Not Playing Favorites: An Experiment on Parental Fairness Preferences

James Berry, Rebecca Dizon-Ross, and Maulik Jagnani
FEBRUARY 2020

Becker Friedman Institute

FOR ECONOMICS AT **UCHICAGO**

Not Playing Favorites: An Experiment on Parental Fairness Preferences
James Berry, Rebecca Dizon-Ross, and Maulik Jagnani
February 2020
JEL No. I20,J13

ABSTRACT

We conduct a lab-in-the-field experiment to identify parents' preferences for investing in their children. The experiment exogenously varied the short-run returns to educational investments to identify how much parents care about maximizing total household earnings, minimizing cross-sibling inequality in "outcomes" (child-level earnings), and minimizing cross-sibling inequality in "inputs" (child-level investments). We show that while parents place some weight on maximizing earnings, they also display a strong preference for equality in inputs, forgoing roughly 40% of their potential earnings or 90% of a day's wage to equalize inputs. We find no evidence that parents care about equalizing outcomes.

James Berry
Department of Economics
University of Delaware
jimberry@udel.edu

Rebecca Dizon-Ross
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
rdr@chicagobooth.edu

Maulik Jagnani
Economic Growth Center
Department of Economics
Yale University
27 Hillhouse Avenue
New Haven, CT 06511
maulik.jagnani@yale.edu

1   Introduction

What are parents' preferences for allocating resources across their children? Although a large and influential literature has documented the profound importance of parental investments for children's outcomes (e.g., Cunha and Heckman, 2007; Cunha et al., 2006), our current understanding of parents' preferences is limited. On the one hand, parents could treat investments in their children as standard investment goods, investing to maximize returns but potentially generating inequality across siblings. On the other hand, parents may be averse to cross-sibling inequality, an aversion which could take multiple forms. They could be averse to inequality in their children's *outcomes* (the amount the children ultimately earn). Alternatively, they could be averse to inequality in *inputs* (the amount spent on each child's schooling), an idea absent from the literature on parents' preferences for investment (e.g., Behrman et al., 1982) but with ties to the fairness and bequests literatures (e.g., Andreoni and Bernheim, 2009).[1] Each of these preferences have different predictions for behavior and ultimately, for policy.

Estimating parental preferences is challenging for three main reasons. First, to estimate preferences from observed behavior, one needs to know the full perceived production function mapping inputs to outcomes. For example, if a parent invests more in her high-ability child than her low-ability child, she could be choosing what she thinks is the returns-maximizing action, or she could be balancing a preference for maximizing returns with a preference for, say, equalizing inputs. Second, it is hard to distinguish a preference for equality from risk aversion. If parents choose to invest more equally across their children than returns maximization without uncertainty would predict, it could reflect a preference for equal inputs, or it could reflect risk-averse parents hedging their investments in the face of uncertainty. The final challenge is identification: one needs multiple exogenous shocks to investment returns to separately identify the weights parents put on different types of preferences.

To overcome these challenges, we designed a lab-in-the-field experiment to identify parents' preferences. Our experiment sampled parents in rural Malawi with two children enrolled

---

[1]Consistent with this idea, raw educational expenditure data from our experimental sample suggest that many parents spend exactly equal amounts on their different children (Figure A.1). American Time Use Survey data also show that many parents spend roughly equal time with each of their children (Price, 2008).

in grades 5 to 7. We first asked each child to take a test. Based on his or her test score, each child received a direct monetary payment. The child's monetary payment is our measure of his or her *outcome.* (In the broader literature on parents' investments, the child's future income is typically thought of as the outcome.) The experiment varied the child-specific payment functions mapping test scores to payments, providing exogenous variation that we use for identification.

Before the test, parents received an *input* to divide between their children: tutoring before the test. Each parent received 10 lottery tickets that she could allocate across her children. One of these 10 tickets was randomly chosen, and the child whose ticket was selected received one hour of tutoring. Therefore, a parent could choose which of her children would receive tutoring by allocating all of her tickets to that child. To measure the perceived production function, we elicited parents' beliefs about each of their children's test scores without tutoring, and their expectations of how much each child's test score would increase if he or she received tutoring.

This setup yields a clean prediction about parents' behavior: a returns-maximizing parent will give all of her tickets to the child she thinks provides greater returns to the investment, and none to the other. She should *only* deviate from an "all-or-nothing" allocation if she is averse to inequality in the (expected) inputs she gives her children or to inequality in her children's expected outcomes. Importantly, as we discuss later, this sharp prediction holds even if there is uncertainty and parents are risk-averse, thus allowing us to sidestep production uncertainty as a confound for identification.

To identify preferences, we varied the child-specific payment functions in ways that have qualitatively different predictions depending on parents' preferences for returns maximization, equality in inputs, and equality in outcomes. For example, certain payment functions provided the same payment per test score point to both children, whereas others delivered a higher payment to the perceived higher-performer.[2] If parents are returns-maximizers, increasing the payment-per-test-score-point for the perceived higher-performing sibling would cause parents to give that child more inputs (lottery tickets) because his or her expected

---

[2]We used the "strategy method" and elicited parents' choices under every pair of child-specific payment functions used in the experiment before randomly selecting one pair of payment functions to implement for each parent.

payment gains from tutoring have increased. In contrast, if parents are averse to inequality in outcomes, they would do the opposite: since the perceived lower-performing sibling's expected payments have decreased relative to his or her sibling's, parents who want to equalize their children's expected payments would reallocate inputs to the perceived lower-performer to help her catch up to her sibling.

Our headline result is that parents display a quantitatively important preference for equality in inputs, which causes them to leave substantial payments on the table. We first test and reject the hypothesis that parents care only about maximizing returns. Only 45% of allocations were "all-or-nothing," assigning all tickets to one child and none to the other, which is the prediction for returns-maximizers in our lottery environment. We next establish that parents are averse to input inequality in particular, and that this preference represents the primary reason they deviated from returns maximization. Parents chose to equalize inputs in roughly 35% of their choices. Even when we substantially increased the expected payment gains from maximizing returns—offering 10 times higher returns per point to one child or the other—at least 30% of parents still equalized. This preference for equal inputs meaningfully decreases earnings: the average parent earned roughly 40% less than the maximum possible amount.[3] Using a structural model to identify parents' preferences, we find that parents are willing to give up an average of 1,300 MWK—90% of an adult's daily wage—in order to invest equally in our experiment.[4]

Our experimental design also allows us to estimate the preference weights parents place on returns maximization and inequality aversion over outcomes. We find that parents place some weight on maximizing returns, as they consistently invest more in a child when the returns to tutoring for that child increase. By contrast, we fail to find any evidence that parents care about equalizing outcomes. Our most direct evidence against inequality aversion over outcomes comes from payment functions that exogenously delivered a lump sum payment to one child only. If parents are inequality averse over outcomes, they should respond to the lump sum by giving more tickets to the child's sibling to help close the gap in

---

[3]These statistics represent forgone expected earnings as a percent of the expected earnings parents controlled; in all scenarios, children received base expected payments that were inframarginal to parents' allocations.

[4]At the time of the study, the exchange rate was 715 MWK : 1 USD.

outcomes. However, we find no evidence that parents respond in this way.[5]

We perform several supplementary analyses to rule out potential confounds to interpretation. We show input equalization in our experiment does not reflect a lack of understanding of how to maximize returns. We also present evidence that parents did not diverge from "all-or-nothing" allocations due to uncertainty about children's performance or indifference about which child should receive the tutoring. To mitigate concerns about demand effects (i.e., that parents did not want to equalize but did it because they thought surveyors wanted them to), we used high monetary stakes, enabled by the fact that we conducted our experiment in a low-income setting. Finally, we ensured children in our experiment were not aware their parents played a role in tutoring allocation. This allays concerns that input inequality might be more visible to children in our experiment than elsewhere.

Our findings have important policy implications. Many policies, such as gifted and talented programs and remedial education programs, may only target one sibling within a given household. The programs can thus have spillover effects, either positive or negative, on non-targeted siblings. Our results suggest that policymakers may be able to promote positive spillovers by leveraging parents' aversion to inequality in inputs. For example, if a gifted and talented program often only affects one child per household, providing parents with information about the high level of *inputs* provided to one of their children through the program may encourage parents to spend more on their other children to mitigate the inequality. Our findings thus suggest policy approaches that could be rigorously tested in the future.

This paper contributes to two main bodies of work in economics: the literature on parental investment in children and the literature on preferences for fairness. Within the research on parental investments, our primary contribution is to a classic economics literature that characterizes parents' preferences for investing in their children and the balance between returns-maximization and inequality aversion (e.g., Behrman et al., 1982, 1986; Pitt et al., 1990). These papers rely primarily on functional form assumptions for identification; we

---

[5]We follow the literature (e.g., Behrman et al. (1982)) and define inequality aversion over outcomes as inequality aversion over *income*, not consumption. In contrast, if parents are inequality averse over consumption, they could maximize returns at the investment stage and redistribute income after the experiment. However, we show that most parents do not reallocate earnings *ex post*.

contribute by using exogenous variation to uncover parents' preferences. In addition, the economics literature has always tested for aversion to inequality in *outcomes*; our paper is the first empirical economics paper to test for an aversion to inequality in inputs, which we show is the dominant preference in our setting.[6] In contrast, other disciplines often conceptualize parents' preferences for fairness as preferences for equal treatment or equality in inputs (see Trivers, 1974 in biology and Hertwig et al., 2002 in psychology). We build on this work by devising a clean test for inequality aversion over inputs and quantifying its role relative to other parental concerns.

Our work relates to two other strands within the literature on parental investment. First, several recent papers examine how parents' investments depend on their children's baseline endowments, and whether parents prefer to reinforce these endowments by investing more in their higher-endowment children, or to compensate by investing less. The findings are mixed, with responses ranging from reinforcing, to zero, to compensatory (see Almond and Mazumder, 2013, for a review). These papers, however, generally do not identify parents' preferences themselves, as investments reflect the interaction between preferences and the (unobserved) perceived production function. It is important to note that our result that parents are inequality averse over inputs is not inconsistent with studies finding that parents spend differently on their children. Indeed, we also find differential spending across children, but that parents balance their desire to spend unequally to maximize returns with their desire to equalize inputs, causing their choices to be *more* equal, not necessarily fully equal.

Second, we relate to papers showing that parents often equally divide bequests among their children.[7] However, bequests differ from investments in several important ways. Unlike investments, equally dividing bequests may have no efficiency costs. Investment-based concepts of inputs and outcomes inequality also do not carry over to the bequests domain.

We also contribute to the extensive literature that has examined fairness preferences and has shown that people are averse to inequalities.[8] The vast majority of this literature

---

[6]From a theoretical perspective, Farmer and Tiefenthaler (1995) outlines different potential concepts of fairness in intrahousehold allocation of resources, with explicit consideration for parents' preferences for equality in both outcomes and inputs.

[7]See, e.g., Bernheim and Severinov (2003); Menchik (1980); Wilhelm (1996).

[8]Key contributions include Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Fehr and Schmidt, 1999; Forsythe et al., 1994.

examines preferences over the distribution of final payoffs. In contrast, our study also examines preferences over the distribution of inputs to production, which generate a different type of inefficiency. While preferences for redistributing final payoffs may generate inefficiencies through moral hazard (since individuals do not keep the full fruits of their labor), preferences for equality at the production stage instead generate inefficiencies by directly causing misallocation in inputs.

Within the fairness literature, we tie most closely to three streams. First, we relate to recent papers that examine whether, in risky situations, people care about equalizing payoffs from an *ex ante* perspective (i.e., equalizing expected payoffs before the risk is realized) or from an *ex post* perspective (i.e., after risk is realized). These papers find evidence for both *ex ante* and *ex post* preferences, with *ex ante* preferences more prevalent upfront (Andreoni et al., 2018; Brock et al., 2013; Cappelen et al., 2013; Krawczyk and Le Lec, 2010); for example, the modal experimental participant in Andreoni et al. (2018) chose an *ex ante* fair allocation *ex ante* and an *ex post* fair allocation *ex post*. In contrast, we focus on *ex ante* preferences only and examine what type of *ex ante* fairness people care about, distinguishing between notions of inputs and outcomes equality.[9] The prior work also tests for *ex ante* fairness concerns in lab settings; we extend this work to a high-stakes field setting. Second, our work relates to papers documenting the widespread use of a "50-50" norm for dividing monetary rewards and costs (Agrawal, 2002; Andreoni and Bernheim, 2009; Dasgupta and Tao, 1998). These papers primarily examine settings where there is no efficiency cost of adopting one sharing rule or another; in contrast, we show that the 50-50 norm extends to a production setting where dividing inputs equally has large potential efficiency costs. Third, we relate to papers examining preferences about the redistribution of income *ex post*, which find that some people believe that inequalities in *earned* income are fair even if inequalities in *unearned* income are not (e.g., Cherry et al., 2002; Fahr and Irlenbusch, 2000; Jakiela, 2015; Konow, 2003). This attitude may help explain why parents in our experiment do not equalize outcomes (i.e., expected earnings), which may be seen as earned income.

The rest of the paper is organized as follows. Section 2 presents a conceptual framework

---

[9]In the previous papers, all individuals are equally skilled at translating inputs to outcomes and hence there are not multiple notions of *ex ante* fairness, whereas for us, *ex ante* equality in inputs and outcomes diverge since some children are better at translating inputs into outcomes.

of parents' preferences. Section 3 describes the experimental design, and Section 4 presents the results. Section 5 concludes.

2   Conceptual Framework

Consider a utility maximizing parent with two children indexed by $i \in [1, 2]$. The parent is choosing the level of educational investment in each child, x_1 and x_2. Investments in child i weakly increase his or her present discounted lifetime earnings, with the relationship determined by the earnings function: $R_i(x_i) \equiv R(x_i, a_i, \varepsilon_i)$. Child i's earnings also depend on her "endowment" or ability a_i, which we define as her earnings potential or earnings when $x_i = 0$, and on ε_i, a mean-0 noise term capturing uncertainty in the production of earnings. When parents choose x_1 and x_2, they know a_1 and a_2, as well as the joint distribution of ε_1 and ε_2 but not their realizations.

Parents' utility functions incorporate three main quantities: i) total expected household earnings (representing returns-maximization), ii) the absolute gap between children's expected earnings (entering negatively, representing inequality aversion in outcomes),[10] and iii) the absolute gap between the inputs allocated to each child (entering negatively, representing inequality aversion in inputs). We also incorporate a preference for spending more on one child or the other (e.g., a preference for a son over a daughter). The utility function is a weighted sum of these quantities:[11]

$$
\begin{aligned}
U(x_1, x_2) = & \lambda \left[ER_1(x_1) + ER_1(x_2) \right] \\
& - \alpha \left| ER_1(x_1) - ER_2(x_2) \right| \\
& - \beta \left| x_1 - x_2 \right| \\
& + \gamma(z_1, z_2)(x_1 - x_2)
\end{aligned}
\tag{1}
$$

---

[10]Note that this formulation, in which parents care about $|ER_1(x_1) - ER_2(x_2)|$, captures aversion to inequality in *ex ante expected* earnings. An alternate formulation might stipulate that parents are averse only to inequalities in the final *ex post* distribution of outcomes ($E|R_1(x_1) - R_2(x_2)|$). We adopt the *ex ante* formulation since we study the decisions parents make from the *ex ante* perspective and because the literature generally finds that people care about *ex ante* fairness from an *ex ante* perspective (Andreoni et al., 2018; Brock et al., 2013; Cappelen et al., 2013; Krawczyk and Le Lec, 2010).

[11]This weighted-sum formulation embeds two important assumptions: separability across the different terms, and linearity in the inequality aversion terms. We make these assumptions for simplicity and (later) ease of estimation, but our results are robust to relaxing them. The formulation also assumes the returns maximization term is linear in earnings; this more substantively affects certain results, and we relax it below.

λ is the weight on returns maximization and α and β are the weights on inequality aversion of outcomes and of inputs. λ, α, and β are all weakly positive. $\gamma(z_1, z_2)$ is the parents' relative preference for investing in child 1 relative to child 2; it depends on a vector of each child's characteristics z_i (e.g., gender, age) and can be positive or negative.

The parent's problem is to choose x_1 and x_2 to maximize $U(x_1, x_2)$ subject to the budget constraint $x_1 + x_2 \leq y$, with y denoting the total educational budget.

Comparative Statics. To develop the basic predictions from the model, we consider cases in which the parent places full weight on one of the components and zero on the others. In these cases, the first-order conditions yield the following intuitive predictions:

Returns Maximization $(\lambda > 0, \alpha = 0, \beta = 0, \gamma = 0)$. The parent invests to equalize the returns to investment across children: $\frac{\partial ER_1(x_1)}{\partial x_1} = \frac{\partial ER_2(x_2)}{\partial x_2}$. If there is no interior solution, she gives all of her inputs to one child. Depending on the complementarity between x and the children's endowments, she could invest more in the higher- or lower-endowment child.

Inequality Aversion in Outcomes $(\lambda = 0, \alpha > 0, \beta = 0, \gamma = 0)$. The parent invests first in the lower-endowment child so that that child's earnings "catch up" with those of the higher-endowment child. At a sufficiently high y, she may begin to invest a positive amount in the higher-endowment child since that child's earnings may then begin to lag behind the lower-endowment child's earnings. If investments and endowments are complements, she will always invest more in total in the lower-endowment child.

Inequality Aversion in Inputs $(\lambda = 0, \alpha = 0, \beta > 0, \gamma = 0)$. The parent seeks to equalize x_1 and x_2, regardless of the shape of $R_1(x_1)$ and $R_2(x_2)$.

Child-specific Preferences $(\lambda = 0, \alpha = 0, \beta = 0, \gamma \neq 0)$. The parent gives more inputs to child 1 if $\gamma(z_1, z_2) > 0$ and to child 2 if $\gamma(z_1, z_2) < 0$.

Identification of Equation (1). We aim to estimate λ, α, and β from observed data on x_1 and x_2. To do so, we need to know the full perceived earnings functions, $ER_1(x_1)$ and $ER_x(x_2)$, and observe how investments respond to multiple exogenous "shocks" to these functions (e.g., shocks that increase the returns for the lower-endowment child relative to the higher-endowment child).

We also want our identification strategy to be robust to different functional forms for

utility. Although for simplicity we model the returns-maximization term as linear in total earnings, our strategy should be robust to nonlinearity (e.g., if the first term in equation (1) became $\lambda E u \left(R_1(x_1) + R_2(x_2) \right)$ for some concave $u(\cdot)$). Nonlinearity complicates identification because it means that returns-maximizing parents are risk averse, and risk aversion may "look like" inequality aversion in the data: risk-averse parents may choose more equal investments than the investments that maximize expected household earnings even if the parents are pure returns-maximizers.

One can sidestep this identification challenge by considering how parents choose to allocate "probabilistic investments." Consider a case where a parent is choosing how to allocate a single free binary investment (in our case: one indivisible tutoring class) between her children. Since the investment is binary, $x_i \in \{0, 1\}$. To determine who receives the investment, the parent receives n lottery tickets to allocate between her children, where each ticket has the same chance of being chosen ($\frac{1}{n}$), exactly one ticket is chosen per household, and the child whose ticket is chosen receives the binary investment.

The lottery setup yields a sharp test for inequality aversion: parents who do not care about inequality aversion ($\alpha = 0, \beta = 0$) will choose "all-or-nothing" allocations of lottery tickets, giving all tickets to one child or the other. If parents choose "split" allocations, giving positive allocations to both children, it must be that either $\alpha \neq 0$ or $\beta \neq 0$, or that it is the knife's edge (and hence empirically unlikely) case that the parent is *exactly* indifferent about which child receives the binary investment. The prediction that split allocations generally imply inequality aversion holds even with risk aversion and uncertainty in the production of earnings. The prediction follows from the fact that, with $\alpha = 0, \beta = 0$, expected utility is linear in lottery tickets, but also holds with many preferences that are non-linear in probability (e.g., prospect theory preferences; Kahneman et al., 2016). See Appendix B for additional discussion. The intuition for this all-or-nothing prediction is straightforward: since only one child ultimately receives the binary investment, if parents are not averse to inequality, they should simply choose which child they want to receive it and give all tickets to that child.

The all-or-nothing test for inequality aversion is a test for caring about a specific type

of inequality: inequality in *ex ante* or *expected* inputs and outcomes.[12] If some parents are only averse to inequalities in the final *ex post* distribution of inputs or outcomes (i.e., if a parent cares only about whether her children ultimately *receive* unequal inputs) but not about the *ex ante* expected input distribution, our test will categorize them as not averse to inequality. In that way, the test is conservative or shaded towards the "null" of the standard returns-maximizing model.[13]

3 Design and Procedures

Our goal is to identify the preference weights λ, α, and β. We use a lab-in-the-field experiment to shock children's short-run earnings and estimate these preferences.

We recruited roughly 300 parents with two children enrolled in grade 5 through grade 7 in government schools in southern Malawi during December of 2017. Both children were asked to take a math test,[14] and a monetary payment was delivered directly to the children based on their test scores. The monetary payment is our measure of the *outcome* (the $R(\cdot)$ function), which, in the literature, generally represents the earnings from education. Parents were given an input: 10 lottery tickets to be allocated across their children, where one ticket would be chosen per household and the lottery winner would receive one hour of tutoring on the material covered on the test. Thus—if they wanted to—parents could guarantee which of their children would receive tutoring by giving all tickets to that child. We measured parents' beliefs of each child's test score without the tutoring and of the amount each child's test score would increase with tutoring. We denote Child L and Child H as the child who the parent believed would have lower and higher test scores without tutoring, respectively.[15]

---

[12]On the inputs side, this means that the third term of equation (1) becomes $-\beta|Ex_1 - Ex_2|$.

[13]That said, it is not always the case that a parent who chooses to split expected inputs would also choose to split actual inputs; for example, if the input has convex returns (and therefore lower total returns when split across children), the parent might split expected inputs but not actual inputs. See, e.g., Andreoni et al. (2018); Brock et al. (2013); Cappelen et al. (2013); Krawczyk and Le Lec (2010) for further discussion of the distinction between *ex ante* and *ex post* notions of fairness. Our notion of aversion to inequality in *ex ante* expected inputs can be seen as a form of "procedural" fairness (e.g., Krawczyk and Le Lec 2010), wherein people care not just about the *ex post* outcome but also the procedure that produced it. A classic example of similar preferences comes from Machina (1989), who observes that a parent may be indifferent about which of her children receive an indivisible good but may strictly prefer to randomize which receives it.

[14]We focused on math to increase the reliability of the test and improve parents' ability to guess their children's scores.

[15]If parents' beliefs about both their children's test scores were equal, we arbitrarily defined Child L as the child whose first name comes first alphabetically.

Our identification comes from exogenously varying the payment functions that mapped child-specific scores to payments. For example, one payment function awarded both children the same payment per test-score-point, whereas another gave one child a higher payment than his or her sibling. We describe the payment functions in detail in Section 3.2 below. In order to maximize statistical power, we used the "strategy method": we presented each parent with five scenarios for what the payment functions might be, each parent chose allocations under each of the five scenarios, and then one scenario was randomly chosen to be implemented for each parent. Lottery tickets were assigned to both children based on their parents' allocation for that scenario.[16]

The elicitation used real stakes. Children were selected for (and received) real tutoring based on the lottery tickets chosen by their parents. Children then took the test and received cash payments based on their test scores and the randomly selected payment functions.

Because cash is potentially transferable within the household, the use of a cash reward could in theory bias us towards finding returns-maximization;[17] our experiment is thus conservative in rejecting the null of the "standard model" of returns maximization. To investigate whether the cash is in fact transferred within the households, we conduct a follow-up survey after the experiment and find limited evidence that it is, thus mitigating concerns about this bias, as described in Section 4.3.

3.1 Experimental Procedures

Figure A.2 presents a visual representation of the process for the experiment. Roughly a week prior to the experiment, participating parents completed a survey that elicited de-

---

[16]This adds a second layer of uncertainty: in addition to there being a lottery within scenarios about which ticket would be chosen, there is also a lottery across scenarios determining which scenario would be chosen. Since inequality aversion in our model is over *expected* inputs and outcomes, this raises a question regarding the level at which parents evaluate the expectation. We assume that parents "narrowly bracket" and try to equalize expected inputs and outcomes *within* scenarios; Exley and Kessler (2019) show that people generally narrowly bracket equity concerns. This assumption is conservative for estimating inequality aversion: if instead parents try to minimize expected inputs and outcomes *across* scenarios, that would bias us away from detecting inequality aversion.

[17]In particular, parents who are averse to inequalities in *consumption* but not earnings might returns maximize at the investment stage and then reallocate *ex post*. Since we aim to identify inequality aversion over earnings, not over consumption, the ability to reallocate *ex post* does not itself create bias. The bias would only arise if it is easier for parents to reallocate *ex post* in our experiment than other settings, thus decreasing the likelihood that they equalize outcomes *ex ante* in our experiment because they know they can easily equalize *ex post*.

mographic information, investments, and attitudes regarding their children's education. At this time, surveyors also described the math test that the children would take and measured parents' beliefs about each of their children's expected test score without tutoring and expected test score gains from tutoring, as well as the certainty of their beliefs. On the day of the experiment, parents were reminded of their beliefs and given a chance to change their responses if they wished. Surveyors then described the experimental design and conducted the experiment.

Because the experiment involves a number of steps, we took multiple measures during the explanation of the design to ensure understanding. We began the experimental design explanation with a "placebo" lottery designed to verify whether parents understood how to maximize monetary returns in a lottery environment; for the very few parents who did not understand, we then explained how to do so. Next, we gave parents a detailed overview of the full experimental design, walking them through two "practice" (hypothetical) scenarios; the practice scenarios used different payment functions than those used in the actual experiment but were explained in the same way as the experimental scenarios. After the practice scenarios, surveyors also walked parents through a sample "scenario lottery" to explain how the strategy method worked in this context. During the lottery explanation, surveyors also told parents that their children would not be told that their parents had any influence over which child received tutoring. Children were told only that tutoring was allocated by a random lottery.

For both the practice scenarios and those used in the experiment, surveyors explained the procedure as follows.[18] They began by describing the payment function for that scenario. They then walked parents through two visual aids, one graphical and one table-based. These displayed, for each ticket allocation the parents could choose, the expected payments for each child as well as (in the graphical version) the total expected payments across both children. Surveyors drew the graphs based on instructions from their tablets, which used parents' beliefs and the specific scenario to calculate expected payoffs under each ticket allocation. The graph clearly displayed the total expected payments as well as child-specific expected payments for each potential allocation (allowing the parent to observe the returns-

---

[18]See Appendix D for sample scripts and visual aids for Scenario 1.

12

maximizing and outcomes-inequality-minimizing allocations), but did not display anything about the gap in inputs.

As part of the explanations, surveyors told parents how to allocate their tickets if they wanted to maximize returns (total expected experimental earnings) or equalize outcomes (child-specific experimental earnings) or inputs in each scenario. Although this explanation could be seen as leading, we prioritized parental understanding of the experiment; we wanted to ensure that departures from returns maximization did not reflect poor understanding. Parents were also told that they could simply choose which child they wanted to receive tutoring by allocating all tickets to that child. Finally, after each practice scenario, surveyors asked parents questions to test their understanding of how to allocate tickets to achieve each of the goals above; no such questions were asked after the real experimental scenarios.

After the practice scenarios and explanation of the lottery, parents made their actual experimental allocations. Surveyors explained each scenario following the procedure described above, and parents made their selections.[19] We then conducted the lottery using the following steps: (1) surveyors' tablets randomly selected one scenario; (2) surveyors assigned the 10 tickets to children based on their parents' allocation for that scenario; and (3) the parents were asked to pick a ticket based on the ticket allocation for the selected scenario. For example, if the parent had allocated five tickets to each child for the selected scenario, the surveyor entered the initials of each child on five out of the ten tickets, and asked the parent to pick a ticket out of a hat to select the winner of the lottery.

Once the experiment was completed, a short survey was administered to gauge parental understanding about the experiment and address confounds. The "winning" child was then provided an hour of tutoring, after which all children took the test, with all children at a given school taking the test at the same time. Immediately after the test, cash earnings were delivered directly to the children in individual envelopes.[20]

---

[19]To ensure parent understanding, when surveyors explained a new scenario, they would not just explain the payment function but also specify what had changed relative to the payment function from the previous scenario. To keep this explanation uniform and easy for surveyors to master, we did not randomize the order in which scenarios were presented. However, because the payment function parameters do not move montonically with the order in which they were presented, it is unlikely that order effects drive our results.

[20]To deliver the cash, all the children and the remaining parents gathered in a classroom after the test. Surveyors called children up to the front of the room to receive their cash envelope. The amounts were not announced publicly.

3.2 Scenarios and Predictions

We now describe the payment functions for each scenario in the experiment and how they allow us to identify the parameters of the parental utility function specified in Section 2. The payment function for child $i \in \{L, H\}$ from household k in scenario j can be expressed as the sum of a lump-sum transfer B_{ij} plus a reward of C_{ij} per point on the test:

$$P_{ijk} = B_{ij} + C_{ij}(TestScore_{ik} - Threshold_k)$$

P_{ijk} is the child's payment if she receives a score of $TestScore_{ik}$ on the test (test scores are percentage scores out of 100). $Threshold_k$ equals parent k's belief about child L's test score without tutoring, rounded down to the nearest 10. We use the same threshold for both children in household k; we reward performance above this value in order to implement steep payment functions while keeping total payments reasonable. We suppress the k index going forward.

Figure 1 presents the five payment function scenarios used in the experiment as well as the predictions for how parents would allocate their tickets in each scenario if their utility functions only weighted (a) returns-maximization, (b) inequality aversion over outcomes, or (c) inequality aversion over inputs. We generate the predictions as follows. In each scenario, the returns-maximizing strategy is to give all tickets to the child with larger expected payment gains from tutoring, calculated as the parent's belief about each child's test score gains from tutoring (denoted R_i) multiplied by the child's scenario-specific payment function slope C_{ij}. Thus, of the payment function parameters, only C_{ij} matters, not the lump sum amount B_{ij}. In contrast, the outcomes-inequality-minimizing strategy is to minimize the cross-child expected payment gap, which depends on both B_{ij} and C_{ij} because both determine expected payments. Finally, the inputs-inequality-minimizing strategy is to equate tickets across children regardless of B_{ij} or C_{ij}.

Figure 1 shows the predictions for each scenario separately, but only for ease of illustration: our analysis primarily uses *cross-scenario* variation for identification. The cross-scenario (i.e., within-child) variation originates from the experimental variation in the payment functions, and controls for the endogenous child- and parent-level factors affecting

14

Figure 1: Payment Function Scenarios and Predictions

| Payment Function Scenarios Child i's payment $= B_i + C_i(TestScore_i - Threshold^1)$ | | | | | Predictions: Would parents give more tickets to Child L, Child H, or give equally to both? | | |
| | Child L[2] | | Child H | | | | |
Scenario	B_L	C_L	B_H	C_H	Returns Maximization[3]	Inequality Aversion of Outcomes	Inequality Aversion of Inputs
1. Base Case	0	10	0	10	If $R_L < R_H$: H If $R_L = R_H$: No prediction If $R_L > R_H$: L	If $R_L < R_H$: L If $R_L = R_H$: L[4] If $R_L > R_H$: Depends on parameters[5]	Equal
2. Higher Returns to Child H	0	10	0	100	H[6]	L	Equal
3. Higher Returns to Child L	0	100	0	10	L[7]	H[8]	Equal
4. Lump Sum to Child L	1000	10	0	10	If $R_L < R_H$: H If $R_L = R_H$: No prediction If $R_L > R_H$: L	H	Equal
5. Higher Returns to Child L & Lump Sum to Child H	0	100	6000	10	L[6]	L	Equal

Notes: 1. Threshold is the parent's belief about Child L's test score without tutoring, rounded down to the nearest 10.
2. Child L (lower-performing child) defined as the one whom the parent perceived would have a (weakly) lower test score without tutoring.
3. R_i = parent's belief about child i's score without tutoring - belief about child i's score without tutoring.
66% of parents believed $R_L < R_H$, 19% believed $R_L = R_H$, and 14% believed $R_L > R_H$.
4. Assumes test score for Child L without tutoring is strictly less than for Child H without tutoring. If equal, no prediction.
5. Defining S_i as the test score without tutoring, tickets to $L = 10(S_H - S_L + R_H)/(R_L + R_H)$ (unless that falls outside of the 0-10 range).
6. For 96% of parents.
7. For 95% of parents.
8. For 98% of parents.

choices, such as parents' beliefs about the benefits for tutoring for each of their children, R_i. In addition, although the predictions shown in Figure 1 assume that the only value of the tutoring to parents is the short-run monetary payments their children will receive (and thus that parents ignore any potential long-run or non-monetary benefits of tutoring), this assumption is again only for expositional purposes.[21] The cross-scenario predictions we take to the data are robust to parents believing that tutoring has additional value, as they difference out any considerations that are constant across scenarios.[22]

Note that there is one important hypothesis test we perform that does not depend on cross-scenario variation: the test of whether parents only choose all-or-nothing allocations (a test for pure returns maximization). The test does not require cross-scenario variation because it holds regardless of parents' beliefs about their children's test scores or the additional value to tutoring, obviating the need for cross-scenario variation to control for beliefs.

In our *Base Case* scenario (Scenario 1), the payment functions for both children are the same: both children receive MWK 10 worth of rewards for each test score point above the threshold ($C_L = C_H = 10$), with no lump-sum transfers ($B_L = B_H = 0$). The other four scenarios are variations on the base case designed to identify the utility function parameters.

We designed two scenarios to yield opposite predictions for returns-maximization and inequality aversion over outcomes, thus letting us test which preference parents weigh more heavily on average. Scenario 2, *Higher Returns to Child H*, gives Child H a ten times higher per-point reward than Child L: C_H is 100 MWK while C_L is 10 MWK. Neither child receives a lump sum. Increasing C_H has two effects. First, it increases the returns to receiving tutoring for Child H relative to Child L enough that, for 96% of households in our sample, the returns-maximizing strategy (for experimental earnings) is to give all inputs to Child H. Second, it increases Child H's expected payments for any ticket allocation that the parent could choose. As a result, the outcomes-inequality-minimizing choice (for experimental earnings) is to give all inputs to Child L. Scenario 3, *Higher Returns to Child L*, exchanges the payment functions

---

[21]The assumption is also consistent with survey data: Over 92% parents said they were only thinking about the short-run monetary payment returns to tutoring when making their experimental ticket allocations.

[22]We also assume away child-specific preferences when discussing these predictions. Because these preferences are constant across scenarios, they are also differenced out in cross-scenario comparisons.

used in *Higher Returns to Child H* between Child L and Child H. Relative to *Higher Returns to Child H*, returns-maximizing parents would now reallocate to Child L, whereas outcomes-inequality-minimizing parents would do the opposite.

Our next scenario, *Lump Sum to Child L* (Scenario 4), was designed specifically to test for inequality aversion over outcomes. Relative to the *Base Case*, *Lump Sum to Child L* delivers a lump sum transfer, B_L, of MWK 1000 to Child L while delivering no lump sum transfer to Child H. For both children, the per-point rewards C_i remain the same as in the *Base Case*: 10 MWK per point. Increasing Child L's lump sum transfer, B_L, should not change his or her expected cash rewards from receiving tutoring, $R_L C_L$, and thus does not affect parents' returns-maximizing choices; nor does it affect their input-equalizing choices. However, lump sum transfers do affect the outcomes-equalizing choice: since giving a lump sum transfer to one child increases his or her expected payments, an outcomes-equalizing parent would respond by reallocating inputs to that child's sibling to increase the sibling's expected payments.

In both the *Higher Returns to Child H* and *Higher Returns to Child L* scenarios, both the outcomes-inequality-minimizing and the returns-maximizing choices are all-or-nothing allocations. However, in both cases, they are *different* all-or-nothing allocations (i.e., the child a returns-maximizing parent would choose is the opposite of the child an outcomes-inequality-minimizing parent would choose). Thus, if parents' preferences place positive weight on both returns maximization and inequality aversion over outcomes, it could cause parents to choose split allocations.[23]

To shed light on whether split allocations reflect inequality aversion over inputs, we introduce a final scenario, *Higher Returns to Child L & Lump Sum to Child H* (Scenario 5), that adjusts the payment functions so that the child a returns-maximizing parent would choose is the *same* as the child a outcomes-inequality-minimizing parent would choose (both are Child L). As a result, if we see similar numbers of parents choosing split allocations in that scenario, it is unlikely that the splitting reflects a balance between inequality aversion over outcomes and returns maximization; instead, it suggests parents care directly about

---

[23]While this is technically only true in the case where the utility function is nonlinear in the various terms instead of linear as in equation (1), since with linear terms the solution will be at a corner, we want our predictions to be robust to nonlinearity.

minimizing inequality in inputs.[24]

Another distinctive prediction of inequality aversion over inputs—testable using any of our scenarios—is that there should be excess mass (relative to a smooth distribution) in the density of choices at the equal allocation point. No other theories should produce this excess mass, since all other factors (e.g., outcomes inequality or expected returns) are smooth through the equal-allocation point.

Figure 2 depicts the scenarios and identifying variation graphically. For each scenario, there are two bars, with the left bar showing the case where parents give all 10 tickets to Child H and the right bar showing the case where parents give all 10 tickets to Child L. The height of each bar represents expected household earnings in that case (i.e., the sum of Child L's and Child H's expected earnings), averaged across all households in the sample.

Figure 2: Expected Earnings Vary Within and Across Scenarios
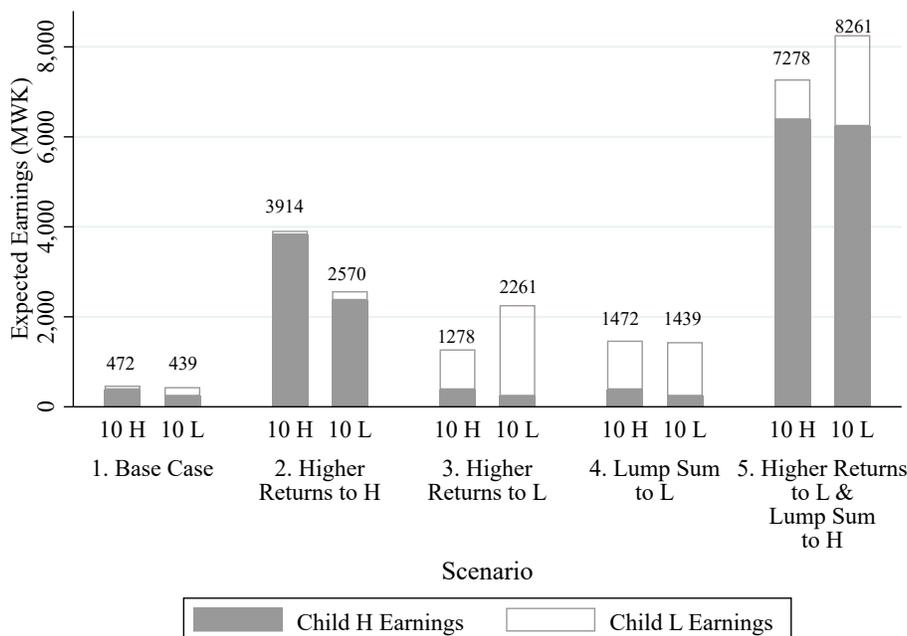
Notes: This figure presents, for each scenario, the mean perceived expected experimental earnings for Child L and Child H in the cases when all ten tickets are given to Child L and when all ten tickets are given to Child H.

There are a few takeaways from Figure 2. First, in the scenarios where the payment per

---

[24]Note that this prediction depends on parents not having important child-specific preferences for Child H, or beliefs that Child H has higher non-monetary or additional benefits from tutoring than Child L. To assess this possibility, we also test the prediction excluding parents who appear to have preferences for Child H, as noted in footnote 29 of Section 4.3.

test score point C_i varies substantially across children (Scenarios 2, 3, and 5), there are large within-scenario differences in expected household earnings depending on whom the parent gives tickets to. For example, in Scenario 2 (*Higher Returns to Child H*), parents on average would sacrifice MWK 1,344, or 96% of an adult's daily wage, if they allocate all of their tickets to Child L instead of Child H.[25] Thus, it is costly for parents to not maximize returns. Second, there are meaningful differences across scenarios in terms of the total earnings, earnings distribution, and which child is the higher-return child. Third, although in several scenarios the average parent cannot perfectly equalize Child L's and Child H's expected payments, parents' allocations can still make meaningful headway towards equality. For example, in scenario 1, the average parent will decrease payment inequality between Child L and H by 88%, from 292 MWK to 35 MWK, if she allocates all tickets to Child L instead of Child H. Across all scenarios, the average parent will decrease payment inequality by 50% if she always chooses the outcome-inequality-minimizing allocation instead of the opposite. Finally, much of our identification requires that parents believe that tutoring increases test scores. The figure implicitly shows that parents *on average* believe this—if not, the left and right bars for each scenario would have the same height. In the next subsection we also show that nearly all parents believe this as well.

Our experiment identifies preferences based on parents' *beliefs* about their children's scores, which other work shows are often inaccurate (Banerji et al., 2017; Dizon-Ross, 2019). Fortunately, whether parents' beliefs are accurate is not important for our identification: we are identifying parents' preferences conditional on their beliefs. Potential inaccuracies should only affect our estimation if parents' beliefs distributions are uncertain and if that uncertainty affects their allocations; as we discuss later in Section 4.3, this does not appear to be the case here.

## 3.3 Validation of Method and Respondent Understanding

Since our experimental design involves a number of steps, it is important to validate that parents indeed understood the set-up and how to maximize returns. We present multiple

---

[25]The daily wage for adults was estimated to be 1,400 MWK by field staff at the time of the study. Gross Domestic Product per capita in Malawi was 260,501 MWK, or 364 USD in 2017 (World Bank, 2020). Assuming a working year of 275 days (based on a six-day workweek minus time off), this is equivalent to 947 MWK, or 1.32 USD per work day, so our experimental stakes are even higher relative to that benchmark.

pieces of evidence that they did. First, in the placebo lottery, we asked parents to allocate 10 lottery tickets between two (hypothetical) prizes just like they did with real stakes in the main experiment; however, here the two prizes were monetary prizes to be given directly to the parent: 50 MWK (0.07 USD) or 100 MWK. As shown in Figure 3, 97% (280/289) of parents allocated 100% of the tickets to the larger 100 MWK prize, suggesting that the vast majority of parents understand how to maximize returns in a lottery.

Figure 3: Nearly All Parents Maximized Returns in the Placebo Lottery
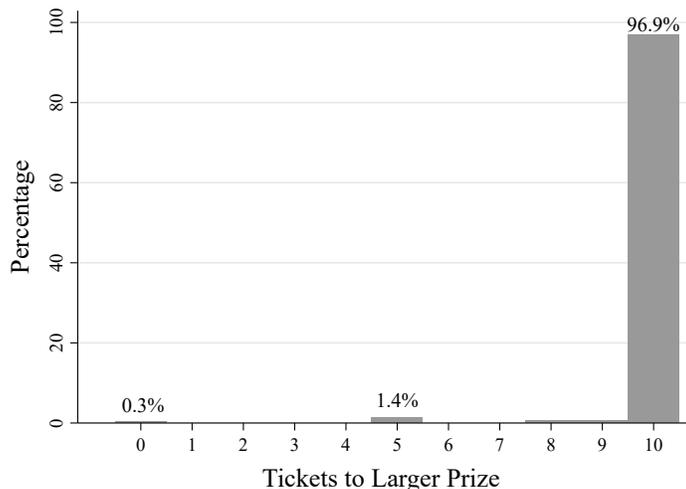
Second, for each scenario, after we elicited ticket allocations, we asked parents about the rationale behind their chosen allocation. In Appendix Table A.1, we show that the stated rationales correlate well with the actual choices. Thus, parents appear to have correctly understood the experimental setup.

3.4 Summary Statistics

Table 1 reports selected summary statistics from our sample. Roughly 85% of all respondents are female. Fifty-five percent of the child sample is female. The average parent in our sample spends roughly MWK 8400/year (11.75 USD) on education for each of her children. On average, parents believe that Child H's score on the test will be 11 percentage points (pp) higher than Child L's. Our experiment depends on parents perceiving that the test score returns to tutoring are positive; critically, 99% of parents thought it would have positive returns for at least one of their children and 93% thought it would have positive returns for both. Parents generally believed that tutoring would have higher score returns

for Child H than Child L, increasing Child L's score by 11pp and Child H's by 18pp, on average.

Table 1: Summary Statistics

A. Respondent Characteristics:		
Completed Grade 8		0.33
		(0.47)
Female		0.85
		(0.36)
Believed tutoring had positive returns for at least one child		0.99
		(0.08)
Believed tutoring had positive returns for both children		0.93
		(0.26)
B. Child Characteristics:	**Child L**	**Child H**
Grade 5	0.40	0.38
	(0.49)	(0.49)
Grade 6	0.32	0.36
	(0.47)	(0.48)
Grade 7	0.27	0.26
	(0.45)	(0.44)
Female	0.55	0.56
	(0.50)	(0.50)
Annual household education expenditure (MWK)	8412	8372
	(9316)	(9020)
Parent believed child had positive returns to tutoring	0.96	0.96
	(0.20)	(0.19)
Parent's belief of score without tutoring (out of 100)	53.14	67.87
	(13.26)	(13.48)
Parent's belief of score with tutoring (out of 100)	64.43	82.43
	(14.02)	(12.76)
Parent believed child had strictly higher returns to tutoring than sibling	0.14	0.66
	(0.35)	(0.47)
Math test score (out of 100)	41.91	44.14
	(25.16)	(24.23)
Received tutoring	0.43	0.57
	(0.50)	(0.50)
Total Households	289	

Notes: This table presents baseline summary statistics for the sample. All statistics are proportions unless otherwise indicated. Standard deviations are in parentheses. "Returns" to tutoring is defined as the difference between beliefs of the test scores with and without tutoring.

4 Experimental Results

This section presents the main results of the experiment. We begin by analyzing the raw data and comparing the results across our experimental scenarios to provide qualitative evidence on the preference parameters. We then estimate the magnitudes of the preference parameters using structural estimation. Finally, we present evidence against possible confounds and rule out alternative explanations for our findings.

4.1 Qualitative Evidence for Each Type of Preference

Returns Maximization: We first test whether parents only value maximizing returns. If so, they should allocate their tickets in an all-or-nothing fashion. Instead, Figure 4 shows that only 44% of allocations were all-or nothing. In 56% of allocations, both children received non-zero tickets. Thus, parents appear to not be pure returns-maximizers but rather to also care about equality and fairness.

Figure 4: Parents Often Did Not Maximize Returns in the Experiment
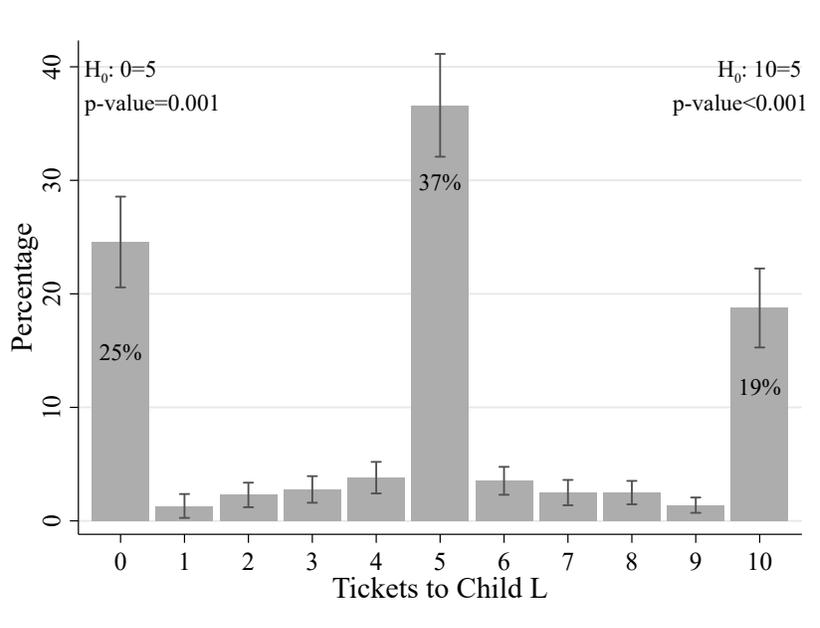
Notes: This figure presents the allocation of experimental lottery tickets across children, pooled across scenarios 1-5, and 95% confidence intervals.

Although parents are not *pure* returns-maximizers, Figure 5 presents evidence that parents do still place positive weight on returns maximization. Panel (a) of Figure 5 compares Scenario 2 (*Higher Returns to Child H*) and Scenario 3 (*Higher Returns to Child L*). Switch-

ing from Scenario 2 to Scenario 3 increases the payment per test score point for Child L relative to Child H. As a result, returns maximization suggests parents should reallocate tickets to Child L, as indicated by the "Returns Maximization" arrow. In contrast, increasing the payment per test score point for Child L also effectively makes Child L richer; if parents were inequality averse over outcomes, they would thus reallocate in the opposite direction, as indicated by the "IA Outcomes" arrow. Consistent with returns-maximization, we find that parents increase their allocations to Child L by a statistically significant 1.6 tickets (Figure 6), with over 20 percentage points (pp) fewer parents giving all of their tickets to Child H, and 10 pp more parents giving all tickets to Child L (Figure 5(a)).

Figure 5: Consistent with Returns Maximization, When the Per-Point Reward for a Child Increased, Parents Shifted Inputs Toward That Child
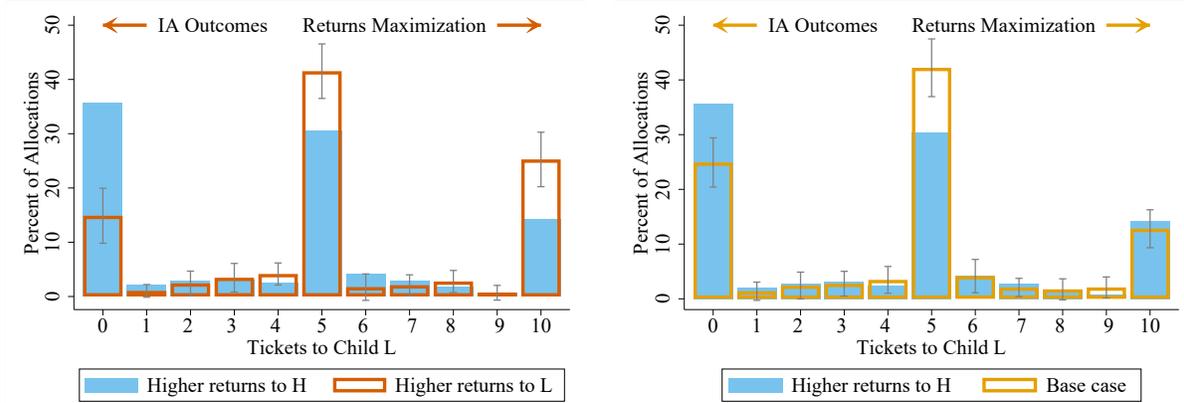

(a) Scenario 2 (*Higher Returns to Child H*) vs. Scenario 3 (*Higher Returns to Child L*)

(b) Scenario 2 (*Higher Returns to Child H*) vs. Scenario 1 (*Base Case*)

Notes: Each of these figures present the allocation of lottery tickets for two separate scenarios. The arrow labeled "Returns Maximization" shows the direction that returns maximization predicts allocations would move when going from the solid scenario to the outlined scenario, and the arrow labeled "IA Outcomes" shows the direction that inequality aversion over outcomes predicts allocations would move when going from the solid to the outlined scenario.

Although these shifts in allocations are meaningful, it is important to note that they are still quite muted relative to a pure returns-maximizing response. If parents were maximizing experimental earnings, the share of parents giving all tickets to Child L in Figure 5(a) should go from 0 to 100%, not 15% to 25%. The fact that parents deviate from returns maximization may be particularly surprising given that the potential costs are large. For example, Child L's average expected experimental gains from tutoring are a whole 983 MWK—70% of a day's

Figure 6: The Tickets Allocated to Child L Varied Across Scenarios
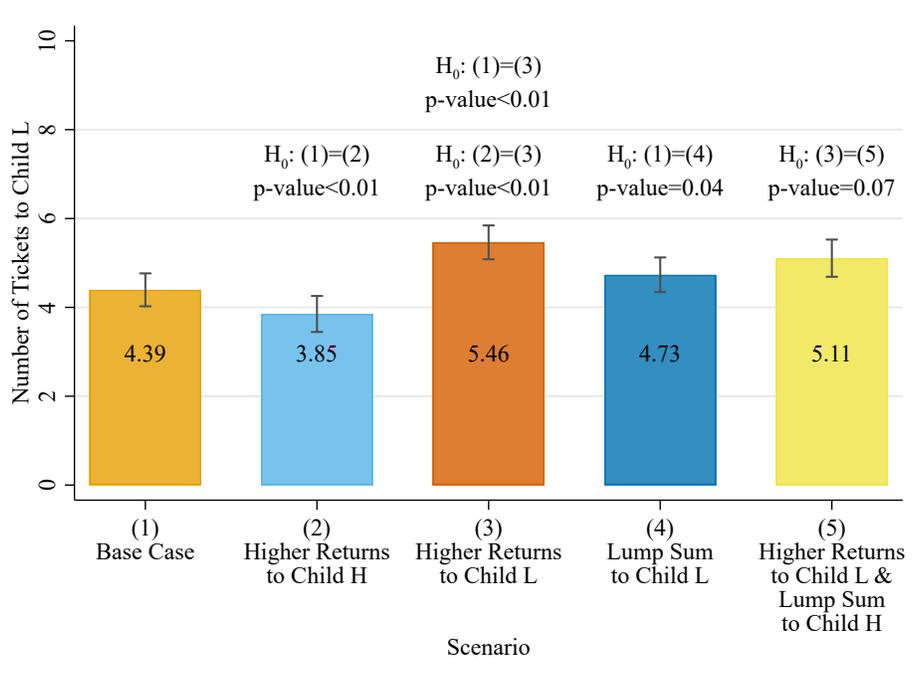
Notes: This figure presents average lottery tickets allocated (out of 10) to Child L, and 95% confidence intervals. P-values correspond to the null hypotheses of equal ticket allocation between given scenarios.

wage—higher than Child H's in Scenario 2, and the differences are even larger in Scenario 3. This suggests that parents also place high weight on preference components other than returns maximization, a point we return to below.

Panel (b) of Figure 5 presents a similar comparison between Scenarios 3 (*Higher Returns to Child L*) and 1 (*Base Case*). The magnitudes of the shifts are smaller due to the smaller variation across payment functions, but the takeaway is the same: parents place positive weight on maximizing returns and, on average, their desire to maximize returns dominates any potential desire to equalize outcomes.

Inequality Aversion over Outcomes. Next, we test whether parents place positive weight on inequality aversion over outcomes ($\beta > 0$). Relative to the *Base Case* (Scenario 1), the *Lump Sum to Child L* scenario (Scenario 4) delivers a lump sum to Child L without changing the per-point rewards for either child. The only theory that predicts parents will react to this change is inequality aversion in outcomes, which predicts that parents would reallocate towards Child H. Note that the lump sum is large enough—at MWK 1000 (1.40 USD), or 70% of a day's wage —that if parents did care about inequality in outcomes, we

would expect to see meaningful responses. However, in aggregate, we fail to find evidence for reallocation towards Child H (Figure 7, Panel (a)). If anything, parents, on average, slightly reallocate in the opposite direction. Turning to individual-level changes, equal numbers of parents reallocate to and away from Child L, suggesting the reallocations primarily represent noise (Figure A.3). Figure 7(b) shows similar findings from comparing *Higher Returns to Child L* (Scenario 3) and *Higher Returns to Child L & Lump Sum to Child H* (Scenario 5). These scenarios were included in the design for other reasons but the only difference between them is again the lump sum.

Figure 7: Inconsistent With Inequality Aversion in Outcomes, When One Child Received a Lump Sum, Parents Did Not Reallocate to His or Her Sibling.
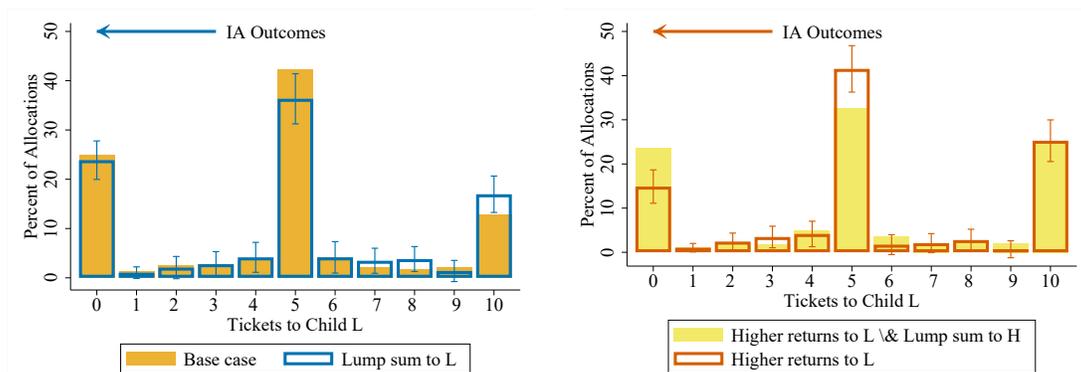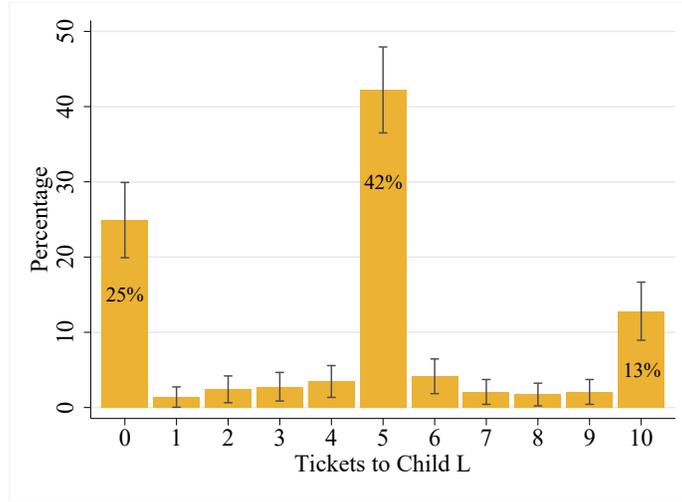

(a) Scenario 1 (*Base Case*) vs. Scenario 4 (*Lump Sum to Child L*)

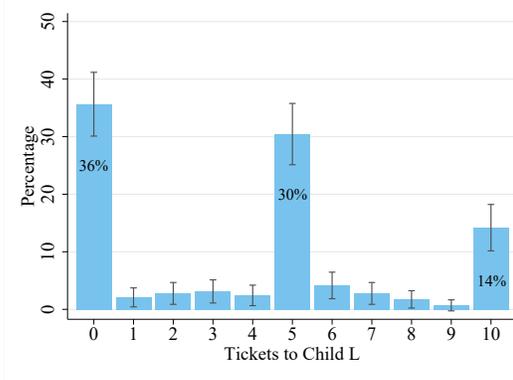(b) Scenario 5 (*Higher Returns to Child L & Lump Sum to Child H*) vs. Scenario 3 (*Higher Returns to Child L*)

Notes: Each of these figures present the allocation of lottery tickets for two separate scenarios. The arrow "IA Outcomes" shows the direction that inequality aversion over outcomes predicts allocations would move when going from the solid to the outlined scenario.

Inequality Aversion over Inputs. Taken together, the evidence presented so far suggests that, first, parents value both returns-maximization and *some* form of equality, and second, they do not value equality in outcomes. This suggests that they likely value equality in inputs. We now substantiate this conclusion with additional analysis. In particular, inequality aversion over inputs is the only theory that predicts excess mass at 50%. Consistent with this, visual inspection of the data in Figure 4 shows a notable spike at equal allocation, with parents choosing exactly equal inputs in roughly 37% of the scenarios. Moreover, we can easily reject that the ticket distribution is smooth around the 5/5 point. Figure 8 shows
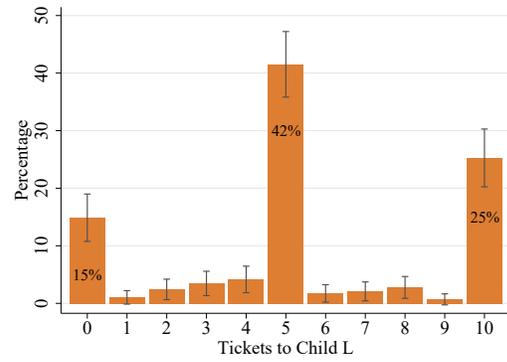
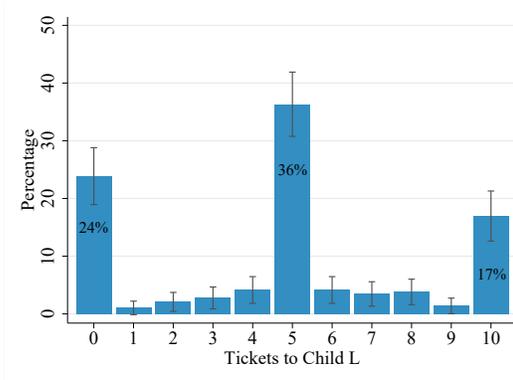Figure 8: Parents Equalized Inputs in All Scenarios
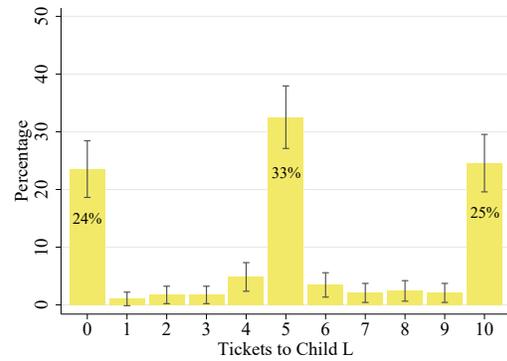

(a) Scenario 1 (*Base Case*)

(b) Scenario 2 (*Higher Returns to Child H*)

(c) Scenario 3 (*Higher Returns to Child L*)

(d) Scenario 4 (*Lump Sum to Child L*)

(e) Scenario 5 (*Higher Returns to Child L & Lump Sum to Child H*)

that a substantial share of parents choose equal inputs in every single scenario, even the scenarios where there are large differences in expected earnings across children. The desire to equate also appears to be widespread across parents, with 58% of parents equalizing at least once. Section 4.3 discusses and rules out other potential explanations for equal splitting and other threats to interpretation of our results.

Child-Specific Preferences   Interestingly, many of parents' choices diverge from either the returns-maximizing, outcome-inequality-minimizing, or input-inequality-minimizing choices, suggesting that parents may also care about something else. For example, 24% percent of parents in Scenario 5 allocated all of their tickets to Child H despite the fact that both returns maximization and inequality aversion over outcomes suggest they should give all to Child L.[26] Appendix Figure A.5 provides evidence that this behavior represents child-specific preferences that vary across parents. Parents who allocated all tickets to Child H or Child L in the *Base Case* scenario (which featured symmetric payment functions) were substantially more likely to allocate all tickets to that same child in all the other scenarios, explaining most of the deviations we described above from the other strategies. Note that the "child-specific preferences" that we are identifying here incorporate both actual child-specific preferences (e.g., preferences for a certain gender) and also beliefs about cross-child differences in the non-experimental benefits of tutoring (e.g., non-monetary or long-run benefits). Both of these factors could lead parents to prefer one child over the other and do not vary across scenarios.

Forgone Earnings from Equalizing Inputs   Finally, we show that parents' deviations from returns maximization have significant earnings implications. For each family × scenario, we calculate the sum of expected earnings for both children under parents' chosen ticket allocations in the experiment ("chosen earnings") and compare those with expected earnings if parents had instead chosen to maximize the sum of their children's payments ("returns-maximizing earnings") or minimize the sum of their children's payments ("returns-minimizing"). Figure 9 then plots forgone earnings (returns-maximizing earnings minus chosen earnings), both in absolute terms and as a percent of "potential earnings" (returns-

---

[26]Technically these predictions for Scenario 5 only hold for 95% of parents but the same 24% statistic holds in that 95% sample as well.

maximizing minus returns-minimizing earnings). On average, across scenarios, parents for-
went roughly 40% of their potential (non-inframarginal) earnings. The experimental stakes
are substantial; the forgone expected payment amounts are correspondingly large, especially
for Scenarios 2, 3, and 5 where we exogenously varied the returns to tutoring across chil-
dren.

Figure 9: Parents Forgo Substantial Experimental Earnings

(a) Forgone Earnings

(b) Forgone Earnings as Percent of Potential
Earnings
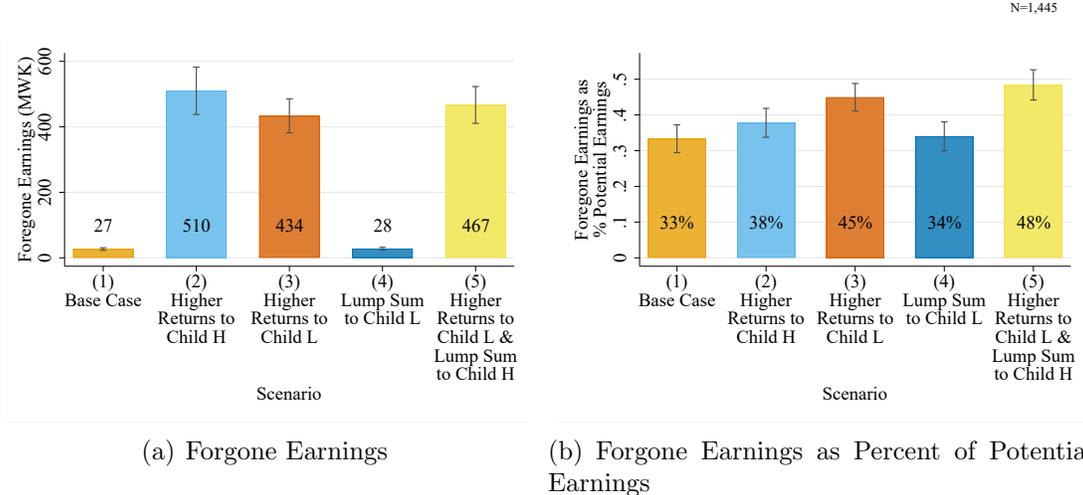
Notes: Panel (a) presents income forgone by parents and 95% confidence intervals. Panel (b) presents forgone
income, and 95% confidence intervals, as a percentage of potential earnings. Forgone earnings are calculated
as the difference between expected income if parents were maximizing returns and expected income based
on their preferred lottery ticket allocation.

4.2 The Value of Equal Allocation

This section uses our experimental results to estimate parents' preference weights from
Section 2: λ, α, β, and γ. We use a structural approach to numerically estimate these
parameters. This strategy is identified from variation across parents' potential choices in ex-
pected inputs and payments. The majority of this variation comes from our experimentally-
generated cross-scenario variation in the payment functions. However, since parents' beliefs
were one input into the payment functions, there are potential concerns about endogeneity.
For example, in Scenario 1, households who perceive the same gain from tutoring for both
children will not forgo earnings if they equalize inputs, while parents who perceive different
gains from tutoring across their children will. As a result, we also present results that take
advantage of a control-function approach to isolate the experimental variation.

We use a mixed logit regression model to estimate parents' preference parameters. The mixed logit allows the preference parameters to vary across the population and avoids the independence of irrelevant alternatives assumption (IIA) entailed by a simpler conditional logit approach. Following equation (1), we assume that parent i has the following utility in scenario j from choosing ticket allocation k:

$$u_{ijk} = \lambda_i TotalPay_{ijk} - \alpha_i OutcomeInequality_{ijk} - \beta_i InputInequality_{ijk}$$
$$+ \gamma_i InputsToChildLvsH_{ijk} + \varepsilon_{ijk} \tag{2}$$

where $TotalPay_{ijk}$ is the total expected combined earnings across both children under allocation k, measured in 100's of MWK (the returns-maximization term); $OutcomeInequality_{ijk}$ is the absolute difference between parent i's children's expected earnings under allocation k measured in 100's of MWK (the inequality aversion in outcomes term); $InputInequality_{ijk}$ is the absolute difference in expected inputs (lottery tickets) between parent i's children under allocation k (the inequality aversion in inputs term); and $InputsToChildLvsH_{ijk}$ is the number of tickets given to Child L relative to Child H in allocation k (the child-specific preference term; note that since γ_i can be positive or negative, parents can prefer either of their children).

We also estimate variations on this specification. Because the earlier reduced form evidence suggests that the primary utility cost from not giving equal inputs is binary instead of continuous, we model the utility cost associated with unequal inputs as binary in addition to the continuous measure defined above. In the binary specification, $InputInequality_{ijk}$ is a dummy that equals 1 if allocation k does not give the two children an equal number of tickets and 0 otherwise. In addition, to facilitate interpretation, in an alternative specification, we model $TotalPay_{ijk}$ as the log of the total expected combined earnings across both children under allocation k.

We allow for the preference parameters $(\lambda, \alpha, \beta, \gamma)$ to vary for each parent i; we assume that each preference parameter is distributed normally with a standard deviation estimated using the regression model. We also allow for correlations across all preference parameters and again estimate the correlations within the regression procedure. Finally, the error term

ε_{ijk} is assumed to be Type I extreme value, independent across i, j, and k.[27] In each scenario j, parent i is assumed to choose the allocation k with the highest utility.

To correct for potential endogeneity, we also adopt the control function approach of Petrin and Train (2010). In particular, in the first stage, we estimate two regressions, regressing the dependent variables $TotalPay_{ijk}$ and $OutcomeInequality_{ijk}$ on our exogenous instruments (indicators for the scenario $\times$ ticket allocation, τ_{jk}) and on the other endogenous regressors from equation (2). We denote $\hat{\eta}_{ijk}$ and $\hat{\mu}_{ijk}$ as the residuals from those regressions. In the second stage, we then include these residuals linearly in the utility function, allowing them to enter with normally-distributed random coefficients ρ_i and τ_i:

$$u_{ijk} = \lambda_i TotalPay_{ijk} - \alpha_i OutcomeInequality_{ijk} - \beta_i InputInequality_{ijk}$$
$$+ \gamma_i InputsToChildLvsH_{ijk} + \rho_i \hat{\eta}_{ijk} + \tau_i \hat{\mu}_{ijk} + \varepsilon'_{ijk} \tag{3}$$

Intuitively, this approach leverages the fact that conditional on the linear control function $\rho_i \hat{\eta}_{ijk} + \tau_i \hat{\mu}_{ijk}$—which captures all of the potentially-endogenous components of $TotalPay_{ijk}$ and $OutcomeInequality_{ijk}$—$TotalPay_{ijk}$ and $OutcomeInequality_{ijk}$ should no longer be correlated with the unobserved component of the decision ε'_{ijk}.

The mixed logit estimates of the means of the parameter distributions are shown in Column 1 of Table 2. Consistent with the Section 4.1 results, we find that parents are more likely to choose a ticket allocation when the monetary returns (total expected payments) associated with that allocation increase. We also find no evidence of aversion to inequality in outcomes: the absolute difference in child-level earnings does not influence choices, with the coefficient not statistically significant, small in magnitude, and wrong-signed. However, parents have a strong preference for equalizing inputs: they are significantly more likely to

---

[27]Allowing each choice to have a separate logit error may seem strange here given that there is a relatively natural numeric ordering between the allocations. Indeed, if the options were "fully ordered" in the sense that, if a parent ranked her preferences, her first choice would always be adjacent to her second choice in the ordering, then this specification would be very unreasonable. However, although choices take on numeric values here, they are not in fact fully-ordered. For example, if a parent's first choice would be to give all tickets to child 1, it does not mean her second choice is necessarily to give 9 tickets to child 1; her second choice might instead be to split 5/5 because she has a high utility from splitting, or to give all tickets to child 2 because she likes to make all-or-nothing allocations. This means that allowing different choices to have separate logit errors is more plausible here than in settings where the choices are fully-ordered, i.e., where knowing a parent's first choice means we know her second choice.

Table 2: Mixed Logit Estimates Suggest Parents Place High Weight on Equalizing Inputs

	Dependent Variable: Chose Ticket Allocation				
	(1)	(2)	(3)	(4)	(5)
Household earnings ('00)	0.2471***	0.1527***	0.2297***	0.1937***	
	(0.0557)	(0.0545)	(0.0454)	(0.0475)	
Log household earnings					7.0107***
					(1.6224)
Absolute difference in inputs	-0.3645***	-0.3095***			
	(0.0613)	(0.0690)			
Inputs not equally split (0/1)			-2.9763***	-2.9031***	-3.0869***
			(0.2921)	(0.2967)	(0.2988)
Gap between children's earnings ('00)	0.0347	0.0773*	0.0108	0.0297	-0.0313
	(0.0353)	(0.0426)	(0.0292)	(0.0300)	(0.0412)
Tickets to Child L	-0.0831	-0.1788**	-0.1398**	-0.1812***	-0.1503**
	(0.0643)	(0.0807)	(0.0684)	(0.0600)	(0.0727)
WTP for 1 unit lower input inequality ('00)	1.48	2.03			
WTP for equal inputs ('00)			12.96	16.57	
WTP for equal inputs (percent)					0.44
Control Function?	NO	YES	NO	YES	YES
Observations	15,895	15,895	15,895	15,895	15,895

Notes: This table presents estimates of the predictors of parents' choice of ticket allocations. Each observation is a parent × scenario × ticket allocation. Willingness to pay is calculated by dividing the "Absolute difference in inputs" (Columns 1 and 2) and "Inputs not equally split" (Columns 3, 4 and 5) coefficients by the "Household earnings" (Columns 1, 2, 3 and 4) and "Log Household earnings" (Column 5) coefficients. Columns 2, 4, and 5 are estimated using a control function approach where control functions are used for "Household earnings" and "Gap between children's earnings" and the instruments are dummies for the scenario × ticket allocation. Standard errors are in parentheses, clustered at the household level. * denotes significance at 0.10; ** at 0.05; and *** at 0.01.

pick choices that have smaller input gaps between children. To interpret the magnitude, note that the coefficient estimates are only identified up to a scale factor.[28] It is thus useful to scale all coefficients relative to the (sample-average) $TotalPay$ coefficient, λ; $\frac{\beta}{\lambda}$ then gives the amount of total household earnings a parent would be willing to forgo to decrease the gap in inputs between her children by 1 lottery ticket. Our estimate of $\frac{\beta}{\lambda}$ is quite large, implying that, on average, parents are willing to give up 148 MWK or roughly 0.20 USD in expected household earnings for each 1 ticket decrease in input inequality.

Including control functions to address endogeneity does not change the conclusion that parents have a high willingness to pay (WTP) for input equality. Column 2 of Table 2 shows that parents are willing to give up 203 MWK or roughly 0.28 USD in expected household earnings for each 1 ticket decrease in input inequality, somewhat larger than our baseline willingness to pay estimate. The similarity suggests that the primary variation identifying the mixed logit model is the experimental variation.

To fit the trends visible in the raw data more closely, in column 3 we estimate a variant of equation (2) where we replace the continuous $InputInequality$ term with a binary term for whether the allocation equally split inputs. We find that parents' mean WTP to avoid unequal inputs is $1,296$ MWK or roughly 1.81 USD, a substantial amount equal to roughly 92% of the daily adult wage in our setting and 15% of per-child annual education spending. As before, our coefficients of interest remain qualitatively similar after we include the control functions (column 4, Table 2): Parents' mean WTP to avoid unequal inputs is $1,657$ MWK or roughly 2.32 USD.

For ease of interpretation, in column 5 we estimate another version of equation (3) where we model $TotalPay_{ijk}$ as the log of the total expected combined earnings across both children under allocation k. We find that parents' mean willingness to pay (WTP) to avoid unequal inputs is 44% of total expected combined earnings. This willingness-to-pay estimate is remarkably similar to the forgone earnings estimate from our reduced-form analysis, where we find on average parents forwent 40% of their potential earnings across scenarios.

In Figure A.4 we show the distribution of the parent-level coefficients for willingness to pay to equalize inputs: 80% of parents have a positive WTP to equally split inputs, and 65%

---

[28]The default scaling of the logit coefficient is relative to the the variance of the error term.

have a positive WTP to decrease the absolute gap in inputs. Reassuringly, we also show that the parent-level preference to equalize inputs correlates with survey measures capturing inequality in *non-experimental* inputs between children. Parents with higher willingness to pay to equalize inputs, as measured by our experiment, have smaller inequalities between their children in educational expenditures (Appendix Table A.2 Panel A), time spent by the mother on each child (Appendix Table A.2 Panel B), and time spent by the father on each child (Appendix Table A.2 Panel C).

4.3 Robustness Checks

In this section, we present evidence against possible confounds and rule out plausible alternative explanations. First, we discuss why demand effects are unlikely to explain our findings. Second, we rule out alternative explanations that may explain why parents equalize inputs. Finally, we present evidence against other explanations for why parents did not equalize outcomes.

Demand Effects. Our experimental design, in which surveyors observe all of parents' decisions, has the potential for demand effects. The use of real stakes for all choices, the standard approach to address demand effects, helps assuage this concern. Indeed, de Quidt et al. (2018) provide evidence that demand effects are modest with incentivized choices. The stakes in our experiment were also substantial: the average gap between the returns-maximizing and returns-minimizing total payments was roughly 700 MWK (0.98 USD) or 50% of the daily wage for adults in the area. In addition, our experimental results are consistent with data from a different experiment, conducted in the same area in Malawi, where the potential for demand effects was lower (Dizon-Ross, 2019). While the focus of Dizon-Ross (2019) is information frictions and the study was not designed to fully distinguish between the preferences we study here, the data suggest that parents prefer to split inputs under a setup with less scope for surveyor demand effects. In the Dizon-Ross (2019) setting, parents also allocated lottery tickets across their children, with the prize again an educational investment. However, unlike in our experiment, the surveyors did not instruct parents how to returns-maximize, equalize inputs, or equalize outcomes; and yet parents still equalized inputs as much as they could. That experiment also used an even higher-valued prize, and

had other features (such as using an odd number of tickets) that should mitigate demand effects further. The consistency of our results with Dizon-Ross (2019) thus suggests that demand effects do not explain our results. See Appendix C for further discussion.

Alternative Explanations for Equalization of Inputs. We next consider whether uncertainty in beliefs, indifference, lack of understanding, or an attempt to balance returns maximization with inequality aversion over outcomes could drive the equal input allocations we observe in the data. We conclude that none do.

Did parents choose to split their allocations due to beliefs uncertainty? Neoclassical risk aversion through concave utility should not cause splitting in our experiment. However, there could still be behavioral channels through which risk aversion might cause parents to choose split allocations. To address this, we perform a heterogeneity analysis based on baseline measures of uncertainty in parents' beliefs. We fail to find evidence that more uncertain parents equalize more; see Appendix Figure A.6.

We can also provide more direct evidence: After the experiment we asked parents whether they would have allocated differently if they were certain about the scores their children would receive with and without tutoring. Only two out of 289 parents replied in the affirmative. Thus, uncertainty does not seem to be causing split allocations here.

Does indifference explain why parents split their tickets evenly? If parents are completely indifferent about which of their children receives tutoring, they could split their tickets evenly even if they are pure returns-maximizers. However, our results show that 19% of parents equate in all scenarios, which should not reflect indifference: if a given parent is indifferent between her children in the base case, then when we offer 10x higher returns per point to Child H (scenario 2) or Child L (scenario 3), the parent should no longer be indifferent, but many parents continue to equalize. Moreover, indifference is a knife's edge and hence empirically improbable case.

Did parents split their tickets because they did not understand the setup? Section 3.3 already showed that parents do understand the setup. Heterogeneity analysis based on parental education provides additional support for this conclusion. Since one would expect more-educated parents to better understand the design, if evenly splitting tickets represented lack of understanding, we would expect it to be more prevalent among the less-educated.

In contrast, relative to parents with below-median education, we find that parents with above-median education are equally likely to maximize returns, and 11% *more* likely to split their tickets evenly (Appendix Table A.3). The structural analysis yields similar conclusions: more-educated parents are willing to pay MWK 400 more to equalize inputs (Appendix Table A.2, Panel D). Although preferences themselves may vary by parental education, these results suggest that misunderstanding of the setup is not the reason that parents split their tickets in our experiment.

Did parents split their tickets to balance inequality aversion over outcomes with returns-maximization? We rule out this possibility by showing that parents still choose a substantial share of split allocations in scenarios where the returns-maximizing and outcomes-inequality-minimizing allocations are the same. In the *Higher Returns to Child L & Lump Sum to Child H* scenario, both returns maximization and inequality aversion over outcomes dictate that almost all parents should allocate all tickets to Child L. The fact that we still see 33% of parents choosing equal inputs in that scenario provides further evidence that equalizing inputs reflects an aversion to input inequality, not a desire to balance returns-maximization against inequality aversion in inputs.[29]

Figure 10 presents additional evidence from other scenarios, limiting to people with $R_L < R_H$ to have only one ticket allocation prediction per strategy per scenario. The left bar pools all parent × scenario observations where inequality aversion in outcomes and returns maximization have different predictions (and thus splitting tickets could represent a balance between the two forces) while the right bar shows all parent × scenario observations where inequality aversion in outcomes and returns maximization have the same prediction (and thus splitting tickets would not represent a balance between the two). Parents equalize in nearly as many scenarios in the right subfigures as the left subfigures, suggesting that the vast majority of "splitting" represents an aversion to inequality in inputs. We find similar results when we include the full sample of households but limit the scenarios to those in which the predictions are uniform across nearly all households (Scenarios 2, 3 and 5) (Figure A.7).

---

[29]This does not address the potential that parents are instead balancing child-specific preferences with the other desires, but the spike at equal allocation makes that explanation unlikely, as does the fact that we get similar results when excluding parents who appear to have a child-specific preference for Child H.

Figure 10: Many Parents Equalized Inputs Even When Returns Maximization and Inequality Aversion of Outcomes Had the Same Prediction
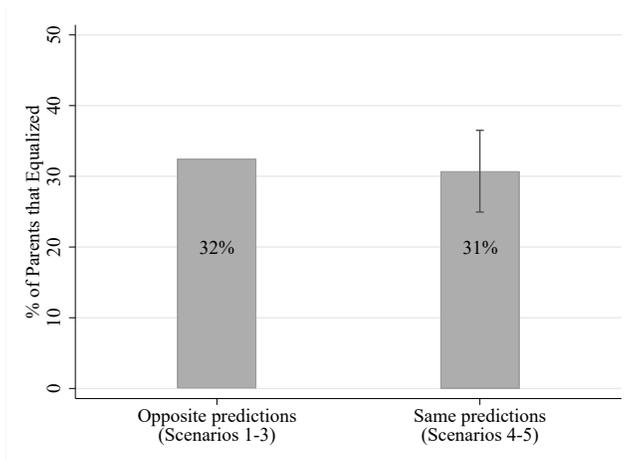
Notes: This figure shows the percentage of parents who equalized tickets across their children, with the left bar depicting the scenarios where inequality aversion over outcomes and returns maximization have opposite predictions (Scenarios 1, 2 and 3) and the right bar depicting the scenarios where IAO and RM have the same prediction (Scenarios 4 and 5). The figure restricts the sample to the 66% of households who believed tutoring yields higher test score gains for Child H ($R_L < R_H$) since the predictions for inequality aversion over outcomes and returns maximization are uniform across that sample. The bracket on the right bar depicts the 95% confidence interval for a test that the difference between the bars is 0; the heights of the two bars are not statistically different from each other.

Absence of Outcomes Equalization. Finally, we discuss several alternative explanations (aside from parents not being inequality averse over outcomes) for why parents did not equalize outcomes within our experiment. Our results are driven neither by parents reallocating earnings after the experiment, by parents equalizing outcomes across scenarios, nor by the fact that parents cannnot perfectly equalize within scenarios.

Did parents not equalize outcomes because they could reallocate children's earnings after the experiment? Our experiment uses cash rewards, which parents could potentially equalize *ex post*. However, if parents were planning to equalize these rewards *ex post*, the natural response would be to maximize returns *ex ante*, thus maximizing the total rewards available for *ex post* reallocation, which is not what we see.

To further address this possibility, we conducted a follow-up survey between August and October 2019 in which we separately interviewed parents and children from our original sample. We located and surveyed 259 (out of 289) parents and 392 (out of 578) children (Table A.4). We asked both children and parents about whether the parents reallocated children's earnings after the experiment. 77% of children with positive earnings indicated

36

that they kept their full earnings; the reports from parents were similar. Importantly, only 2% of children said that their parents gave their earnings to their sibling, and only 5% of parents said that they took earnings from one of their children because they wanted to *ex post* equalize outcomes. Further evidence against *ex post* reallocation comes from the fact that we fail to find significant evidence that parents took earnings from their higher-earning children more often than their lower-earning children. Parents report that they allowed similar shares of high and low earnings to keep their earnings (74% of high earners and 79% of low earners), and the difference is not statistically significant (p-value = 0.19). Overall, these data show that the potential for reallocation is an unlikely explanation for why parents did not equalize expected outcomes in the experiment.[30]

Did parents equalize outcomes across scenarios rather than within scenarios? To test this potential explanation for why we see no outcomes equalization within scenario, after we elicited all ticket allocations, we asked parents if they had equalized outcomes across scenarios. Reassuringly, only 19% of parents said they equalized outcomes across scenarios, and our results look very similar when we exclude these respondents.

Did parents not equalize outcomes because they could not perfectly do so? If parents only care about equalizing outcomes when they can *perfectly* do so (i.e., if their utility term governing inequality aversion over outcomes is $-\alpha\mathbb{1}\{ER(x_1) \neq ER(x_2)\}$), it could bias us away from finding evidence of inequality aversion in outcomes because, in many of our scenarios, it was not possible for parents to perfectly equalize expected outcomes. To test this idea, we present summary statistics from Scenario 1 (the scenario which had the majority of cases in which parents could perfectly equalize outcomes), separately by whether the parent had the option to perfectly equalize experimental earnings. Although the option to perfectly equalize depends on potentially endogenous variation in parents' beliefs about their children's scores, we view the analysis as suggestive. Appendix Table A.5 shows that the percentage of choices in which parents minimize outcomes inequality is not significantly different across parent types and, if anything, is smaller among parents who had the option to perfectly equalize. Thus, not being able to perfectly equalize appears not to explain our

---

[30]Similarly, in the follow-up survey we also asked parents if after the experiment they treated either of their children differently than usual, either by giving them more or less of something than usual. Only 9 (8%) out of the 110 parents we posed this question to replied in the affirmative.

findings.

5   Conclusion

Our experiment provides the first evidence that parents have a quantitatively important preference for equalizing the inputs they invest in their children. In order to identify parents' preferences, we experimentally shock the short-run returns to educational investments to identify the degree to which parents care about (a) maximizing total returns, (b) minimizing cross-sibling inequality in outcomes, and (c) minimizing cross-sibling inequality in inputs. We find that parents care about both maximizing returns and minimizing inequality in inputs, but find no evidence for aversion to inequality in child-level outcomes or earnings. Parents' aversion to inequality in inputs is quantitatively important, causing parents to forgo roughly 40% of their potential experimental earnings.

Because the effects of policies depend on individuals' behavioral responses, our estimates have important implications for policy design. In particular, when policies target children, parents' preferences for investment can create important spillovers to non-targeted children within the household. Our results imply, for example, that parents may respond to policies that deliver inputs to one child in the household by providing more inputs to non-targeted children to try to mitigate the input inequality. The impacts are likely to depend on whether parents know what level of inputs the policy is providing, and thus giving information to parents about the inputs provided may encourage parental spending on non-targeted children. These results have relevance to numerous policies that target a subset of children within the household, such as programs based on academic achievement, gender, or age.

Our findings open up several directions for future work. One direction is to investigate the implications for policy design by using the preference parameters estimated here to develop and evaluate "optimized" policies that account for parents' behavioral responses. Another important area is to examine the extent to which the preferences we identify here over investments with short-run returns extend to investments with larger, longer-term returns, such as long-run schooling decisions. One specific issue to investigate is whether parents who are averse to inequality in *expected* inputs also prefer to equalize actual inputs when inputs are divisible. Another issue to explore is how the cost of choosing unequal

inputs scales with the stakes of the decision. If the cost is fixed, then inequality aversion would be more important for small-stakes than large-stakes decisions. In contrast, if the cost scales in proportion with the returns, then inequality aversion would be important for large-stakes decisions as well. In that case, our structural estimates imply that parents are willing to give up 44% of total investment returns to equalize, a huge cost for large-stakes decisions. A final area for future work is to determine the frame or bracket within which parents normally equalize. Do they equalize within short or long time horizons, or within specific investment domains (e.g., health vs. education) or across all investments? In our experiment, parents narrowly bracket, and Exley and Kessler (2019) suggest that narrow bracketing of equity concerns is a widespread phenomenon. The more narrowly parents bracket outside of experimental settings, the higher the efficiency cost of equalizing inputs is likely to be.

References

Agrawal, P. (2002). Incentives, Risk, and Agency Costs in the Choice of Contractual Arrangements in Agriculture. *Review of Development Economics 6*(3), 460–477.

Almond, D. and B. Mazumder (2013). Fetal Origins and Parental Responses. *Annual Review of Economics 5*(1), 37–56.

Andreoni, J., D. Aydin, B. Barton, B. D. Bernheim, and J. Naecker (2018). When Fair Isn't Fair: Understanding Choice Reversals Involving Social Preferences. *NBER Working Paper 25257*.

Andreoni, J. and B. D. Bernheim (2009). Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica 77*(5), 1607–1636.

Banerji, R., J. Berry, and M. Shotland (2017). The impact of maternal literacy and participation programs: Evidence from a randomized evaluation in india. *American Economic Journal: Applied Economics 9*(4), 303–337.

Behrman, J. R., R. A. Pollak, and P. Taubman (1982). Parental Preferences and Provision for Progeny. *Journal of Political Economy 90*(1), 52–73.

Behrman, J. R., R. A. Pollak, and P. Taubman (1986). Do parents favor boys? *International Economic Review 27*(1), 33–54.

Bernheim, B. D. and S. Severinov (2003). Bequests as Signals: An Explanation for the Equal Division Puzzle. *Journal of Political Economy 111*(4), 733–764.

Bolton, G. E. and A. Ockenfels (2000). ERC: A Theory of Equity, Reciprocity, and Vompetition. *American Economic Review 90*(1), 166–193.

Brock, J. M., A. Lange, and E. Y. Ozbay (2013). Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review 103*(1), 413–437.

Cappelen, A. W., K. O. Moene, E. Ø. Sørensen, and B. Tungodden (2013). Just Luck : An Experimental Study of Risk Taking and Fairness. *American Economic Review 124*(4), 1398–1413.

Charness, G. and M. Rabin (2002). Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics 117*(3), 817–869.

Cherry, T. L., P. Frykblom, and J. F. Shogren (2002). Hardnose the Dictator. *The American Economic Review 92*(4), 1218–1221.

Cunha, F. and J. Heckman (2007). The Technology of Skill Formation. *American Economic Review Papers and Proceedings 97*(2), 31–47.

Cunha, F., J. J. Heckman, L. Lochner, and D. V. Masterov (2006). Interpreting the Evidence on Life Cycle Skill Formation. *Handbook of the Economics of Education 1*(6), 697–812.

Dasgupta, S. and Z. Tao (1998). Contractual Incompleteness and the Optimality of Equity Joint Ventures. *Journal of Economic Behavior and Organization 37*(4), 391–413.

de Quidt, J., J. Haushofer, and C. Roth (2018). Measuring and bounding experimenter demand. *American Economic Review 108*(11), 3266–3302.

Dizon-Ross, R. (2019). Parents' beliefs about their children's academic ability: Implications for educational investments. *American Economic Review 109*(8), 2728–2765.

Exley, C. L. and J. B. Kessler (2019). Equity concerns are narrowly framed. *NBER Working Paper 25326*.

Fahr, R. and B. Irlenbusch (2000). Fairness as a constraint on trust in reciprocity: Earned property rights in a reciprocal exchange experiment. *Economics Letters 66*(3), 275–282.

Farmer, A. and J. Tiefenthaler (1995). Fairness concepts and the intrahousehold allocation of resources. *Journal of Development Economics 47*(2), 179–189.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics 114*(3), 817–868.

Forsythe, R., J. L. Horowitz, N. E. Savin, and M. Sefton (1994). Fairness in simple bargaining experiments.

Hertwig, R., J. N. Davis, and F. J. Sulloway (2002). Parental investment: how an equity motive can produce inequality. *Psychological bulletin 128*(5), 728.

Jakiela, P. (2015). How Fair Shares Compare: Experimental Evidence from Two Cultures. *Journal of Economic Behavior and Organization 118*, 40–54.

Kahneman, D., A. Tversky, B. Y. D. Kahneman, and A. Tversky (2016). Prospect Theory: An Analysis of Decision under Risk Linked. *Econometrica 47*(2), 263–292.

Konow, J. (2003). Which is the fairest one of all? a positive analysis of justice theories. *Journal of Economic Literature 41*(4), 1188–1239.

Krawczyk, M. and F. Le Lec (2010). 'Give me a chance!' an experiment in social decision under risk. *Experimental Economics 13*, 500–511.

Machina, M. J. (1989). Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty. *Journal of Economic Literature 27*(4), 1622–1668.

Menchik, P. L. (1980). Primogeniture, Equal Sharing, and the U.S. Distribution of Wealth. *The Quarterly Journal of Economics 94*(2), 299–316.

Petrin, A. and K. Train (2010). Control function approach to endogeneity in consumer choice models. *Journal of Marketing Research 47*, 2–13.

Pitt, M. M., M. R. Rosenzweig, and M. N. Hassan (1990). Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries. *The American Economic Review 80*(5), 1139–1156.

Price, J. (2008). Parent-child quality time does birth order matter? *Journal of human resources 43*(1), 240–265.

Trivers, R. L. (1974). Parent-offspring conflict. *Integrative and Comparative Biology 14*(1), 249–264.

Wilhelm, B. M. (1996). Bequest Behavior and the Effect of Heirs' Earnings: Testing the Altruistic Model of Bequests. *American Economic Review 86*(4), 874–892.

World Bank (2020). World Bank Open Data. https://data.worldbank.org/. Accessed January 10, 2020.
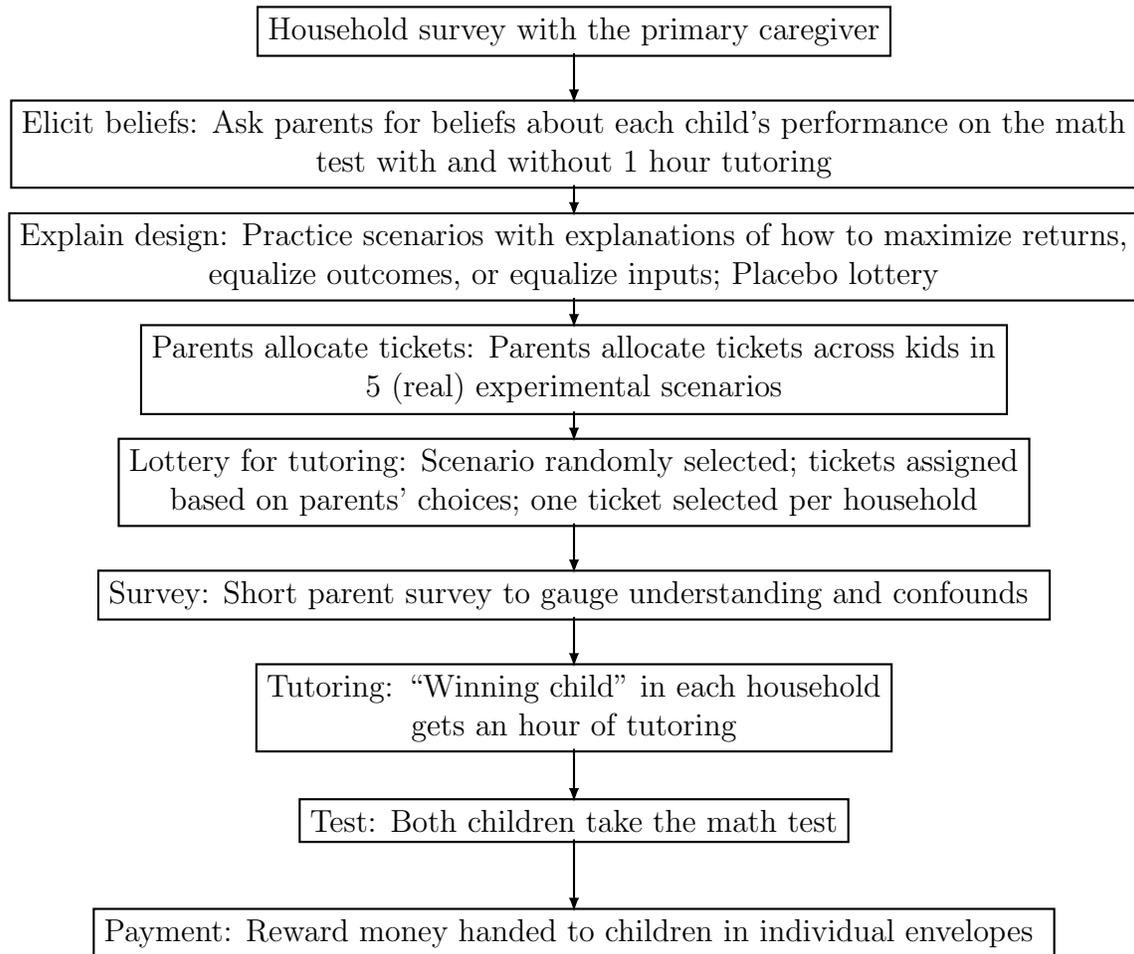
A   Appendix Tables and Figures

Appendix Figure A.1: The Share of Baseline Educational Expenditures on the Perceived Lower-Performing Child Has a Spike at 50%
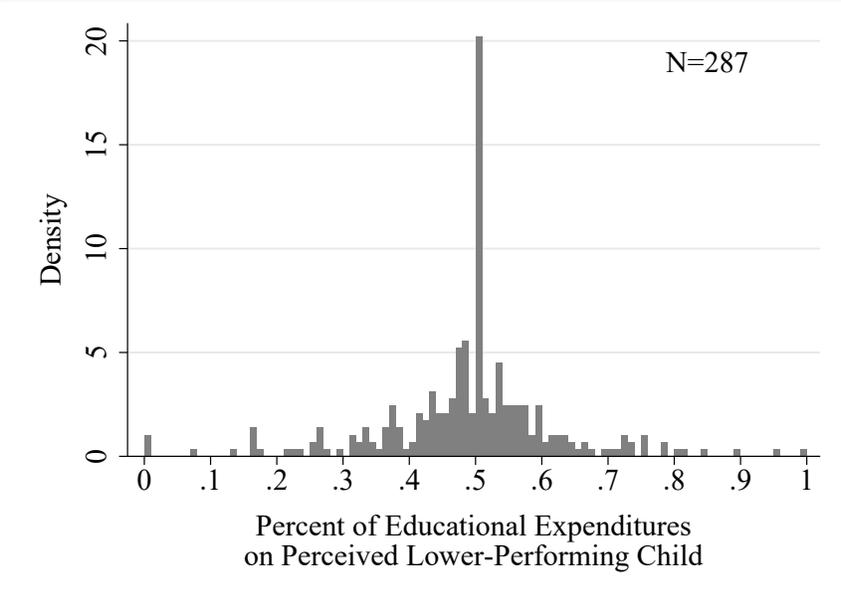

Notes: This figure shows the distribution of the percent of educational expenditures on the perceived lower-performing child, as a share of the total educational expenditures on both sampled children in the study sample.

Appendix Figure A.2: Sequence of Events

Household survey with the primary caregiver

↓

Elicit beliefs: Ask parents for beliefs about each child's performance on the math test with and without 1 hour tutoring

↓

Explain design: Practice scenarios with explanations of how to maximize returns, equalize outcomes, or equalize inputs; Placebo lottery

↓

Parents allocate tickets: Parents allocate tickets across kids in 5 (real) experimental scenarios

↓

Lottery for tutoring: Scenario randomly selected; tickets assigned based on parents' choices; one ticket selected per household

↓

Survey: Short parent survey to gauge understanding and confounds

↓

Tutoring: "Winning child" in each household gets an hour of tutoring

↓

Test: Both children take the math test

↓

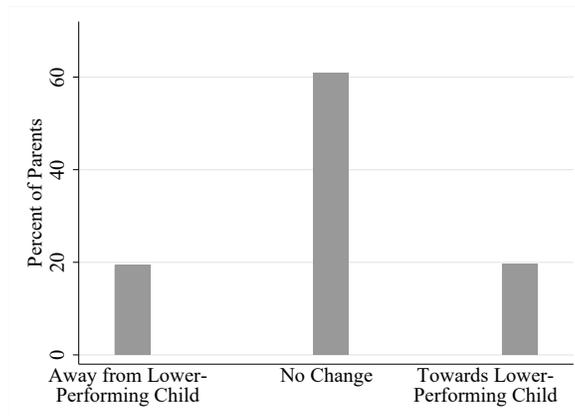Payment: Reward money handed to children in individual envelopes

Appendix Figure A.3: Inconsistent With Inequality Aversion in Outcomes, When One Child Received a Lump Sum, Parents Did Not Reallocate to His/Her Sibling.
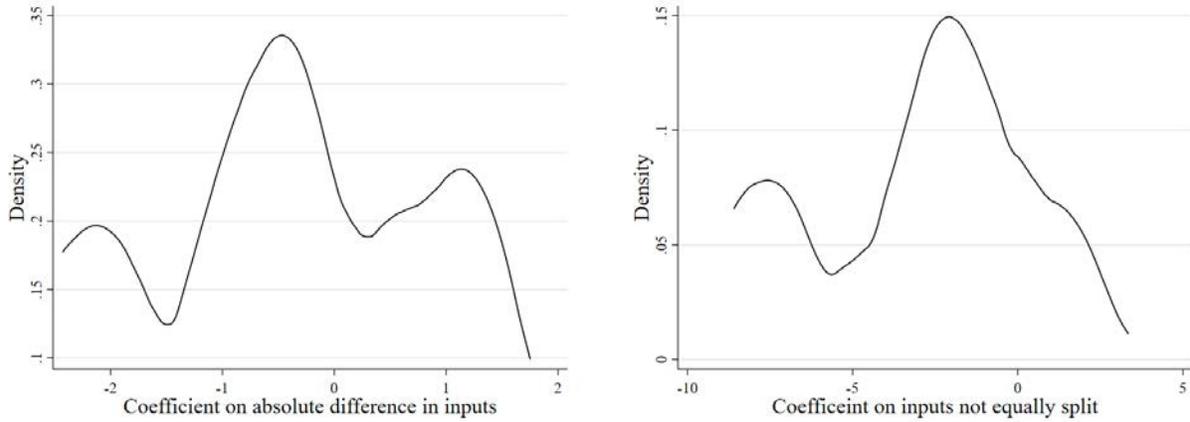
Individual Parent-level Changes
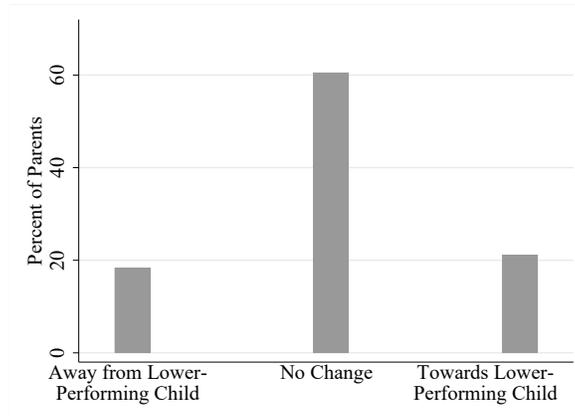
(a) Parent-level Changes From 1. *Base Case* to 4. *Lump Sum to Child L*

(b) Parent-level Changes From 5. *Higher Returns to Child L & Lump Sum to Child H* to 4. *Higher Returns to Child L*

Appendix Figure A.4: Distribution of Parents' Preferences for Equalizing Inputs


(a) Distribution of parents' preferences to decrease absolute gap in inputs

(b) Distribution of parents' preferences to equally split inputs

Note: Panel (a) presents the distribution of the coefficient on "Absolute difference in inputs" from Column 1 of Table 2. Panel (a) presents the distribution of the coefficient on "Inputs not equally split" from Column 3 of Table 2. We use an Epanechnikov kernel function, with bandwidths of 0.36 and 0.92 for Panels (a) and (b), respectively.

Appendix Figure A.5: Evidence for Child-Specific Preferences

Parents Who Gave More Tickets to Child H in Scenario 1 Continue to Allocate More to Child H in Scenarios 2-5
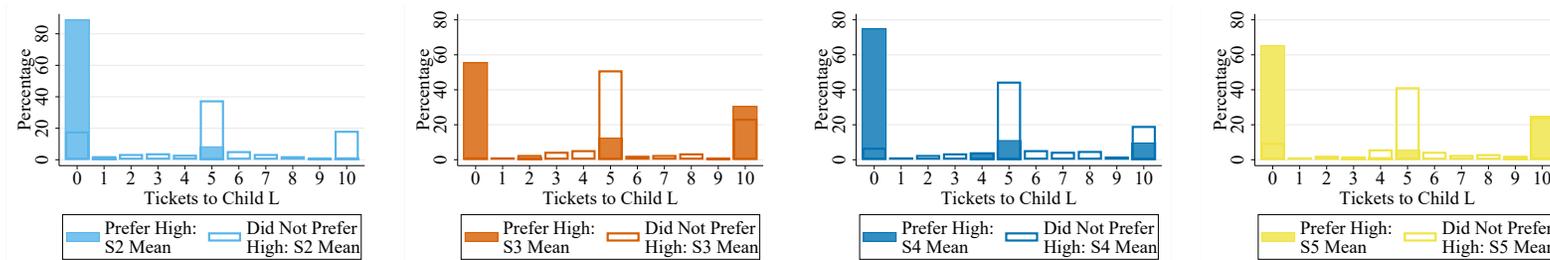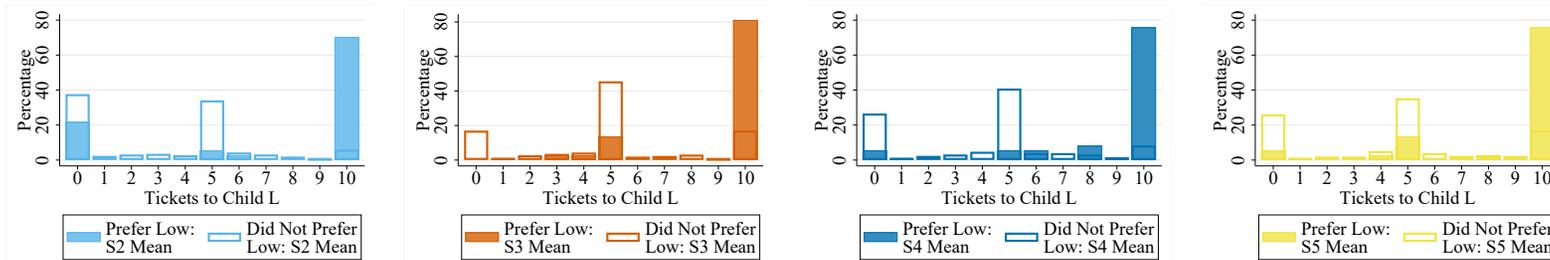

(a) Scenario 2 *(Higher Returns to Child H)*

(b) Scenario 3 *(Higher Returns to Child L)*

(c) Scenario 4 *(Lump Sum to Child L)*

(d) Scenario 5 *(Higher Returns to Child L & Lump Sum to Child H)*

Parents Who Gave More Tickets to Child L in Scenario 1 Continue to Allocate More to Child H in Scenarios 2-5

(e) Scenario 2 *(Higher Returns to Child H)*

(f) Scenario 3 *(Higher Returns to Child L)*

(g) Scenario 4 *(Lump Sum to Child L)*

(h) Scenario 5 *(Higher Returns to Child L & Lump Sum to Child H)*

Notes: This figure presents allocation of lottery tickets across children for Scenarios 2-5. Panels (a) through (d) present allocations separately for parents who gave strictly more tickets to Child H in Scenario 1 (2,387 parents, labeled as "Prefer High" in the figures) and those who did not (792 parents, labeled as "Did Not Prefer High" in the figures). Panels (e) through (h) present allocations separately for parents who gave strictly more tickets to Child L in Scenario 1 (2,772 parents, labeled as "Prefer Low" in the figures) and those who did not (407 parents, labeled as "Did Not Prefer Low" in the figures).

Appendix Figure A.6: More Uncertain Parents Did Not Equalize Inputs More

*Heterogeneity in allocations by parents' baseline measure of uncertainty*
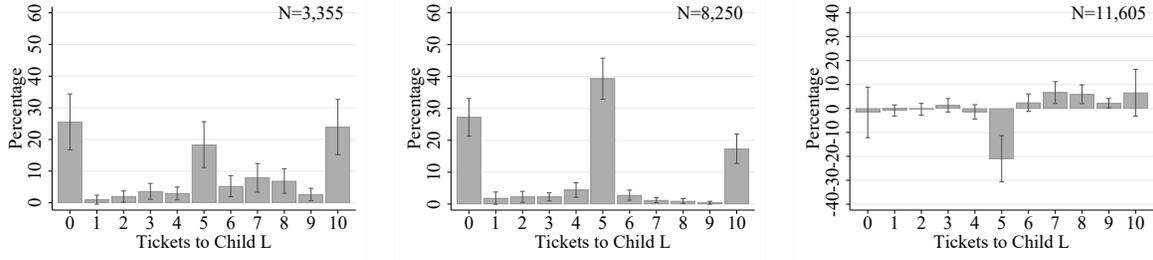
(a) Parents who are uncertain about their beliefs

(b) Parents who are certain about their beliefs

(c) Difference in means (uncertain - certain)

Heterogeneity in allocations by whether parents changed beliefs between baseline survey and experiment
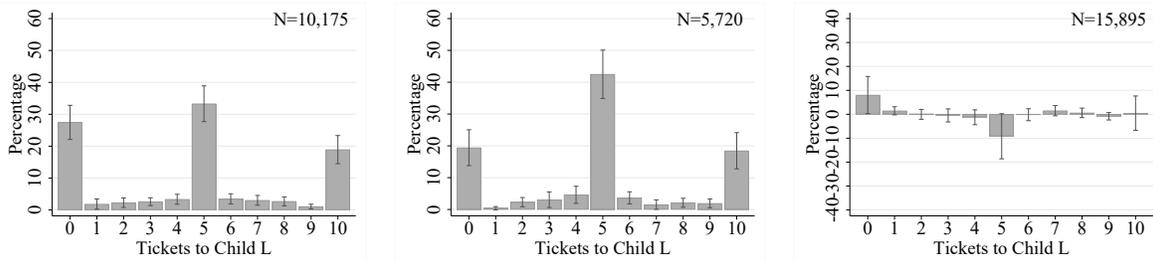
(d) Parents who changed their beliefs

(e) Parents who did not change their beliefs

(f) Difference in means (changed - did not change)

Heterogeneity in allocations by difference between parents' beliefs and actual scores
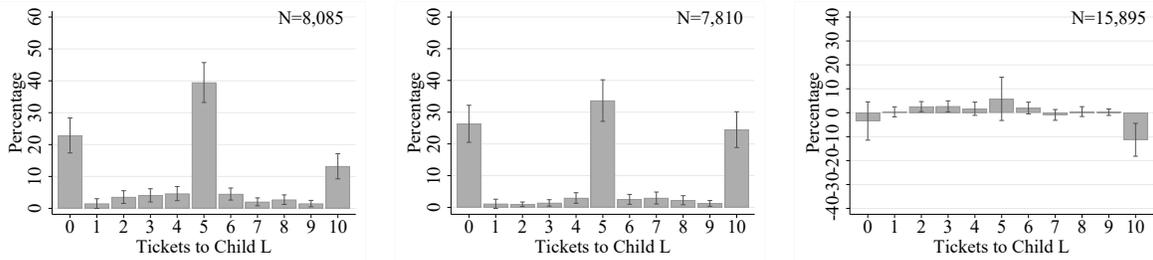
(g) Difference between beliefs and actual scores above median

(h) Difference between beliefs and actual scores below median

(i) Difference in means (Difference in means (above-median - below-median)

Notes: This figure presents allocation of lottery tickets across children for Scenarios 2-5, and 95% confidence intervals for each scenario, separately by beliefs uncertainty. We use three measures of uncertainty. Panels (a)-(c) use parents' stated uncertainty about their beliefs during the baseline survey. However, since we allowed parents to adjust their beliefs at the experimental visit but only measured beliefs uncertainty in the baseline survey baseline, this measure may not perfectly capture uncertainty at the time of the experimental allocations. We thus use two additional proxies for uncertainty: Whether parents changed their beliefs between the baseline survey and experimental visit, and whether the absolute value of the gap between the parents' beliefs and their children's true scores was above-median.

47

Appendix Figure A.7: Many Parents Equalize Inputs Even When Returns Maximization and Inequality Aversion of Outcomes Have the Same Prediction
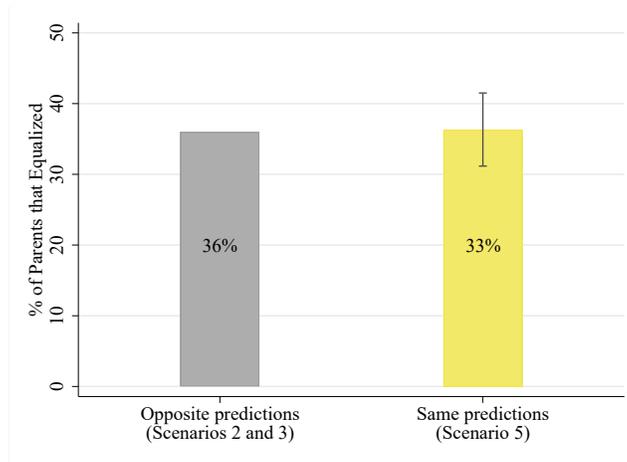
Notes: These figures presents the percentage of parents who equalized tickets across their two children. It only displays the scenarios for which, for at least 94% of the sample, there is only one ticket allocation prediction per strategy. The left bar show all scenarios where inequality aversion over outcomes (IAO) and returns maximization (RM) have opposite predictions (Scenarios 2 and 3). The right bar show the scenario where inequality aversion over outcomes (IAO) and returns maximization (RM) have the same prediction (Scenario 5). The bracket on the right bar shows that the two bars are not statistically different from each other.

Appendix Table A.1: Parents' Stated Reasons Predict Their Actual Choices

	(1) Equalize Inputs	(2) Maximize Returns	(3) Equalize Outcomes
Prefer Equalizing Inputs	0.79***	-0.14***	-0.17***
	(0.05)	(0.04)	(0.04)
Prefer Maximizing Returns	-0.13***	0.36***	-0.16***
	(0.05)	(0.06)	(0.04)
Prefer Equalizing Outcomes	0.08	-0.03	0.00
	(0.09)	(0.06)	(0.06)
Prefer Child L	0.02	-0.02	0.03
	(0.05)	(0.05)	(0.04)
Prefer Child H	0.02	-0.02	-0.04
	(0.05)	(0.05)	(0.04)
Scenario FE	Yes	Yes	Yes
Household FE	Yes	Yes	Yes
Observations	1445	1445	1445
R^2	0.886	0.468	0.298

Notes: The dependent variables represent dummies for the actual choices parents made in a given scenario. The independent variables represent the parents' stated reasons for making those choices. Standard errors in parentheses, clustered at the household level. * denotes significance at 0.10; ** at 0.05; and *** at 0.01.

Appendix Table A.2: Estimated Preference for Equal Inputs Correlates with More Equal Allocations of Spending and Time and Higher Parental Education

	(1) Coefficient on Absolute Gap in Inputs	(2) Coefficient on Inputs Not Equally Split
A. Household Expenditure		
Above-Median Abs. Gap in Expenditures	0.29**	0.61
	(0.14)	(0.37)
Constant	-0.55***	-3.27***
	(0.10)	(0.28)
Observations	288	288
R^2	0.013	0.009
B. Mother's Time Use		
Mother's Time Not Equally Split	0.32**	0.90**
	(0.16)	(0.40)
Constant	-0.52***	-3.35***
	(0.10)	(0.27)
Observations	251	251
R^2	0.015	0.019
C. Father's Time Use		
Father's Time Not Equally Split	0.39**	0.99**
	(0.19)	(0.49)
Constant	-0.59***	-3.39***
	(0.12)	(0.33)
Observations	175	175
R^2	0.022	0.022
D. Parental Education		
Grade 8 Completed	-0.31**	-0.92**
	(0.16)	(0.40)
Constant	-0.30***	-2.66***
	(0.09)	(0.23)
Observations	289	289
R^2	0.014	0.019

Notes: Panel A presents heterogeneity by the difference in children's share of household expenditure, which is measured by a binary variable that takes value 1 if the absolute gap in expenditure between children is above the median. Panel B presents heterogeneity by the difference in children's share of mother's time, which is measured by a binary variable that takes value 1 if mother's time is not equally split. Panel C presents Panel A presents heterogeneity by the difference in children's share of father's time, which is measured by a binary variable that takes value 1 if father's time is not equally split. Panel D presents heterogeneity by parents' education, captured by a binary variable that takes the value 1 if the primary caregiver has at least completed Grade 8. Standard errors are in parentheses, clustered at the household level. * denotes significance at 0.10; ** at 0.05; and *** at 0.01.   50

Appendix Table A.3: More Educated Parents Are More Likely to Equalize Inputs

	(1) Child L Tickets	(2) Equalize Inputs	(3) Maximize Returns	(4) Equalize Outcomes
Caregiver > Class 8	0.42	0.11**	-0.02	-0.01
	(0.31)	(0.05)	(0.04)	(0.03)
Constant	4.57***	0.33***	0.28***	0.13***
	(0.19)	(0.03)	(0.02)	(0.01)
Observations	1445	1445	867	867
R^2	0.003	0.012	0.000	0.000

Notes: Regression is at the parent × scenario level. Table shows heterogeneity in parents' choices by parental education. "Caregiver > Class 8" is a binary variable that takes the value 1 if the primary caregiver has at least completed Grade 8; 0 otherwise. The dependent variables represent the parents' choices (in particular, for columns (1)-(4), the dependent variables represent the number of tickets given to Child L, a dummy for whether the parent equalized inputs, a dummy for whether the parent chose the returns-maximizing strategy for experimental earnings, and a dummy for whether the parent chose the outcomes-equalizing strategy for experimental earnings). Standard errors in parentheses, clustered at the household level. * denotes significance at 0.10; ** at 0.05; and *** at 0.01.

Appendix Table A.4: Most Children Kept Their Earnings After the Experiment

Total Households			289		
Total Parents Surveyed			259		
Total Children Surveyed			392		

A. Parent Survey	Child L	Child H	High Earner	Low Earner	Equal Earner
Children with non-zero earnings	135	163	210	84	4
Children who kept earnings	111 (82%)	113 (69%)	154 (74%)	66 (79%)	4 (100%)
Children who did not get to keep earnings	24 (18%)	50 (31%)	56 (26%)	18 (21%)	0 (0%)
Reasons children did not keep earnings					
Parents didn't trust child with money	8 (6%)	16 (10%)	20 (9%)	4 (4%)	
Parents wanted to equalize outcomes	5 (4%)	4 (3%)	5 (2%)	4 (5%)	
Other reasons	11 (8%)	30 (18%)	31 (15%)	10 (12%)	

B. Child Survey	Child L	Child H	High Earner	Low Earner	Equal Earner
Children surveyed	190	202	193	196	3
Children with non-zero earnings	125	150	193	79	3
Children who kept earnings	94 (75%)	116 (77%)	143 (74%)	64 (81%)	3 (100%)
Children who did not get to keep earnings	31 (25%)	34 (23%)	50 (26%)	15 (19%)	0 (0%)
Reasons children did not keep earnings					
Parents took earnings	6 (5%)	8 (5%)	10 (5%)	4 (5%)	
Parents gave it to sibling	2 (2%)	2 (2%)	4 (2%)	0 (0%)	
Children with other reasons	23 (18%)	24 (16%)	36 (19%)	11 (14%)	

Notes: This table presents parents' and children's responses from the follow-up survey where they were asked questions about reallocation of children's earnings after the experiment.

Appendix Table A.5: No Evidence of Aversion to Inequality in Outcomes Even Among Parents Who Can Perfectly Equalize Outcomes

Scenario 1 (Base Case) choices by whether parents could perfectly equalize outcomes

	Whether parents' IAO choice perfectly equalized outcomes		
	Yes	No	P-value (Yes=No)
	(1)	(2)	(3)
IAI (% of parents)	0.47	0.39	0.20
	(0.50)	(0.49)	
RM (% of parents)	0.25	0.21	0.46
	(0.43)	(0.41)	
IAO (% of parents)	0.08	0.11	0.29
	(0.27)	(0.32)	
Observations	106	183	

Notes: This table presents the proportion of parents who equalized inputs, maximized returns, and equalized outcomes for Scenario 1, summarized separately by whether the parent had the option to perfectly equalize outcomes. The P-value reported in column 3 tests for a difference in means between columns 1 and 2. There were two parents for which the returns-maximizing and outcomes-equalizing allocations were the same — these parents were categorized in the "equalized outcomes" category.

B   Section 2 Proofs

In Section 2, we posit that, if parents do not care about inequality aversion (i.e., if $\alpha = 0$ and $\beta = 0$), then parents should only choose all-or-nothing allocations of inputs. In this section, we first show this idea more rigorously when parents have expected-utility preferences. We then explore the robustness of the prediction even when parents have other types of preferences such as prospect theory preferences.

Expected Utility

To see that parents with expected-utility preferences and no inequality aversion ($\alpha = 0, \beta = 0$) would choose all-or-nothing allocations, we first rewrite equation (1) in the case where $\alpha = 0, \beta = 0$, and the first term is a nonlinear expected-utility term:

$$U(x_1, x_2) = \lambda Eu\left(R_1(x_1) + R_2(x_2)\right) + \gamma(z_1, z_2)\left(x_1 - x_2\right) \tag{4}$$

Now define t_i as the fraction of tickets allocated by the parent to child i and rewrite equation (4) as a function of the new choice variable t_i:

$$U(t_1, t_2) = \lambda\left[t_1 Eu\left(R_1(1) + R_2(0)\right) + t_2 Eu\left(R_1(0) + R_2(1)\right)\right] + \gamma(z_1, z_2)\left(t_1 - t_2\right) \tag{5}$$

The first returns-maximizing term represents a weighted sum of the parent's expected utility if child 1 received the binary investment and her expected utility if child 2 received the binary investment, weighted by the probability that each event occurs, t_1 and t_2. Utility is thus linear in t_1 and t_2, and hence parents will in general choose "all-or-nothing" corner solutions, setting either t_1 or t_2 to 0. The only case in which parents with utility function (5) would not choose corner solutions is if they are exactly indifferent between which child receives the tutoring (i.e., $Eu\left(R_1(1) + R_2(0)\right) + \gamma(z_1, z_2) = Eu\left(R_1(0) + R_2(1)\right) - \gamma(z_1, z_2)$) which is a knife's edge case that is unlikely to hold in practice.

Non-Expected-Utility Preferences (e.g., Prospect Theory)

Define $\pi(p)$ as the "probability weighting" function or the *actual* probability weight a parent places on an event whose true probabililty of occurrence is p. A sufficient condition for the prediction that parents who are not averse to inequality will not choose split allocations to hold is that $\pi(p) + \pi(1 - p) \leq \pi(0) + \pi(1)$. Intuitively, this condition means that parents understand that choosing split allocations does not increase their chance of winning the lottery.

Most non-expected-utility preferences that the literature has posited will satisfy this condition. For example, if people prefer certainty (as prospect theory suggests), which

means that $\pi(1)$ is particularly large, then the prediction should go through. If people have probability weighting functions that are concave at the bottom and convex at the top (as again suggested by prospect theory) then that is also fine for the prediction as long as the function is rotationally symmetric. In our specific setting, we show that the $\pi(p) + \pi(1-p) \leq \pi(0) + \pi(1)$ condition likely holds by conducting a "placebo lottery" allocation, which suggests that parents likely understand that choosing split allocations does not increase their chances of winning the lottery.

C Using Data From A Different Experiment To Address Potential Confounds

Here, we bring in data from a different experiment conducted in the same area in Malawi in 2012 (Dizon-Ross, 2019). The data come from the control group of the experiment, so we can think of the data as representing "baseline" allocations. These data describe the results of a lottery similar to the one conducted here: parents allocated lottery tickets between two of their children, and the child whose ticket was chosen won a prize. However, there are several differences between the Dizon-Ross (2019) lottery and ours.

First, the prize was something parents could not *ex post* equalize. Instead, it was a scholarship to secondary school. Most parents in this setting could not afford secondary school on their own. The earnings return to secondary school would also not be realized for many years, until the children were out of the house and adults in the labor market (at the time, the children were in 2nd through 7th grade). Thus, this is a setting where parents could not *ex post* equalize any more easily than they could in non-experimental settings.

Second, the Dizon-Ross (2019) experiment had several features that decreased the potential for demand effects. Unlike in our experiment, the enumerator script for the Dizon-Ross (2019) experiment did not instruct parents how to returns-maximize, equalize inputs, or equalize outcomes. In addition, the total number of tickets was an odd number (9).[31] Thus, it is unlikely that parents felt pressured by enumerators to choose equal allocations, since they were not capable of doing so. However, we can still use the data from that experiment to test for inequality aversion in inputs. In particular, a parent inequality averse in inputs should choose the most equal allocation—4 to one child and 5 to the other. In contrast, a parent who does not care about equality should still assign all tickets in all-or-nothing fashion. As a result, if we see excess mass at the 4/5 choice, that provides evidence of inequality aversion in inputs, as returns-maximization would be all-or-nothing and inequality aversion over outcomes would be smooth through the 4/5 point. Finally, the value of the Dizon-Ross (2019) prize was also larger than ours.

Figure C.1(a) shows the absolute value of the gap in tickets between children from the Dizon-Ross (2019) experiment. Seventy-five percent of parents chose as equal an allocation as possible, whereas only 12% chose an all-or-nothing allocation. Panel (b) shows the number of tickets that were given to the child perceived as lower-performing. Among those who chose the most equal allocation, 4 times as many chose to allocate 1 more ticket to their higher-performing child than lower-performing child, suggesting that, again, in that setting,

---

[31]That study used an odd number because, when piloting for the experiment, we found that parents primarily chose equal allocations when given an even number of tickets. For that project, allowing equal allocation was not desirable.

parents were not truly indifferent or did not misunderstand the returns.[32] Rather, they were simply averse to splitting inputs unequally.

Appendix Figure C.1: Many Parents Also Equalize Inputs As Much As Possible In Another Setting
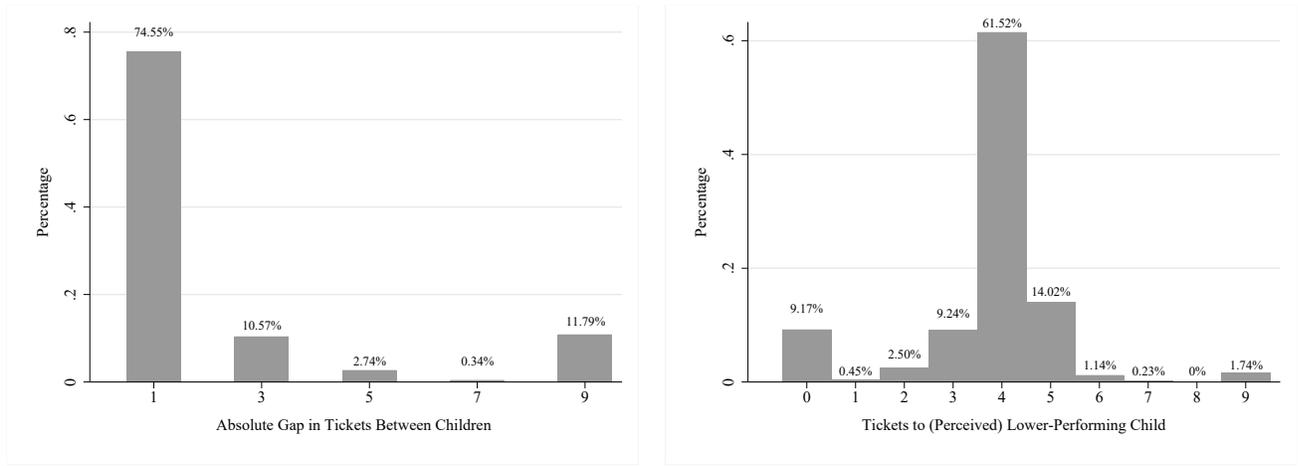
(a) Absolute Gap in Tickets   (b) Tickets Allocated to Lower-Performing Child

Notes: Data from the control group from the Dizon-Ross (2019) experiment. Panel A shows the distribution of the absolute gap between the number of tickets allocated to a parents' two children in a setting where parents were asked to allocate 9 tickets between their two children. Panel B shows the number of tickets allocated to the child the parent perceived was lower-performing child. Here, one out of every 100 households was randomly selected and the child whose name was on the selected ticket received a scholarship for four years of government school fees. In the cases where parents believed both children performed equally, we randomly select which child is designated as the "lower-perfoming child."

---

[32]In that setting, the vast majority of parents believed the earnings return to the lottery prize would be higher for higher-performing children.

D   Additional Detail on Experimental Scenarios

In the experiment, the scenarios were presented to respondents in the following order:

1. Base Case
2. Lump Sum to Child L
3. Higher Returns to Child H
4. Higher-Returns to Child L & Lump Sum to Child H
5. Higher Returns to Child L

This order was optimized during piloting to facilitate respondent understanding of the payment functions. For example, this order of payment functions changes the lump sum parameter B and the slope parameter C separately before changing them together.

The remainder of this appendix presents the script and visual aids that surveyors used to explain Scenario 1 (the Base Case) to respondents.

NOTE TO READER: "Child A" is referred to in the text as "Child L" or the "lower-performing child." "Child B" is referred to in the text as "Child H" or the "higher-performing child."

Scenario 1 Script

OVERVIEW

Here's your first scenario. With this scenario, both children get 10 MWK for every point scored over 40 on the test.

	Scenario 1					
	Beliefs w/o T	Beliefs w T	Scenario	Payoff w/o T	Payoff w T	# Tickets o/f 10
Child A	50	60	10*(TS-40)	100	200	
Child B	70	90	10*(TS-40)	300	500	

So, if Child A gets 50 points and Child B gets 70 points, with this scenario, Child A would get a reward worth (50-40) points X 10 MWK per point = 100 MWK, and Child B would get a reward worth (70-40) points X 10 MWK per point = 300 MWK. So, the expected reward for each child depends on the score they receive, but with this scenario, both children get 10 MWK for each point above 40 scored.

Without tutoring, you expected Child A to score 50 on the test; if they do in fact score 50, then Child A would get a prize worth 10*(50-40) = 100 MWK. With tutoring, you expected Child A to get a score of 60. If he/she did score 60, he/she will receive a prize worth 10*(60-40) = 200 MWK. So, then the more tickets you give to Child A, the higher chance you move them from a prize worth 100 MWK to a prize worth 200 MWK.

Similarly, without tutoring, you expected Child B to score 70 on the test, which means that Child B would get a prize worth 10*(70-40) = 300 MWK. With tutoring, you expected Child B to get a score of 90. With this reward scenario, he/she will receive 10 * (90-40) = 500 MWK. So, then the more tickets you give to Child B, the higher chance you move them from a prize worth 300 MWK to a prize worth 500 MWK.

Because your ticket allocation can make a big difference on which child gets tutoring, and because only one scenario is randomly selected by the computer, you should think of each scenario as a standalone scenario and evaluate it in isolation, pretending that that scenario is the scenario selected by the computer and thinking what you want to happen in that case.

BAR CHART

This bar chart is another way to see the information in the table. It shows how expected reward for Child A and Child B and total rewards depend on how you allocated tickets.

[RA INSTRUCTIONS] Draw a graph with y axis labeled rewards and x axis with space for 11 bars. Below the x-axis, create a row labeled "Tickets to Child A" and write the numbers 0-10 from left to right. Under this, create a row labeled "Tickets to Child B" and write the numbers 10-0 from left to right]

With this scenario, if you allocate 0 tickets to Child A and 10 tickets to Child B, Child A's expected reward is 100 and Child B's expected reward is 500. The total reward for this allocation is 600. [RA INSTRUCTIONS: Draw a bar for Child A from y=0 to y=100 and shade this in lightly. Label this region 100. Draw a bar for Child B that starts from y=100 to y=600 and do not shade in. Label this region 500. Label top of bar 600.]

If you allocate 10 tickets to Child A and 0 tickets to Child B, Child A's expected reward is 200 and Child B's expected reward is 300.The total reward for this allocation is 500. [RA INSTRUCTIONS: Draw a bar for Child A from y=0 to y=200 and shade this in lightly. Label this region 200. Draw a bar for Child B that starts from y=200 to y=500 and do not shade in. Label this region 300. Label top of bar 500.]

ALLOCATIONS TABLE

For instance, if Child A is allocated 0 tickets and Child B is allocated 10 tickets, expected reward for Child A is 100, while expected reward for Child B is 500. [RA INSTRUCTIONS: point to first row of visual aid table in which Child A is allocated 0 tickets]

If Child A is allocated 10 tickets and Child B is allocated 0 tickets, expected reward for Child A is 200, while expected reward for Child B is 300. [RA INSTRUCTIONS: point to last row of visual aid table in which Child A is allocated 10 tickets]

If you would like to maximize Child A's reward, you should allocate all tickets to Child A [RA INSTRUCTIONS: draw an arrow to last row of table in which Child A gets 10 tickets and label "highest reward for Child A"]

If you would like to maximize Child B's reward, you should allocate all tickets to Child B [RA INSTRUCTIONS: draw arrow to first row in table in which Child B gets 10 tickets and label "highest reward for Child B"]

[RA CHECK] *If one child has a higher score gain from tutoring than the other*: If you would like to maximize the total reward amount received by Child A and Child B combined, you would give the tutoring to [Child with higher returns to tutoring] because s/he is the one whose expected reward would increase more with tutoring. To do that, you would allocate all the tickets to [Child with higher returns to tutoring] [RA INSTRUCTIONS: draw an arrow to the row in which [Child with higher returns to tutoring] gets 10 tickets and label "highest total reward".].

[RA CHECK] *If one child does not have a higher score gain from tutoring than the other*: If you want to maximize the total reward amount received by Child A and Child B combined, it wouldn't matter how you allocate the tickets because they all have the same total expected reward -- the only thing that differs across allocations is the split between Child A and Child B, not the total.

If you would like to give both children an equal opportunity to get tutoring, you should allocate tickets equally; 5 tickets to Child A and 5 tickets to Child B. [RA INSTRUCTIONS: draw arrow to row in which Child A and Child B each get 5 tickets and label "equal tickets"].

[RA CHECK*] If Child A's maximum expected reward is greater than or equal to Child B's minimum expected reward:* If you would like to give both children as close to the same expected reward as possible, you should allocate more tickets to Child A than to Child B. [RA INSTRUCTIONS: draw an arrow to right half of table in which Child A gets more tickets than Child B and label "most equal rewards"]

[RA CHECK] *If Child A's maximum expected reward is less than Child B's minimum expected reward*: If you would like to give both children as close to the same expected reward as possible, you should allocate 10 tickets to Child A and 0 tickets to Child B. [RA INSTRUCTIONS: Draw an arrow to last row in which Child A gets 10 tickets and label "most equal rewards"]
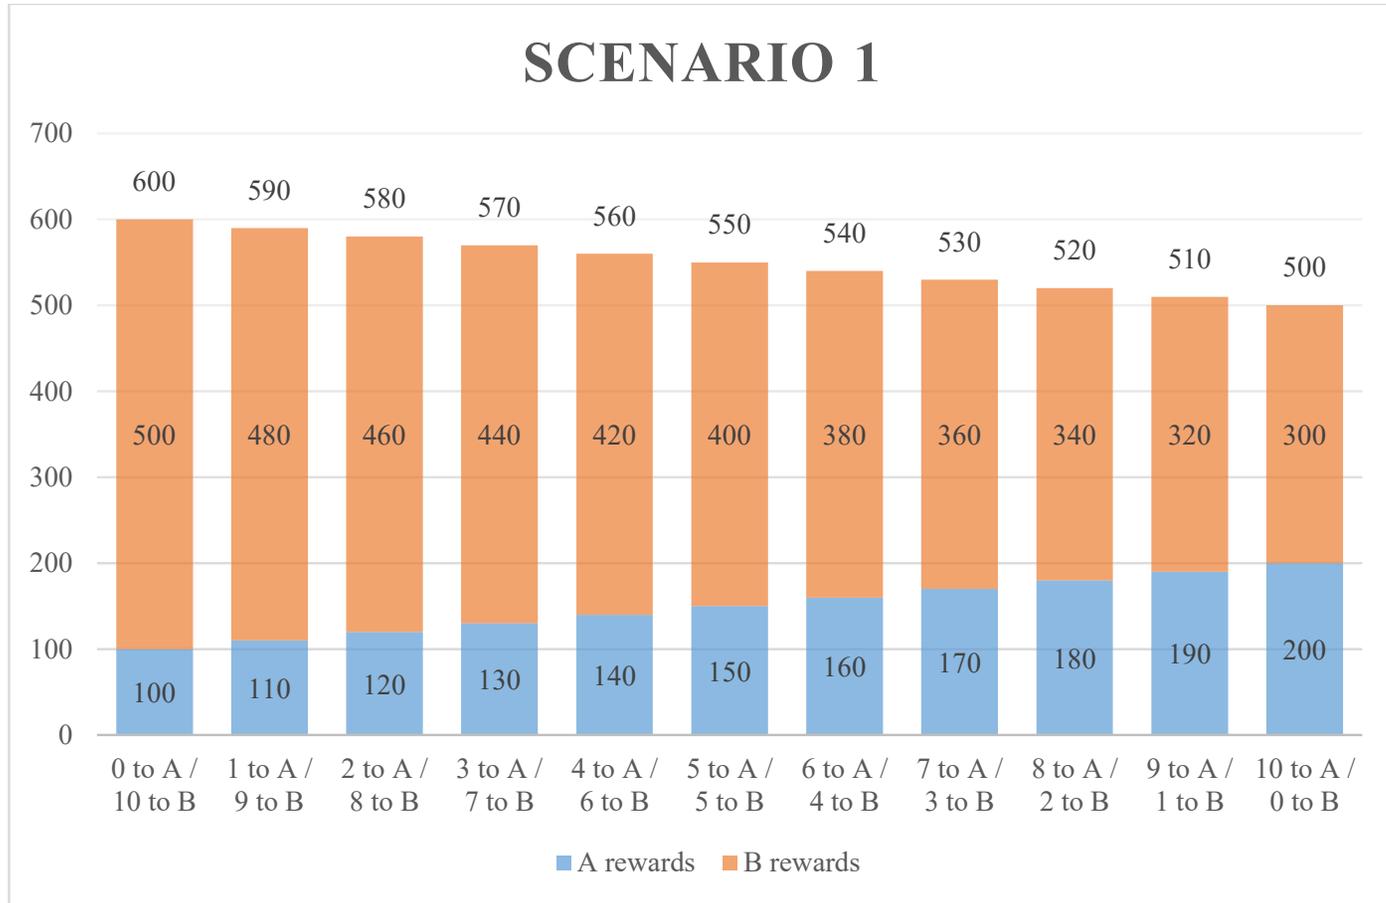
Do you have any questions?

TICKET ALLOCATION

So, here are 10 tickets. We'll ask you to divide these 10 lottery tickets between your children. One out of the 10 lottery tickets will be randomly selected by you, and the child whose ticket it is will receive one hour of tutoring. If that chosen lottery ticket belongs to "Child A", he/she will receive tutoring, otherwise, " Child B" will receive tutoring. Thus, the child with a larger allocation of lottery tickets has a higher chance of receiving tutoring.

Please allocate these 10 lottery tickets across your two children.

Scenario 1: Visual Aid 2

Tickets to Child A	0	1	2	3	4	5	6	7	8	9	10
Child A's Expected Reward	100	110	120	130	140	150	160	170	180	190	200
Child B's Expected Reward	500	480	460	440	420	400	380	360	340	320	300
Tickets to Child B	10	9	8	7	6	5	4	3	2	1	0