt}). \tag{7}$$

Under these assumptions, the optimal price for preschool j, p_{jt}, equals its marginal cost mc_{jt} plus a markup η_{jt} given by

$$\eta_{jt} = \left[-\sum_i w_i \frac{dp^c_{ijt}}{dp_{jt}} \frac{\partial s_{ijt}}{\partial p^c_{ijt}}\right]^{-1} \left(\sum_i w_i \frac{dp^c_{ijt}}{dp_{jt}} \frac{\partial s_{ijt}}{\partial p^c_{ijt}} (p^p_{ijt}(p_{jt}, y_{it}) - p_{jt}) + \sum_i w_i s_{ijt}(\boldsymbol{p}_t) \frac{dp^p_{ijt}}{dp_{jt}}\right). \tag{8}$$

[17] Figure B2 shows that capacity constraints are unlikely to bind under several conservative measures of effective capacity. If constraints do bind, treating constrained centers as available may understate latent demand and therefore subsequent entry or capacity investment. The model may consequently understate the extent to which the benefits and cost-effectiveness of the counterfactual policies grow over time.

[18] Although some centers are nonprofits, the majority of preschools in the data are for-profit. In addition, without further assumptions, I cannot distinguish preschools with low marginal costs from preschools with an altruistic objective that places weight on enrollment. I therefore treat all centers as pure profit maximizers.

The expression above highlights the impact of the early education subsidies on market power. Following the institutional setting in Pennsylvania, providers are allowed to charge top-ups that guarantee a weakly positive $p^p_{ijt}(p_{jt}, y_{it}) - p_{jt}$. If quality-tiered reimbursements $addon_t(r_{jt})$ were set to very generous levels, this term could be pushed to strictly positive values. The price sensitivity term $\frac{\partial s_{ijt}}{\partial p^c_{ijt}}$ is always negative. Therefore, abstracting from demand responses, greater financial support from the state leads high-quality centers serving more subsidized students to lower their markups, as public funds substitute for higher prices. However, note that $addon_t(r_{jt})$ can also change the mix of marginal consumers and, by reducing top-ups, mitigate the connection between posted prices and consumer prices for subsidized families.

4.2 Entry, Exit, and Quality Upgrades

State space, actions, and timing. The stage game takes the set and quality of operating centers as given. I now endogenize these through providers' entry, exit, and quality decisions. Let $\mathcal{M}_{xjt}$ collect the public payoff-relevant state variables for provider j at time t, including family demographics, provider characteristics and demand and cost shocks, and the spatial distribution of drive times. Provider j's own state is $x_{jt} \in \{\text{Out}, \text{STAR } k^4_{k=1}\}$. Every period, an incumbent chooses an action from $\mathcal{A}(\mathcal{M}_{xjt}) = \{\text{Out}, \text{STAR } k^4_{k=1}\}$. Potential entrants are tied to a specific location and choose between remaining out and entering at STAR 1.[19] Figure D6 reports the frequency of each action by state in the data. I assume that information, decisions, and payoffs follow the timing below:

1. $\mathcal{M}_{xjt}$ is realized.
2. Fixed costs of operating in the current state x, $C(x_{jt}; \boldsymbol{\psi}^C)$, are paid.
3. Private action-specific payoff shocks $\epsilon(a)$ are realized. The state for the next period is chosen, potentially resulting in a scrap value or a transition or entry cost $W(a_{jt}, \mathcal{M}_{xjt}; \boldsymbol{\psi}^W)$.
4. Variable profits $\pi(\mathcal{M}_{xjt})$ from stage-game price competition are realized.
5. The next-period state is determined by j's chosen action a_{jt}, j's competitors' actions, and aggregate state transitions.

Denoting providers' discount factor as β, the flow payoff from step 3 of period t through step

[19] This differs from spatial competition models such as Seim (2006) and Caoui et al. (2026). For child care centers, suggestive evidence supports the assumption that providers are more tied to a specific neighborhood than firms in the video-rental and supermarket industries studied in those papers. Appendix D.3 details how these location-specific potential entrants are constructed.

2 of period $t+1$, $\Pi(a_{jt}, \mathcal{M}_{xjt}) + \psi^{\epsilon}\epsilon(a_{jt})$, is defined as

$$\Pi(a_{jt}, \mathcal{M}_{xjt}) + \psi^{\epsilon}\epsilon(a_{jt}) = \pi(\mathcal{M}_{xjt}) - W(a_{jt}, \mathcal{M}_{xjt}; \boldsymbol{\psi}^{W}) - \beta C(a_{jt}; \boldsymbol{\psi}^{C}) + \psi^{\epsilon}\epsilon(a_{jt}). \tag{9}$$

Equilibrium Concept. I focus on stationary Markov-Perfect Bayesian Nash Equilibria (MPE). I denote the strategy of a firm in state $\mathcal{M}_{xjt}$ with payoff shocks ϵ_{jt} as $\mathcal{P}_j(\mathcal{M}_{xjt}, \epsilon_{jt})$, and the corresponding strategy profile as $\mathcal{P} = \{\mathcal{P}_j(\mathcal{M}_{xjt}, \epsilon_{jt})\}_{j\in\mathcal{J}}$.[20] The ex ante value function of preschool j in state $\mathcal{M}_{xjt}$ given a strategy profile $\mathcal{P}$ can be written as:

$$V^{\mathcal{P}}(\mathcal{M}_{xjt}; \boldsymbol{\psi}) = \mathbb{E}_{\epsilon_t} \max_{a_{jt}} \{v^{\mathcal{P}}(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) + \psi^{\epsilon}\epsilon(a_{jt})\}, \tag{10}$$

where $v^{\mathcal{P}}(a, \mathcal{M}_{xjt}; \boldsymbol{\psi})$ is the choice-specific value function corresponding to action a in state $\mathcal{M}_{xjt}$, defined as

$$v^{\mathcal{P}}(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) = \Pi(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) + \beta \int V^{\mathcal{P}}(\boldsymbol{\mathcal{M}}_{\tilde{\boldsymbol{x}}\boldsymbol{jt+1}}; \boldsymbol{\psi}) g(\boldsymbol{\mathcal{M}}_{\tilde{\boldsymbol{x}}\boldsymbol{jt+1}} | a, \mathcal{P}_{-j}, \mathcal{M}_{xjt}) d\boldsymbol{\mathcal{M}}_{\tilde{\boldsymbol{x}}\boldsymbol{jt+1}}. \tag{11}$$

The second term corresponds to the future value of taking action a in state $\mathcal{M}_{xjt}$ given a strategy profile $\mathcal{P}$ and is denoted by $FV(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}, \mathcal{P})$. Making the dependence on dynamic parameters explicit, the choice-specific value function can be written as

$$v^{\mathcal{P}}(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) = \Pi(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) + FV(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}, \mathcal{P}).$$

The transition kernel $g(\boldsymbol{\mathcal{M}}_{\tilde{\boldsymbol{x}}\boldsymbol{jt+1}} | a_{jt}, \mathcal{P}_{-j}, \mathcal{M}_{xjt})$ incorporates transitions of exogenous variables relevant for profits, such as demographics and i.i.d. transitory demand and cost shocks, as well as expectations over other players' actions. An MPE is a strategy profile $\mathcal{P}^*$ satisfying

$$\mathcal{P}^*(\mathcal{M}_{xjt}, \epsilon_{jt}) = \arg\max_{a_{jt}} \{v^{\mathcal{P}^*}(a, \mathcal{M}_{xjt}; \boldsymbol{\psi}) + \psi^{\epsilon}\epsilon(a_{jt})\}.$$

Players' equilibrium conditional choice probabilities (CCPs) $\mathbb{P}$ are obtained by integrating $\mathcal{P}^*$ over private payoff shocks. Assuming a Type I extreme-value distribution for the action-specific payoff shocks ϵ, the CCP corresponding to provider j's best response to the strategies

[20] The focus on stationary equilibria requires that players' strategies are not indexed by t. While the environment in which centers compete evolves over time, the relevant information on time-varying factors is already captured in $\mathcal{M}_{xjt}$ and players' expectations. The exceptions are the changes in the policy environment described in Appendix A, for which expectations are not modeled. Instead, I assume that centers believe that the next period's policy environment will continue forever.

$\mathcal{P}$ can be expressed as

$$\mathbb{P}^{\mathcal{P}}(a_{jt}|\mathcal{M}_{xjt};\boldsymbol{\psi}) = \frac{\exp\left(\frac{v^{\mathcal{P}}(a,\mathcal{M}_{xjt};\boldsymbol{\psi})}{\psi^{\epsilon}}\right)}{\sum_{a'}\exp\left(\frac{v^{\mathcal{P}}(a',\mathcal{M}_{xjt};\boldsymbol{\psi})}{\psi^{\epsilon}}\right)} = \frac{\exp\left(\frac{\Pi(a,\mathcal{M}_{xjt};\boldsymbol{\psi})+FV(a,\mathcal{M}_{xjt};\boldsymbol{\psi},\mathcal{P})}{\psi^{\epsilon}}\right)}{\sum_{a'}\exp\left(\frac{\Pi(a',\mathcal{M}_{xjt};\boldsymbol{\psi})+FV(a',\mathcal{M}_{xjt};\boldsymbol{\psi},\mathcal{P})}{\psi^{\epsilon}}\right)}. \tag{12}$$

Curse of dimensionality. The model outlined above is difficult to solve in practice due to the dimensionality of the state space. Information available to providers includes multiple continuous variables, such as demographics and providers' unobservable characteristics, making the state space high-dimensional. In addition, several urban markets feature numerous players, further complicating the game. This complexity makes the value function $V^{\mathcal{P}}$ and transition kernel g intractable. I use a recursive procedure that iterates on centers' strategies, discretizes the state space, and approximates the value function to make the computation tractable. Section 5.2 details the solution method and estimation approach.

5 Estimation

Estimation proceeds in two steps. First, I estimate families' preschool preferences and providers' marginal costs from the within-year equilibrium in enrollment and prices, following the differentiated-products demand approach outlined in Conlon and Gortmaker (2020), but incorporating space as part of product characteristics. These estimates map each local market state and policy environment into equilibrium prices, enrollment, and provider flow profits. Second, I combine these flow profits with observed provider transitions to estimate the fixed operating, entry, and quality-investment costs governing providers' dynamic decisions. The central computational challenge is therefore to preserve the rich household heterogeneity and spatial competition of the stage game while estimating the dynamic model. This section describes the estimation and computational approach used to address that challenge. See Appendices C and D for details.

5.1 Estimation and Identification of the Stage Game

Preference and Marginal-Cost Parameters. The stage game parameters $\boldsymbol{\theta}$ consist of heterogeneous preferences for preschool in general, β_i^I, and for highly rated preschools, β_i^r, as well as heterogeneous price and distance sensitivities, α_i and λ_i. I parameterize the intercept as a function of income to capture heterogeneous needs for formal early education and variation in outside options across demographic groups: $\beta_i^I = \beta_0^I + \beta_y^I \log(y_{it})$. Motivated by Miravete et al. (2023), I define price sensitivities as $\alpha_i = -\exp\left(\alpha_0 + \alpha_y \log y_{it} + \alpha_e e_{it} + \sigma_p v_i^p\right)$.

I allow distance sensitivities to vary with income, $\lambda_i = -\exp(\lambda_0 + \lambda_y \log y_{it})$, and specify preferences for quality ratings as $\beta_i^r = \beta_0^r + \beta_y^r \log y_{it}$.[21] Unobserved preferences v^p are drawn from a standard normal distribution. Lastly, the demand parameters include fixed unobserved quality $\boldsymbol{\xi}$ and transitory demand shocks $\boldsymbol{\Delta\xi}$. Marginal cost is assumed to depend linearly on observed preschool characteristics $\boldsymbol{X}_{jt}$ and unobserved residual cost components. The observed preschool characteristics include the center's STAR rating at time t and neighborhood characteristics such as demographics and density. As with the residuals in the common component of utility, I decompose unobserved cost shocks into a year fixed effect ω_t, a preschool fixed effect ω_j, and a transitory component $\Delta\omega_{jt}$:

$$mc_{jt} = \omega_t + \omega_j + X_{jt}'\boldsymbol{\gamma} + \Delta\omega_{jt}. \tag{13}$$

Estimator and Identification. I estimate the stage-game parameters by GMM, following Berry et al. (1995) and Conlon and Gortmaker (2020). With center fixed effects $(\boldsymbol{\xi}, \boldsymbol{\omega})$ absorbing permanent demand and cost heterogeneity, identification requires instruments for endogenous prices and market shares that are orthogonal to the transitory shocks $(\boldsymbol{\Delta\xi}, \boldsymbol{\Delta\omega})$. Following Berry et al. (1995), I construct instruments from competitors' characteristics and from cost shifters such as tract-level wage data for teachers and administrators in child-service occupations. Although market structure is endogenous in this model, the timing assumption makes period-t market structure predetermined with respect to the transitory demand and cost shocks. Competitors' characteristics, excluding prices, can therefore serve as instruments in the stage-game estimation; see Appendix C.3 for details. Having constructed these excluded instruments Z, and denoting exogenous covariates included in demand by X_D and supply by X_S, and the resulting instrument matrices by $Z_D = [X_D, Z]$ and $Z_S = [X_S, Z]$, the GMM moments are:

$$g^{IV}(\boldsymbol{\theta}) = \frac{1}{N} \begin{bmatrix} Z_D' & 0 \\ 0 & Z_S' \end{bmatrix} \begin{bmatrix} \Delta\xi(\boldsymbol{\theta}) \\ \Delta\omega(\boldsymbol{\theta}) \end{bmatrix}.$$

Additional Moments. Distinctive features of the data provide additional moments that can be used in estimation. First, centers report not only their aggregate market shares, but also how their enrollment is distributed between subsidized and unsubsidized children. Because subsidy status is based primarily on income, and income distributions differ by geography and educational attainment, these data provide useful variation for estimating

[21] I have experimented with specifications including a random coefficient on rating v^r and found qualitatively very similar results.

the demographic-specific preference parameters such as $\beta_y^I, \alpha_y, \alpha_e$. I also use microdata on families' choices from the NSECE to form additional moments linking demographics to the characteristics of early education arrangements to inform the estimation of λ_y. Appendix C.5 lists the micromoments, and their implied values in the data and at the estimated parameters.

Computational Approach. Large urban markets, sometimes with more than one hundred providers, make demand estimation computationally intensive. To accurately integrate over household heterogeneity, I draw 10,000 Halton household types per market. Each type is defined by location, income, education, and random coefficients $(\ell(i), y_i, e_i, \boldsymbol{v}_i)$. To make this computation feasible, following Lewis et al. (2021), I stack markets into Torch tensors and exploit fast array operations and GPU acceleration.[22]

5.2 Estimation of the Dynamic Parameters

Solution Approach. Estimating the dynamic parameters requires addressing the high dimensionality of the state space generated by spatial demand and imperfect competition. Following Sweeting (2013); Arcidiacono et al. (2013), I approximate the value function governing entry, exit, and quality decisions as a linear combination of basis functions that summarize each center's competitive environment. Intuitively, these basis functions must adequately summarize the local competitive environment each center faces. I begin with more than 900 candidate basis functions, including estimated demand and cost primitives, own characteristics, local demographics, and nearby competition aggregated over distance bands. Because flow profits arise from a rich Nash–Bertrand pricing game, these handcrafted bases approximate stage-game profits relatively poorly (out-of-sample $R^2 \leq 0.70$). I cover in more detail the specific architecture and training of the neural network in Appendix D.1.2. This motivates a second step in which I use a neural network with an architecture informed by Nash–Bertrand competition to approximate variable profits over observed and simulated market states and transform the initial variables into a smaller set of nonlinear basis functions. The resulting basis functions achieve substantially better predictive performance for variable profits (out-of-sample $R^2 \geq 0.96$). As stage-game profits are the key input into providers' flow payoffs in Equation 9 and eventually in their value function, this motivates the use of these bases to approximate the value function.

[22] I use Float64 precision to meet the contraction-mapping tolerance recommended by Conlon and Gortmaker (2020). Although Torch supports automatic differentiation, I compute the GMM gradient analytically to reduce memory requirements.

Estimation. The dynamic parameters $\boldsymbol{\psi}$ govern the transition-cost function W, the fixed operating-cost function C, and the scale ψ^{ϵ} of private action-specific payoff shocks. I parameterize the transition-cost matrix with a common downgrade cost W^{Dwn}, capturing the low propensity of providers to move down the quality ladder. Centers face no transition cost of staying at the same quality rating. Upgrade costs are allowed to be origin-destination specific. I estimate a constant operating fixed cost common across providers, whose payoff contribution is normalized to be $-\exp(\psi_0^C)$, and a common payoff-shock scale parameter ψ_0^{ϵ}.

Firms are forward-looking over their own entry, exit, and quality choices. They are also strategic: their beliefs about the future profitability of an action depend on competitors' anticipated choices. To keep these expectations tractable, I impose an approximation akin to spatial obliviousness. When computing expectations over next-period payoff states, $\Phi(\mathcal{M}_{t+1})$, I use simulated future states in which only a restricted set of nearby competitors are allowed to change their operating states, while more distant centers are held fixed. This approximation focuses the simulation on the strategic responses most relevant for a provider's future profitability and is natural in early education because household demand is spatial and competition is essentially local. Demographic state variables evolve according to estimated AR(1) processes, so firms forecast future local demand conditions from current demographics. Policy beliefs are fixed: firms take the current subsidy rule as the policy environment governing future payoffs and do not anticipate later policy reforms. Thus, firms are forward-looking over local market evolution but ignore future policy changes.

I estimate the parameters by maximum likelihood, with each likelihood evaluation nesting an approximate value-function iteration, following Arcidiacono et al. (2013). Intuitively, for a candidate parameter vector, the estimator first solves for the coefficients of the approximate value function so that centers' projected Bellman equations are satisfied. It then computes the likelihood of observed entry, exit, and quality choices using the continuation values implied by this approximation. The dynamic parameters are updated, which in turn implies a new Bellman equation, and the procedure is iterated until convergence. As in the demand estimation, I stack market-year-state observations into batched Torch tensor objects, which allow the nested likelihood evaluation and value-function and choice-probability calculations to run on a GPU with automatic differentiation for the dynamic optimization. Appendix D.1.1 provides additional details on the solution algorithm.

Dynamic games can have multiple equilibria, which can make iterative estimators that update players' strategies unstable or prevent convergence (Pesendorfer and Schmidt-Dengler, 2010). To avoid such instability, I follow Aguirregabiria and Mira (2010); Sweeting (2013) and hold beliefs about competitors' actions fixed at their first-stage estimates, $\hat{\mathbb{P}}$. This fixed-belief approach avoids repeatedly updating rivals' strategies during estimation, making likelihood evaluation stable while still imposing optimality of each provider's own choices given those beliefs. In counterfactuals, however, I solve for a full equilibrium, allowing all providers' strategies to adjust.

6 Estimation Results

This section reports the model estimates. I first discuss parental preferences and provider marginal costs, then use them to characterize centers' stage-game profits and how those profits vary with nearby demand and competition, and with providers' observed and unobserved quality. I then report estimates of entry costs, fixed operating costs, and the investment costs of moving across quality ratings.

6.1 Demand Estimation

Heterogeneity in Demand for Preschools. Panel A of Table 4 reports the stage-game estimates. With center and year fixed effects, the linear demand specification contains only a limited set of time-varying characteristics. The baseline taste for a one-level increase in STAR rating, β_0^r, is small and negative. This does not imply that higher-rated centers face lower overall demand: Appendix Figure C6 shows that centers observed at higher ratings also have larger estimated unobserved demand components, ξ_j, consistent with positive selection on persistent center attributes. Moreover, preference heterogeneity is substantial. First, higher-income families value quality more: β_y^r is positive and significant, offsetting β_0^r at roughly half a standard deviation above mean income. These estimates clarify the descriptive finding in Section 3 that demand rises following a rating upgrade. The model decomposes this response into distinct channels: for subsidized families, the higher add-on associated with an upgrade lowers effective prices, while higher-income families respond more strongly to the rating itself. Second, richer and college-educated families are less price sensitive, as captured by α_y and α_e, while $\sigma_p = 0.447$ indicates substantial unobserved heterogeneity in price sensitivity. Third, families dislike travel time, with higher-income families substantially less sensitive to distance, consistent with survey results.

Turning to costs, higher quality raises marginal costs even with center and year fixed effects: upgrading by one rating raises daily costs by \$0.57 and median marginal cost increases from \$28.4 for STAR 1 centers to \$31.8 for STAR 4 centers. Costs are also higher for centers serving a larger share of younger children, which may reflect cross-subsidization across age groups that I abstract from in the supply model. The supply-side statistics in Panel B are

Table 4. Stage-Game Parameter Estimates

	Parameter	Estimate	Std. Err.
A. Estimates			
Linear			
STAR rating	β_0^r	-0.15	(0.012)
Nonlinear			
Price	α_0	-1.118	(0.066)
Distance	λ_0	-0.69	(0.019)
Intercept (Income)	β_y^I	2.052	(0.151)
STAR rating (Income)	β_y^r	0.306	(0.011)
Price (Income)	α_y	-0.064	(0.016)
Price (Educ)	α_e	-0.347	(0.030)
Distance (Income)	λ_y	-0.527	(0.016)
Random Coefficients			
Price	σ_p	0.447	(0.048)
Supply			
STAR rating	γ^r	0.57	(0.049)
log Center Age	γ^{age}	-1.983	(0.100)
Share Infant/Toddler	$\gamma^{infant/toddler}$	0.731	(0.343)
Tract log Teacher Wage	γ^{wage}	0.064	(0.054)
Tract log Rent	γ^{rent}	0.241	(0.186)
Tract log Density	$\gamma^{density}$	0.098	(0.322)

B. Selected Model Implications				
Supply Statistics				
Median Own Price Elasticity	-6.069			
Median Lerner Index	0.154			
Median Marginal Cost ($)	Quality Rating			
	STAR 1	STAR 2	STAR 3	STAR 4
	28.429	28.785	30.545	31.793

Notes: Estimation uses data from all geographic markets from 2010 to 2018, totaling 26,995 center–year observations. The specification includes center and year fixed effects. Log Center Age is measured as $\log(1+\min\{5, \text{years in operation since 2010}\})$ to capture potentially higher costs for newer centers. Standard errors are reported in parentheses for parameter estimates in Panel A.

also reasonable: the median own-price elasticity is -6.1, and almost the entire distribution lies below -1 (Figure C5), consistent with profit-maximizing pricing in a market with price-sensitive demand and thin margins.

The estimated travel-time preferences imply that both demand and competition are highly local. From the perspective of families, the value of an additional center declines rapidly with distance and is negligible beyond 30 minutes of travel time (Appendix Figure C7). Similarly, diversion ratios between centers more than 5 km apart are less than one quarter of those between centers within 500 m. Appendix C.6 provides further details.

Table 5. Stage-Game Profit Drivers

Variable	Estimate	Std. Err.
Competitor count (within 5,000 m)	-0.118	(0.007)
Competitor count (beyond 5,000 m)	0.044	(0.010)
Competitor sum exp(ξ) (within 5,000 m)	-0.024	(0.002)
STAR 4	0.171	(0.003)
STAR 3	0.018	(0.003)
STAR 2	-0.124	(0.003)
Unobserved quality (ξ)	1.506	(0.004)
Unobserved cost (ω)	-0.619	(0.005)
Children aged 0–5, college-educated household, income \$100k+	0.092	(0.004)
Observations	361,047	

Notes: The sample includes many simulated markets where providers' operating status and quality are randomly perturbed and demographics are randomly changed, hence the large number of preschool-year observations compared to the estimation sample. Each cell reports the coefficient from the regression $\pi_{jt} = \beta x_{jt} + \mu_j + \phi_t + \epsilon_{jt}$, where π_{jt} denotes profits from the stage-game estimates. Continuous regressors and profits are standardized, and the regression includes year and preschool fixed effects.

The model fits the demographic micro-moments closely (Table C2); in particular, it predicts preschool-level subsidized enrollment shares with an average relative error of only 7%.

The preference and cost estimates imply stage-game profits consistent with spatial competition. To make these relationships transparent, I construct perturbed copies of observed market states by randomly altering providers' ratings and operating status and the demographic composition of nearby neighborhoods, and then re-solve the Nash–Bertrand pricing game in each configuration. Table 5 reports regressions of the resulting profits on a selected set of variables. A one s.d. increase in the number of competitors within five kilometers lowers profits by 0.12 s.d., while greater unobserved quality among nearby competitors also lowers profits. Conversely, a one s.d. increase in the number of young children from college-educated households earning above \$100,000 nearby raises profits by 0.09 s.d.

The table also clarifies providers' dynamic investment incentives. Relative to STAR 1, operating at STAR 4 increases profits by 0.17 s.d., while gains at lower ratings are smaller and are negative at STAR 2. Returns to quality investment therefore materialize primarily near the top of the rating ladder. Reaching these higher tiers requires upfront investment, and providers must evaluate current costs against higher variable profits over subsequent years. These trade-offs form the basis of the dynamic incentives in Section 4.

Table 6. Dynamic Parameter Estimates

Name	Parameter	Estimate	Std. Err.
A. Fixed-Cost Parameters			
FC Intercept	ψ_0^C	-0.518	(0.058)
B. Transition Parameters			
Entry $0 \to 1$ Intercept	W_0^{01}	2.594	(0.248)
Upgrade $1 \to 2$ Intercept	W_0^{12}	1.975	(0.165)
Upgrade $1 \to 3$ Intercept	W_0^{13}	4.260	(0.459)
Upgrade $1 \to 4$ Intercept	W_0^{14}	3.828	(0.369)
Upgrade $2 \to 3$ Intercept	W_0^{23}	2.026	(0.226)
Upgrade $2 \to 4$ Intercept	W_0^{24}	2.455	(0.317)
Upgrade $3 \to 4$ Intercept	W_0^{34}	1.288	(0.179)
Downgrade Intercept	W^{Dwn}	2.698	(0.304)
C. Shock Parameters			
Std. Dev. Intercept	ψ_0^ϵ	0.794	(0.070)

Notes: Estimation uses data from a subset of geographic markets over the years 2010–2017. Bootstrap standard errors are reported in parentheses. These standard errors reflect only uncertainty from the dynamic estimation and do not account for uncertainty in the demand estimates or the variable-profit approximation. In estimation, the discount factor is set to $\beta = 0.9$, and scrap values are assumed to be zero. Estimates are scaled by profits expressed in thousands of dollars per day; annual values can therefore be recovered by multiplying by 180,000, assuming 180 operating days per year.

6.2 Dynamic Estimation

Dynamic Parameters. Table 6 reports the dynamic estimates. First, quality ratings operate as a ladder: upgrades are costly, particularly when skipping levels, consistent with quality improvements taking several years to build. Second, the estimates imply annual fixed costs of slightly over $\exp(-0.518) \times 180 \times \$1000 \approx \$107{,}000$, consistent with the relatively large exit rate, particularly among STAR 1 centers, as reported in Appendix Figure D6. These fixed costs amount to roughly 30% of annual revenue for a center serving 40 preschoolers at a daily price of $50. Combining these fixed costs with the median marginal cost of a highly rated provider, about $30 per child-day, implies annual costs of roughly $8,000 per child. This is close to industry estimates: $7,876 per child for Pennsylvania public programs operating STAR 3–4 centers and $9,543 per preschooler in Head Start (Friedman-Krauss et al., 2018). The model fits observed transitions well (Figure D6), reproducing the main patterns: exit declines with quality, most providers remain at their current rating, and upgrades typically occur one level at a time. Estimated entry costs are around $2.594 \times 180 \times$

$1000 $\approx$ $467,000, comparable to state grants supporting center openings.[23]

7 Counterfactuals

This section shows how providers' short-run price responses and medium-run adjustments in market structure shape the incidence and cost-effectiveness of early education policies. First, comparing quality-neutral and quality-targeted subsidies with a no-subsidy scenario, I show how these policies reshape the distribution of provider market power across neighborhoods and quality tiers. Over the medium run, quality-targeted subsidies have spatially differentiated effects on market structure, increasing access to high-quality preschool particularly in low-income neighborhoods. These short-run and dynamic responses interact: incumbent markups fall most where nearby entry expands most strongly. Next, I simulate plausible preschool expansion scenarios. I compare policies that consist of increasing eligibility for demand subsidies with (quality-targeted) and without (quality-neutral) quality-tiered reimbursements, and policies that fix eligibility but vary quality-tiered reimbursements. I find that these policies expand preschool supply, increasing cost-effectiveness in consumer-surplus terms over time as parents' nearby options grow. The consumer-surplus measure favors quality-neutral demand expansions, while the educational-benefit measure favors quality-targeted expansions that match low-income children to high-quality providers.

7.1 Equilibrium Mechanisms

Subsidies influence how providers set prices across neighborhoods and quality tiers. To highlight the interaction between provider market power, early education subsidies, and changing market structure, I implement three counterfactual scenarios. In the first, no subsidies are available. In the second, I implement a quality-neutral subsidy: demand subsidies follow the same means-tested, income-based schedule as in Pennsylvania, but there are no quality-tiered reimbursements. In the third, I implement a quality-targeted subsidy: in addition to the same income-based subsidies, high-quality providers receive quality-tiered add-on rates for serving subsidized children.

First, the distribution of markups across neighborhoods and provider characteristics highlights how consumer preferences shape provider market power across space. Figure 4 shows that markups in 2010 increase with neighborhood income in all three policy scenarios. Absent quality-tiered add-ons (left and middle panels), they are also generally higher for higher-

[23] For comparison, following the American Rescue Plan, New York State awarded between $398,000 and half a million dollars to grantees opening a center (NY Division of Child Care Services, 2022).

Figure 4. Markups across neighborhoods, quality ratings, and subsidy regimes

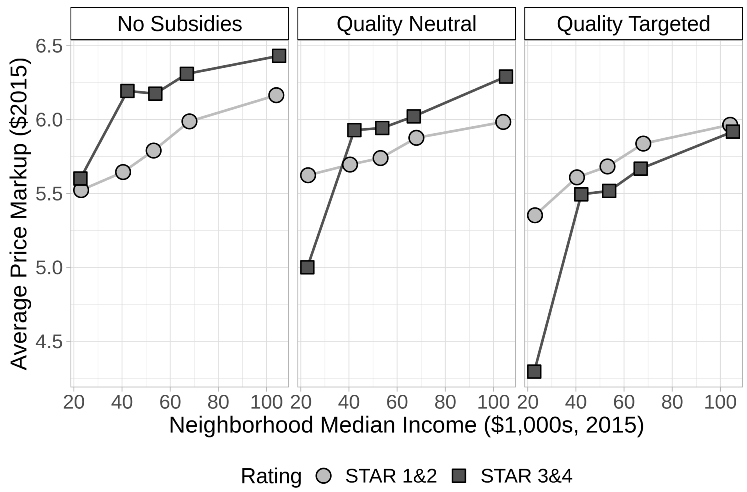

Notes: The figure plots average price markups by neighborhood median household income, quality rating, and subsidy regime. Neighborhoods are grouped into five income bins defined using the 2010 distribution of block-group median household income. Points are plotted at the average income of centers in each bin. Markups are measured in 2015 dollars and are averaged over ten simulation runs before aggregation.

quality providers. This pattern is consistent with the demand estimates, which show that richer households and households in which the mother is college-educated are less price sensitive.

Comparing the first two panels illustrates how demand subsidies affect market power. Relative to the no-subsidy benchmark (left), quality-neutral subsidies (middle) reduce markups for both low- and high-quality preschools. These average effects mask substantial heterogeneity in centers' responses, as highlighted in Appendix Figure E10. Indeed, subsidies generate two opposing forces. On the one hand, a positive demand shock from a reduction in consumers' price sensitivity may lead centers to increase markups. On the other hand, the entry from the outside option of highly price-sensitive consumers can make residual demand more elastic and induce some centers to lower prices. Both positive and negative markup responses are largest in low-income neighborhoods, where subsidies have the greatest impact on demand. Quality-tiered add-ons (right) compress markups further, particularly among high-quality providers serving low-income neighborhoods.

Quality-targeted subsidies improve access to high-quality providers over time. To illustrate how market-structure adjustment shapes the spatial distribution of policy benefits, I focus on one subdivision of the Pittsburgh metropolitan area. This market, one of three that together cover the metropolitan area, includes the part of Pittsburgh between the Allegheny and Monongahela rivers, as well as eastern and southern suburbs, and exhibits

Figure 5. Subsidies and High-Quality Access

(a) Quality-Neutral (b) Quality-Targeted (c) Correlation with Income

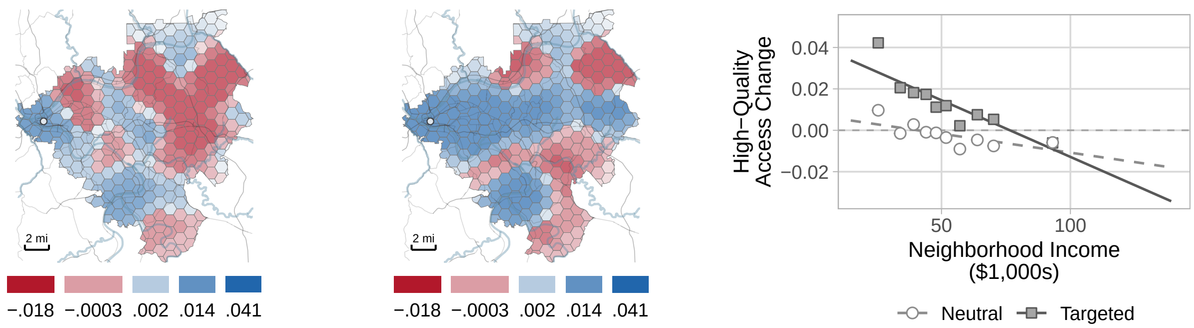

Notes: Panels A and B report differences in access to high-quality centers under quality-neutral subsidies (demand subsidies without an add-on) and quality-targeted subsidies (demand subsidies with an add-on), relative to a baseline scenario without subsidies, in 2017. Panel C reports the relationship between differences in access and neighborhood income. Access is computed using the two-step floating catchment area approach applied to the early education sector in Davis et al. (2019).

substantial variation in household income, as illustrated in Appendix Figure E11. I use the model to simulate how the quality-neutral and quality-targeted subsidy designs reshape access to high-quality preschool across neighborhoods within this market.

Figure 5 panel (a) shows that quality-neutral demand subsidies generate relatively modest gains in access to high-quality preschool in 2017. In contrast, panel (b) shows that quality-targeted subsidies produce larger and more spatially dispersed improvements, despite some localized declines. These gains extend into lower-income parts of the market, including areas south of Pittsburgh in the Monongahela Valley and neighborhoods toward the eastern city limits. By increasing reimbursements for high-quality preschools serving lower-income families, quality-tiered subsidies allow providers in these areas to invest in quality without passing the full cost of upgrading on to families, whose willingness to pay for higher quality may otherwise be insufficient to support such investments.

Panel (c) quantifies the relationship between neighborhood income and the access gains shown in the maps. Under quality-targeted subsidies, a $10,000 increase in neighborhood income is associated with a roughly 0.005 smaller gain in high-quality access, compared with about 0.002 under quality-neutral subsidies. At neighborhood incomes around $25,000, quality-targeted subsidies increase access by roughly 0.03, equivalent to 55% of mean access in Pittsburgh East under the no-subsidy counterfactual. The corresponding effect of quality-neutral subsidies is close to zero.

Together, these counterfactuals show that subsidy design shapes both immediate affordability and longer-run supply: quality-tiered add-ons direct entry and upgrading toward lower-

Figure 6. Changes in Markups and Competitor Entry

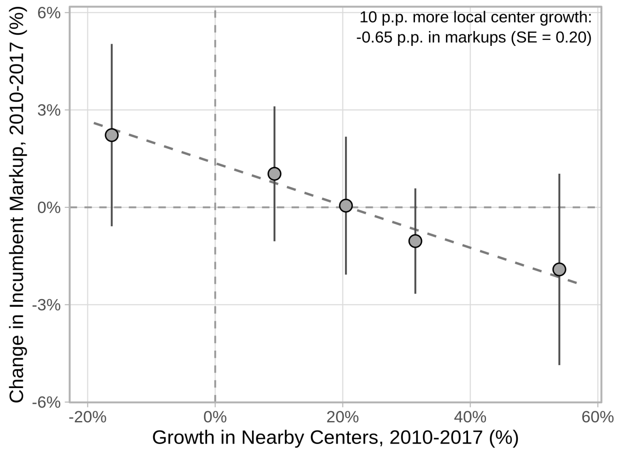

Notes: The sample includes centers operating in both 2010 and 2017 under both the quality-neutral and no-subsidy counterfactuals. Axes report $100 \times \Delta^{Sub}_{\text{No Sub}} \log \frac{X_{2017}}{X_{2010}}$, where X stands for markups on the y-axis and number of centers within 5 km on the x-axis. Outcomes are averaged across 10 simulation runs before aggregation.

income neighborhoods, materially reshaping the spatial distribution of high-quality access over seven years.

Policies that promote entry reduce provider market power. Beyond improving affordability directly, preschool subsidies can make formal early education more accessible over the medium run by inducing entry, particularly in areas where subsidized families reside. Figure 6 plots the difference in local center growth between the quality-neutral and no-subsidy counterfactuals against the corresponding difference in incumbent markup growth from 2010 to 2017. The sample is restricted to centers operating in both years under both counterfactuals. Most binscatter points lie to the right of zero, indicating that nearby center growth is greater under quality-neutral subsidies than in the absence of subsidies. More importantly, the downward-sloping relationship shows that places where subsidies generate larger increases in nearby competition also experience larger relative declines in incumbent markups. This competitive response constitutes an additional dynamic benefit of the policy: subsidy-induced expansion in local supply disciplines incumbent pricing. Because these lower markups apply to all consumers, the gains extend beyond subsidized families to unsubsidized families as well.

7.2 Policy Expansions and Cost-Effectiveness in the Short and Medium Run

I consider a range of policy experiments along two margins: the breadth of subsidy eligibility and the level of quality-tiered reimbursements. Along the eligibility margin, I compare expansions with and without quality-tiered reimbursements. I then hold eligibility fixed at Pennsylvania's benchmark threshold of approximately $48,000 and vary quality-tiered reimbursements. These counterfactuals separate the effects of broader subsidy coverage from stronger incentives for high-quality provision and trace how each policy's cost-effectiveness varies with generosity.

Varying the eligibility threshold for demand subsidies. I first consider expansions of subsidy eligibility to progressively higher-income households, either with or without quality-tiered reimbursements. While less encompassing than universal preschool, these policies represent plausible extensions of existing subsidy programs and preserve the central role of private providers as the main delivery channel for subsidized early education (NSECE Project Team, 2014). Figure E7 summarizes their design in the case with quality-tiered reimbursements. Eligibility expands through progressively higher income thresholds. Family copayments rise smoothly with income before plateauing at approximately $20 per day.

Table 7 reports equilibrium outcomes across subsidy eligibility thresholds relative to a policy with an eligibility threshold of 200% of the federal poverty line, or approximately $48,000, which prevailed in Pennsylvania throughout the sample period. Panel A reports short-run enrollment effects, overall and by STAR-rating group. Eliminating subsidies reduces overall enrollment by almost 37% (column 1). This decline reflects the increase in out-of-pocket prices faced by formerly subsidized families, who are also particularly price sensitive. Enrollment losses exceed 53% at STAR 3–4 centers, in part because of the elimination of the quality-tiered reimbursements that help low-income families afford these higher-cost, higher-priced providers.

Holding quality-tiered reimbursements fixed, lowering the threshold from $48,000 to $24,000 (column 4) reduces overall enrollment by 17.4%, showing that subsidies sustain enrollment among families with incomes between $24,000 and $48,000. At the same $24,000 threshold, quality-targeted subsidies (column 4) generate about 6% more enrollment overall and 41% more enrollment at STAR 3–4 centers than quality-neutral subsidies (column 2). Raising the eligibility threshold to $72,000 has markedly different effects depending on whether subsidies remain quality-targeted. Under quality-neutral subsidies (column 3), overall enrollment remains close to the Pennsylvania benchmark, but its composition shifts: STAR 1–2 en-

Table 7. Market Structure and Market Expansion Effects

	No Subsidy	Quality Neutral		Quality Targeted	
Eligibility Threshold ($1,000s)	0	24	72	24	72
	(1)	(2)	(3)	(4)	(5)
A. Enrollment share relative to $48k with 2018 add-on, 2010					
Overall	0.631	0.777	0.987	0.826	1.113
STAR 1–2	0.716	0.903	1.153	0.871	1.055
STAR 3–4	0.461	0.524	0.655	0.738	1.227
B. Enrollment share relative to $48k with 2018 add-on, 2017					
Overall	0.599 [0.582, 0.609]	0.743 [0.719, 0.761]	0.993 [0.974, 1.013]	0.790 [0.782, 0.804]	1.150 [1.130, 1.173]
STAR 1–2	0.696 [0.639, 0.764]	0.924 [0.820, 0.999]	1.296 [1.188, 1.387]	0.891 [0.830, 0.990]	1.026 [0.882, 1.159]
STAR 3–4	0.436 [0.366, 0.505]	0.433 [0.347, 0.547]	0.476 [0.405, 0.631]	0.621 [0.528, 0.702]	1.386 [1.193, 1.663]
C. Number of centers relative to $48k with 2018 add-on, 2017					
Overall	0.788 [0.752, 0.827]	0.777 [0.745, 0.809]	0.923 [0.878, 0.953]	0.867 [0.848, 0.894]	1.113 [1.031, 1.184]
STAR 1–2	0.841 [0.758, 0.929]	0.870 [0.801, 0.914]	1.051 [0.959, 1.116]	0.902 [0.859, 0.953]	1.060 [0.957, 1.159]
STAR 3–4	0.675 [0.598, 0.734]	0.574 [0.531, 0.632]	0.642 [0.564, 0.746]	0.793 [0.717, 0.864]	1.235 [1.150, 1.430]

Notes: Columns report subsidy eligibility thresholds in thousands of 2015 dollars, separately for quality-neutral and quality-targeted subsidies. All entries are normalized relative to the omitted Pennsylvania benchmark, which combines a $48k eligibility threshold with the 2018 quality-tiered add-on schedule; the corresponding benchmark value is 1.000. "Quality-neutral" columns (2) and (3) correspond to policy environments without quality-tiered reimbursements, while "quality-targeted" columns (4) and (5) have those reimbursements in place. Panel A reports enrollment shares in 2010 relative to this benchmark. Panel B reports relative enrollment shares in 2017, and Panel C reports the relative number of centers in 2017. Rows report effects overall and separately for STAR 1–2 and STAR 3–4 providers. For the 2017 panels, bracketed values report the minimum and maximum across ten simulation runs.

rollment increases by 15%, while STAR 3–4 enrollment falls by 35%. When quality-tiered reimbursements are retained (column 5), overall enrollment increases by 11%, with gains of 6% at STAR 1–2 centers and 23% at STAR 3–4 centers. At the same $72,000 threshold, quality targeting therefore nearly doubles high-quality enrollment relative to the quality-neutral policy. Together, these comparisons show that demand subsidies expand participation, while quality-tiered reimbursements direct enrollment toward high-quality providers.

Panel B shows how medium-run adjustments in entry, exit, and quality amplify these short-run effects. By 2017, enrollment without subsidies falls to only 60% of the Pennsylvania benchmark, compared with 63% on impact. Under the $72,000 expansions, the aggregate enrollment effects persist, but supply adjustment sharpens the divergence across quality tiers. With quality-neutral subsidies, STAR 1–2 enrollment rises to almost 30% above the benchmark, while STAR 3–4 enrollment falls by 53%. With quality-targeted subsidies, by contrast, overall enrollment rises by 15% and STAR 3–4 enrollment by 39%. Dynamic responses therefore magnify the extent to which quality-tiered reimbursements direct market expansion toward high-quality provision.

Panel C shows the market-structure adjustment underlying these medium-run enrollment patterns. The immediate enrollment losses from subsidy removal are compounded by center closures: by 2017, there are 21% fewer centers overall and 33% fewer high-quality centers. At the $72,000 threshold, quality-neutral subsidies leave the total number of centers 8% below the Pennsylvania benchmark and the number of STAR 3–4 centers 36% below it. With quality-tiered reimbursements, by contrast, the total number of centers rises by 11% and the number of STAR 3–4 centers by 24%. Thus, at the same eligibility threshold, quality targeting generates almost twice as many high-quality centers as the quality-neutral policy, while leaving the number of lower-quality centers nearly unchanged. These differences reflect providers' entry and exit decisions and incumbents' quality investments.

Increasing quality-tiered reimbursements generates spillovers to unsubsidized families. I next hold demand-subsidy eligibility fixed at the Pennsylvania benchmark of approximately $48,000 and vary the add-on rates paid to high-quality providers serving subsidized children. This exercise isolates the effects of strengthening quality targeting while holding the breadth of subsidy coverage fixed.

To measure the resulting benefits to families, I compute consumer surplus in counterfactual Ψ and year t as

$$CS_{gt}(\Psi) = \sum_m N_{mt} \sum_{i \in g} \frac{1}{|\alpha_i|} w_i \log \left(1 + \sum_{j \in \mathcal{J}_{mt}(\Psi)} \exp U_{ijt}(\Psi) \right), \tag{14}$$

where $\mathcal{J}_{mt}(\Psi)$ is the set of centers active under policy environment Ψ in year t and $U_{ijt}(\Psi)$ is the utility that parent i in demographic group g derives from choosing center j. This measure captures changes in prices, quality, and the set of available providers.

Figure 7 reports consumer-surplus gains from increasingly generous quality-tiered reimbursements separately for subsidized and unsubsidized families. For subsidized families, more generous add-ons raise provider reimbursements and reduce top-ups, generating direct gains. In the short run, these families therefore capture most of the policy benefits, while gains to unsubsidized families arise only through equilibrium price responses and remain modest. Over time, however, providers respond through entry and quality investment, expanding the availability of high-quality options. As a result, unsubsidized higher-income families also experience substantial consumer-surplus gains in the medium run.

Figure 7. Benefits of Quality-Tiered Reimbursements for Families

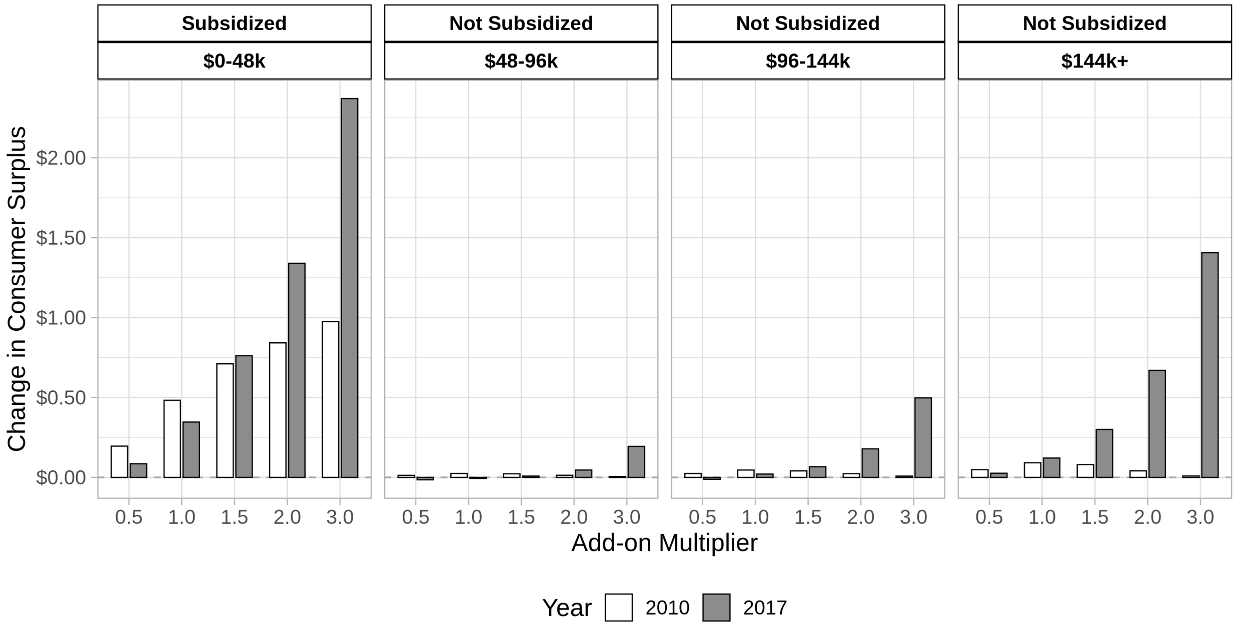

Notes: This figure shows the average consumer-surplus gains experienced by subsidized and unsubsidized families in 2010 and 2017 across levels of quality-tiered reimbursement generosity.

Cost-effectiveness of expansions in the short and medium run. Finally, I evaluate the cost-effectiveness of these policy designs at varying levels of generosity using two marginal value of public funds (MVPF) measures, which relate policy benefits to government expenditures.[24] The first uses the consumer-surplus measure defined above to capture families' valuation of changes in affordability, access, and quality. Following evaluations of early-education policies including Hendren and Sprung-Keyser (2020); Kline and Walters (2016); Cascio (2023), I construct a second MVPF based on children's educational benefits. I calibrate these benefits using estimates of the effect of Head Start on kindergarten test scores for low-income children from Kline and Walters (2016), assume that gains decline with family income and reach zero above $100,000, and assign STAR 1&2 preschools half the benefit of STAR 3&4 preschools. For this second measure, government expenditures are net of the additional tax revenue generated by children's higher future earnings. See Appendix E.2.3 for details.

Figure 8 reports consumer-surplus MVPFs across the three policy designs in the short and medium run. Panel (a) varies the eligibility threshold under quality-neutral subsidies, panel (b) varies eligibility under quality-targeted subsidies, and panel (c) varies the generosity of quality-tiered reimbursements while holding eligibility fixed.

[24] These calculations abstract from the distortionary costs of financing the policies through taxation.

Figure 8. Consumer Surplus MVPFs Across Policies and Over Time

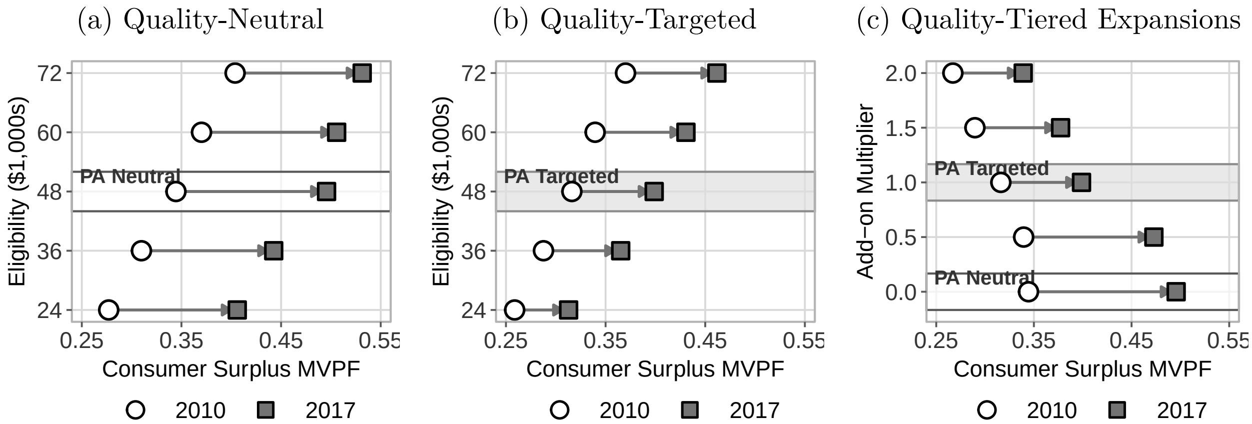

Notes: This figure reports the consumer surplus marginal value of public funds (MVPF) for different expansion policies relative to a no-subsidy scenario. Policy generosities are measured by eligibility thresholds (panels (a) and (b)) and by quality-tiered add-on multipliers (panel (c)), and the MVPF is shown on the x-axis. White outlined rows labeled "PA Neutral" correspond to a policy with Pennsylvania's eligibility threshold but without quality-tiered add-ons. Gray rows labeled "PA Targeted" combine Pennsylvania's eligibility threshold with its 2018 quality-tiered add-on levels.

All policies generate positive consumer-surplus gains relative to no subsidies, reflecting both immediate reductions in families' out-of-pocket costs and, over time, expansions in their choice sets. Panels (a) and (b) show that raising the eligibility threshold for quality-neutral and quality-targeted subsidies from $24,000 to $72,000 increases cost-effectiveness. Two forces contribute to this pattern. On the numerator, broader eligibility extends subsidies to families with stronger demand for preschool and lower price sensitivity, generating larger consumer-surplus gains. On the denominator, because family copayments rise with income, higher-income families require smaller government subsidies. Together, these forces increase consumer-surplus gains per dollar of public spending.

Comparing panels (a) and (b), adding quality-tiered reimbursements to demand subsidies reduces the consumer-surplus MVPF. Under quality-targeted policies, providers receive additional payments for serving subsidized children at the highest rating tiers, but the demand estimates imply that low-income families do not value these higher ratings more strongly, and some of the additional reimbursements accrue to providers as profits rather than to families as consumer surplus. Panel (c) reinforces this mechanism: as quality-tiered reimbursements become more generous, a larger share of the additional payments accrues to providers, reducing the consumer-surplus MVPF. Thus, from the perspective of families' revealed preferences, the most cost-effective policy expands demand subsidies without quality-tiered reimbursements.

Across all three panels, consumer-surplus MVPFs are substantially higher in 2017 than in

2010. Providers' dynamic responses, including entry and quality investment, expand families' choice sets and amplify policy benefits over time.

Consumer surplus reflects families' revealed preferences, but may not capture the educational benefits that motivate preschool policy. I therefore turn to the educational-benefit MVPF. Figure 9 reports this measure for the same three policy designs. All MVPFs exceed 1 under the chosen calibration, implying that educational gains exceed policy costs. The MVPF associated with Pennsylvania's policies is around 1.8. Cost-effectiveness is nevertheless non-monotonic in generosity across all three designs. In panels (a) and (b), broader eligibility initially reaches children with large preschool benefits but eventually extends subsidies to children with smaller gains. In panel (c), more generous quality-tiered reimbursements initially shift subsidized children toward high-quality providers, but at higher levels an increasing share of the additional payments accrues to providers with little further effect on families' choices, causing the MVPF to decline.

In contrast to the consumer-surplus results, quality-targeted policies in panels (b) and (c) generate higher educational-benefit MVPFs than quality-neutral policies in panel (a). By shifting low-income children toward high-quality providers, these policies generate larger educational gains under the chosen calibration. Dynamic supply responses, however, do not uniformly increase educational-benefit MVPFs: they amplify gains under generous quality-tiered reimbursements in panel (c) but can reduce them under more restrictive eligibility policies in panels (a) and (b). This divergence reflects the different welfare objects: an expansion in nearby options can raise consumer surplus while inducing families to choose closer, lower-quality providers, thereby reducing educational benefits.

Figure 9. Educational-Benefit MVPFs Across Policies and Over Time

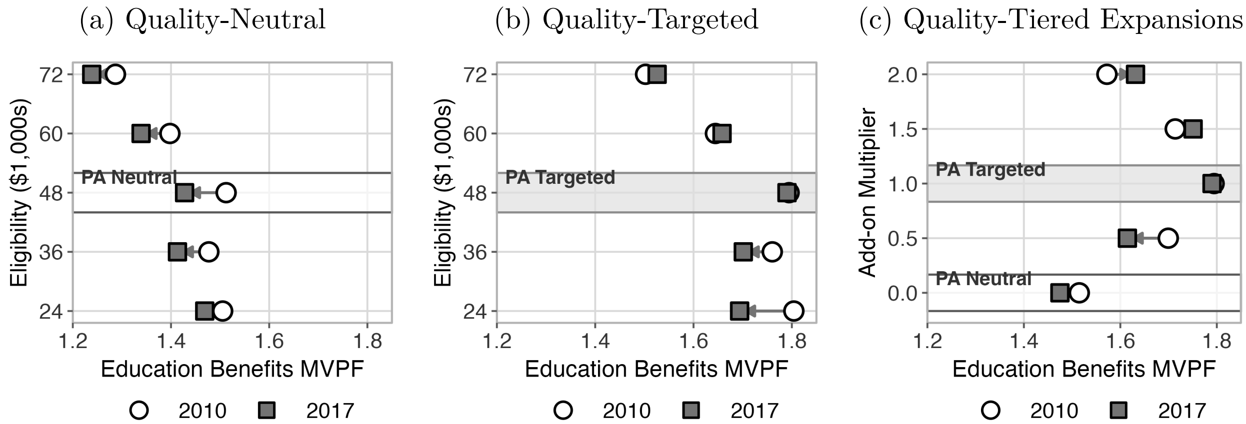

Notes: This figure reports the educational-benefit marginal value of public funds (MVPF) relative to a no-subsidy scenario. Appendix E.2.3 describes how these MVPFs are calculated. Policy generosities are measured by eligibility thresholds —panels (a) and (b)— and by quality-tiered add-on multipliers —panel (c). White outlined rows labeled "PA Neutral" correspond to Pennsylvania's eligibility threshold but without quality-tiered add-ons. Gray rows labeled "PA Targeted" combine Pennsylvania's eligibility threshold with its 2018 quality-tiered add-on levels.

8 Conclusion

High-quality preschool holds the promise of reducing educational inequality before kindergarten, generating sustained policy interest in its expansion. Yet limited empirical guidance exists on how to design cost-effective expansion policies when preschool is predominantly delivered through private early education markets. I develop a dynamic model of the preschool market capturing providers' key adjustment margins: entry, exit, quality investment, and pricing. Counterfactual simulations show that these adjustments materially shape policy incidence and cost-effectiveness: over time, entry and quality investment reshape geographic access to high-quality preschool, amplify enrollment gains from subsidy expansions, and generate benefits for subsidized and unsubsidized families. More broadly, the analysis complements cost-effectiveness calculations based on ex post program evaluations by accounting for policy-induced changes in the private delivery system and hence in take-up and program costs.

References

Aguirregabiria, V. and P. Mira (2010): "Stability of Equilibria and Properties of Sequential Estimators: Problems and Solutions," Unpublished manuscript.

Allende, C. (2019): "Competition Under Social Interactions and the Design of Education Policies," Unpublished manuscript.

Arcidiacono, P., P. Bayer, F. A. Bugni, and J. James (2013): "Approximating High-Dimensional Dynamic Models: Sieve Value Function Iteration," *Advances in Econometrics*, 31, 45–95.

Armona, L. and S. Cao (2022): "Redesigning federal student aid in higher education," Unpublished manuscript.

Baker, M., J. Gruber, and K. Milligan (2008): "Universal child care, maternal labor supply, and family well-being," *Journal of Political Economy*, 116, 709–745.

——— (2019): "The long-run impacts of a universal child care program," *American Economic Journal: Economic Policy*, 11, 1–26.

Bassok, D., M. Fitzpatrick, and S. Loeb (2014): "Does state preschool crowd-out private provision? The impact of universal preschool on the childcare sector in Oklahoma and Georgia," *Journal of Urban Economics*, 83, 18–33.

Berlinski, S., M. M. Ferreyra, L. Flabbi, and J. D. Martin (2020): "Child care markets, Parental labor supply, and child development," *Parental Labor Supply, and Child Development.*

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890.

Berry, S. T. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, 242–262.

Borowsky, J., J. H. Brown, E. E. Davis, C. Gibbs, C. M. Herbst, A. Sojourner, E. Tekin, and M. J. Wiswall (2022): "An equilibrium model of the impact of increased public investment in early childhood education," Unpublished manuscript.

Brown, J. H. (2018): "Does public pre-k have unintended consequences on the child care market for infants and toddlers?" Unpublished manuscript.

Cannon, J. S., G. L. Zellman, L. A. Karoly, and H. L. Schwartz (2017): *Quality rating and improvement systems for early care and education programs: Making the second generation better*, RAND.

Caoui, E. H., B. Hollenbeck, M. Osborne, and E. Page (2026): "The impact of dollar store expansion on local market structure and food access," *Quantitative Marketing and Economics*, 24, 7.

Cascio, E. U. (2015): "The promises and pitfalls of universal early education," *IZA World of Labor.*

——— (2023): "Does universal preschool hit the target?: Program access and preschool impacts," *Journal of Human Resources*, 58, 1–42.

Cascio, E. U. and D. W. Schanzenbach (2013): "The impacts of expanding access to high-quality preschool education," *Brookings Papers on Economic Activity*, 2013, 127–192.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024): "On binscatter," *American Economic Review*, 114, 1488–1514.

Collard-Wexler, A. (2013): "Demand fluctuations in the ready-mix concrete industry," *Econometrica*, 81, 1003–1037.

Conlon, C. and J. Gortmaker (2020): "Best practices for differentiated products demand estimation with pyblp," *The RAND Journal of Economics*, 51, 1108–1161.

Davis, E. E., W. F. Lee, and A. Sojourner (2019): "Family-centered measures of access to early care and education," *Early Childhood Research Quarterly*, 47, 472–486.

Davis, E. E. and A. Sojourner (2021): "Increasing Federal Investment in Children's Early Care and Education to Raise Quality, Access, and Affordability," *Washington. DC: Hamilton Project.*

Department of the Treasury (2021): *The Economics of Child Care Supply in the United States*, United States Department of the Treasury.

Dinerstein, M., C. Neilson, and S. Otero (2020): *The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy*, Princeton University, Industrial Relations Section.

Duncan, G. J. and A. J. Sojourner (2013): "Can intensive early childhood intervention programs eliminate income-based cognitive and achievement gaps?" *Journal of human resources*, 48, 945–968.

Durkin, K., M. W. Lipsey, D. C. Farran, and S. E. Wiesen (2022): "Effects of a statewide pre-kindergarten program on children's achievement and behavior through sixth grade." *Developmental Psychology.*

Elango, S., J. L. García, J. J. Heckman, and A. Hojman (2015): "Early childhood education," in *Economics of Means-Tested Transfer Programs in the United States, Volume 2*, University of Chicago Press, 235–297.

Friedman-Krauss, A. H., K. S. Barnett, G. Weisenfeld, and N. DiCrecchio (2018): "The State of Preschool 2017: State Preschool Yearbook." *National Institute for Early Education Research.*

Gray-Lobe, G., P. A. Pathak, and C. R. Walters (2023): "The long-term effects of universal preschool in Boston," *The Quarterly Journal of Economics*, 138, 363–411.

Grunewald, R. and P. Davis (2011): "Hardly child's play," *Federal Reserve Bank of Minneapolis.*

Heckman, J., R. Pinto, and P. Savelyev (2013): "Understanding the mechanisms through which an influential early childhood program boosted adult outcomes," *American Economic Review*, 103, 2052–86.

Hendren, N. and B. Sprung-Keyser (2020): "A unified welfare analysis of government policies," *The Quarterly Journal of Economics*, 135, 1209–1318.

Hotz, V. J. and M. Wiswall (2019): "Child care and child care policy: existing policies, their effects, and reforms," *The ANNALS of the American Academy of Political and Social Science*, 686, 310–338.

Hotz, V. J. and M. Xiao (2011): "The impact of regulations on the supply and quality of care in child care markets," *American Economic Review*, 101, 1775–1805.

Keys, T. D., G. Farkas, M. R. Burchinal, G. J. Duncan, D. L. Vandell, W. Li, E. A. Ruzek, and C. Howes (2013): "Preschool center quality and school readiness: Quality effects and variation by demographic and child characteristics," *Child development*, 84, 1171–1190.

Kline, P. and C. R. Walters (2016): "Evaluating public programs with close substitutes: The case of Head Start," *The Quarterly Journal of Economics*, 131, 1795–1848.

Lee, W. F., A. Sojourner, E. E. Davis, and J. Borowsky (2025): "Effects of Child-Care Vouchers on Price, Quantity, and Provider Turnover in Private-Care Markets," *Journal of Human Capital*, 19, 1–30.

Lewis, G., B. Ozaltun, and G. Zervas (2021): "Maximum Likelihood Estimation of Differentiated Products Demand Systems," Unpublished manuscript.

Ludwig, J. and D. L. Miller (2007): "Does Head Start improve children's life chances? Evidence from a regression discontinuity design," *The Quarterly journal of economics*, 122, 159–208.

Miravete, E. J., K. Seim, and J. Thurk (2023): *Elasticity and curvature of discrete choice demand models*, Centre for Economic Policy Research London, UK.

Neilson, C. (2013): "Targeted vouchers, competition among schools, and the academic achievement of poor students," Unpublished manuscript.

NSECE Project Team (2014): "Characteristics of Center-based Early Care and Education Programs: Initial Findings from the National Survey of Early Care and Education," OPRE Report No. 2014-73a, Office of Planning, Research and Evaluation, Administration for Children and Families, U.S. Department of Health and Human Services, available at: https://www.acf.hhs.gov/opre/report/characteristics-center-based-early-care-and-education-programs-initial-findings.

NY Division of Child Care Services (2022): "Invest in New York - Child Care Deserts Grant for New Providers," https://ocfs.ny.gov/programs/childcare/deserts/#awardees, accessed: 2022-11-01.

OPRE (2016): "Fact Sheet: How Far Are Early Care and Education Arrangements from Children's Homes?" *OPRE Report*, No. 2016-10.

Pesendorfer, M. and P. Schmidt-Dengler (2010): "Sequential estimation of dynamic discrete games: A comment," *Econometrica*, 78, 833–842.

Polyakova, M. and S. P. Ryan (2023): "Market Power and Redistribution: Evidence from the Affordable Care Act," Unpublished manuscript.

Ryan, S. P. (2012): "The costs of environmental regulation in a concentrated industry," *Econometrica*, 80, 1019–1061.

Seim, K. (2006): "An empirical model of firm entry with endogenous product-type choices," *The RAND Journal of Economics*, 37, 619–640.

Singleton, J. D. (2019): "Incentives and the supply of effective charter schools," *American Economic Review*, 109, 2568–2612.

Sweeting, A. (2013): "Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry," *Econometrica*, 81, 1763–1803.

Tebaldi, P. (2025): "Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca," *Review of Economic Studies*, 92, 586–620.

U.S. Department of Education (2014): "U.S. Department of Education National Center for Education Statistics, Digest of Education Statistics," https://nces.ed.gov/programs/digest/d19/tables/dt19_202.60.asp, accessed: 2022-12-30.

Waldfogel, J. (2003): "Preference externalities: An empirical study of who benefits whom in differentiated-product markets," *The RAND Journal of Economics*, 34, 557–568.

Appendix

A Institutional Details

A.1 Definitions for Pennsylvania Early Education Subsidies

Private rate p. Daily amount charged by the provider to private-pay families. Sometimes referred to as Converted Payment Rate (CPR) (Pennsylvania Office of Child Development and Early Learning, 2024a).

Maximum child care allowance $MCCA$. Highest amount the Department will pay for child care services provided to families eligible for subsidized child care (55 Pa. Code § 3042.3). From 55 Pa. Code § 3042.14(d): "The eligibility agency may not pay child care costs that exceed the maximum child care allowance minus the family copayment for the type of care the child received from the provider, except when the Department provides tiered-reimbursement to providers that are eligible based on their participation in the Department's Quality Rating and Improvement System."

Copayment $copay$. Weekly amount a family pays for subsidized child care (55 Pa. Code § 3042.3). Copayments depend on family size and income, $copay(y_i)$.

Subsidy sub. Amount paid by the eligibility agency to the provider. Under the provider agreements, the base subsidy is the lesser of the provider's private rate and the MCCA, minus the family copayment, plus any quality-tiered reimbursement for which the provider is eligible.

Top-up $topup$. Difference between a provider's private-pay rate and the subsidized child care payment rate. Providers are allowed to charge this difference to subsidized parents under 55 Pa. Code § 3042.14(d).

Quality-tiered reimbursement, or add-on rate, $addon$. Amount added to a provider's payment rate for meeting additional quality standards (55 Pa. Code § 3042.3). Add-on rates increase with Keystone STAR ratings and reduce any remaining gap between the subsidy payment and the provider's private-pay rate. They may raise subsidy payments above the MCCA maximum (55 Pa. Code § 3042.14(b)) and, for high-quality providers, above private-pay rates (Pennsylvania Office of Child Development and Early Learning, 2024b, Section 4.3.3).

A.2 Implied Subsidy Formulas

These definitions imply:

$$
\begin{aligned}
sub_{ij} &= \min\{p_j, MCCA_j\} - copay(y_i) + addon(r_j) \\
topup_j &= \max\{0, p_j - (copay(y_i) + sub_{ij})\} \\
&= \max\{0, \max\{0, p_j - MCCA_j\} - addon_j\} \\
&= \max\{0, p_j - MCCA_j - addon_j\} \\
p_{ij}^{cons} &= copay(y_i) + topup_j \text{ if } y_i \leq \bar{y}, \text{otherwise } p_j \\
&= copay(y_i) + \max\{0, p_j - MCCA_j - addon_j\} \text{ if } y_i \leq \bar{y}, \text{otherwise } p_j \\
p_{ij}^{prov} &= sub_{ij} + copay(y_i) + topup_j \text{ if } y_i \leq \bar{y}, \text{otherwise } p_j \\
&= \max\{p_j, \min\{p_j, MCCA_j\} + addon_j\} \text{ if } y_i \leq \bar{y}, \text{otherwise } p_j \\
p_{ij}^{prov} - p_j &= \max\{0, addon_j - \max\{0, p_j - MCCA_j\}\} \text{ if } y_i \leq \bar{y}, \text{otherwise } 0
\end{aligned}
$$

For differentiability, the max and min operators are replaced in the code with differentiable "smoothmax" and "smoothmin" operators.

A.3 Descriptive Statistics on Policy Variables

Figure A1. Supply Side: Add-On Rates by Quality Rating

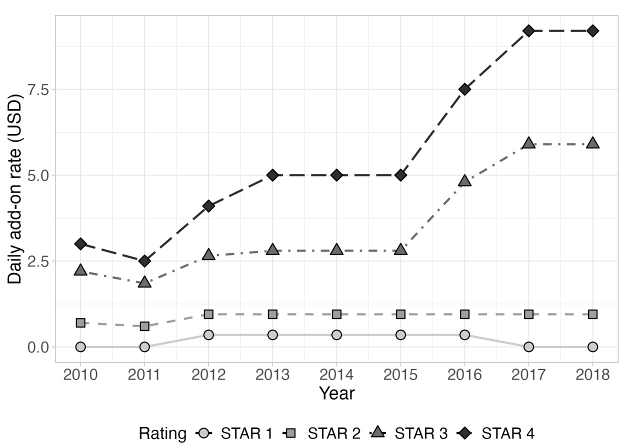

Notes: This figure reports the quality-tiered reimbursement schedule for Keystone STARS providers. Providers at higher ratings receive a daily add-on rate for each subsidized child served. Rates are reported for 2010–2018 and are measured in nominal dollars. The data come from FOIA requests to the Pennsylvania Department of Human Services.

Figure A1 shows the evolution of add-on rates over the period of the sample. The policy gradually becomes more generous and more targeted to the two highest quality tiers. Fig-

Figure A2. Components of the Child Care Works Demand-Side Subsidy

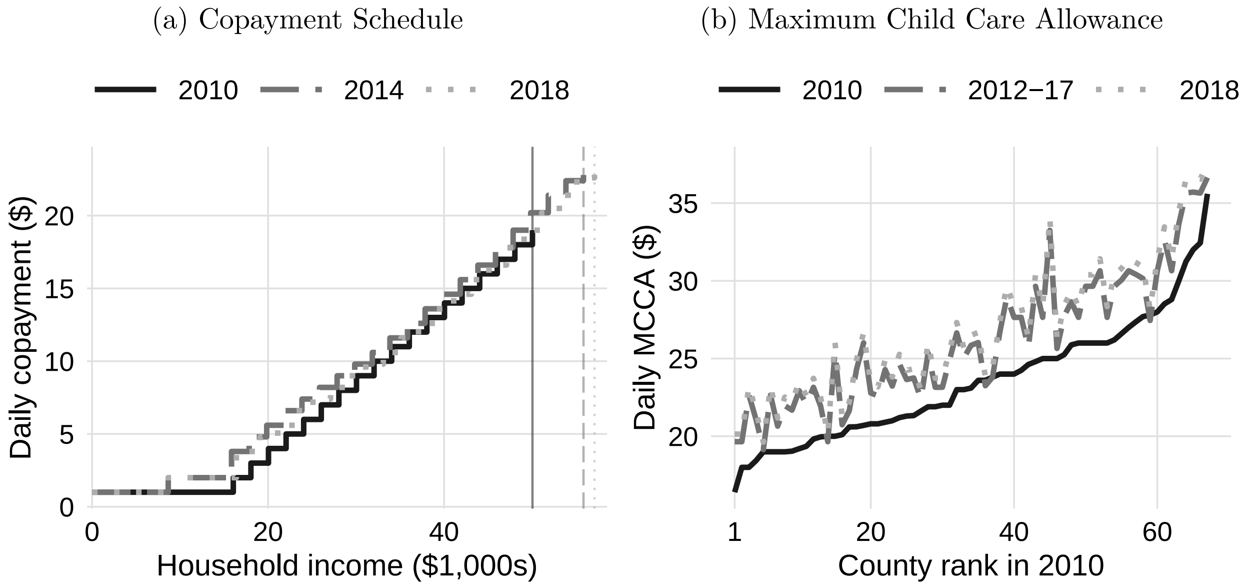

Notes: The figure describes the two main policy schedules entering the Child Care Works demand-side subsidy. The left panel plots the daily family copayment schedule for a household size of four in 2010, 2014, and 2018; vertical lines mark the corresponding income eligibility cutoffs. Copayments are a function of household size and income. The right panel plots county-level daily MCCA values in 2010, 2012–2017, and 2018, ranking counties by their 2010 MCCA. MCCA values vary by county and were updated discretely over the sample period. Copayment schedules come from Pennsylvania Bulletins recovered for 2010–2018, and MCCA data come from FOIA requests to the Pennsylvania Department of Human Services.

ure A2 shows the remaining components of Pennsylvania early education policies. The left panel displays copay schedules across 3 selected years, 2010, 2014 and 2018. The variation across years in copay schedules and eligibility thresholds implies subsidy-driven price variation across years for low-income households. The right panel shows maximum care allowance, which determines the local ceiling in the generosity of subsidized rates. Again, the figure emphasizes variations across time and across counties that the model will leverage in estimation.

B Additional Descriptive Facts

B.1 Centers Descriptive Facts

Geographic access measure definition. To account for differences in the local population of preschool-aged children, I follow Davis et al. (2019) and construct a demand-adjusted measure of supply, which I refer to as access. For each preschool j with capacity C_j, and for each spatial unit ℓ with a population N_ℓ of preschool-aged children, I define slots per capita at the school (S_j) and spatial-unit (A_ℓ) levels as follows:

$$S_j = \frac{C_j}{\sum_{\ell,\, t_{\ell j}<20 \text{ min}} w(t_{\ell j}) N_\ell} \quad \text{and} \quad A_\ell = \sum_{j,\, t_{\ell j}<20 \text{ min}} w(t_{\ell j}) S_j,$$

where $t_{\ell j}$ is the driving time between spatial unit ℓ and preschool j, and $w(t_{\ell j}) = \exp\left(-t_{\ell j}^4/1000\right)$ is a Gaussian decay weight. S_j constitutes an upper bound on the true number of seats per child because licensed capacity C_j reflects the size of the building rather than the actual number of available spots.

Figure B1 shows the distribution of the fraction of preschoolers among enrollees in a center. Consistent with the modeling assumption that centers make their decisions based on this age-group, for the vast majority of centers, preschoolers are most of their enrollment. Moreover, prior research suggests that preschool-aged is the profit-making age group that cross-subsidizes other younger children, further lending support to our conduct assumption (Brown, 2018). Figure B2 shows the distribution of enrollment relative to classroom capacity in the centers, lending credibility to the modeling assumption that capacity constraints are not binding: the vast majority of center-years observations operate far below 90% of their classroom capacity. This finding is consistent with prior studies of the early education sector in the U.S. (Blau, 2001; Herbst, 2023).

Figure B1. Share of Preschoolers in Centers

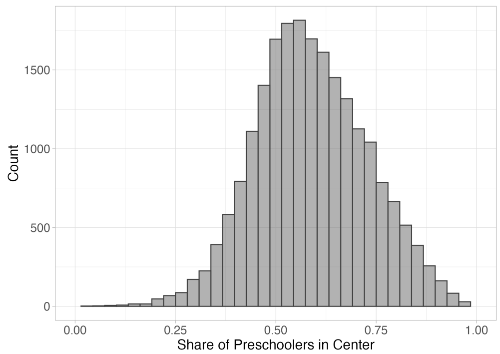

Notes: The figure displays the distribution of the fraction of enrolled children who belong to the preschool age-group category, defined as 3- and 4-year-olds. The remaining children are mostly infants and toddlers, aged one to two. For most centers, preschoolers constitute the majority of their enrollment, consistent with a model where centers' decisions are determined by the market for this age-group.

Figure B2. Centers and Classroom Capacity Constraints

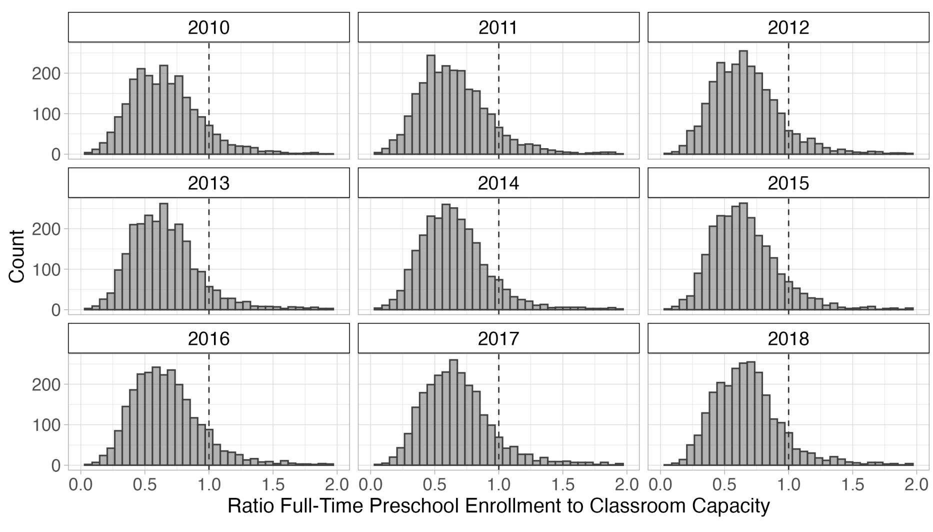

Notes: The figure displays the distribution of the ratio of preschool enrollment to a capacity constraint based on the number of classrooms in the center. The vertical line marks the statutory benchmark implied by Pennsylvania classroom-size requirements. The distribution suggests that in any given year, only a limited share of centers operates close to this classroom-capacity benchmark.

B.2 Stylized Facts

Figure B3 illustrates that, after controlling for the fact that higher-income neighborhoods tend to be less dense, access is positively correlated with both the share of college-educated mothers and median income in the census tract.

Figure B3. Access to Preschool and Tract Demographics

(a) College Educated Women (b) Household Income

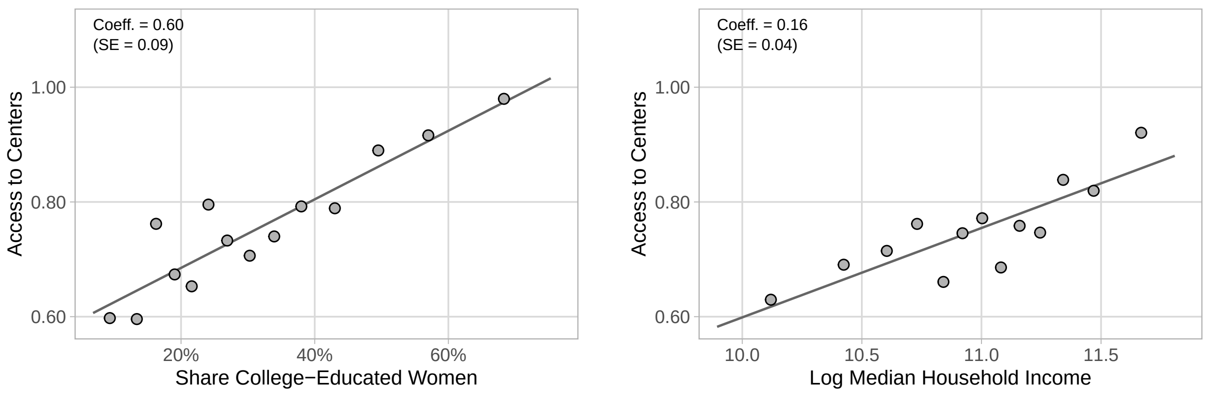

Notes: The figure relates access to preschool to neighborhood demographics. Both panels control for log density and weight census tracts by the number of preschool-aged children.

Figure B4 illustrates that beyond access, enrollment also increases with the share of college educated women in a neighborhood. A similar pattern holds for household median income, as illustrated on Figure B4b.

Figure B4. Formal Preschool Enrollment and Tract Demographics

(a) College Ratio (b) Household Income

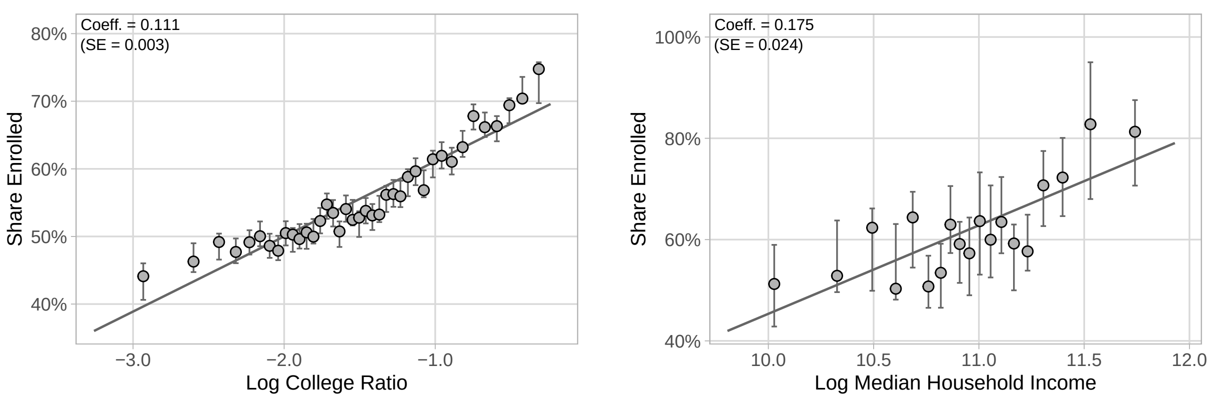

Notes: The figure relates formal preschool enrollment to neighborhood demographics. The left panel plots the share of 3- and 4-year-olds enrolled in school against the log ratio of college-educated women in the census tract. The right panel plots the same enrollment measure against log median household income in 2018. Dots report binned averages, vertical bars report confidence intervals, and the fitted line summarizes the relationship across bins.

Prices, Enrollment, and Upgrading. Table 1 shows that highly rated centers enroll more students and charge higher prices than lower-quality preschools. To connect these cross-sectional differences to centers' quality investments, I use the panel to compare the evolution of prices and enrollment for centers that start at STAR 1 or STAR 2 in 2014 and reach different ratings by 2018. I estimate regressions of the form

$$\Delta_{2014}^{2018} Y_{jc} = \beta_3 \mathbb{1}\{r_{j,2018} = \text{STAR } 3\} + \beta_4 \mathbb{1}\{r_{j,2018} = \text{STAR } 4\} + \phi_c + \epsilon_j, \tag{15}$$

where $\Delta_{2014}^{2018} Y_{jc}$ is the change in either daily price, full-time-equivalent enrollment, or private-pay enrollment of preschool j operating in county c between 2014 and 2018. The omitted group is centers that remain at STAR 1 or STAR 2 in 2018. All specifications include county fixed effects ϕ_c.

The results are presented in Table B1. Centers that reach STAR 4 increase daily prices by \$3.16 more than centers that remain at STAR 1 or STAR 2. This difference corresponds to roughly 9 percent of the average 2018 daily price of STAR 1 centers. Price increases are accompanied by enrollment growth: centers that reach STAR 4 add 5.84 more full-time-equivalent preschoolers than centers that remain at the lower ratings. The point estimate for private-pay enrollment is also positive, at 1.13 more children, although it is less precise once county fixed effects are included. This pattern is directionally consistent with enrollment gains not being driven only by subsidized children who are less exposed to posted prices. The STAR 3 estimates are also positive across all three outcomes. These longer-run panel patterns complement the Berry-style demand evidence in the main text: centers that upgrade subsequently raise prices and expand enrollment, which is consistent with parents valuing higher observed quality over multi-year horizons.

Table B1. Price and Enrollment Changes of STAR 1 and STAR 2 Rated Preschools

Model:	Δ Daily Price (1)	Δ F.T.E Enrollment (2)	Δ F.T.E Private Pay Enrollment (3)
A. Estimates			
STAR 3 in 2018	2.748**	2.512***	1.102
	(1.059)	(0.8318)	(0.7926)
STAR 4 in 2018	3.157***	5.840***	1.125
	(0.6226)	(1.912)	(0.7944)
B. Fixed Effects			
County fixed effects	Yes	Yes	Yes
C. Fit Statistics			
Observations	2,142	2,142	2,142

Notes: The sample includes centers rated STAR 1 or STAR 2 in 2014 and observed in 2018. The omitted category is centers that remain at STAR 1 or STAR 2 in 2018. Outcomes are changes between 2014 and 2018. All specifications include county fixed effects. County-clustered standard errors are in parentheses. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

C Stage-game Model Supplements

C.1 Markets

For computational reasons, I partition Pennsylvania into multiple geographically distinct markets for demand estimation. Market boundaries are constructed by taking unions of school districts. The markets are intentionally large, so that each market plausibly contains the relevant choice set for residents in the corresponding urban area and minimizes arbitrary boundaries. For the two largest cities in Pennsylvania, Philadelphia and Pittsburgh, I further divide the urban area into submarkets, using prominent geographic features such as rivers and major highways whenever possible to define boundaries. The resulting market map is shown in Figure C1. Table C1 below provides descriptive statistics of the markets used in the estimation of the stage-game model.

Figure C1. Geographic Preschool Markets

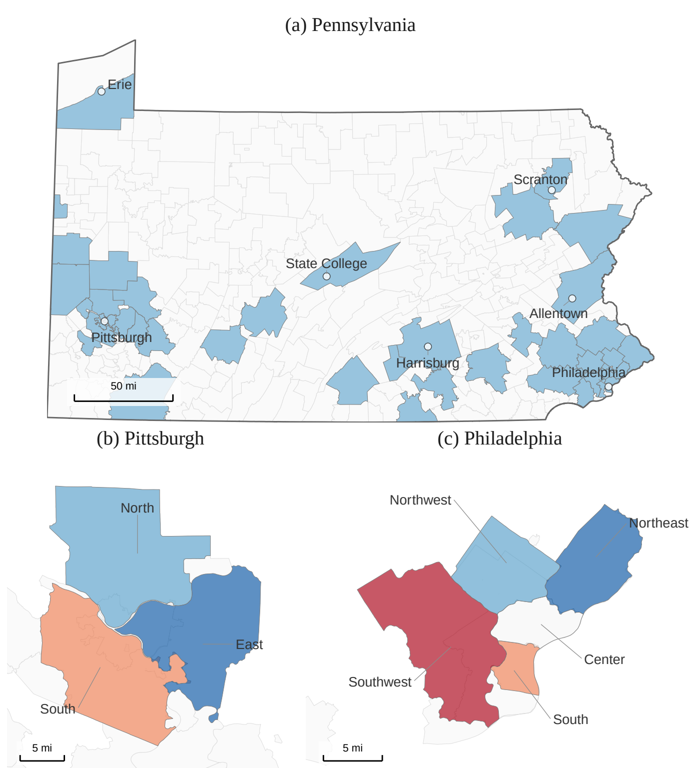

Notes: The figure displays the geographic markets used in the demand estimation. Shaded areas correspond to preschool markets. The bottom panels zoom into the Philadelphia and Pittsburgh metropolitan areas, where large urban markets are split into submarkets.

Table C1. Market Characteristics in 2018

Market	Children	Sub. (%)	Income ($000)	Centers	Price ($)	S1	S2	S3	S4	BGs
Abington Hatboro Horsham	8,626	35.6	119.0	68	49.0	39	6	7	16	262
Allentown Bethlehem Easton	14,919	47.3	81.8	201	34.4	95	42	29	35	444
Altoona	2,355	57.5	61.2	35	29.0	11	8	1	15	101
Canonsburg Peters	2,254	31.9	123.7	22	45.1	12	1	2	7	61
Carlisle	1,389	43.8	87.0	14	34.8	6	0	1	7	37
Central Bucks North Penn	7,803	29.6	127.2	88	46.6	45	9	8	26	202
Chambersburg Shippensburg	2,470	47.0	79.8	22	29.7	14	4	1	3	58
Chester	8,960	41.8	111.1	79	45.4	43	15	5	16	306
Erie	5,833	58.4	65.7	69	35.1	14	15	12	28	217
Freeport Burrell New Kensington	3,037	47.5	86.7	32	34.4	16	10	4	2	121
Greensburg	2,538	48.5	78.0	24	36.1	7	6	2	9	98
Hanover	2,180	43.9	79.8	16	32.2	10	2	1	3	61
Harrisburg	11,883	47.9	86.9	160	39.3	76	40	18	26	325
Hermitage	896	61.9	56.2	12	40.7	3	1	4	4	50
Johnstown	1,960	61.6	58.9	22	27.1	11	4	4	3	102
Lancaster	9,888	44.0	84.0	83	40.0	31	9	9	34	245
Laurel Highlands	2,311	62.8	56.6	24	30.1	10	7	3	4	86
Levittown	9,497	36.4	119.4	102	46.8	44	20	16	22	290
Mars Butler Seneca	3,743	37.6	114.2	36	39.0	23	4	3	6	99
Moon Sewickley	1,739	35.5	108.7	11	43.3	5	1	2	3	55
New Castle	1,985	57.3	61.1	15	35.4	3	3	5	4	78
Norristown	7,712	32.0	128.6	82	48.5	48	11	8	15	226
Penn Trafford Norwin	2,851	42.0	93.8	24	35.8	16	1	2	5	92
Philadelphia Center	17,912	70.5	50.3	172	38.4	111	25	23	13	500
Philadelphia Northeast	18,300	58.9	68.5	194	38.4	116	35	19	24	448
Philadelphia Northwest	16,527	58.5	72.5	222	38.8	143	37	19	23	520
Philadelphia South	9,447	59.8	76.5	128	43.9	83	11	16	18	387
Philadelphia Southwest	22,681	52.1	92.1	375	40.2	239	60	32	44	808
Pittsburgh East	10,797	56.3	73.1	161	38.8	85	24	14	38	562
Pittsburgh North	8,476	42.1	104.7	69	42.6	40	9	4	16	321
Pittsburgh South	11,607	51.0	82.0	101	41.5	67	12	4	18	545
Pottstown Phoenixville	7,446	32.8	115.9	83	44.4	29	14	9	31	179
Reading	7,239	54.2	72.2	64	37.2	26	17	6	15	187
Rochester Aliquippa	4,325	47.2	84.0	48	33.1	24	11	8	5	160
Scranton	4,590	55.3	66.2	60	33.4	18	19	8	15	181
State College	1,896	46.6	104.3	34	41.3	13	6	6	9	82
Stroudsburg Pocono	3,374	45.9	86.2	55	32.2	44	0	4	7	96
West Chester Great Valley Downingtown	8,924	29.1	137.7	100	50.9	50	19	6	25	211
Wilkes Barre	5,860	54.3	67.6	72	30.5	33	17	13	9	230
York	6,432	49.3	82.0	54	34.0	16	18	7	13	201

Notes: The table reports market-level descriptive statistics for 2018. Children are 3- and 4-year-olds. Average household income is reported in thousands of 2015 dollars. Columns S1-S4 report the breakdown of centers by STAR rating. BGs denotes census block groups.

C.2 Provider Data Pre-processing

Although unusually rich for early education data, the Pennsylvania provider data contain occasional missing prices and substantial missing enrollment among STAR 1 providers, which were not required to report enrollment. Because these providers are prevalent and contribute to the competitive environment faced by nearby centers, I impute missing prices and market shares rather than drop them from demand estimation.[25] I train gradient-boosted regression models on observed prices and enrollment and use them to predict missing outcomes. Figure C2 shows that both models perform well out of sample.

Figure C2. Price and Enrollment Imputation, Out-of-Sample Fit

(a) Price (b) Enrollment

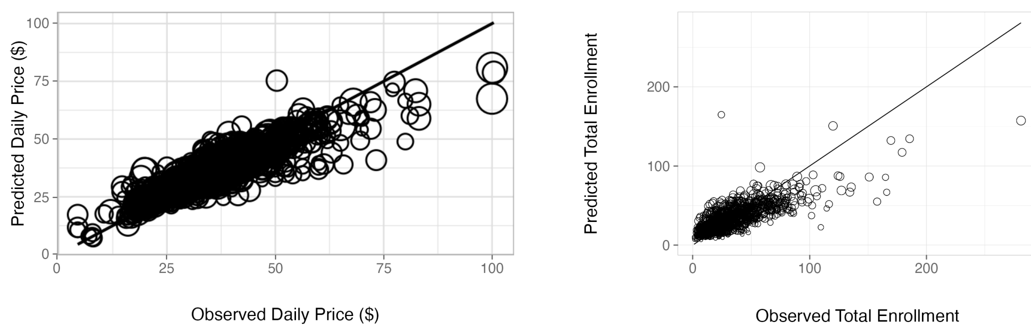

Notes: The figure compares observed outcomes with predictions from XGBoost models evaluated on randomly held-out 10 percent test samples. Panel (a) evaluates the specification used to impute daily prices for providers never observed with a price. Panel (b) evaluates the specification used to impute total preschool enrollment and excludes observations for which observed or predicted enrollment exceeds 300. Point size denotes provider capacity, the diagonal line indicates exact fit, and panel titles report out-of-sample R^2.

C.3 Instrument Construction

Following Berry et al. (1995); Conlon and Gortmaker (2020), I construct instruments combining competitors' characteristics ("BLP instruments") with demographics of consumers located near each provider. Competitor characteristics include nearby center counts, capacity, classrooms, STAR ratings, chain ownership, name-based center-type indicators, and a predicted price based on the center's name. Although STAR ratings are endogenous in the dynamic model, the timing assumption makes them exogenous for stage-game estimation.

[25] An earlier version grouped providers with missing shares into a single category with a common transitory unobserved shock, $\Delta\xi^{miss}$. Besides complicating computation, this approach failed to exploit information about providers' likely enrollment.

Figure C3. First-Stage Strength of BLP-PCA Instruments

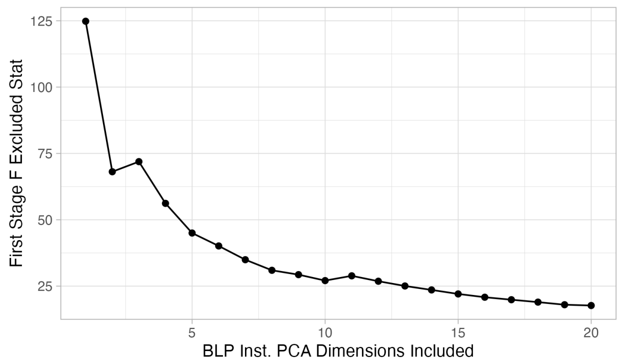

Notes: The figure reports excluded F-statistics for the BLP-PCA instruments. The full model regresses posted daily prices on exogenous product characteristics, center and year fixed effects, and the number of BLP-PCA instruments shown on the x-axis. The restricted model keeps the same exogenous product characteristics and fixed effects but excludes the PCA instruments.

I aggregate these characteristics using inverse-distance sums, Gaussian distance-decay averages and sums, and discrete distance bands. I apply the same distance-weighting schemes to local demographics, including block-group income, preschool-age population, racial composition, college education, and work status, and also include tract-level wages for teachers and administrators in child-service occupations as cost shifters.

These variables generate a high-dimensional instrument matrix. I residualize the instruments with respect to center and year fixed effects and apply PCA to the residualized matrix, including interactions between BLP and demographic instruments. The resulting orthogonal components retain much of the variation in the original instruments while reducing their number and strengthening the first stage. Figure C3 reports the excluded first-stage F-statistic as a function of the number of PCA instruments and shows that the procedure yields a reasonably strong first stage even with center and year fixed effects.

C.4 Simulation of Preschool Consumers

I construct the distribution of preschool-aged children using ACS data at the census-block-group level. Because counts of three- and four-year-olds are unavailable, I apply each tract's share of three- and four-year-olds among children under five to the under-five population of its constituent block groups. For each block group and year, I combine the resulting preschool-age population with the share of college-educated women and the household-income histogram. To construct education-specific income distributions, I use PUMA-year distributions of log household income by education as priors, updating their means using the block-group income histogram while retaining the PUMA-year conditional variances.

Given the size of the markets, I depart from Neilson (2013) and approximate each market-year consumer distribution with 10,000 equally weighted Halton draws. The first coordinate assigns a block group in proportion to its population of three- and four-year-olds, the second assigns college education using the corresponding block-group share, and the third draws income from the resulting block-group- and education-specific lognormal distribution. I divide annual income by 180 to express it on the same daily scale as preschool prices and convert it to 2015 dollars. Depending on the specification, additional Halton coordinates generate the standard-normal draws for the random coefficients. The resulting simulated types $(\ell_i, e_i, y_i, \boldsymbol{v}_i)$ approximate the full spatial and demographic support without enumerating it.

C.5 Model Fit

Table C2 reports the fit of the micro-moments used to inform the estimation of the stage-game parameters governing heterogeneous preferences across demographics. Data to form the micro-moments come from the NSECE as described in Section 2 and from the main preschool panel data set, exploiting the feature that not only aggregate market shares but also enrollment by payment streams are observed.

Table C2. Stage-Game Model Fit

Name	Moment	Data	Model
A. Targeted Moments			
Avg. Distance (College)*	$\mathbb{E}[d_{ij} \mid j \neq 0 \,\&\, educ_i = 1]$	8.42	8.42
Avg. Distance (Non College)*	$\mathbb{E}[d_{ij} \mid j \neq 0 \,\&\, educ_i = 0]$	5.59	5.59
College (cond. Enrol.)*	$\mathbb{E}[educ_i \mid j \neq 0]$	0.47	0.47
Subsidy cond. STAR 4 Enrol.	$\mathbb{E}[\mathbf{1}\{y_i < 2FPL\} \mid r_j = 4]$	0.39	0.39
Subsidy cond. Enrol. and price above MCCA†	$\mathbb{E}[\mathbf{1}\{y_i < 2FPL\} \mid j \neq 0 \,\&\, p_j > MCCA_j]$	0.41	0.44
Subsidy times Price (cond. Enrol. and price above MCCA)†	$\mathbb{E}[p_j \mathbf{1}\{y_i < 2FPL\} \mid j \neq 0 \,\&\, p_j > MCCA_j]$	15.38	16.41
Subsidized Enrollment Error, Relative†	$\frac{\mathbb{E}[N_m \sum_i \mathbf{1}\{y_i < 2FPL\} w_{im} s_{ijm}] - n_{jm}^{sub,obs}}{\mathbb{E}[n_{jm}^{sub,obs}]}$	0.00	-0.07
B. Implied Statistics			
Covariance (price, subsidy)†	$\mathbb{C}(p_j, \mathbf{1}\{y_i < 2FPL\} \mid j \neq 0 \,\&\, p_j > MCCA_j)$	-0.62	-0.68

Notes: The data column reports empirical moments and the model column reports their simulated counterparts at the estimated parameters. MCCA denotes the state's maximum child care allowance. * National Survey of Early Care and Education (NSECE), † Pennsylvania Office of Child Development and Early Learning (OCDEL).

C.6 Stage-Game Estimates: Additional Figures

Figure C4. Distribution of Household Price Sensitivity

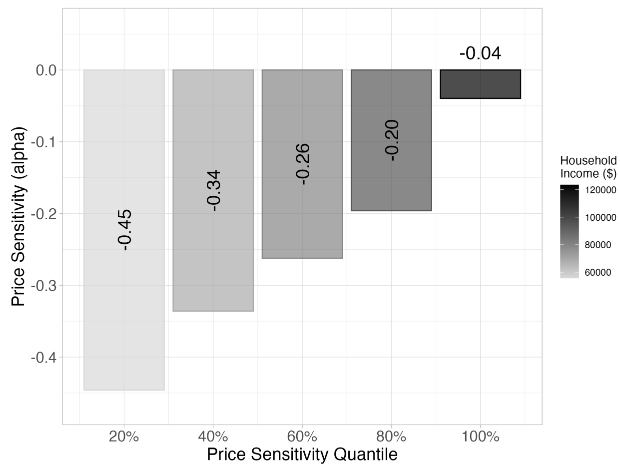

Notes: The figure reports quantiles of household-level price sensitivity α_i. More negative values indicate greater sensitivity to higher consumer prices. Shading reports the average household income within each price-sensitivity bin, in 2015 dollars.

Figure C5. Supply Statistics by STAR Rating, 2018

(a) Marginal Cost

(b) Markup

(c) Lerner Index

(d) Own-price Elasticity

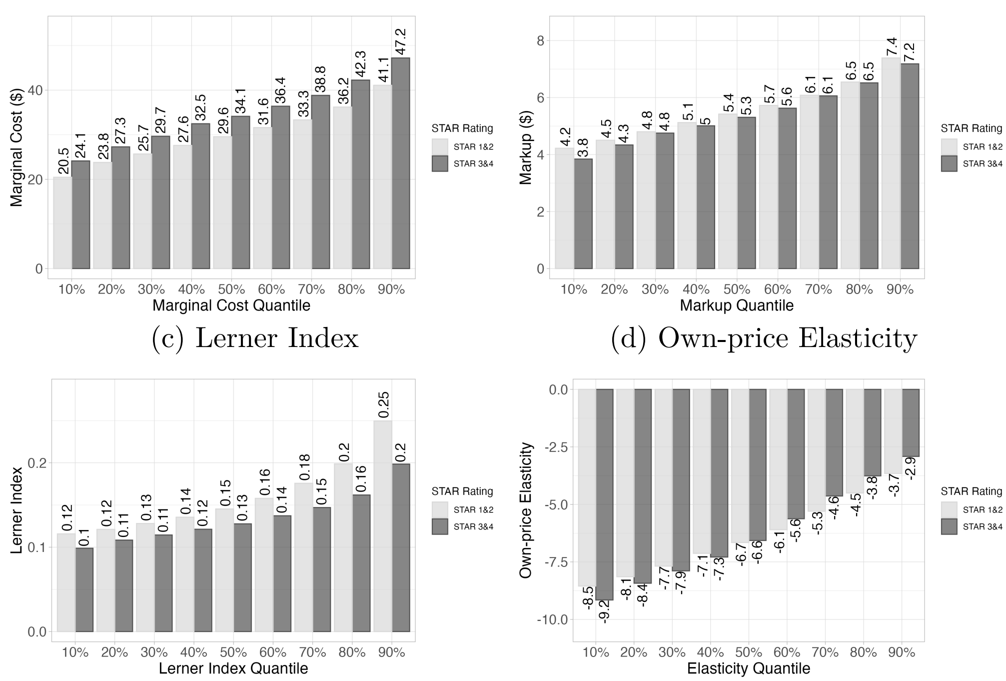

Notes: Bars report deciles of estimated provider-level supply statistics in 2018, separately for STAR 1&2 and STAR 3&4 centers. Marginal costs and markups are in 2015 dollars per day.

Figure C6. Unobserved Center Characteristics by STAR Rating

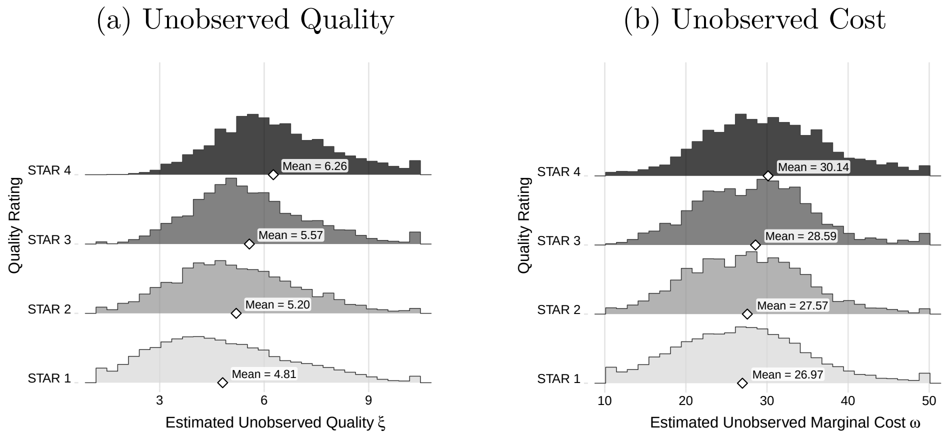

Notes: Panel (a) reports the distribution of $\xi_j + \xi_t$ and panel (b) of $\omega_j + \omega_t$ by STAR rating across all centers–year observations in the demand estimation sample.

Implications of Spatial Competition. Estimated coefficients in Table 4 show distaste for travel time, especially for low-income households. This negative coefficient on distance has two key consequences for the model. On the demand side, parents only value centers that are close enough to their home locations. To see this, I compute a parent's willingness to pay to keep a preschool in the choice set, following Ho (2006); Conlon and Mortimer (2021):

$$WTP_i(j) = \frac{1}{|\alpha_i|}\left(\log\left(1 + \sum_{k\in\mathcal{J}} \exp U_{ik}\right) - \log\left(1 + \sum_{k\in\mathcal{J}\setminus j} \exp U_{ik}\right)\right) \tag{16}$$

I then average willingness to pay across distance bands around a family's home location. Figure C7a shows the resulting averages by increments of 10 minutes of travel time. The willingness to pay for a nearby preschool is about \$4 per household-year, reflecting the availability of multiple close substitutes.[26] Because of parents' distaste for travel time, this willingness to pay falls quickly with distance: parents are almost indifferent to the removal of a center more than 30 minutes away. This spatial structure of demand in turn impacts how preschools compete. The first column of Appendix Table D1 shows that stage-game profits exhibit the patterns of a spatial oligopoly: within a preschool, profits increase with unobserved quality, are higher at highly rated tiers, and increase with the presence of nearby high-income consumers, but fall with the presence of nearby competitors, especially those with higher unobserved quality. I compute price diversion ratios, which measure how demand

[26] These estimates do not reflect the value of removing access to preschool altogether, but only the welfare implication of a one-unit change in parents' choice set. Because of the nonlinearity of the logarithm, removing additional preschools would increasingly harm parents.

shifts between centers when one raises its price, as $D_{jk} = -\frac{\partial s_k}{\partial p_j} / \frac{\partial s_j}{\partial p_j}$ (Conlon and Mortimer, 2021). As shown in Figure C7b, the intensity of competition between two centers declines steeply with distance. Centers more than 5 km apart compete very little.

Figure C7. Role of Distance for Demand and Supply

(a) Willingness to Pay (b) Diversion ratios

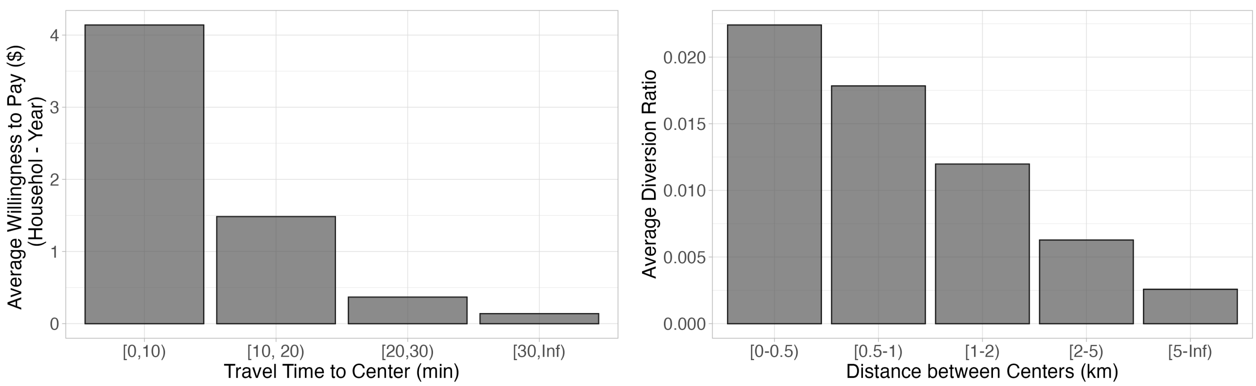

Notes: Panel (a) reports the average annual willingness to pay to keep a center in a household's choice set, by travel-time bins from the household to the center. Values are in 2015 dollars and annualized to a 180-day preschool year. Panel (b) reports average price diversion ratios between distinct centers, by distance bins in kilometers; own diversions are excluded.

C.7 Stage-Game Price Equilibrium Algorithm

Having estimated the stage-game parameters, I can compute equilibrium prices in the Nash–Bertrand pricing game. Although the existence of an equilibrium in a differentiated-products Nash–Bertrand game is not guaranteed, standard numerical practice over a reasonable range of price elasticities typically yields a stable equilibrium price vector (Morrow and Skerlos, 2011; Conlon and Gortmaker, 2020). I use a standard fixed-point approach to solve for a new stage-game equilibrium, sometimes described as "BLP markup" iteration in the differentiated-demand literature. Given demand parameters $\boldsymbol{\theta}$, a policy environment Ψ, a vector of estimated marginal costs $\boldsymbol{mc}$, and an initial price vector $\boldsymbol{p}^0$, the procedure follows this iterative scheme:

$$F_\Psi : \boldsymbol{p}^h \to \boldsymbol{p}^{h+1} = \boldsymbol{\eta}(\boldsymbol{p}^h; \boldsymbol{\theta}, \Psi) + \boldsymbol{mc}$$

I augment this standard fixed-point iteration with Anderson acceleration (Anderson, 1965), which greatly reduces the number of iterations required for convergence. In the more generous demand-subsidy scenarios, convergence may be hindered by numerical instability caused by highly price-insensitive consumers who are shielded from most price variation by subsidies. To address this issue, I use a homotopy approach in which I successively solve for the fixed point of the map $F_{\Psi(\tau)}$ as $\tau \to 1$, where τ gradually introduces the wedge that subsidies create between provider and consumer prices. More precisely, $F_{\Psi(\tau)}$ corresponds to the mapping in which $\frac{\partial p_{ij}^{\text{cons}}(p;\Psi)}{\partial p_j}$ is replaced by

$$D_{ij,\tau}(p) = \tau \frac{\partial p_{ij}^{\text{cons}}(p;\Psi)}{\partial p_j} + (1-\tau), \qquad \tau \in [0,1].$$

I stop this procedure either when τ reaches 1 or when the fixed point for $\tau < 1$ fails to converge. I typically obtain convergence for values of τ very close to one. Thus, I find either an equilibrium corresponding to the exact policy environment or an equilibrium corresponding to an environment very close to the target.

C.8 Computational Implementation

As Table C1 shows, markets contain many providers and simulated families. I implement the core demand and pricing operations in Torch using GPU-accelerated tensors. Following the masking-and-stacking approach in Lewis et al. (2021), I batch consumer draws and pad each market's choice set to $J^{\max}$ products, using masks to exclude padded products, before stacking M markets for parallel evaluation. Computations therefore operate on $[I_b, J^{\max}, M]$ tensors, where $I_b < I$ is the number of consumer draws in batch b, $J^{\max}$ is the largest choice set among the stacked markets, and M is the number of markets evaluated in parallel. To reduce memory requirements further, I compute the required derivatives analytically rather than carrying automatic-differentiation graphs.

D Dynamic Model Supplements

D.1 Solution Approach

I develop a state-reduction and profit-approximation framework for dynamic spatial competition games. In contrast to oblivious or partially oblivious approaches, which obtain tractability by restricting firms' strategies to depend on their own state and an aggregate summary of the industry (Weintraub et al., 2008; Ifrach and Weintraub, 2017), this setting requires local statistics: in spatial oligopoly games, firms' profitability depends on nearby competitors, local demographics, local cost conditions, and subsidy generosity. Moreover, the relevant low-dimensional summary might not be known ex ante because stage-game profits are shaped by an equilibrium that depends on complex interactions between local demand and competition. As I show below, even carefully engineered measures of local conditions constructed by the researcher, as in Sweeting (2013), might not adequately represent the stage game. The approach therefore combines spatial aggregation of local demand and competition conditions with an economics-informed neural network that imposes basic pricing and profit restrictions while learning a low-dimensional representation of the stage game for use in the dynamic problem.

D.1.1 Dynamic Game Solution Algorithm

Solution. For a given candidate dynamic parameter vector $\boldsymbol{\psi}$, solving the model requires three objects: a conjectured profile of conditional choice probabilities $\mathcal{P}$, a set of simulated next-period states $\{\mathcal{M}'_h(\mathcal{M}, a)\}_{h=1}^{H}$, and a vector of sieve basis functions $\Phi(\mathcal{M})$ defined over the state space. The basis functions are used to project the value function onto a low-dimensional sieve space, while the simulated next states are used to approximate transition expectations. In particular, for each current state-action pair $(\mathcal{M}, a)$, I approximate expectations over next-period states under $\mathcal{P}$ using importance-sampling weights induced by rivals' conjectured choice probabilities. Denoting these weights by $\omega_h(\mathcal{M}, a; \mathcal{P})$, with $\sum_{h=1}^{H} \omega_h(\mathcal{M}, a; \mathcal{P}) = 1$, the expected basis vector is

$$\mathbb{E}_{\mathcal{P}}[\Phi(\mathcal{M}') \mid \mathcal{M}, a] \approx \sum_{h=1}^{H} \omega_h(\mathcal{M}, a; \mathcal{P})\, \Phi(\mathcal{M}'_h(\mathcal{M}, a)) \, .$$

Conditional on these objects, I approximate the value function with a linear sieve,

$$V(\mathcal{M}; \boldsymbol{\psi}, \mathcal{P}) \approx \Phi(\mathcal{M})' \rho(\boldsymbol{\psi}; \mathcal{P}),$$

where $\rho(\boldsymbol{\psi}; \mathcal{P})$ denotes the vector of sieve coefficients. The corresponding Bellman operator is

$$(T_{\boldsymbol{\psi},\mathcal{P}} V)(\mathcal{M}) = \gamma \psi^{\epsilon} + \psi^{\epsilon} \log \sum_{a \in \mathcal{A}(\mathcal{M})} \exp\left(\frac{\Pi(a, \mathcal{M}; \boldsymbol{\psi}) + \beta\, \mathbb{E}_{\mathcal{P}}[\Phi(\mathcal{M}') \mid \mathcal{M}, a]'\, \rho(\boldsymbol{\psi}; \mathcal{P})}{\psi^{\epsilon}} \right)$$

where γ is the Euler-Mascheroni constant. For a fixed conjecture $\mathcal{P}$, I solve the agent's dynamic problem by iterating on the projected Bellman equation,

$$\rho^{\text{new}}(\boldsymbol{\psi}; \mathcal{P}) = (\Phi'\Phi)^{-1} \Phi' T_{\boldsymbol{\psi},\mathcal{P}}(\Phi \rho(\boldsymbol{\psi}; \mathcal{P})) \, ,$$

which yields the continuation-value approximation

$$\mathbb{E}_{\mathcal{P}}[V(\mathcal{M}'; \boldsymbol{\psi}, \mathcal{P}) \mid \mathcal{M}, a] \approx \mathbb{E}_{\mathcal{P}}[\Phi(\mathcal{M}') \mid \mathcal{M}, a]'\, \rho(\boldsymbol{\psi}; \mathcal{P}).$$

Given the converged sieve coefficients, define the action-specific continuation value

$$FV(a, \mathcal{M}; \boldsymbol{\psi}, \mathcal{P}) = \beta\, \mathbb{E}_{\mathcal{P}}[\Phi(\mathcal{M}') \mid \mathcal{M}, a]'\, \rho(\boldsymbol{\psi}; \mathcal{P}).$$

Combining this continuation value with the action-specific flow payoff gives the implied choice probabilities

$$P(a \mid \mathcal{M}; \boldsymbol{\psi}, \mathcal{P}) = \frac{\exp([\Pi(a, \mathcal{M}; \boldsymbol{\psi}) + FV(a, \mathcal{M}; \boldsymbol{\psi}, \mathcal{P})] / \psi^{\epsilon})}{\sum_{a' \in \mathcal{A}(\mathcal{M})} \exp([\Pi(a', \mathcal{M}; \boldsymbol{\psi}) + FV(a', \mathcal{M}; \boldsymbol{\psi}, \mathcal{P})] / \psi^{\epsilon})}.$$

The variable-profit component of Π is common across actions and therefore cancels from these probabilities. For entrant states, only the subset of feasible actions is included in the denominator, with infeasible actions assigned probability zero by construction.

The equilibrium is therefore computed as a fixed point in conditional choice probabilities. Starting from an initial guess $\mathcal{P}$, I construct importance-sampling transition expectations, solve the projected Bellman equation for $\rho(\boldsymbol{\psi}; \mathcal{P})$, update the implied probabilities $P(\cdot \mid \mathcal{M}; \boldsymbol{\psi}, \mathcal{P})$, and iterate until the CCPs converge in the sup norm. In the implementation, both the sieve coefficients and the CCPs are warm-started from the previous outer iteration, and the CCP update is damped for numerical stability.

Estimation. To evaluate the dynamic model for a candidate parameter vector $\boldsymbol{\psi}$, I take as given two objects from the first stage: the estimated rivals' conditional choice probabilities $\hat{\mathcal{P}}_0$, which determine the transition expectations, and the vector of sieve basis functions $\Phi(\mathcal{M})$. Conditional on these first-stage objects, I approximate the value function as

$$V(\mathcal{M}; \boldsymbol{\psi}, \hat{\mathcal{P}}_0) \approx \Phi(\mathcal{M})' \rho(\boldsymbol{\psi}; \hat{\mathcal{P}}_0),$$

and compute the expected basis vector using the same importance-sampling construction as in the solution step, but fixing rivals' behavior at $\hat{\mathcal{P}}_0$:

$$\mathbb{E}_{\hat{\mathcal{P}}_0}[\Phi(\mathcal{M}') \mid \mathcal{M}, a].$$

Thus, when $\boldsymbol{\psi}$ changes, the sieve coefficients $\rho(\boldsymbol{\psi}; \hat{\mathcal{P}}_0)$ are updated, while the transition expectations remain pinned down by the first-stage estimate $\hat{\mathcal{P}}_0$.

For each $\boldsymbol{\psi}$, I recover $\rho(\boldsymbol{\psi}; \hat{\mathcal{P}}_0)$ by iterating on

$$\rho^{\text{new}}(\boldsymbol{\psi}; \hat{\mathcal{P}}_0) = (\Phi'\Phi)^{-1} \Phi' T_{\psi, \hat{\mathcal{P}}_0}\left(\Phi \rho(\boldsymbol{\psi}; \hat{\mathcal{P}}_0)\right),$$

where the Bellman operator $T_{\psi,\hat{\mathcal{P}}_0}$ is defined as above with $\mathcal{P}$ replaced by $\hat{\mathcal{P}}_0$. Given the converged coefficients, the action-specific continuation values are

$$FV(a, \mathcal{M}; \boldsymbol{\psi}, \hat{\mathcal{P}}_0) = \beta\, \mathbb{E}_{\hat{\mathcal{P}}_0}[\Phi(\mathcal{M}') \mid \mathcal{M}, a]'\, \rho(\boldsymbol{\psi}; \hat{\mathcal{P}}_0),$$

which, together with the action-specific flow payoffs, imply conditional choice probabilities

$$P(a \mid \mathcal{M}; \boldsymbol{\psi}, \hat{\mathcal{P}}_0) = \frac{\exp\Big(\Big[\Pi(a, \mathcal{M}; \boldsymbol{\psi}) + FV(a, \mathcal{M}; \boldsymbol{\psi}, \hat{\mathcal{P}}_0)\Big] / \psi^{\epsilon}\Big)}{\sum_{a' \in \mathcal{A}(\mathcal{M})} \exp\Big(\Big[\Pi(a', \mathcal{M}; \boldsymbol{\psi}) + FV(a', \mathcal{M}; \boldsymbol{\psi}, \hat{\mathcal{P}}_0)\Big] / \psi^{\epsilon}\Big)}.$$

For entrant states, only the subset of feasible actions is included in the denominator, with infeasible actions assigned probability zero by construction.

Let a_n denote the observed action in observation n. The sample negative log likelihood is

$$\mathcal{L}(\boldsymbol{\psi}) = -\frac{1}{N} \sum_{n=1}^{N} \log P(a_n \mid \mathcal{M}_n; \boldsymbol{\psi}, \hat{\mathcal{P}}_0).$$

Estimation therefore nests a sieve approximation to the Bellman fixed point inside the likelihood-based estimation of the structural parameters, conditional on the first-stage estimate $\hat{\mathcal{P}}_0$.

D.1.2 Neural Network for Profit Approximation and Basis Construction

Initial representation of the spatial stage-game equilibrium. The nature of competition in the spatial equilibrium model implies a complex state structure. Preschool profitability depends on local demand and supply conditions, and local demand in turn depends on subsidies, the income distribution, and household preferences. On the supply side, preschools are differentiated by observed and unobserved quality, star rating, unobserved costs, and cost shifters. Together, these forces determine the profitability of operating a preschool in a given location, so the approximation must encode a rich local market environment.

I summarize the local stage-game demand-and-supply environment facing center j in market m and year t by a covariate vector x_{jmt}. The construction is similar in spirit to Sweeting (2013): first-stage demand variables are combined with market characteristics to form a tractable state representation for the dynamic problem. There are two additional challenges in this setting. First, competition is inherently local, so no common state summarizing market-level aggregate variables would adequately represent the local environment facing each preschool. Second, the model comprises a pricing game with potentially dozens or hundreds of players in a given market-year, which poses additional computational hurdles. The vector x_{jmt} contains four broad groups of variables: estimated demand and cost primitives $(\xi_{jmt}, \omega_{jmt}, \delta_{jmt}, mc_{jmt})$, center characteristics, surrounding demographics, and local competition measures. Demographic and competition variables are aggregated around each center over concentric distance bands using either sums or inverse-distance-weighted averages. Specifically, for any local variable Z_{lt}, I compute ring-specific sums and inverse-distance-weighted averages,

$$\text{Sum}^r_{jmt}(Z) = \sum_{l:\, d_{lj} \in B_r} Z_{lt}, \qquad \text{Avg}^r_{jmt}(Z) = \sum_{l:\, d_{lj} \in B_r} \omega^r_{lj} Z_{lt},$$

where

$$\omega^r_{lj} = \frac{d_{lj}^{-1}}{\sum_{l':\, d_{l'j} \in B_r} d_{l'j}^{-1}}.$$

I use the same distance-band logic for local competition, aggregating nearby centers' counts, capacities, quality composition, and neighboring demand and cost fundamentals. The resulting vector x_{jmt} therefore provides a reduced-form encoding of the local stage-game equilibrium. Adding interactions between quality-rating indicators and the main demographic, competition, center, and demand blocks yields a total of 933 covariates.

Training Data Construction. To expand the support of market environments used to train the profit approximation, I augment the realized panel with simulated market states. For each market m, year t, and draw s, I perturb provider operating states by allowing incumbents to switch quality tiers or exit and potential entrants to enter at STAR 1. I redraw provider demand and cost shocks, $(\Delta\xi^s_{jmt}, \Delta\omega^s_{jmt})$, from a truncated empirical bivariate normal distribution estimated from the first-stage residuals, and perturb local demographic primitives within bounded ranges around their observed values, including income-location parameters, income dispersion, education shares, and the population of 3- and 4-year-olds. I then recompute the associated quadrature weights and income-bin variables and solve for Nash–Bertrand equilibrium prices as described in Section C. Repeating this procedure across markets, years, and draws generates a large set of counterfactual stage-game environments, expanding the support of x_{jmt} and profits used to train the neural network. This data-augmentation approach is similar to Sweeting (2013), with the additional requirement of solving for equilibrium prices in each simulated market.

Motivation for a Nonlinear Model. I assess the informativeness of the preliminary basis vector $x \in \mathbb{R}^K$ by predicting firm profits with a LASSO regression, selecting the ℓ_1 penalty by grid search. The resulting model achieves an out-of-sample R^2 of about 0.69 (Figure

Figure D1. LASSO Model Profits Out-of-sample Fit

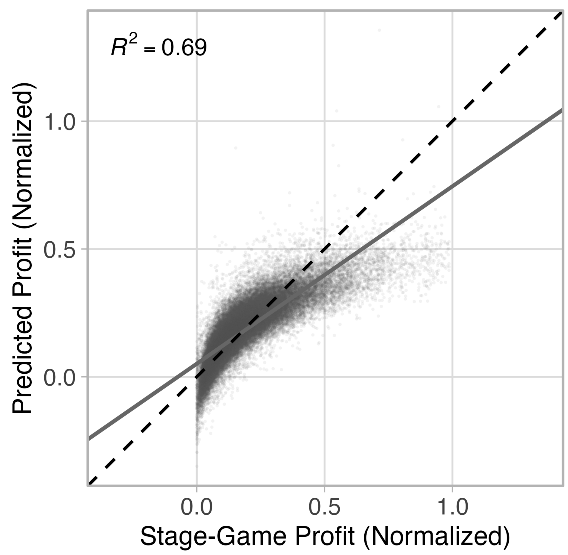

Notes: The figure compares realized and LASSO-predicted profits in the held-out test sample. The penalty parameter is selected via cross-validation. The dashed line is the 45-degree line, and the solid line is the linear fit.

D1), indicating that the hand-crafted basis functions are informative but leave substantial variation in profits unexplained. Because prediction errors may bias the dynamic estimates by missing interactions among competition, demand, firm characteristics, and policy, I construct improved basis functions Φ using supervised dimension reduction through a neural network whose architecture incorporates restrictions implied by oligopolistic competition.

Model Architecture. To improve on the LASSO benchmark, I use a neural-network encoder as a nonlinear analogue to PCA (Hinton and Salakhutdinov, 2006; Goodfellow et al., 2016), combined with supervised training to improve generalization (Le et al., 2018). Generalization is particularly important because the learned representation must approximate not only profits but also the value functions governing entry, exit, and quality decisions. Figure D2 summarizes the architecture. A shared encoder maps the covariates x into a low-dimensional representation Φ, which is concatenated with x through a skip connection for the price and quantity heads. The architecture imposes economic restrictions: predicted price exceeds marginal cost, $\hat{p} = mc + \text{softplus}(\cdot)$, and predicted quantity is positive. Profits are predicted both linearly from Φ and structurally as $\hat{\pi}^{PQ} = \hat{q}(\hat{p} - mc)$, encouraging Φ to retain information relevant for oligopolistic profits. Finally, a decoder reconstructs the original inputs x from Φ, ensuring that the latent representation also preserves information potentially relevant for value functions.

For each observation i, the network takes inputs (x_i, mc_i) and produces predictions $\hat{p}_i$, $\hat{q}_i$, $\hat{\pi}_i^{\text{lin}}$, and $\hat{x}_i$. In addition, the architecture defines a profit measure implied by the price and quantity predictions:

$$\hat{\pi}_i^{PQ} = \hat{q}_i(\hat{p}_i - mc_i). \tag{17}$$

Let p_i denote observed price, q_i observed quantity (enrollment), π_i observed profit, and x_i the original covariates. The model is trained using a weighted sum of six mean-squared-error loss components:

$$\mathcal{L}_p = \frac{1}{N}\sum_{i=1}^{N}(\hat{p}_i - p_i)^2, \qquad \mathcal{L}_q = \frac{1}{N}\sum_{i=1}^{N}(\hat{q}_i - q_i)^2, \qquad \mathcal{L}_{\pi,PQ} = \frac{1}{N}\sum_{i=1}^{N}(\hat{\pi}_i^{PQ} - \pi_i)^2,$$
$$\mathcal{L}_{\pi,\text{lin}} = \frac{1}{N}\sum_{i=1}^{N}(\hat{\pi}_i^{\text{lin}} - \pi_i)^2, \qquad \mathcal{L}_{\text{cons}} = \frac{1}{N}\sum_{i=1}^{N}(\hat{\pi}_i^{PQ} - \hat{\pi}_i^{\text{lin}})^2, \qquad \mathcal{L}_{\text{dec}} = \frac{1}{N}\sum_{i=1}^{N}\|\hat{x}_i - x_i\|^2.$$

The overall training objective is

$$\mathcal{L} = \lambda_p\mathcal{L}_p + \lambda_q\mathcal{L}_q + \lambda_{\pi,PQ}\mathcal{L}_{\pi,PQ} + \lambda_{\pi,\text{lin}}\mathcal{L}_{\pi,\text{lin}} + \lambda_{\text{cons}}\mathcal{L}_{\text{cons}} + \lambda_{\text{dec}}\mathcal{L}_{\text{dec}}.$$

The weights $\lambda_p, \lambda_q, \lambda_{\pi,PQ}, \lambda_{\pi,\text{lin}}, \lambda_{\text{cons}}, \lambda_{\text{dec}}$ are allowed to vary over training. In particular, the training schedule initially places substantial weight on the decoder loss to learn a stable latent representation, and later shifts weight toward the profit and consistency losses.

Figure D2. Architecture of the Economics-Informed Neural Network Model

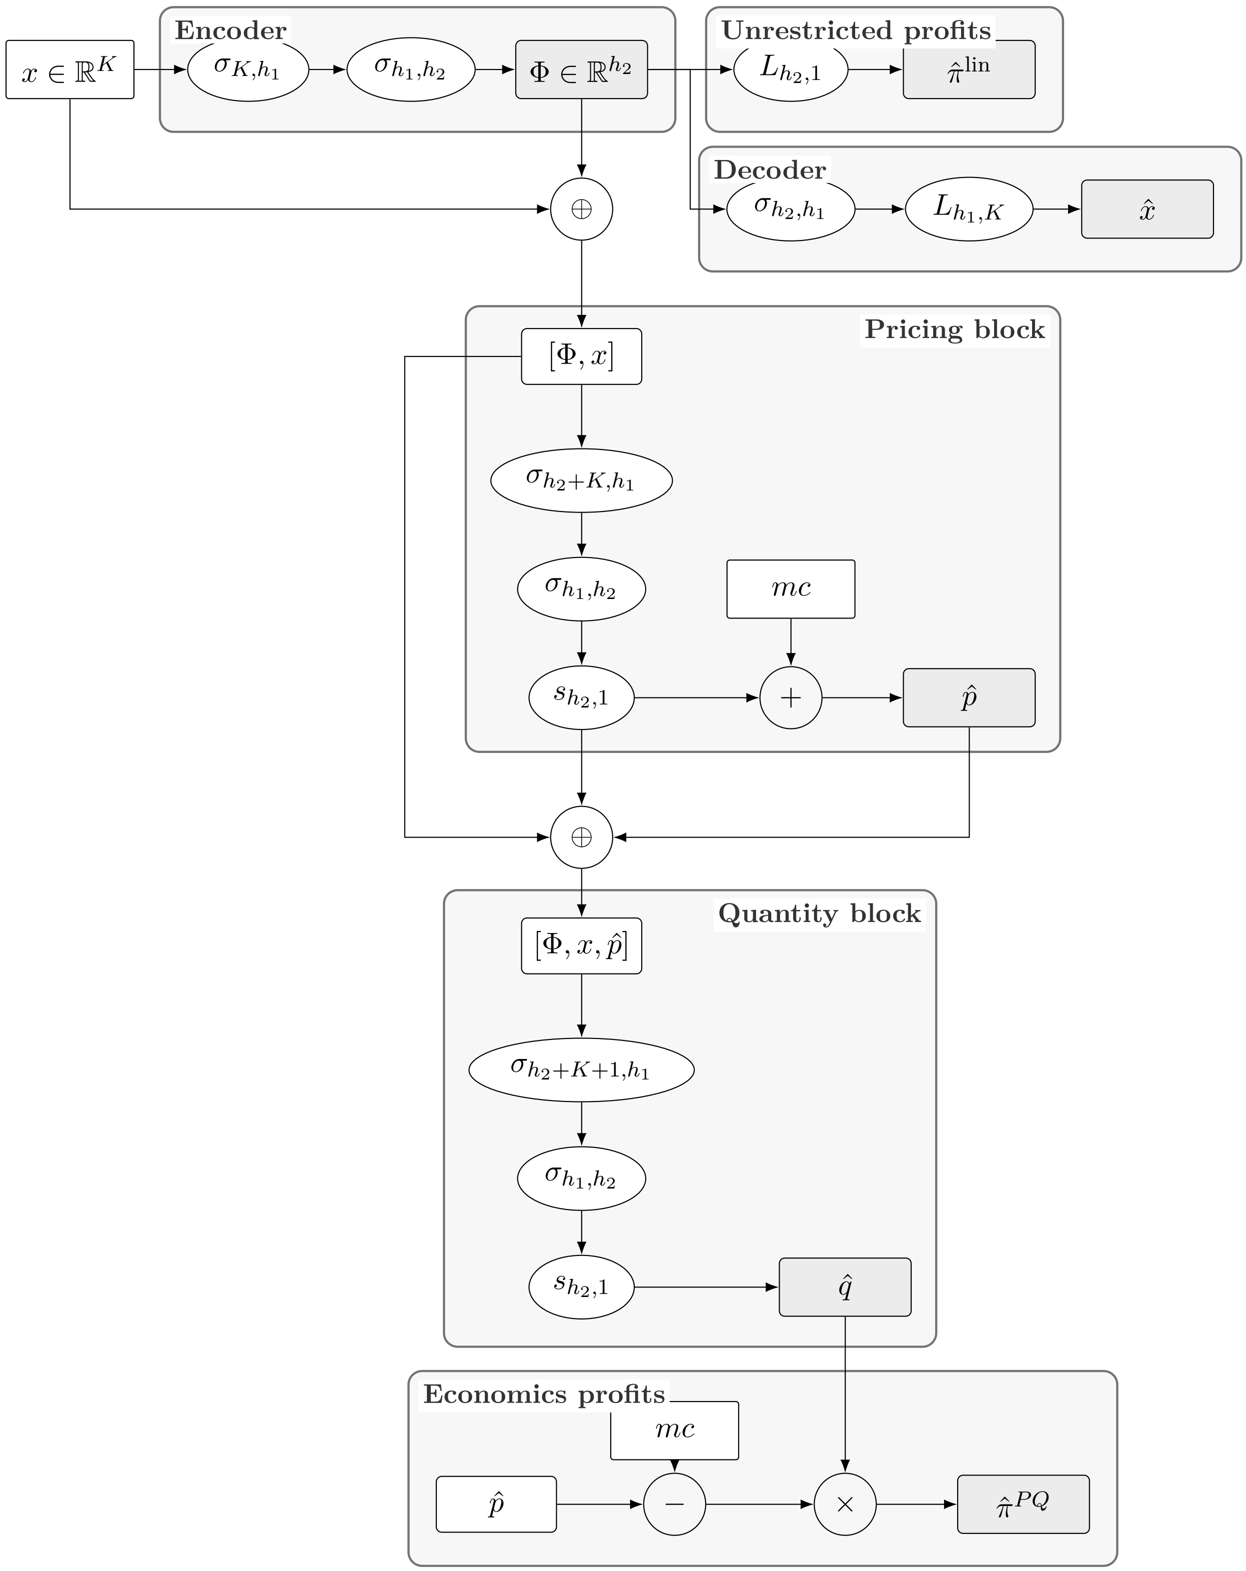

Notes: L denotes linear layers, σ denotes ReLU activations, and s denotes softplus activations. h_1 and h_2 denote the widths of the hidden layers.

Neural Net Fit.

Figure D3. Neural Net Model Profits Out-of-sample Fit

(a) Econ-informed Model π^{PQ}, $R^2 = 0.98$ (b) Linear Model $\pi^{Lin}(\Phi)$, $R^2 = 0.97$

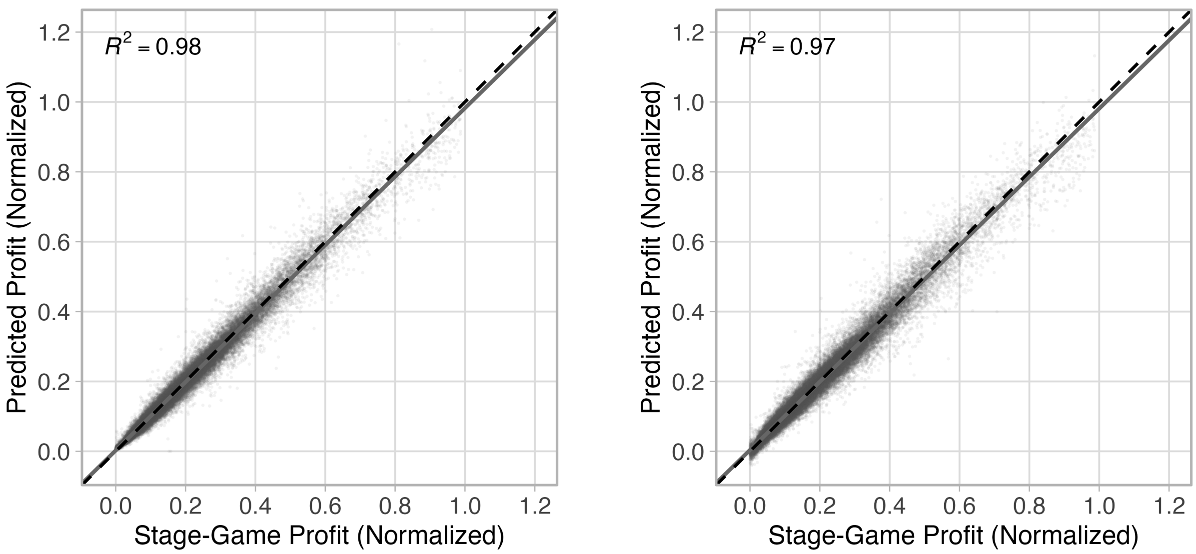

Notes: The figure compares realized and predicted profits in the held-out test sample. Panel (a) constructs profits from the network's price and enrollment predictions, $\hat{\pi}^{PQ} = \hat{q}(\hat{p} - mc)$. Panel (b) predicts profits linearly from the learned representation, $\hat{\pi}^{Lin}(\Phi)$. Dashed lines are 45-degree lines; solid lines are linear fits.

Table D1. Neural Network Moment Fit

Variable	Stage game profits	NNet profits
Competitor count (within 5,000m)	-3.926***	-4.027***
	(0.237)	(0.235)
Competitor count (beyond 5,000m)	6.190***	6.232***
	(1.347)	(1.334)
Competitor sum ξ (0–500m)	-7.256***	-7.345***
	(0.393)	(0.390)
STAR 4	22.759***	20.960***
	(0.390)	(0.386)
STAR 3	2.422***	1.909***
	(0.389)	(0.385)
STAR 2	-16.527***	-15.687***
	(0.355)	(0.351)
Unobserved quality (ξ)	202.419***	199.914***
	(0.550)	(0.545)
Unobserved cost (ω)	-82.062***	-80.934***
	(0.692)	(0.685)
Children age 0–5, college, hh. income 100k+	11.993***	12.227***
	(0.511)	(0.506)
Preschool and year fixed effects	Yes	Yes
Observations	361,848	361,848

Notes: Each cell reports the coefficient from the regression $\pi^k_{jt} = \beta x_{jt} + \mu_j + \phi_t + \epsilon_{jt}$, where profit type k denotes either profits from the stage-game demand-and-supply estimates ("Stage Game") or profits predicted by the neural network ("NNet").

D.2 Expectations

D.2.1 Demographics

I model the evolution of demographic state variables as independent log-linear autoregressive processes estimated on panel data.

For each block-group-level state variable $X_{gt} \in \{\mu^{\text{clg}}_{gt}, \mu^{\text{nonclg}}_{gt}, s^{\text{clg}}_{gt}, N^{3\text{–}4}_{gt}\}$, where g indexes block groups and t indexes years, I estimate:

$$\log X_{gt} = \alpha_g + \lambda_t + \rho_X \log X_{g,t-1} + \varepsilon_{gt}. \tag{18}$$

Here, α_g denotes a block-group fixed effect, λ_t denotes a year fixed effect, and ρ_X is the persistence parameter for variable X.

Similarly, for each PUMA-level dispersion variable $Z_{pt} \in \{\sigma^{\text{clg}}_{pt}, \sigma^{\text{nonclg}}_{pt}\}$, where p indexes PUMAs, I estimate:

$$\log Z_{pt} = \alpha_p + \lambda_t + \rho_Z \log Z_{p,t-1} + u_{pt}. \tag{19}$$

In this specification, α_p denotes a PUMA fixed effect and λ_t again denotes a year fixed effect.

All equations are estimated separately by ordinary least squares using the panel observations for 2010–2017, with standard errors clustered at the geographic level of the panel unit. I present the estimated coefficients below:

Table D2. Demographic State Laws of Motion

Logged state variable	Persistence ρ	Block-group FE	PUMA FE	Year FE	R^2	Observations
A. Block-Group States						
Mean household income, college	0.5486^{***} (0.0076)	Yes	No	Yes	0.91192	66,697
Mean household income, noncollege	0.6772^{***} (0.0061)	Yes	No	Yes	0.94925	66,697
College share	0.6408^{***} (0.0068)	Yes	No	Yes	0.91206	63,689
Population ages 3–4	0.6322^{***} (0.0052)	Yes	No	Yes	0.79679	59,649
B. PUMA States						
SD household income, college	-0.1274^{***} (0.0454)	No	Yes	Yes	0.40380	592
SD household income, noncollege	-0.1625^{***} (0.0481)	No	Yes	Yes	0.50816	592

Notes: Each row reports the autoregressive coefficient from a separate AR(1) model for the logged demographic state variable listed in the first column. Standard errors are reported in parentheses and clustered at the block-group level in Panel A and the PUMA level in Panel B. All models include year fixed effects and the geographic fixed effects indicated in the table. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.2.2 Competitors: Spatially Oblivious Model

I assume that providers only consider geographically close competitors when forming their expectations of their value function next period. This assumption stems from the consequences of spatial oligopolistic competition displayed in Figure C7, namely that centers far away from each other compete less intensely, and that demand is highly local. In practice, when computing the weights $\omega_h(\mathcal{M}, a; \mathcal{P})$ on simulated next states h, I only consider in h reconfigurations of either the closest 5 competitors, or competitors within a 1,000m radius, whichever yields a larger number of rivals. Providers further away in the market are treated as if their endogenous state (rating, operating status) stays fixed at its value implied by $\mathcal{M}$. This assumption allows me to sample densely alternative states in the state space corresponding to those where relevant close rivals make a move.

D.3 Potential Entrants Simulation

For each market, I set the number of potential entrants in every year equal to twice the largest annual number of entrants observed in that market. The fact that the number of simulated potential entrants is arbitrary is a feature of many entry models (Singleton, 2019; Caoui et al., 2026). I could assume more potential entrants who ended up not entering, which would translate into a higher estimated entry cost. Observed entrants are part of this total, and I simulate only the remaining candidates. To determine their locations, I estimate the tract-level Poisson model

$$N_{cmt}^{\text{Entry}} \mid X_{cmt} \sim \text{Poisson}(\lambda_{cmt}), \tag{20}$$

$$\log \lambda_{cmt} = \beta_1 \log(\text{ChildrenUnder5}_{cmt}) + \beta_2 \log(\text{Density}_{cmt}) + \gamma_m + \delta_t, \tag{21}$$

where c, m, and t index tracts, markets, and years, respectively, while γ_m and δ_t denote market and year fixed effects. I combine the resulting predicted entry intensity with a measure that favors locations with limited existing preschool access, so that counterfactuals may potentially change the spatial distribution of preschool supply, should the policies make entry in these low-realized access locations profitable:

$$\omega_{cmt} = 0.25 \frac{\widehat{\lambda}_{cmt}}{\sum_{c'} \widehat{\lambda}_{c'mt}} + 0.75 \frac{\exp(-A_{cmt}/0.5)}{\sum_{c'} \exp(-A_{c'mt}/0.5)}, \tag{22}$$

where A_{cmt} is preschool capacity per child aged three or four accessible within 20 minutes of tract c. Thus, simulated locations resemble tracts that historically attract entry while deliberately oversampling areas with limited existing access. Candidate tracts are drawn with replacement, excluding tracts containing an observed entrant in that year.

Figure D4. Simulated Entrants in Pittsburgh East

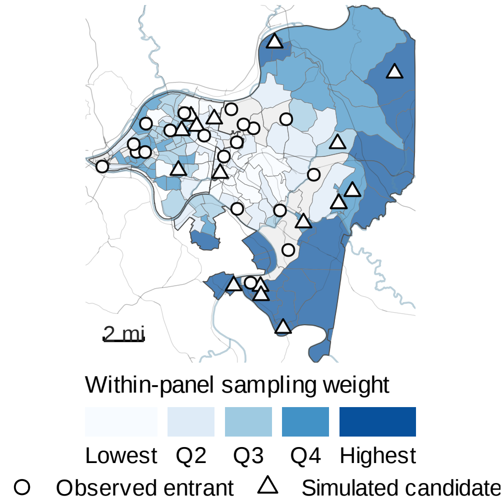

Notes: The figure plots observed entrants and the potential simulated entrants for Pittsburgh East in 2010 that resulted from the procedure described above.

D.4 First-Stage Conditional Choice Probabilities Estimation.

I estimate first-stage conditional choice probabilities (CCPs) as a flexible policy function mapping the payoff-relevant state of each preschool into action probabilities. For center i in market m and year t, let s_{imt} denote the observed state, summarized by the covariate vector x_{imt}, and let $a_{imt} \in \{0, 1, 2, 3, 4\}$ denote the chosen action, corresponding to exit and the four STAR ratings. The neural network outputs

$$P(a_{imt} = a \mid s_{imt}) = P(a_{imt} = a \mid x_{imt}), \qquad a \in \{0, 1, 2, 3, 4\}.$$

The model is trained on the realized panel of observed actions to minimizing

$$\mathcal{L}_{\text{CCP}} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{a=0}^{4} \mathbf{1}\{a_i = a\} \log P(a_i = a \mid x_i),$$

with dropout regularization and early stopping based on validation performance. I evaluate out-of-sample fit by comparing the predicted probabilities assigned to each action across true action classes and, in particular, by examining the predicted probability of remaining in the current state separately for movers and stayers. This first-stage neural-network CCP estimator provides a flexible reduced-form approximation to firms' policy rules in a high-dimensional spatial state space. Figure D5 shows that the model adequately reproduces key out-of-sample transition patterns in the data: the transition matrix is diagonal-heavy except for upgrades by one STAR rating, consistent with the QRIS design of encouraging providers to build quality over time by progressively climbing the rating ladder. Exit ("action 1" in Figure D5) is rare except among potential entrants who stay out and, to a lesser extent, STAR 1 providers.

Figure D5. Conditional Choice Probabilities Out-of-sample Fit

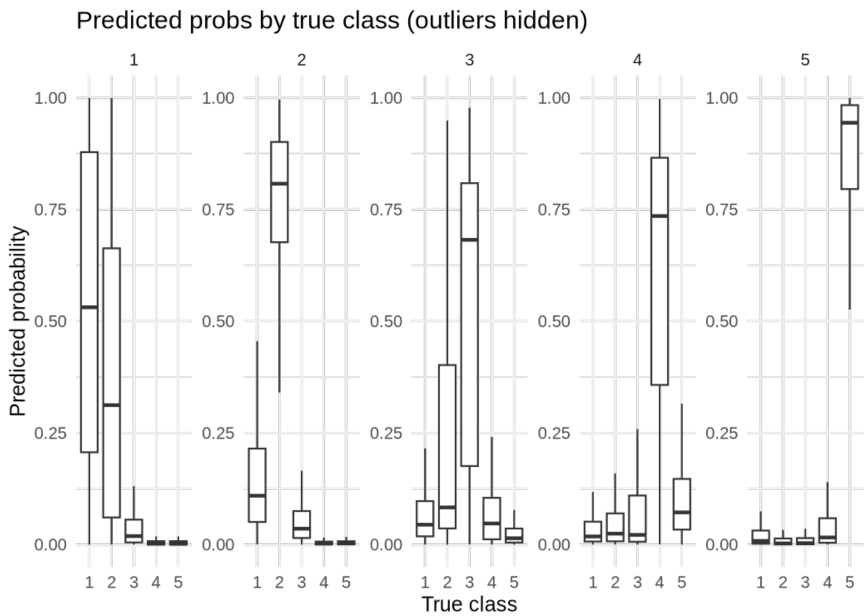

Notes: The figure evaluates the first-stage neural-network CCP model in the held-out test sample. Each panel conditions on the realized action and reports the distribution of predicted probabilities assigned to each of the five possible actions. Boxplot outliers are omitted from the display.

D.5 Dynamic Model Fit

Figure D6. Dynamic Model Fit

(a) True Conditional Choice Probabilities (b) Fitted Conditional Choice Probabilities

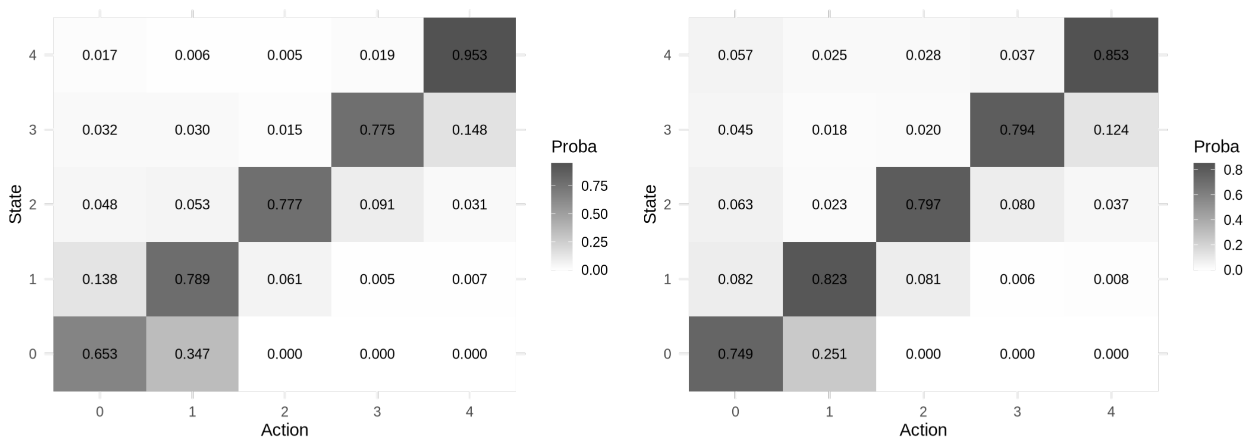

Notes: This figure shows the fit of the dynamic model. On the y-axis, state represents whether a preschool is outside ("0") or operating at one of the 4 possible STAR ratings. On the x-axis figures the action chosen among exit ("0") or one of the 4 STAR ratings.

E Counterfactual Supplements

E.1 Counterfactual Solution Procedure

I solve each counterfactual by combining the estimated stage-game and dynamic parameters with a policy-specific neural-network approximation of stage-game profits and a forward simulation of market structure. For computational reasons, I do not solve for the invariant state distribution. Instead, following Sweeting (2013), I initialize each simulation at the observed 2010 market structure and simulate its evolution through 2017. This approach is consistent with the perspective of a policymaker introducing a policy into an existing market. Let Ψ denote a policy environment, including the subsidy eligibility rule, copayment schedule, provider reimbursement schedule, and quality-tiered add-ons. For each Ψ, I take as given the estimated stage-game parameters $\hat{\boldsymbol{\theta}}$, the estimated dynamic parameters $\hat{\boldsymbol{\psi}}$, the first-stage CCP estimates $\hat{\mathcal{P}}_0$, and the initial market states $\{\mathcal{M}_{m,0}\}_{m=1}^{M}$.

1. Market-state augmentation and stage-game approximation. I first construct a collection of training states $\{\mathcal{M}_s^{\text{train}}\}_{s=1}^{S}$ by perturbing providers' operating status, quality ratings, demand and cost shocks, and local demographic conditions around their observed values. For each observed or perturbed state, I solve the stage-game Nash–Bertrand equilibrium under policy environment Ψ:

$$\left(p_\Psi^*(\mathcal{M}_s^{\text{train}}), q_\Psi^*(\mathcal{M}_s^{\text{train}}), \pi_\Psi^*(\mathcal{M}_s^{\text{train}})\right) = \text{StageGameEq}\left(\mathcal{M}_s^{\text{train}}, \Psi; \hat{\boldsymbol{\theta}}\right).$$

This produces a counterfactual-specific training sample that maps local market conditions into equilibrium prices, enrollment, and provider profits.

I use this sample to train a counterfactual-specific neural network described in Figure D2. Given market covariates $x(\mathcal{M}, \Psi)$, the network produces

$$x(\mathcal{M}, \Psi) \longmapsto (\hat{p}_\Psi(\mathcal{M}), \hat{q}_\Psi(\mathcal{M}), \hat{\pi}_\Psi(\mathcal{M}), \Phi(\mathcal{M})),$$

where $\hat{\pi}_\Psi(\mathcal{M})$ approximates stage-game profits and $\Phi(\mathcal{M})$ is the low-dimensional representation used to approximate the dynamic value function. In the quality-tiered reimbursement counterfactuals, I train a single network jointly across add-on multipliers and include the policy environment among its conditioning variables. Joint training both expands the training sample and supplies variation along the policy dimension.

2. Simulated next-period states. For each approximation state-action pair $(\mathcal{M}, a)$, I construct a fixed set of possible next-period states,

$$\mathcal{H}(\mathcal{M}, a) = \{\mathcal{M}'_h(\mathcal{M}, a)\}_{h=1}^{H}.$$

The focal provider's action is fixed at a, rivals' actions are drawn from the first-stage CCP model $\hat{\mathcal{P}}_0$, and demographic conditions evolve according to the estimated transition processes. I evaluate the neural network at each simulated next-period state to obtain

$$\{\hat{\pi}_\Psi(\mathcal{M}'_h), \Phi(\mathcal{M}'_h)\}_{h=1}^{H}.$$

The draws in $\mathcal{H}(\mathcal{M}, a)$ are held fixed while solving the counterfactual dynamic equilibrium. For any conjectured counterfactual CCP profile $\mathcal{P}$, normalized importance-sampling weights adjust the probability assigned to each draw:

$$\mathbb{E}_{\mathcal{P}}\left[\Phi(\mathcal{M}') \mid \mathcal{M}, a\right] \approx \sum_{h=1}^{H} \omega_h(\mathcal{M}, a; \mathcal{P}) \Phi\left(\mathcal{M}'_h(\mathcal{M}, a)\right), \qquad \sum_{h=1}^{H} \omega_h(\mathcal{M}, a; \mathcal{P}) = 1.$$

Thus, the support of possible next states remains fixed within a dynamic solution, while its probability distribution changes with the conjectured policy functions.

3. Dynamic equilibrium and forward simulation. For each counterfactual simulation $r = 1, \dots, R$, I initialize the markets at their observed states in 2010,

$$\mathcal{M}_{m,0}^{r,\Psi} = \mathcal{M}_{m,0}.$$

I then simulate the market forward for $t = 0, \dots, T-1$.

At the beginning of period t, I combine the current simulated market states with the pre-computed approximation states and initialize a policy function profile $\mathcal{P}_t^{(0)}$. Given iteration-ℓ CCPs $\mathcal{P}_t^{(\ell)}$, I update the importance-sampling weights and construct

$$\mathbb{E}_{\mathcal{P}_t^{(\ell)}}\left[\Phi(\mathcal{M}') \mid \mathcal{M}, a\right].$$

Conditional on these expectations, I solve the projected Bellman equation for the sieve coefficients

$$\rho_t^{(\ell)} = \rho\left(\hat{\boldsymbol{\psi}}; \mathcal{P}_t^{(\ell)}, \Psi\right).$$

The resulting action-specific continuation values are

$$FV_t^{(\ell)}(a, \mathcal{M}) = \beta \mathbb{E}_{\mathcal{P}_t^{(\ell)}} \left[\Phi(\mathcal{M}') \mid \mathcal{M}, a \right]' \rho_t^{(\ell)}.$$

Combining these continuation values with the action-specific flow payoffs under policy Ψ gives the undamped best-response probabilities

$$\mathcal{P}_{t,\mathrm{BR}}^{(\ell+1)}(a \mid \mathcal{M}) = \frac{\exp\left(\frac{\Pi_\Psi(a,\mathcal{M};\hat{\psi}) + FV_t^{(\ell)}(a,\mathcal{M})}{\hat{\psi}_0^\epsilon} \right)}{\sum_{a' \in \mathcal{A}(\mathcal{M})} \exp\left(\frac{\Pi_\Psi(a',\mathcal{M};\hat{\psi}) + FV_t^{(\ell)}(a',\mathcal{M})}{\hat{\psi}_0^\epsilon} \right)}.$$

I damp the CCP update according to

$$\mathcal{P}_t^{(\ell+1)} = \lambda \mathcal{P}_t^{(\ell)} + (1-\lambda) \mathcal{P}_{t,\mathrm{BR}}^{(\ell+1)},$$

and iterate until

$$\left\| \mathcal{P}_{t,\mathrm{BR}}^{(\ell+1)} - \mathcal{P}_t^{(\ell)} \right\|_\infty < \varepsilon_{\mathcal{P}}.$$

Denote the resulting equilibrium policy function by $\mathcal{P}_t^\Psi$.

I then draw entry, exit, and quality choices for each provider from

$$a_{j,m,t}^{r,\Psi} \sim \mathcal{P}_t^\Psi \left(\cdot \mid \mathcal{M}_{m,t}^{r,\Psi} \right).$$

Combining these actions with draws from the demographic and provider-shock transition processes yields the next-period market state:

$$\mathcal{M}_{m,t+1}^{r,\Psi} = \mathcal{G}\left(\mathcal{M}_{m,t}^{r,\Psi}, a_{m,t}^{r,\Psi}, \nu_{m,t+1}^r; \Psi \right).$$

For each resulting state, I reconstruct the covariates, evaluate $\hat{\pi}_\Psi(\mathcal{M}_{m,t+1}^{r,\Psi})$ and $\Phi(\mathcal{M}_{m,t+1}^{r,\Psi})$, and construct its simulated next-state support. These objects enter the dynamic equilibrium calculation in period $t+1$.

4. Exact equilibrium outcomes. The neural-network approximation is used to evaluate profits and continuation values inside the dynamic problem. It is not used to construct the reported counterfactual outcomes. For every realized market state along each simulated path, I re-solve the exact stage-game pricing equilibrium:

$$\left(p_{m,t}^{r,\Psi}, s_{m,t}^{r,\Psi}, e_{m,t}^{r,\Psi}, \pi_{m,t}^{r,\Psi} \right) = \mathrm{StageGameEq}\left(\mathcal{M}_{m,t}^{r,\Psi}, \Psi; \hat{\boldsymbol{\theta}} \right),$$

where $p^{r,\Psi}_{m,t}$, $s^{r,\Psi}_{m,t}$, and $e^{r,\Psi}_{m,t}$ denote equilibrium posted prices, market shares, and enrollment. Using these equilibria, I also compute markups, provider profits, subsidy expenditures, consumer surplus, and fiscal costs.

Finally, I aggregate market-level outcomes using the relevant population or provider weights and average across the R simulated paths. For an additive outcome Y, this aggregation is

$$\overline{Y}^{\Psi}_t = \frac{1}{R}\sum_{r=1}^{R}\sum_{m=1}^{M} Y\left(\mathcal{M}^{r,\Psi}_{m,t}, \Psi\right).$$

Differences across policy environments therefore reflect both their direct effects on the stage-game pricing and enrollment equilibrium and their indirect effects on the evolution of entry, exit, quality, and local market structure.

E.2 Policy Evaluations

E.2.1 Demand Expansion: Policy Details

Figure E7 below shows consumer-facing prices under alternative income eligibility thresholds, holding the provider's posted price fixed at \$50 per day. I keep quality add-ons to their levels in the data in 2018 as shown in Figure A1. Broader eligibility lowers prices for newly eligible households across all ratings, while the quality add-ons produce larger reductions at higher-rated centers. These differences arise between the alternative eligibility thresholds and disappear below \$24,000, where all four policies provide subsidies, and above \$120,000, where none do.

E.2.2 Quality-Tiered Reimbursement Expansion: Policy Details

Figure E8 below shows consumer-facing prices under alternative add-on multipliers, holding the provider's posted price fixed at \$50 per day. I keep the feature in place in Pennsylvania, 2018, that add-ons apply only to STAR 3 and 4 providers. As such, their generosity does not affect prices at STAR 1 and 2 centers, whereas larger multipliers reduce prices for eligible households attending higher-rated centers. These differences are largest at low incomes, narrow as copayments rise, and disappear above the \$48,000 eligibility threshold.

Figure E7. Consumer Prices by Household Income Eligibility Thresholds

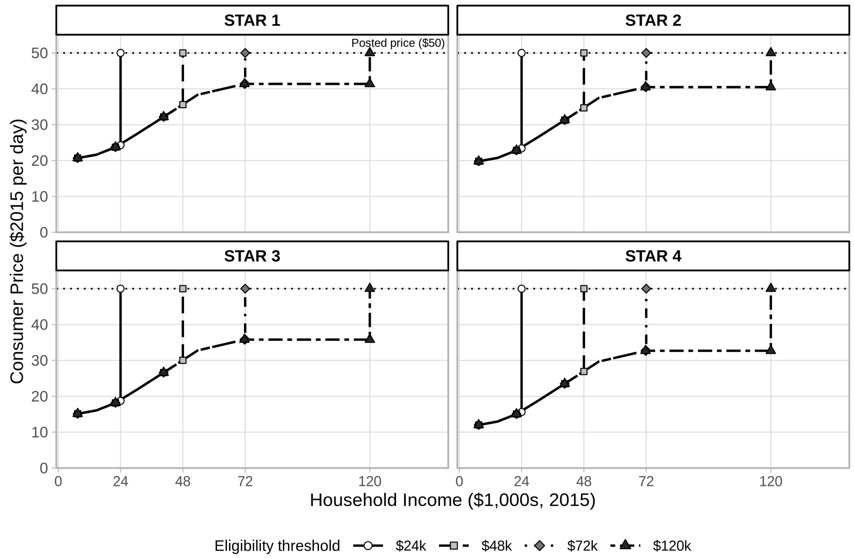

Notes: The figure plots the consumer price for a provider with a fixed posted price of $50. Subsidy-eligible households pay the policy copay plus any top-up when the posted price exceeds the copay, base reimbursement, and STAR-specific add-on. Households above the eligibility threshold pay the full posted price. Vertical lines indicate the income eligibility thresholds under each policy.

Figure E8. Consumer Prices by Add-on Multiplier

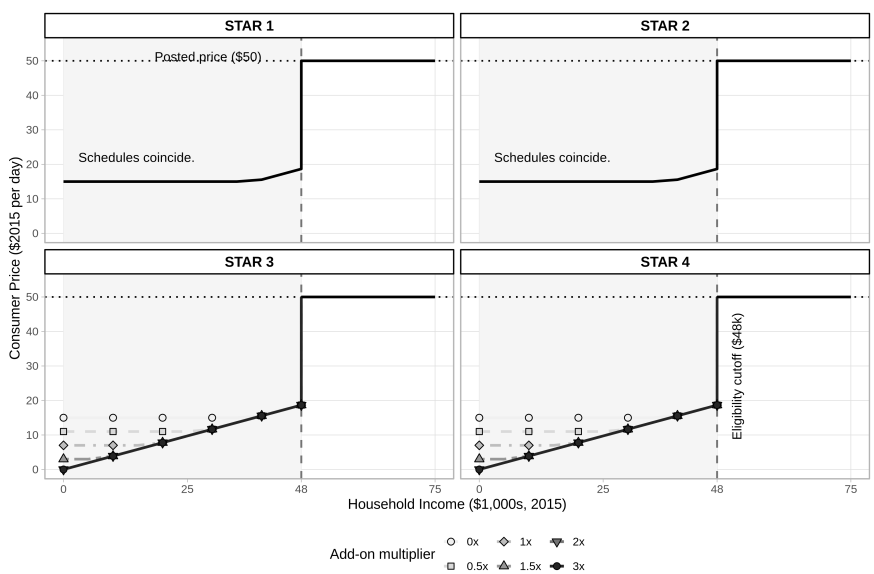

Notes: The figure plots consumer prices for a provider charging a fixed posted price of $50. Subsidy-eligible households pay the scheduled copay plus any top-up when the posted price exceeds the copay, base reimbursement, and STAR-specific add-on. Ineligible households pay the full posted price.

E.2.3 Educational-Benefit Marginal Value of Public Funds Calculations

Policy Comparison Framework. I focus on the potential learning benefits of attending high-quality preschool. These benefits are well documented in the policy-evaluation literature, especially for low-income children, and are a central motivation for preschool subsidies (see Elango et al. (2015) for a review). Focusing on educational-benefit MVPF calculations also facilitates comparisons with previous calculations in the literature, including Hendren and Sprung-Keyser (2020); Kline and Walters (2016); Cascio (2023). As noted by Humphries et al. (2024), effects on children's future adult earnings should be included in MVPF calculations only if parents do not fully internalize those benefits when making early-education arrangement choices. This assumption is supported by prior studies on early education showing that parents exhibit limited responsiveness to structural measures of quality (Council of Economic Advisers, 1997), and more broadly by the literature using revealed preferences in education contexts and documenting limited demand for value-added (Abdulkadiroğlu et al., 2020).[27]

I first estimate learning gains. The treatment effect of interest is the net present value (NPV) of the lifetime earnings gains from attending a preschool of quality q for a child from socioeconomic background y, relative to the outside option of informal early education. I denote this effect by $\tau_{ps}^{NPV}(y,q)$. I decompose these NPV gains into components that can be calibrated from studies relying on quasi-experimental variation: the short-term effect of attending preschool on kindergarten test scores, $\tau_{ps}^{KScore}(y,q)$, the effect of higher kindergarten test scores on adult earnings, τ_{KScore}^{NPV}, and the baseline net present value of adult earnings by income group, $NPV(y)$. Therefore,

$$\tau_{ps}^{NPV}(y,q) = \tau_{ps}^{KScore}(y,q) \times \tau_{KScore}^{NPV} \times NPV(y).$$

Given these effects of preschool attendance on children's future earnings, and assuming a common tax rate κ on income, the educational benefits of a policy scenario Ψ in which the equilibrium implies preschool enrollment $\{N_{yq}\}$ are given by

$$B^c(\Psi) = (1-\kappa)\sum_y \sum_q N_{yq}(\Psi)\tau_{ps}^{NPV}(y,q).$$

[27] Additional benefits may stem from reductions in crime (Elango et al., 2015), spillovers to other children (Williams, 2019), or increased household income driven by labor-supply responses (Humphries et al., 2024). I ignore these benefits in the main analysis, but note that regardless of the outcome considered, providers' equilibrium responses will impact take-up via entry and prices, and thus influence the overall effects of policies that target these alternative outcomes.

In the policy comparisons, I always compute benefits relative to a scenario Ψ_0 without any early-education subsidy. The numerator of the cost-benefit analysis is therefore

$$B^c(\Psi) - B^c(\Psi_0) = (1-\kappa)\sum_y \sum_q (N_{yq}(\Psi) - N_{yq}(\Psi_0))\tau_{ps}^{NPV}(y,q).$$

The denominator is given by the sum of demand- and supply-side subsidies minus fiscal benefits:

$$\begin{aligned} C(\Psi) &= sub^{\mathcal{S}}(\Psi) + sub^{\mathcal{D}}(\Psi) - \kappa \sum_y \sum_q (N_{yq}(\Psi) - N_{yq}(\Psi_0))\tau_{ps}^{NPV}(y,q) \\ &= sub^{\mathcal{S}}(\Psi) + \sum_y \sum_q N_{yq}(\Psi)\overline{sub}_{yq}(\Psi) - \kappa \sum_y \sum_q (N_{yq}(\Psi) - N_{yq}(\Psi_0))\tau_{ps}^{NPV}(y,q). \end{aligned}$$

Here, $sub^{\mathcal{S}}(\Psi)$ denotes potential fixed supply-side subsidies that are not per child (for instance a start-up grant) and $\overline{sub}_{yq}(\Psi)$ is the average government cost per child of income group y enrolled in quality tier q. This cost depends on equilibrium provider prices and realized enrollment. By definition, Ψ_0 satisfies $C(\Psi_0) = 0$ because no subsidies are in place in the baseline scenario.

Calibration. First, to form expected adult earnings by household income y, $NPV(y)$, I use the Opportunity Insights (OI) mobility data (Chetty et al., 2026). Specifically, I convert parent-income percentiles into dollars, match them to Pennsylvania child-income outcomes from the OI county mobility file, and then scale earnings over the life cycle using an age-earnings profile to construct the discounted NPV of adult earnings by household-income background at age 3.5 years.

In the MVPF calculation, this baseline $NPV(y)$ is multiplied by the effect of kindergarten test scores on adult earnings, τ_{KScore}^{NPV}, which I set to 10% based on Chetty et al. (2011). Finally, I calibrate the effect of preschool on kindergarten test scores, $\tau_{ps}^{KScore}(y,q)$, as follows. Kline and Walters (2016) report a local average treatment effect of $LATE_{nh} = 0.37$ for children substituting from an informal arrangement, the outside option in this setting, to Head Start. I use this effect to calibrate $\tau_{ps}^{KScore}(y \leq FPL, q = H)$, where $H \equiv \{\text{STAR 3\&4}\}$, because Head Start provides formal preschool to low-income children and its estimated annual cost is close to the cost estimates for highly rated providers.

An increasing number of studies in both universal and targeted contexts report plausibly causal effects of preschool attendance on kindergarten test scores. However, this evidence comes from multiple contexts, making it difficult to construct a fully satisfactory profile for

Figure E9. Assumptions on $\tau_{ps}^{KScore}(y, q)$ and Resulting $\tau_{ps}^{NPV}(y, q)$

(a) Test Scores (b) NPV

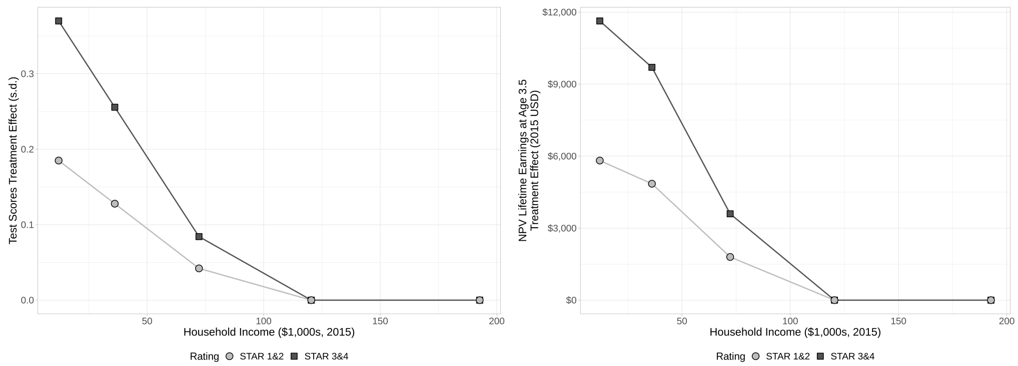

Notes: Panel A shows calibrated test-score effects of attending preschool by household income and preschool quality.

$\tau_{ps}^{KScore}(y, q)$. In the main analysis, I therefore adopt assumptions that reflect two takeaways from the existing evidence: first, learning gains from attending formal preschool relative to the outside option tend to be larger for children from lower-income families (e.g., Elango et al. (2015); Havnes and Mogstad (2015); Cornelissen et al. (2018); Blau (2021)); second, learning gains are more pronounced when the quality of the learning environment is higher (see, e.g., Cascio (2015)). To capture these features in a stylized way, I extrapolate the Kline and Walters (2016) estimate along the income dimension by assuming that (1) $\tau_{ps}^{KScore}(., H)$ decreases linearly with income before plateauing at 0, such that $\tau_{ps}^{KScore}(y \geq \bar{y}, H) = 0$ with $\bar{y}$ = \$100,000, and (2) $\tau_{ps}^{KScore}(y, L) = 0.5\tau_{ps}^{KScore}(y, H) \forall y$. These assumptions could, for instance, be consistent with a model in which cognitive-skill production is determined by the highest-quality arrangement a child attends, rather than the average arrangement, and the quality of parental or available informal care increases with income. Figure E9 represents these assumptions. They rule out negative effects of preschool attendance at any income level, which have been found in some evaluation studies (Havnes and Mogstad, 2015). However, they allow policy comparisons to focus entirely on learning benefits arising from low- and middle-income families' substitution from informal to formal care and across quality tiers, which better summarizes the objective function of policymakers.

E.3 Additional Results

E.3.1 Markup Effects of Introducing Demand Subsidies

Figure E10 describes providers' equilibrium markup responses when demand subsidies are introduced and centers re-optimize posted prices. For each provider j, the markup response is

$$100 \times \frac{\eta_j^{\mathrm{Sub}} - \eta_j^{\mathrm{NoSub}}}{\eta_j^{\mathrm{NoSub}}}.$$

Providers operating in 2010 are ordered from the largest markup reduction to the largest markup increase and divided into 50 equal-count bins, each containing approximately 2% of providers. Markup responses to subsidies are highly heterogeneous, with the largest positive and negative responses concentrated in low-income neighborhoods, where most subsidized families reside. Centers that lower their markups tend to be high-marginal-cost centers competing for subsidized, price-sensitive families. A subset of low-cost centers operating in low-income neighborhoods instead respond to the subsidy-induced demand shift by increasing their markups.

Figure E10. Markup Responses to Demand Subsidies

(a) Neighborhood Income (b) Marginal Cost

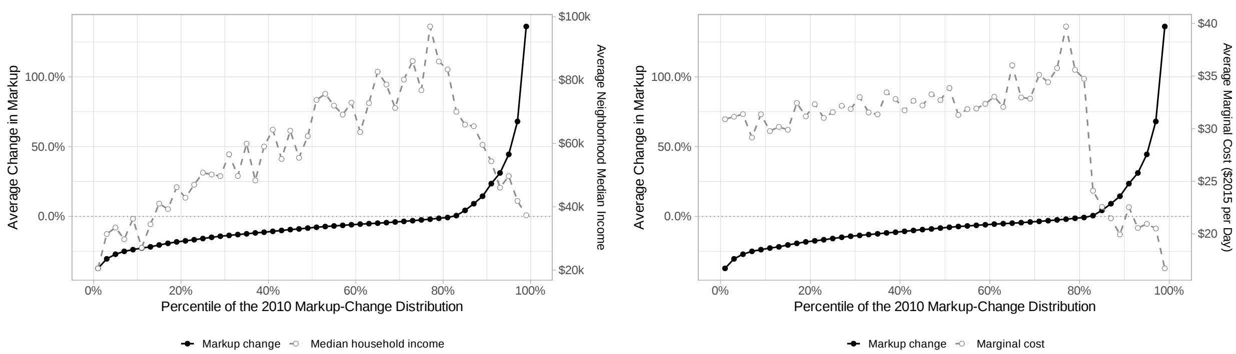

Notes: Panel (a) presents the distribution of markup responses by center block-group income, and panel (b) by center marginal cost. The sample includes all markets in 2010. Income and daily marginal cost are measured in 2015 dollars.

E.3.2 Distributional Effects Across Space

Figure E11 shows the distribution of income in Pittsburgh East, the market used to illustrate the spatial access changes resulting from quality-neutral and quality-targeted subsidies. While Figure 5 in the main text focuses on access to high-quality centers, Figure E12 below displays changes in overall preschool access.

Figure E11. Household Income in Pittsburgh East

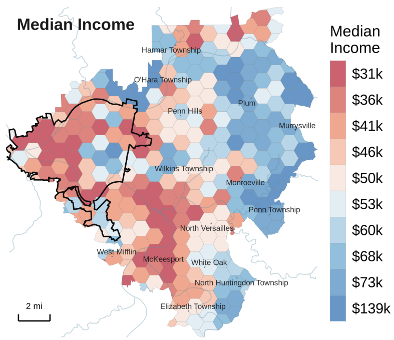

Notes: The figure plots average median income across a hexagonal grid in the Pittsburgh East market.

Figure E12. Subsidies and Access to Centers

(a) Quality-Neutral (b) Quality-Targeted (c) Correlation with Income

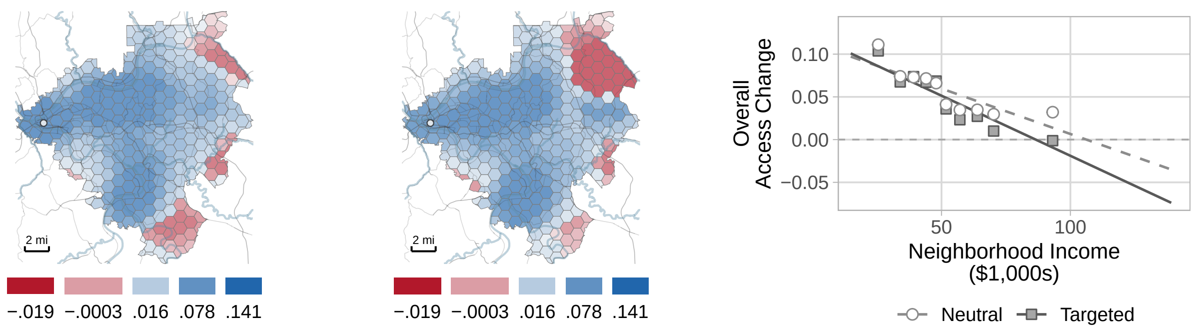

Notes: Panels A and B report differences in access to centers under quality-neutral subsidies (demand subsidies without an add-on) and quality-targeted subsidies (demand subsidies with an add-on), relative to a baseline scenario without subsidies, in 2017. Panel C reports the relationship between differences in access and neighborhood income. Access is computed using the two-step floating catchment area approach applied to the early-education sector in Davis et al. (2019).

E.3.3 Robustness of Educational-Benefit MVPF Calculations

Figure E13. Educational-Benefit MVPF of Pennsylvania Policies at Various Test-Score Effects

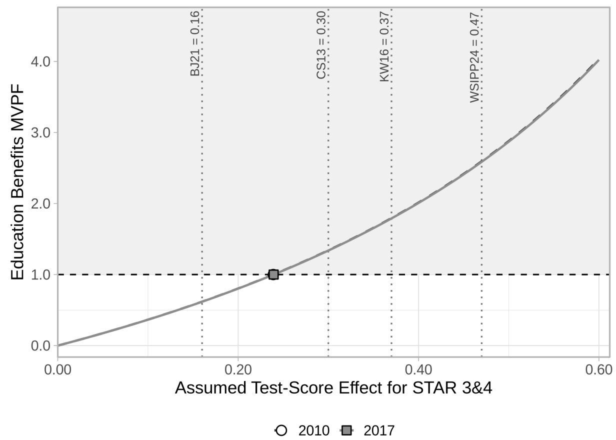

Notes: The figure plots the educational-benefit MVPF of the $48,000 eligibility policy relative to no subsidies as the assumed STAR 3&4 test-score effect varies. STAR 1&2 also adjust to half of the STAR 3&4 test-score effect, and decay with parental income at the same rate as assumed in section E.2.3. Vertical lines indicate test score estimates from BJ21: Barnett and Jung (2021), CS13: Cascio and Schanzenbach (2013), KW16: Kline and Walters (2016), and WSIPP24: Washington State Institute for Public Policy (2024).

F Notation

Table F1. Summary of Notation

Symbol	Description
Indices, types, and sets	
m, t, ℓ	Geographic market, School year, Neighborhood.
j, k	Preschool providers.
i	Family or simulated consumer type.
$\mathcal{L}_m$	Set of neighborhoods contained in market m.
$\mathcal{J}_{mt}$	Set of preschool providers operating in market m in year t.
Policy variables	
$\bar{y}_t$	Household-income threshold for subsidy eligibility.
$copay_t(y)$	Family copayment schedule for subsidy-eligible households.
$MCCA_{m(j)t}$	Maximum Child Care Allowance applying to provider j in year t.
$addon_t(r_{jt})$	Quality-tiered reimbursement for rating r_{jt}.
sub, $topup$	Subsidy to provider, top-up paid by subsidized family.
Stage game: families, prices, and provider profits	
$N_{\ell t},\ N_{mt}$	Number of preschoolers in neighborhood ℓ and market m.
$e_{it} \in \{H, L\}$	Mother's educational attainment: (non-) college educated H (L).
y_{it}	Household income.
U_{ijt}	Utility family i obtains from choosing provider j in year t, net of taste shocks.
s_{ijt}, s_{jt}	Choice probability of consumer i and market share of provider j.
$r_{jt} \in \{1, 2, 3, 4\}$	Keystone STARS quality rating of provider j in year t.
p_{jt}	Posted daily price charged by provider j.
p^c_{ijt}, p^p_{ijt}	Consumer and provider price.
η_{jt}	Provider j's markup over marginal cost, so that $p_{jt} = mc_{jt} + \eta_{jt}$.
π_{jt}	Variable profits generated by the stage-game pricing equilibrium.
mc_{jt}	Daily marginal cost of provider j in year t.
α_i	Consumer type i's price sensitivity.
$\xi_j,\ \xi_t,\ \Delta\xi_{jt}$	Provider fixed effect, time effect, and transitory demand shock.
$\omega_j,\ \omega_t,\ \Delta\omega_{jt}$	Provider fixed effect, time effect, and transitory marginal-cost shock.
$\boldsymbol{\theta}$	Stage-game demand and marginal-cost parameters.
Dynamic provider problem	
$\mathcal{M}_{xjt}$, $\mathcal{M}$	Public payoff-relevant state for provider j with own state x
$\mathcal{H}(\mathcal{M}_{xjt})$	set of simulated potential next states from $\mathcal{M}_{xjt}$
$\boldsymbol{\psi}$	Dynamic parameters: transition costs $\boldsymbol{\psi}^W$, fixed-cost $\boldsymbol{\psi}^C$, scale ψ^{ϵ}.
$W(\mathcal{M}, a; \boldsymbol{\psi})$	Entry or quality-transition cost associated with action a.
$C(\mathcal{M}, a; \boldsymbol{\psi})$	Fixed operating-cost component associated with action a.
$\mathcal{P}$	Profile of provider strategies or conditional choice probabilities.
$\Pi(a, \mathcal{M}; \boldsymbol{\psi})$	Action-specific flow payoff after transition and fixed costs.
$\bar{\Pi}(\mathcal{M}, \mathcal{P})$	Average flow payoff in state $\mathcal{M}$ under strategy profile $\mathcal{P}$
$V(\mathcal{M})$, $FV(a, \mathcal{M})$	Ex-ante value function and action-specific continuation value.
$\Phi(\mathcal{M})$	Basis vector used to approximate the value function.
$\rho(\boldsymbol{\psi}; \mathcal{P})$	Coefficients on the value-function basis.

Appendix References

Abdulkadiroğlu, A., P. A. Pathak, J. Schellenberg, and C. R. Walters (2020): "Do parents value school effectiveness?" *American Economic Review*, 110, 1502–1539.

Anderson, D. G. (1965): "Iterative procedures for nonlinear integral equations," *Journal of the ACM (JACM)*, 12, 547–560.

Barnett, W. S. and K. Jung (2021): "Effects of New Jersey's Abbott preschool program on children's achievement, grade retention, and special education through tenth grade," *Early Childhood Research Quarterly*, 56, 248–259.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890.

Blau, D. M. (2001): *Child Care Problem: An Economic Analysis*, Russell Sage Foundation.

——— (2021): "The effects of universal preschool on child and adult outcomes: A review of recent evidence from Europe with implications for the United States," *Early Childhood Research Quarterly*, 55, 52–63.

Brown, J. H. (2018): "Does public pre-k have unintended consequences on the child care market for infants and toddlers?" Unpublished manuscript.

Caoui, E. H., B. Hollenbeck, M. Osborne, and E. Page (2026): "The impact of dollar store expansion on local market structure and food access," *Quantitative Marketing and Economics*, 24, 7.

Cascio, E. U. (2015): "The promises and pitfalls of universal early education," *IZA World of Labor*.

——— (2023): "Does universal preschool hit the target?: Program access and preschool impacts," *Journal of Human Resources*, 58, 1–42.

Cascio, E. U. and D. W. Schanzenbach (2013): "The impacts of expanding access to high-quality preschool education," *Brookings Papers on Economic Activity*, 2013, 127–192.

Chetty, R., J. N. Friedman, N. Hendren, M. R. Jones, and S. R. Porter (2026): "The opportunity atlas: Mapping the childhood roots of social mobility," *American Economic Review*, 116, 1–51.

Chetty, R., J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan (2011): "How does your kindergarten classroom affect your earnings? Evidence from Project STAR," *The Quarterly journal of economics*, 126, 1593–1660.

Conlon, C. and J. Gortmaker (2020): "Best practices for differentiated products demand estimation with pyblp," *The RAND Journal of Economics*, 51, 1108–1161.

Conlon, C. T. and J. H. Mortimer (2021): "Empirical Properties of Diversion Ratios," *The RAND Journal of Economics*, 52, 693–726.

Cornelissen, T., C. Dustmann, A. Raute, and U. Schönberg (2018): "Who benefits from universal child care? Estimating marginal returns to early child care attendance," *Journal of Political Economy*, 126, 2356–2409.

Council of Economic Advisers (1997): "The Economics of Child Care," White House, Executive Office of the President, accessed May 21, 2026.

Davis, E. E., W. F. Lee, and A. Sojourner (2019): "Family-centered measures of access to early care and education," *Early Childhood Research Quarterly*, 47, 472–486.

Elango, S., J. L. García, J. J. Heckman, and A. Hojman (2015): "Early childhood education," in *Economics of Means-Tested Transfer Programs in the United States, Volume 2*, University of Chicago Press, 235–297.

Goodfellow, I., Y. Bengio, A. Courville, and Y. Bengio (2016): *Deep learning*, vol. 1, MIT press Cambridge.

Havnes, T. and M. Mogstad (2015): "Is universal child care leveling the playing field?" *Journal of public economics*, 127, 100–114.

Hendren, N. and B. Sprung-Keyser (2020): "A unified welfare analysis of government policies," *The Quarterly Journal of Economics*, 135, 1209–1318.

Herbst, C. M. (2023): "Child care in the United States: Markets, policy, and evidence," *Journal of Policy Analysis and Management*, 42, 255–304.

Hinton, G. E. and R. R. Salakhutdinov (2006): "Reducing the dimensionality of data with neural networks," *Science*, 313, 504–507.

Ho, K. (2006): "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market," *Journal of Applied Econometrics*, 21, 1039–1079.

Humphries, J. E., C. Neilson, X. Ye, and S. D. Zimmerman (2024): "Parents' Earnings and the Returns to Universal Pre-Kindergarten," Unpublished manuscript.

Ifrach, B. and G. Y. Weintraub (2017): "A framework for dynamic oligopoly in concentrated industries," *The Review of Economic Studies*, 84, 1106–1150.

Kline, P. and C. R. Walters (2016): "Evaluating public programs with close substitutes: The case of Head Start," *The Quarterly Journal of Economics*, 131, 1795–1848.

Le, L., A. Patterson, and M. White (2018): "Supervised autoencoders: Improving generalization performance with unsupervised regularizers," *Advances in neural information processing systems*, 31.

Lewis, G., B. Ozaltun, and G. Zervas (2021): "Maximum Likelihood Estimation of Differentiated Products Demand Systems," Unpublished manuscript.

Morrow, W. R. and S. J. Skerlos (2011): "Fixed-point approaches to computing Bertrand-Nash equilibrium prices under mixed-logit demand," *Operations research*, 59, 328–345.

Neilson, C. (2013): "Targeted vouchers, competition among schools, and the academic achievement of poor students," Unpublished manuscript.

Pennsylvania Office of Child Development and Early Learning (2024a): "ELRC 1 Authorization Agreement for Child Care Works Subsidized Child Care Provider Agreement," https://nwirelrc.org/wp-content/uploads/2025/03/Combined-Regulated-Provider-Form.pdf, combined Regulated Provider Form, including Appendices A and F.

——— (2024b): "Pennsylvania FFY 2025–2027 Child Care and Development Fund State Plan," https://www.pa.gov/content/dam/copapwp-pagov/en/dhs/documents/services/children/documents/2025-2027-ocdel-ccdf-state-plan.pdf, section 4.3.3.

Singleton, J. D. (2019): "Incentives and the supply of effective charter schools," *American Economic Review*, 109, 2568–2612.

Sweeting, A. (2013): "Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry," *Econometrica*, 81, 1763–1803.

Washington State Institute for Public Policy (2024): "State Early Childhood Education Programs: Universal," https://www.wsipp.wa.gov/BenefitCost/Program/975, benefit-cost methods updated December 2024; literature review updated July 2019; accessed July 21, 2026.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008): "Markov perfect industry dynamics with many firms," *Econometrica*, 76, 1375–1411.

Williams, B. J. (2019): "The spillover benefits of expanding access to preschool," *Economics of Education Review*, 70, 127–143.