(Not) Playing Favorites: An Experiment on Parental Preferences for Educational Investment*

James Berry, Rebecca Dizon-Ross, and Maulik Jagnani

This version: February 11, 2019

Abstract

How do parents choose to allocate investments across children? Do they maximize the returns to their investments (total household earnings), or equalize across their children because of an aversion to cross-sibling inequality? In this paper, we conduct the first experiment that identifies parents' preferences for investing in their children's education. The experiment exogenously varies the short-run returns to educational investments to identify the degree to which parents care about (a) maximizing total household earnings, (b) minimizing cross-sibling inequality in "outcomes" (i.e., child-level earnings), and (c) minimizing cross-sibling inequality in "inputs" (i.e., the investments each child receives). We find that parents care about both maximizing total household earnings and minimizing inequality in inputs. Parents place a high value on equality of inputs, choosing exactly equal inputs 35% of the time and foregoing roughly 40-50% of their potential experimental earnings in order to equalize inputs. By contrast, we find no evidence that parents are averse to inequality in outcomes.

*Berry: Department of Economics, University of Delaware; Dizon-Ross: Booth School of Business, University of Chicago; Jagnani: Charles H. Dyson School of Applied Economics and Management, Cornell University. This study was funded by the J-PAL Post-Secondary School Initiative. The study protocols received approval from the IRBs of IPA and the University of Chicago. We thank Faith Millongo and Marumbo Munyenyembe for leading the fieldwork, and Yashna Nandan for excellent research assistance. We are grateful to Sabrin Beg, Marianne Bertrand, Esther Duflo, Pascaline Dupas, Richard Hornbeck, Seema Jayachandran, Maria Rueda, Laura Schechter, and Jeremy Tobacman for helpful conversations and feedback, and to participants in the Booth Junior group and MIT Development lunch for helpful comments. All errors are our own.

"Be fair and just in terms of the
gifts you offer your children."
– *Prophet Muhammad*

"If you choose not to decide,
you still have made a choice."
– *Rush*

1 Introduction

What are parents' preferences for allocating resources among their children? Although social scientists since Adler (1956) and Freud (1961) have examined this question, a growing realization that parental investments have persistent and profound impacts on their children's outcomes has spurred contemporary interest in the topic (e.g., Cunha and Heckman, 2007; Cunha et al., 2006). However, our current understanding of parents' preferences is limited: Do parents treat investments in their children as standard investment goods, maximizing returns but potentially leading to inequality across siblings? Alternatively, are they averse to cross-sibling inequality? And if parents are inequality averse, do they prefer to equalize the *inputs* they invest in their children, such as expenditures on books and schooling, or to equalize their children's *outcomes*, i.e., the amounts their children ultimately earn?

Understanding these preferences can help governments design better policies that account for households' endogenous responses. For example, consider one of the most prevalent education policies in developing countries: conditional cash transfer programs (CCTs), which make large monetary transfers to households if children attend school. Many of these policies directly target only one child in a household; understanding the nature of the spillovers onto non-targeted children is thus important (Barrera-Osorio et al., 2011). The spillovers depend directly on parents' preferences. If parents want to maximize returns, the policies could have negative spillovers onto non-targeted siblings by increasing the relative returns to investing in the targeted sibling; in contrast, if parents have a preference for equalizing their children's outcomes, then CCTs are more likely to have positive spillovers. Finally, if parents have a preference for equalizing the inputs they invest in their children's education, the CCTs could have negative to ambiguously-signed spillovers, depending on whether the targeted child would already have attended school in the absence of the CCT.[1] To optimize policies for these reactions, we must understand parents' preferences.

This paper reports the results from the first experiment to identify parents' preferences for investing in their children's education. Our lab-in-the-field experiment sampled parents with at least two children between class 5 and class 7 enrolled in government schools in southern Malawi. We first asked each child to take a test. We then delivered monetary payments directly to the children based on their test scores; here, each child's monetary payment is our measure of his/her *outcome* (which, in the broader literature, has represented an individual child's income resulting from education). Before the test, parents were given an *input*: 10 lottery tickets they could allocate across their children, where exactly one of the 10 tickets would be randomly chosen and where the child whose ticket was selected would receive one hour of tutoring. This setup allowed parents to choose which of their children would receive tutoring by allocating all their tickets to that child; thus, parents should only split their tickets between their children if they are averse to inequality and/or making a choice between their children.

Before the experiment, we measured parents' beliefs about each of their children's test scores without tutoring, as well as their beliefs about each of their children's "test score gains" from tutoring (i.e., how much parents expect each of their children's test scores to increase if they received tutoring). These beliefs yield predictions for how parents would allocate their tickets (inputs) under returns-maximization, inequality aversion over inputs, and inequality aversion over outcomes. In particular, returns maximization means the parents want to maximize their children's total payments and so would assign all lottery tickets to the child whose expected payment gain from tutoring (based on both parental beliefs about each child's test score gains from tutoring and the exogenously-given payment functions mapping scores to payments) would be highest. Inequality aversion over inputs means the parents want to give both children an equal number of lottery tickets, and inequality aversion over outcomes means the parents want to minimize the gap between their children's expected payments.

[1] For example, if the targeted child would have attended school in the absence of the CCT, then the CCT would decrease the family's net expenditures on the targeted child, potentially causing negative spillovers if parents are inequality averse over inputs.

The experiment then exogenously varied the child-specific payment functions mapping test scores to payments. For example, under some payment functions, the tutoring improved payments (outcomes) more for children who parents perceive to have higher baseline performance on the test, while in others, the tutoring improved payments more for children with lower baseline performance. These types of shocks enable us to estimate parents' utility parameters. Because cash is potentially transferable within the household, the use of a cash reward biases us towards finding returns-maximization;[2] we can thus think of our experiment as estimating a lower bound on the level of inequality aversion. Given that most options for rewards could be seen as biasing estimation one way or the other, to be conservative, we prefer to be biased against rejecting the null of the "standard model" (returns maximization) rather than towards.[3]

Our headline result is that parents display a quantitatively important aversion to inequality in inputs, which causes them to leave non-trivial payments on the table. We establish this result as follows: We first test and reject the null 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. Second, we establish that parents care about equalizing the inputs they give to their children, and that this preference represents the primary reason that they deviate from returns maximization in our experiment. Parents choose exactly equal inputs in roughly 35% of their choices. Even when we meaningfully increase the expected payment gains from maximizing returns, offering 10x higher returns per point to the child parents perceive is higher-performing or lower-performing, at least 30% of parents still equalize. Third, we fail to find evidence that parents care about equalizing outcomes. Although this (non-)finding could, in theory, partly reflect the fact that parents could *ex post* equalize their children's payments (outcomes), we present evidence from another setting showing that even when parents cannot *ex post* equalize, they appear more averse to inequality in inputs than outcomes. Finally, we find that parents are willing to forego meaningful financial payments to satisfy their preference for equality in inputs. The average parent forwent roughly 40-50%

[2] In particular, outside of our experimental setting, one reason that parents might be inequality averse is limited commitment: parents may not be able to enforce transfers between their children in adulthood, and so cannot *ex post* equalize via transfers. They may, however, be able to perform offsetting transfers with the cash rewards we use in this experiment.

[3] In particular, when designing the experiment, we considered using a reward that parents could not *ex post* equalize, such as non-fungible consumption. However, we were concerned that parents may have highly concave utility in the non-fungible consumption (and/or simply not value it), which could bias us towards finding inequality aversion.

of their potential earnings in each scenario, thus earning only 50-60% as much as they could have.[4] In the scenarios where we exogenously varied the returns to tutoring across children, average foregone earnings are substantial in magnitude, representing 6% of average annual per-child educational spending or 35% of a local adult's daily wage.

We perform various supplementary analyses to rule out potential confounds to interpretation, including providing evidence that equalizing inputs does not reflect a lack of understanding of how to maximize returns. For example, we present data from a "placebo test" where we gave parents lottery tickets to allocate between two monetary prizes (instead of between their children); in that task, over 95% of parents' allocations were all-or-nothing. This suggests that parents do, in fact, know how to maximize returns in a lottery; they just prefer not to do so when it means they have to invest unequally in their children. We also present evidence that parents do not diverge from "all-or-nothing" due to uncertainty about children's performance or indifference about which child should receive the tutoring.

Our work contributes to various strands of literature. First, we contribute to a classic literature in economics that characterizes parents' preferences for investing in their children and their balance between returns-maximization and inequality aversion (e.g., Behrman et al., 1982, 1986; Pitt et al., 1990). This literature has typically relied on functional form assumptions for identification; our project is the first to use exogenous variation to uncover parents' preferences. Moreover, the economics literature has always tested for aversion to inequality in outcomes; our paper is the first empirical economics paper to introduce a desire for equality in inputs, which we demonstrate to be the dominant preference in our setting.[5] 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 incorporating this preference into a utility function framework and quantifying its role relative to other parental concerns.

Second, there has been a recent flurry of articles that investigates 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

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

[5] On the theory side, 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.

compensate by investing less.[6] The results are mixed, with responses ranging from reinforcing, through 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 (normally unobserved) features of the economic environment, such as the perceived production function mapping endowments and investments to returns.[7] But, understanding the preferences themselves is necessary to proactively design policies that account for parents' endogenous responses. Note that our finding that parents are inequality averse over inputs is not inconsistent with the studies finding that parents often spend differently on their children (indeed we find that too). Rather, we find that inequality aversion is a powerful force that parents balance with other concerns (e.g., returns maximization), thus causing their input choices to be *more* equal than they would be without it, but not necessarily fully equal.[8]

Third, this paper relates to a growing body of research estimating individuals' preferences about fairness. Both experimental and observational studies have shown behavioral patterns which are consistent with inequality aversion.[9] A small number of experimental papers have also examined preferences between equality of opportunity or fairness from an ex-ante perspective (inequality aversion in inputs) and inequality aversion from an ex-post perspective (inequality aversion of outcomes), finding disproportionate support for the former (e.g., Andreoni et al., 2016; Cappelen et al., 2014).[10] These studies examine preferences for fairness in the lab among strangers; our work extends this examination to a field setting and to the important and policy-relevant domain of parents' preferences among their own children. Perhaps more related to our study are articles that examine how parents divide estates among their children, documenting the presence of an "equal division puzzle" (Bernheim and Severinov, 2003; Menchik, 1980; Wilhelm, 1996); however, transfer payments do

[6]See, e.g., Abufhele et al. (2017); Adhvaryu and Nyshadham (2016); Akresh et al. (2012); Almond et al. (2009); Ayalew (2005); Bharadwaj et al. (2018, 2013); Datar et al. (2010); Hsin (2012); Leight (2017); Yi et al. (2015).

[7]For example, just because parents spend more on average on their lower-endowment children, this does not mean they care about inequality in outcomes; rather, they could perceive that investments in lower-endowment children yield higher returns.

[8]For example, it could be the case that the perceived returns-maximizing strategy is to spend 10x as much on higher-endowment children but, because of inequality aversion over inputs, parents only spend 2x as much on higher-endowment children.

[9]See, e.g., Andreoni and Bernheim (2009); Forsythe et al. (1994).

[10]Andreoni et al. (2016), study time inconsistency in preferences for fairness, and find that most individuals view fairness from an ex-ante perspective when making decisions ex-ante, and from an ex-post perspective when making decisions ex-post. We abstract away from this finding as we don't allow ex-post transfers in our experiment: Parents were responsible for ex-ante allocation of tickets, but the final reward was delivered directly to the children based on their test score.

not entail the same tradeoffs as educational investments do between efficiency and equity or between inputs and outcomes, and so parents' preferences may be very different in the transfer and investment domains.

The rest of the paper is organized as follows: Section 2 presents a conceptual framework of parents' preferences. Section 3 describes the experimental design and procedures, and Section 4 presents the results. Section 5 concludes.

2 Conceptual Framework

Consider a utility maximizing parent p with two children $i = [1, 2]$. Parent p's problem is to choose the level of investment in each child, x_1 and x_2. Three quantities enter the household's maximization problem: i) total household earnings (representing returns-maximization), ii) the gap between children's earnings (entering negatively, representing inequality aversion in outcomes), and iii) the gap between children's investment (entering negatively, representing inequality aversion in inputs). We assume that the utility function is a weighted sum of these quantities. Utility can therefore be expressed as:

$$
\begin{aligned}
\max_{x_1,x_2} U(x_1, x_2|a_1, a_2) = & \lambda\mathbb{E}\left[R(x_1|a_1) + R(x_2|a_2)\right] \\
& - \alpha\Big|\mathbb{E}\left[R(x_1|a_1) - R(x_2|a_2)\right]\Big| \\
& - \beta\Big|\mathbb{E}\left[x_1 - x_2\right]\Big|
\end{aligned} \tag{1}
$$

where λ is the weight on returns maximization and α and β are the weights on inequality aversion of outcomes and of inputs. These weights are scaled such that each is weakly positive. a_1 and a_2 are the "endowment" of child $i = [1, 2]$, respectively, which we define as a child's "baseline" earnings potential (i.e., earnings potential when $x_i = 0$). $R()$ is the earnings function capturing expected earnings from schooling, which increases in x_i and, for $x_i = 0$, by definition, in a_i. This utility function is maximized subject to the budget constraint

$$x_1 + x_2 \leq y$$

where y is the total educational budget.

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

1. Returns maximization ($\lambda = 1, \alpha = 0, \beta = 0$):
   - If a_i and x_i are complements ($\frac{\partial^2 R_i}{\partial x_i \partial a_i} > 0$), then parents give more inputs to their higher-endowment child ($\frac{\partial x_i^*}{\partial a_i} > 0$).
   - If a_i and x_i are substitutes ($\frac{\partial^2 R_i}{\partial x_i \partial a_i} < 0$), then parents give more inputs to their lower-endowment child ($\frac{\partial x_i^*}{\partial a_i} < 0$).
2. Inequality aversion in outcomes ($\lambda = 0, \alpha = 1, \beta = 0$): Parents invest so that $R(x_1|a_1) = R(x_2|a_2)$.[11]
   - If a_i and x_i are complements ($\frac{\partial^2 R_i}{\partial x_i \partial a_i} > 0$), then parents give more inputs to their lower-endowment child ($\frac{\partial x_i^*}{\partial a_i} < 0$).
   - If a_i and x_i are substitutes ($\frac{\partial^2 R_i}{\partial x_i \partial a_i} < 0$), then parents give more inputs to their lower-endowment child for low y; depending on how strong the substitutability is and how large the baseline earnings gap is, at a sufficiently high y, they may begin to give more to their higher-endowment child since that child's earnings may begin to lag behind the lower-endowment child's earnings.
3. Inequality aversion in inputs ($\lambda = 0, \alpha = 0, \beta = 1$): Parents invest so that $x_1 = x_2$.
   - Parents equalize inputs regardless of complementarity.

The above utility function exhibits what is known in the literature as "equal concern": if the monetary return to investment is equal across children, parents will not prefer to invest in one child over another. We can also incorporate "unequal concern", that is, child-specific preferences, such as a preference for investing in a son over a daughter. Here we model unequal concern as a preference for giving more inputs x_i to the preferred child, rather than for giving the preferred child higher earnings $R(x_i|a_i)$.[12] Define $\gamma(z_1, z_2)$ as the parents' relative preference for investing in child 1 relative to child 2. $\gamma(z_1, z_2)$ can be either positive or negative, meaning that the parent could prefer either of her children, and is allowed to

[11] If perfect equality is not acheiveable, parents will simply minimize $|R(x_1|a_1) = R(x_2|a_2)|$.

[12] This choice is based on data from the experiment that shows that, when there are no returns implications, a significantly higher share of parents choose to give their children unequal inputs than unequal outcomes, suggesting a higher degree of child-specific preference over inputs than earnings.

depend on a vector of each child's characteristics z_i (e.g., gender, age, etc.). The utility function can then be written as follows:

$$\begin{aligned}\max_{x_1,x_2} U(x_1,x_2|a_1,a_2) =& \lambda\mathbb{E}\left[R(x_1|a_1)+R(x_2|a_2)\right] \\ &-\alpha\big|\mathbb{E}\left[R(x_1|a_1)-R(x_2|a_2)\right]\big| \\ &-\beta\big|\mathbb{E}\left[x_1-x_2\right]\big| \\ &+\gamma(z_1,z_2)[x_1-x_2]\end{aligned} \tag{2}$$

Note that, while the addition of $\gamma \neq 0$ may change predictions for which child a parent will give more inputs to for a given R function, it does not change the predictions for how parents should respond to *shocks* to the R function.[13]

For simplicity of exposition and model estimation, the above utility functions are linear in payments and hence do not exhibit risk aversion. While, in general, adding risk aversion could cause input choices to seem more compressed across children than returns maximization without uncertainty would dictate, in our experimental setting that should not be the case because we are using lottery tickets as our input and expected utility is linear in lottery tickets. We discuss this issue further in Section 4.1.1.

3 Design and Procedures

Our goal is to identify the preference weights λ, α, and β. It is very difficult to estimate these preferences using observational data, most notably because it is hard to find a setting with enough separate exogenous shocks to separately identify the parameters, and also because parents' perceptions about the production function are generally not observable. We use an experiment to overcome these challenges. Since it is difficult to generate multiple long-run shocks for identification, we instead experimentally shock the short-run returns to investment.

We recruited 300 parents with at least two children between class 5 and class 7 enrolled

[13] For example, if the complementarity of a and x increases, regardless of γ, a parent who cared about maximizing returns but not equality would increase her relative investments in her higher-endowment child, while a parent who cared only about equality in outcomes would increase her relative investments in her lower-endowment child.

in government schools in southern Malawi. All children were asked to take a math test,[14] and a monetary reward was delivered directly to the children based on their test score. The monetary reward is our measure of the *outcome* (the $R()$ 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. This means that (if they wanted to) parents could guarantee which child would receive tutoring by giving all tickets to one of their children. We measured parents' beliefs about each child's test score without the tutoring and about how much each child's test score would increase with tutoring. Going forward, 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]

We use lottery tickets as the input because it means that a child's expected payment is linear in the number of tickets; linearity is advantageous as it yields clean predictions and facilitates measurement. For example, the lottery setup yields an unambiguous prediction for the returns-maximizing allocation: parents will allocate all their tickets to the child they want to receive tutoring.[16]

In order to maximize our statistical power, we use the "strategy method" to exogenously vary the payment functions that map child-specific scores to payments. Under this method, parents choose their allocations under five scenarios varying the payment function; we then randomly choose one scenario to be implemented for each parent.[17]

For each scenario, we constructed the payment functions for each child as lump-sum transfer B plus a reward of C per point on the test. For child i, this can be expressed

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

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

[16]The linearity also makes it easier to elicit parental beliefs about how expected payments vary with inputs since one only needs to elicit beliefs at two points on the function. If instead parents were splitting say one hour of tutoring between their children, we would have had to elicit beliefs about the concavity in returns for each child over minutes.

[17]This adds a second layer of uncertainty: in addition to there being a lottery within scenario 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 the question of at what level parents evaluate the expectation. We assume that parents "narrowly bracket" and try to equalize expected inputs and outcomes *within* scenario. This is conservative for estimating inequality aversion: If instead, they try to minimize expected inputs and outcomes *across* scenarios, that would bias us away from detecting inequality aversion. This is another reason our experiment can be seen as estimating a lower bound on inequality aversion. We also asked parents survey questions after the experiment which suggested that they generally did not try to equalize across scenarios.

as:

$$P_i = B_i + C_i(Testscore_i - Threshold)$$

where i = {L, H}, P_i is child i's expected payment given expected test score $Testscore_i$, and $Threshold$ equals the parent's belief of child L's test score without tutoring, rounded down to the nearest 10. We chose to only reward performance above this value in order to implement steep payment functions while keeping total payments reasonable. We used 5 different scenarios for the payment functions, depicted in Figure 1. In our first *Base Case* scenario, 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 ($B_L = B_H = 0$).[18] The other 4 scenarios are variations on the base case designed to identify the utility function parameters, as described further in Section 3.3.

The elicitation used real stakes: after ticket allocations were elicited for each scenario, one scenario was randomly selected, and tickets were assigned to both children based on their allocation for that scenario. Children were then asked to take a test, and cash was delivered directly to the children (in individual envelopes) based on their test score and the randomly-selected payment function scenario.

3.1 Process

Roughly a week prior to the experiment, all participating parents were surveyed about demographics and their children's education (e.g., educational expenditures). We also described the math test that their children would take and measured parents' beliefs about each of their children's expected performance 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. We then described the experimental design and conducted the experiment.

Because the design involves a number of steps, we took multiple measures 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

[18] 10 MWK is roughly 1.4 US cents or 0.12% of average annual per-child household educational spending.

so. Next, we gave parents a detailed overview of the full experimental design, walking them through two "practice (hypothetical) scenarios" using different payment functions than those used in the actual experiment, but which we explained in the same way that we explained the experimental scenarios (described in detail below). After the "practice" scenarios, surveyors also walked parents through a sample "scenario lottery" to explain how the strategy method worked.[19]

For both "practice" and "experimental" scenarios, surveyors explained as follows (See Appendix B for Sample scripts and visual aids for Scenario 1.). They first explained the payment functions. 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 for each potential allocation (making the returns-maximizing solution clear), but did not display anything about the gap in expected outcomes or inputs. We decided that it would be infeasible to display total expected payments in addition to the expected outcomes and inputs gap in a clean manner. Thus, we emphasized returns maximization so that we would be conservative in testing the null of the standard returns-maximizing model.

As part of the explanations, surveyors showed parents how they should allocate their tickets if they wanted to maximize returns or equalize outcomes or inputs in each scenario. Although this explanation could be seen as leading, we prioritized being sure that parents fully understood 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 the "practice" scenarios only, surveyors asked parents a set of questions to test understanding of how to allocate tickets to achieve each of the goals above; no such questions were asked after the real experimental scenarios.

[19] In particular, surveyors explained how only one of the scenarios would be chosen, and then the payment function from that scenario would be used and that the tutoring winner would be chosen using lottery tickets assigned to children based on the parents' ticket allocation for that scenario.

After the practice scenario and lottery explanation, parents made their actual experimental allocations. Surveyors explained each scenario in order from 1 to 5, showing them the relevant visual aids and giving the explanations described above, and parents made their selections. 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 the 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 randomly pick a ticket out of a hat to select the winner of the lottery.

Once the experiment was completed, a short-post experiment 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 were asked to take a test and cash rewards were delivered.

3.2 Validation of the Method and Respondent Understanding

Since our experimental design is complex, it is important to validate whether parents indeed understood the set-up and how to maximize returns. We present several arguments and pieces of evidence that they did. First, as a "placebo" test, we asked parents to allocate 10 lottery tickets between two prizes in a hypothetical scenario before the experiment, just like they did in the main experiment; however, here the two prizes were both monetary prizes to be given directly to the parent: 50 MWK (0.07 USD) or 100 MWK. In this task, over 95% parents allocated 100% of the tickets to the 100 MWK prize, suggesting that the vast majority of parents understand how to maximize returns in a lottery.

Second, we examine heterogeneity by parent education. A key pattern we will show in the data is that, in the (non-placebo) experimental lottery, parents often split their tickets relatively evenly instead of maximizing returns. It is important for our interpretation to show that this behavior does not reflect lack of understanding. One would expect more-educated parents to better understand the design; thus, if evenly splitting tickets or other departures from returns maximization represented lack of understanding we would expect them to be more prevalent among the less-educated. However, Appendix Table A.2 shows that less-educated parents are equally likely to maximize returns as more-educated parents, and 12% *less* likely to split their tickets evenly. This suggests that departures from returns maximization are unlikely to be driven by lack of understanding. That said, one caveat to this analysis is that preferences may also be heterogeneous by parental education; however, the heterogeneity analysis still offers some reassurance.

Finally, recall that, for each scenario, surveyors explicitly told parents the expected total household earnings associated with each potential allocation (as well as the expected earnings for each child), as well as indicating which allocation would maximize expected total household earnings or equalize expected outcomes or inputs, given their beliefs. Thus, if parents wanted to maximize returns, they were told how to do so. Note that this strategy, which we adopted to ensure understanding, introduces 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 were also substantial: we paid the highest amounts that our field team thought would not seem outsized to participants. The average gap between the returns-maximizing and returns-minimizing total payments was roughly 700 MWK (3.5 USD) or 50% of the daily wage for adults in the area (roughly 8% of average per-child annual educational expenditures). In Section 4.1.2, we also show similar results in data collected for a different experiment where surveyors did not explicitly explain to parents how to execute various strategies, suggesting that the script did not drive choices here.

3.3 Scenarios and Predictions

We now describe the scenarios we included in the experiment and how they allow us to identify the parameters of the parental utility function specified in Section 2. All scenarios are variations on the *Base Case* scenario (scenario 1), which gave both children 10 MWK per test score point above the threshold, with no lump sum transfers. Figure 1 also summarizes the payment functions, as well as the predictions for how parents would allocate in the various scenarios if their utility function only weighted (a) returns-maximization, (b) inequality aversion over outcomes, or (c) inequality aversion over inputs.

To fix ideas, we rewrite the utility function from equation (2) for our experimental setting. Denote x_i as the tickets (inputs) given to child i ($i \in \{L, H\}$); S_i as child i's expected test score without tutoring; and R_i as child i's expected test score gains to tutoring

$$\begin{aligned} U(x_L, x_H) =& \lambda \left(B_L + C_L(S_L + \frac{x_L}{10} R_L) + B_H + C_H(S_H + \frac{x_H}{10} R_H) \right) \\ & - \alpha \left| (B_L + C_L(S_L + \frac{x_L}{10} R_L)) - (B_H + C_H(S_H + \frac{x_H}{10} R_H) \right| \\ & - \beta \left| x_L - x_H \right| + \gamma \left(x_L - x_H \right) \end{aligned} \tag{3}$$

A key point is that, because of our lottery setting, if α and β are 0 (i.e., if parents do not care about inequality in outcomes or inputs and instead care only about maximizing returns or child-specific preferences), they should allocate their tickets in an "all-or-nothing" fashion, giving all tickets to one child and none to the other. Because expected utility is linear in probabilities, this is also true in the more general case where we allow parents' utility from total household earnings to be concave (not linear).[20] As a result, we can test whether parents exhibit *any* inequality aversion by testing whether parents choose split allocations in any scenario in our lottery environment. Our goal with the additional scenarios is to vary the payment functions to yield qualitatively different predictions for returns-maximization, inequality aversion over outcomes, and inequality aversion over inputs, to allow us to estimate the relative weights.

Our first additional scenario, *Lump Sum to Child L* (scenario 2), was designed 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 (1.37 USD) 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 her expected payment returns to tutoring, as those depend only on her expected test score gain from tutoring and her per-point rewards, C_L. Thus, lump sum transfers do not affect parents' returns-maximizing choices, nor do they affect their input-equalizing choices or child-specific preferences. However, lump sum transfers do affect the outcomes-equalizing

[20] In that case, taking u as a concave function, the first term becomes $\lambda\left(\frac{x_L}{10} Eu(B_L + C_L(S_L + R_L) + B_H + C_H S_H) + \frac{x_H}{10} Eu(B_L + C_L S_L + B_H + C_H(S_H + R_H))\right)$, where the expectation $Eu(\cdot)$ is taken over the risk in parents' beliefs about S_i and R_i. Importantly, this is linear in x_L and x_H.

choice: since giving a lump sum transfer to one child increases her expected payments, an outcomes-equalizing parent would respond by reallocating tickets to that child's sibling to increase the sibling's expected payments.

We designed two other scenarios to yield opposite predictions for returns-maximization and inequality aversion over outcomes, thus letting us test which strategy parents weight more heavily on average. The two scenarios are as follows, with neither having any lump sum transfer. The first (*Higher Returns to Child H*, Scenario 3) gives Child H a ten times higher per-point reward than Child L: C_H is 100 MWK (0.14 USD) while C_L is 10 MWK. Increasing C_H has two effects: 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 is to give all inputs to Child H. However, it also increases Child H's expected payments (outcome) for *any* ticket allocation that the parent could choose. As a result, for all households, the outcomes-inequality-minimizing choice is to give all inputs to Child L. The second scenario, (*Higher Returns to Child L*, Scenario 5), exchanges the payment functions used in *Higher Returns to Child H* between Child L and Child H. That is, there are still no lump sum transfers, but now it is C_L that is 100 MWK while C_H is 10 MWK. For almost all households, this correspondingly reverses the predictions: returns-maximizing parents would now allocate all inputs to Child L, whereas outcomes-inequality-minimizing parents would do the opposite.

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.

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 4) that adjusted the payment functions so that the child a returns-maximizing parent would choose is the *same* as the child an outcomes-inequality-minimizing parent would choose. As a result, if we see a high share of parents choosing split allocations in that scenario, it cannot be that parents are balancing inequality aversion over outcomes against returns maximization; instead, it suggests parents care directly about minimizing inequality in inputs.

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

Identification depends on parents believing that tutoring yields positive test score gains for a child; in the next subsection, we show that this is the case. More broadly, we are identifying preferences based on parents' *beliefs* about their children's scores, which other work shows are often inaccurate (Dizon-Ross, 2019). Fortunately, whether parents' beliefs are accurate is not important for our identification: we are identifying parents' preferences conditional on their beliefs. Thus, the only way that potential inaccuracies should affect our estimation is if parents' beliefs distributions are uncertain and if that uncertainty affects their allocations; however, as we show in Section 4.1.1, uncertainty does not appear to affect the results here.

3.4 Summary Statistics and Prima Facie Evidence

Table 1 reports selected summary statistics from our sample. Nearly all respondents are mothers. Fifty-five percent of the child sample is female. On average, parents believe that Child H's score on the test will be 11 percentage points (pp) higher than Child L's, and that tutoring will have higher score returns for Child H than child L, increasing Child L's score by 11pp and Child H's by 18pp. 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. The average parent in our sample spends roughly MWK 8400/year (11.6 USD) on education for each of her children (Panel D).

We also present prima facie evidence suggesting that parents may have a preference for equalizing their spending on their children. Figure 2 displays expenditures on Child L, measured during the baseline survey, as a share of total combined spending across both children. Consistent with inequality aversion in inputs, there is a notable spike in density at 50%. To examine the external validity of this pattern, we also analyze data from the cross-country Young Lives Survey and show similar patterns there; see Appendix Figure A.2.[21] In our baseline survey, over 85% of parents also state that they spend a roughly equal amount of money on both their children's education (Table 1 Panel A). Of course, these results are suggestive only; there are other potential reasons that parents might spend equally on their children besides inequality aversion. We now turn to our experimental results to provide more definitive evidence on this issue.

4 Experimental Results

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 shed light on the quantitative magnitude of the parameters through a combination of reduced-form and structural analysis.

4.1 Reduced-Form Analysis

We first test – and reject – the null that parents care only about maximizing returns. Because we conducted our experiment in a lottery setting, we have a very clear prediction for returns maximization: unless parents care about some form of equality (inputs or outcomes), they should allocate their tickets in an "all-or-nothing" fashion, giving all tickets to the child who has the highest returns to investment and none to the other. However, as shown in Figure 3, which shows the pooled distribution of ticket allocations across all scenarios, only 45% of allocations were all-or nothing. In 55% of allocations, both children received non-zero tickets. There is one other reason, besides inequality aversion, that parents could split their tickets: in the knife's edge (empirically improbable) case that parents are completely indifferent about which of their children receives the lottery. However, Figure 4 suggests that the reason is not that parents are indifferent between their two children: when we offer 10x higher returns per point to Child H (scenario 3) or Child L (scenario 5), that should

[21] We use the Young Lives Survey because it has data on expenditures on two children across multiple children. As described in the figure notes, the data collection methods for that survey were not fully ideal for this question since they asked about joint spending in each category (e.g., books, school fees) and then had parents apportion that across children. However, we view the data as suggestive.

presumably be enough to break indifference for most parents, but we still find that at least 30% of parents equalize inputs (Figure 4). Moreover, Table 2 shows that 19% of parents equate in all scenarios, which should not reflect indifference: if a given parent is indifferent between her children in one scenario, then if a second scenario meaningfully changes the relative returns to investing in her two children, she should no longer be indifferent unless she does not value payments.

We next use the *Lump Sum to Child L* scenario to test for evidence of inequality aversion over outcomes. We fail to find evidence that parents care about equalizing outcomes. Recall that, relative to the *Base Case* (scenario 1), the *Lump Sum to Child L* scenario (scenario 2) 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. However, in aggregate, we fail to find evidence for reallocation towards Child H (Figure 5). If anything, parents, on net, reallocate in the opposite direction, although the magnitudes are very small.[22] Turning to individual-level changes, equal numbers of parents reallocate to and away from Child L, suggesting the reallocations primarily represent noise. We find similar findings when comparing the *Higher Returns to Child L* and *Higher Returns to Child L & Lump Sum to Child H* scenarios, which were included in the design for other reasons but across which the only difference is in whether one child receives a lump sum.

The variation in outcomes inequality across parents' potential choices is large enough (both absolutely and relative to the variation in total payments) that, if parents did care about inequality in outcomes, we would expect to see meaningful responses. For example, across all scenarios, Figure 6 shows that parents' chosen allocations had 400 MWK (0.55 USD) more unequal expected payments than the outcomes-inequality-minimizing allocation; the daily wage in our setting is roughly 1400 MWK and 400 MWK represents 5% of annual per-child educational spending.[23] However, because we provide cash rewards, there

[22] We are unable to explain the marginal significance of that finding.

[23] An alternative explanation is that parents only care about equalizing outcomes when they can *perfectly* equalize outcomes (i.e., that their utility term capturing inequality aversion over outcomes is $-\alpha\mathbb{1}\{\mathbb{E}R(x_1|a_1) \neq \mathbb{E}R(x_2|a_2)\}$); this could bias us away from finding evidence of inequality aversion in outcomes since, in many of parents' scenarios, even the outcomes-inequality-minimizing choice did not bring the expected outcomes gap between children equal to zero. To test this, we present summary statistics of the chosen allocation for 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 (whether parents had the option to perfectly equalize depends fully on potentially endogenous variation in their beliefs about their children's scores with

is potential that parents could ex post equalize via transfers, which could contribute to the null finding. We provide evidence against this explanation in Section 4.1.2.

We next test whether returns maximization or inequality aversion over outcomes dominate each other on average, presenting the results in Figures 7 and 8. We find that returns maximization dominates. This finding may be unsurprising given the above evidence that parents do not care about equality in outcomes, but is useful both to verify the earlier finding and to provide evidence that parents do in fact care about returns maximization. Panels (a)-(c) of Figure 7 compare scenario 3 (*Higher Returns to Child H*) and scenario 5 (*Higher Returns to Child L*). When switching from *Higher Returns to Child H* to *Higher Returns to Child L*, returns-maximization suggests parents should increase their tickets to Child L whereas inequality aversion over outcomes suggests the opposite. We find that parents increase their allocations to Child L by a statistically significant 1.6 (Figure 8), 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 7(b)).

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 parents likely value equality in inputs; we now present positive evidence that that is indeed the case. First, visual inspection of the data in Figure 3 shows a notable spike at equal allocation, with parents choosing exactly equal inputs in roughly 37% of the scenarios. 58% of parents equalize at least once. Moreover, we can easily reject that the ticket distribution is smooth around the 5/5 point.

Second, we rule out that splitting represents a desire to balance inequality aversion over outcomes with returns-maximization 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

and without tutoring, but we view the analysis as suggestive.). The results in Appendix Table A.1 show that the percentage of choices in which parents minimize outcomes inequality is not significantly different between cases and, if anything, is smaller in the case where parents had the option to perfectly equalize. Thus, not being able to perfectly equalize does not seem to explain why parents do not care about IAO.

inputs reflects an aversion to input inequality, not a desire to balance returns-maximization against inequality aversion in inputs.[24] Figure 9 presents further consistent evidence from other scenarios. The left subfigures (9(a) and 9(c)) pool all parent $\times$ 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 subfigures (9(b) and 9(d)) show all parent $\times$ scenario observations where inequality aversion in outcomes and returns maximization have the same prediction (and thus splitting tickets can only represent inequality aversion in inputs or a positive desire to split).[25] 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.

Interestingly, the results in subfigure 9(b) and 9(d) showing that parents often allocate substantial inputs to the child who is neither the returns-maximizing nor outcomes-inequality-minimizing choice suggests that parents have important child-specific preferences ($\gamma \neq 0$). In Appendix Figures A.3 and A.4 we provide further evidence for the existence of child-specific preferences, and show that the preference varies across parents. Parents that allocated all tickets to Child H or Child L in the *Base Case* scenario (which featured symmetric payment functions) were more likely to allocate all tickets to that same child in the following scenarios.

Finally, we find that parents' deviations from returns maximization have significant monetary implications. For each family $\times$ 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 10 then plots "foregone earnings" ("returns-maximizing earnings" minus "chosen earnings") as a percent of "potential earnings" ("returns-maximizing" minus "returns-minimizing earnings"). On average,

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

[25] The difference between subfigures a/b and subfigures c/d is how the x-axes are ordered and labeled: in (a) and (b), the axes represent the number of tickets given to Child L, whereas in (c) and (d), we order the x-axes by how the alignment fits with the predictions of the various models. For subfigures (c) and (d) variants, we thus omit scenarios where returns-maximization does not yield clean predictions; thus the sample is slightly different than for subfigures (a) and (b).

parents forwent roughly 40-50% of their potential (non-inframarginal) earnings in each scenario, thus earning only 50-60% as much as they could have. The experimental stakes are substantial; the foregone expected payment amounts, shown in Figure 11, are correspondingly large, especially for scenarios 3-5 where we exogenously varied the returns to tutoring across children. In those scenarios, average foregone earnings represents roughly 6% of average annual per-child educational spending or 33% of the adult daily wage in the catchment area.

4.1.1 Uncertainty

Outside of the experimental environment, if parents' utility is concave in total household earnings (i.e., they are risk-averse), then uncertainty in the cross-sibling returns to investment could cause parents to equalize more than one would expect based purely on the production function and the means of their beliefs distributions about their children's returns. That is, risk aversion could cause parents to split their inputs more evenly than they would if there were no uncertainty, potentially causing behavior that looks like inequality aversion in inputs. However, this is not the case in our experimental environment: even if parents are risk averse and there is uncertainty about children's scores, this should not cause parents to diverge from all-or-nothing allocations. This is because our inputs are lottery tickets and expected utility is linear in probability.[26]

However, although neoclassical risk aversion through concave utility should not cause splitting here, there could still potentially be a more behavioral channel through which risk aversion might cause parents to choose split allocations. To address this, we perform heterogeneity analysis based on baseline measures of parents' beliefs uncertainty. We find no evidence that more uncertain parents equalize more; see Appendix Figure A.5.[27] 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, it does not seem that uncertainty is causing parents to choose split allocations here;

[26]To see this, now say that S_i and R_i (parents' beliefs about test scores without tutoring and about test score gains to tutoring) are random variables, and take u as a concave function. The first returns-maximization term of utility function (3) becomes: $\lambda\left(\frac{x_L}{10}Eu(B_L + C_L(S_L + R_L) + B_H + C_H S_H) + \frac{x_H}{10}Eu(B_L + C_L S_L + B_H + C_H(S_H + R_H))\right)$, where the expectation $Eu(\cdot)$ is taken over the distribution of parents' beliefs S_i and R_i. Importantly, this function is also linear in x_L and x_H, producing corner solutions if the weights on the other terms are 0.

[27]If anything, uncertain parents equalize a little less than other parents.

rather they appear to be averse to inequality in inputs.

4.1.2 Additional Evidence Against Inequality Aversion Over Outcomes and Demand Effects

One potential concern with using our experimental setting to shed light on inequality aversion over outcomes is that we use cash rewards, which parents could potentially equalize *ex post.* The first order impact of cash should be to bias us towards returns maximization, which is not what we see. Indeed, if parents' primary aversion to inequality were to inequality in outcomes and they were planning to *ex post* equalize, the most natural response would be to returns maximize during the experiment and *ex post* equalize, not to do what we actually see – which is equalize inputs during the experiment, thereby decreasing the pie for *ex post* equalization. However, there is still a potential concern that, for some reason, the fact that parents could equalize *ex post* caused them to equalize inputs more during the experiment than they would have otherwise.

To address this possibility, 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 them as representing "baseline" allocations. These data describe the results of a lottery conducted with parents where parents allocated lottery tickets between two of their children, and where the child whose ticket was chosen would win a prize. However, there are two main differences between this lottery and our setting.

First, and importantly for this exercise, the prize was a scholarship to secondary school, which most parents could not afford on their own. The earnings return to secondary school would 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 can in non-experimental settings. These data thus allow us to verify whether parents' choices look similar in a setting where they cannot easily ex post equalize.

Second, and less importantly, the total number of tickets was an odd number (9). 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. However, we can still test for inequality aversion in inputs in this odd-ticket-number setting. In particular, a parent who does not care about equality should still assign all tickets in all-or-nothing fashion. A parent inequality averse in inputs should choose the most equal allocation – 4 to one child and 5 to the other. 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.

Figure 12(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, showing that, again, in that setting, parents were not truly indifferent or did not misunderstand the returns.[28] Rather, they were simply averse to splitting unequally.

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. The consistency of the results with ours thus suggests that our results were not caused by demand effects. The fact that the other experiment used a prize with larger value also suggests that our results are robust to the size of the prize.

4.2 The Value of Equal Allocation

This section uses our experimental results to estimate parents' preference weights for equality in inputs and outcomes from Section 2: α and β. We employ several methods to do so. Because any structural approach relies on several assumptions, we first present a reduced-form approach that focuses only on quantifying the trade-off between input equalization and total earnings. We next present a structural approach which allows us to numerically estimate both α and β. Both strategies are identified off of variation in the expected total payments and the expected payment inequality across the choices parents made. The majority of this variation comes from our experimentally-generated cross-scenario variation in the payment

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

functions. However, since parents' beliefs were one input into the payment functions, there are potential concerns about endogeneity. For example, households that have small cross-sibling differences in the perceived gains from tutoring will have relatively smaller potential (perceived) foregone earnings as a result of equalizing than parents with larger cross-sibling differences. As a result, for both the reduced-form and structural approaches, we also present instrumental variable approaches that isolate the experimental variation.

Reduced-Form Approach Our earlier analysis suggests that parents have a preference for equality in inputs but not in outcomes.[29] Our primary interest is thus in shedding light on the weight they place on inputs inequality. Since utility functions are only defined up to a normalization, if we reweight the preference term on returns-maximization (γ) to 1, then the weight on inequality aversion over inputs β essentially expresses how much total household earnings parents are willing to give up to choose more equal inputs.

In this section, we use reduced form analysis to shed light on the tradeoff parents make between total earnings and inputs equality. To implement the approach, we compute, for each household and scenario, the difference between earnings from the returns-maximizing choice and earnings from the input-equalizing choice. This measure can be thought of as potential foregone earnings from equalizing inputs. We then regress a binary variable indicating whether the household chose to equalize inputs in that scenario on the potential foregone earnings measure:

$$Equalized_{ij} = d_0 + d_1 * Foregone_{ij} + \varepsilon_{ij},$$

where $Equalized_{ij}$ indicates whether respondent i equalized inputs in scenario j, and $Foregone_{ij}$ represents potential foregone earnings for that household and scenario. Note that this analysis focuses only on the binary choice of whether to equalize completely; in the structural analysis we allow for parents to care about a continuous measure of inputs equality.

Figure 13 displays the results graphically by depicting the fraction of respondents who equalized inputs by bins of foregone earnings. As shown in the figure, this fraction is generally decreasing in foregone earning. Column (1) of Table 3 presents the OLS regression results.

[29] We present a formal test of their weight on equality in outcomes in the structural analysis.

This column suggests that parents must forego substantial income to be induced to stop equalizing: an additional MWK 1000 of foregone earnings are associated with just a 5 percentage point (and not statistically significant) decrease in the likelihood of equalizing inputs across children. Column (2) uses an instrumental variable (IV) strategy, using a dummy for the scenario as an instrument for $Foregone_{ij}$. The IV estimates again imply a substantial willingness to pay for equal inputs: an additional MWK 1000 (1.38 USD or 12% of annual per-child educational expenditures) of foregone earnings leads to a 10-percentage point decrease in the likelihood of a household equalizing inputs across children (significant at the 1-percent level). We also depict the IV strategy graphically in Figure 13, by creating a variable that equals the average potential foregone earnings in the sample for each scenario and plotting the share of households who equalized by each value of predicted foregone earnings. Again, the graph shows a strong downward relationship.

Overall, the analysis suggests that, while parents do trade off income for equality, at least some parents require a large amount of compensation (in the form of extra household earnings) to be willing to deviate from equal input allocation. This suggests that their utility weight on inequality in inputs, β, is large. We now move to the structural approach to provide a numeric estimate of β.

Structural Approach We next estimate parents' preference parameters using a mixed logit regression model. We adopt the mixed logit regression model so that we can allow the preference parameters to vary across the population, as well as avoid the independence of irrelevant alternatives assumption (IIA) that would be entailed by using a simpler conditional logit approach. Following equation (2), we assume that parent i has the following underlying preferences in scenario j from choosing ticket allocation k:

$$\begin{aligned} u_{ijk} =& \lambda_i TotalPay_{ijk} - \alpha_i OutcomeInequality_{ijk} - \beta_i InputInequality_{ijk} \\ &+ \gamma_i InputsToChildLvsH_{ijk} + \varepsilon_{ijk} \end{aligned} \tag{4}$$

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, $\mathbb{E}\left[R(x_1) + R(x_2)\right]$); $OutcomeInequality_{ijk}$ 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, $|\mathbb{E}[R(x_1|a_1) - R(x_2|a_2)]\,|$); $InputInequality_{ijk}$ is the absolute difference in inputs (lottery tickets) between parent i's children under allocation k (the inequality aversion in inputs term, $\left|\mathbb{E}[x_1 - x_2]\,\right|$); and $InputsToChildLvsH_{ijk}$ is the number of tickets given to child L relative to child H in allocation k (the "unequal concern" term, $\mathbb{E}[x_2 - x_1]$; note that since γ_i can be positive or negative, parents can prefer either of their children). 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 through the estimation procedure. We also allow for correlations across all preference parameters, and again estimate the correlations within the estimation procedure. Finally, the error term ε_{ijk} is assumed to be type I extreme value, independent across i, j, and k.[30] In each scenario j, parent i is assumed to choose the allocation k with the highest utility.

A key caveat to bear in mind when interpreting the analysis is that the estimation uses all of the variation in earnings from potential allocations, not just the experimentally-induced variation. However, as shown later in this section, isolating the experimentally-induced variation within a similar specification does not meaningfully change the results.

The mixed logit estimates of the means of the parameter distributions are shown in column 1 of Table 4. 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 strong preferences for equalizing inputs: they are significantly more likely to 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.[31] It is thus useful

[30]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.

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

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 forego 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 330 MWK or roughly 0.45 USD in expected household earnings for each 1 ticket decrease in input inequality. For ease of interpretation (and to fit the trends visible in the raw data more closely), in column 2 we estimate a variant of equation (5) where we replace the continuous $InputInequality$ term with a binary term for whether the allocation equally split inputs. We find that parents' mean willingness to pay (WTP) to avoid unequal inputs is $1,640$ MWK or roughly 2.6 USD, a substantial amount equal to rougly 117% of daily wage in our setting and 20% of per-child annual education spending.

Figure 14 shows the results graphically. In particular, we estimate a mixed logit model using the following utility specification which includes fixed effects τ_k for all 11 potential ticket allocation choice (10/0, 9/1, 8/2, etc.)

$$u_{ijk} = \lambda_i TotalPay_{ijk} + \beta_i OutcomeInequality_{ijk} + \tau_k + \varepsilon_{ijk} \tag{5}$$

The inclusion of τ_k causes the $InputInequality_{ijk}$ and $TixToChildLvsH_{ijk}$ terms to drop out due to multicollinearity; thus we can think of the τ_k as capturing average preferences for a given allocation, incorporating how unequal that allocation is and how many tickets that allocation gives to the child who is (on average) preferred.[32] We then plot the fixed effects τ_k normalized by sample-average λ, thus showing the average WTP to move from the least-preferred allocation (which was to give 1 ticket to Child L) to any other allocation, conditional on $TotalPay$ and $OutcomeInequality$. We find that the WTP to move to equal inputs from any other allocation is sizeable, and that the WTP for all split allocations decrease smoothly as one moves away from the equal input point. In addition, there are large WTP estimates for the all-or-nothing allocations, demonstrating strong child-specific preferences conditional on returns; however, the WTP for equal inputs is the highest on average by a meaningful margin (at least 1000 MWK) and the difference is highly statistically significant.

[32]For analysis tractability, the τ_k coefficients are modeled as fixed not random coefficients.

One concern with our mixed logit specification is that it uses all of the variation in earnings across potential allocations. To address this issue, we use an instrumental-variables strategy to isolate experimental variation in earnings, similar to that presented in the reduced-form analysis at the beginning of this section. We implement this approach by estimating equation (5) as a linear probability model (LPM) and compare OLS estimates with IV estimates, using the experimental variation as instruments.[33] An LPM specification of multinomial choice is not a particularly realistic model, but provides us with a simple and transparent way to compare estimates which are identified off of all of the variation identifying our logit model, with those identified just with the experimental variation. In particular, we estimate the following regression:

$$\begin{aligned} Prob(AllocationChosen) =& \lambda TotalPay_{ijk} - \alpha OutcomeInequality_{ijk} - \beta InputInequality_{ijk} \\ &+ \gamma InputsToChildLvsH_{ijk} + \tau_{ij} + \mu_{ijk} \end{aligned} \tag{6}$$

Because the choice is made at the parent $\times$ scenario level, we include parent $\times$ scenario fixed effects, τ_{ij}, in both specifications. The IV specification then instruments for $TotalPay_{ijk}$ and $PayInequality_{ijk}$ using indicators for the scenario $\times$ ticket allocation, τ_{jk}. Columns 3 through 4 show that the IV and OLS estimates are quantitatively similar to one another (and also qualitatively consistent with the conditional logit results, if not directly comparable, since they are estimating marginal effects, not latent utility coefficients). This suggests that the primary variation driving the identification of the mixed logit model is the experimental variation and that our high-level conclusion – that parents have a high willingness to pay to avoid investing unequally in their children – would not change if we incorporated instruments into the estimation.

5 Conclusion

Our experiment provides the first evidence that parents have a quantitatively important preference for equalizing the inputs they invest in their children's education. In order to identify these preferences, we experimentally shock the short-run returns to educational

[33] A more direct way to assess the role of endogeneity would be to estimate a mixed logit model with instruments but we are as of yet unaware of a computational method for doing so that is tractable here.

investments to identify the degree to which parents care about (a) maximizing returns, (b) minimizing cross-sibling inequality in "outcomes" (i.e., the amount their children earn), and (c) minimizing cross-sibling inequality in "inputs" (i.e., the inputs each child receives). 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 earnings (outcomes). Parents' aversion to inequality in inputs is quantitatively important, causing parents to forego roughly 40-50% of their potential earnings.

Understanding parents' preferences for investment is important for policy design. The effects of public policies depend on individuals' behavioral responses, which can either enhance or undo the policies' intended impacts. One important area for future work will be to use the preference parameters estimated here to develop "optimized" policies that take current policies and optimize them for parents' behavioral responses; one could then compare the efficacy of the optimized and baseline policies.

References

Abufhele, A., J. Behrman, and D. Bravo (2017). Parental preferences and allocations of investments in children's learning and health within families. *Social Science and Medicine 194*, 76–86.

Adhvaryu, A. and A. Nyshadham (2016). Endowments at Birth and Parents' Investments in Children. *Economic Journal 126*(593), 781–820.

Adler, A. (1956). *The Individual Psychology of Alfred Adler*. New York: Harper Torchbooks.

Akresh, R., E. Bagby, D. de Walque, and H. Kazianga (2012). Child Ability and Household Human Capital Investment Decisions in Burkina Faso. *Economic Development and Cultural Change 61*(1), 157–186.

Almond, D., L. Edlund, and M. Palme (2009). Chernobyl's Subclinical Legacy: Prenatal Exposure to Radioactive Fallout and School Outcomes in Sweden. *Quarterly Journal of Economics 124*(4), 1729–1772.

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 (2016). When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences. *Mimeo*.

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.

Ayalew, T. (2005). Parental Preference, Heterogeneity, and Human Capital Inequality. *Economic Development and Cultural Change 53*(2), 381–407.

Barrera-Osorio, F., M. Bertrand, L. L. Linden, and F. Perez-Calle (2011). Improving the Design of Conditional Cash Transfer Programs: Evidence from a Randomized Evaluation in Colombia. *American Economic Journal: Applied Economics 3*(April), 167–195.

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*, 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.

Bharadwaj, P., J. P. Eberhard, and C. A. Neilson (2018). Health at Birth, Parental Investments, and Academic Outcomes. *Journal of Labor Economics 36*(2), 349–394.

Bharadwaj, P., K. V. Loken, and C. Neilson (2013). Early Life Health Interventions and Academic Achievement. *American Economic Review 103*(5), 1862–1891.

Cappelen, A. W., K. O. Moene, E. Ø. Sørensen, and B. Tungodden (2014). Just Luck : An Experimental Study of Risk Taking and Fairness. *American Economic Review 124*(4), 1398–1413.

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.

Datar, A., M. R. Kilburn, and D. S. Loughran (2010). Endowments and Parental Investments in Infancy and Early Childhood. *Demography 47*(1), 145–162.

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

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

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

Freud, S. (1961). *Civilization and Its Discontents.* W.W.Norton and Company, Inc.

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

Hsin, A. (2012). Is Biology Destiny? Birth Weight and Differential Parental Treatment. *Demography 49*(4), 1385–1405.

Leight, J. (2017). Sibling Rivalry: Endowment and Intrahousehold Allocation in Gansu Province, China. *Economic Development and Cultural Change 65*(3), 457–493.

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

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*, 1139–1156.

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.

Yi, J., J. J. Heckman, J. Zhang, and G. Conti (2015). Early Health Shocks, Intra-household Resource Allocation and Child Outcomes. *Economic Journal 125*(588), F347–F371.

Tables and Figures

Figures

Figure 1: Scenarios and Predictions

Scenario					Predictions: Would parents give more tickets to child L, child H, or give equally to both?		
	Payment Function Parameters Child i's payment in a given scenario is: $P_i = B_i + C_i(TestScore_i - Threshold)$ where i = {L, H} and Threshold = the parent's belief about child L's test score without tutoring, rounded down to the nearest 10						
	Child L (Lower-performing child)		**Child H** (Higher-performing child)				
Scenario	**B$_L$**	**C$_L$**	**B$_H$**	**C$_H$**	**Returns Maximization** Note: R_i = parent's belief about (child i's score with tutoring) - (child i's score without tutoring)	**Inequality Aversion of Outcomes**	**Inequality Aversion of Inputs**
1. Base Case	0	10	0	10	If $R_L = R_H$: No prediction If $R_L < R_H$: H If $R_L > R_H$: L	If $R_L = R_H$: L[1] If $R_L < R_H$: L If $R_L > R_H$: Depends on parameters[2]	Equal
2. Lump Sum to Child L Lump sum transfer to child L	1000	10	0	10	If $R_L = R_H$: No prediction If $R_L < R_H$: H If $R_L > R_H$: L	H	Equal
3. Higher Returns to Child H Higher returns to child H, who has higher expected income without tutoring	0	10	0	100	H (For 96% of parents)	L	Equal
4. Higher Returns to Child L & Lump Sum to Child H Higher returns to child L, who has lower expected income without tutoring	0	100	6000	10	L (For 95% of parents)	L	Equal
5. Higher Returns to Child L Higher returns to child L, who almost always has higher expected income without tutoring	0	100	0	10	L (For 95% of parents)	H (For 98% of parents)	Equal

Notes: Child L (lower-performing child) defined as the one who the parent perceived would have a (weakly) lower test score without tutoring.
1. Assumes test score for child L without tutoring is strictly less than for child H without tutoring. If equal, no prediction.
2. Defining S_i as the test score without tutoring, one can solve that tickets to $L = 10(S_H - S_L + R_H)/(R_L + R_H)$ (unless that quantity falls outside of the 0-10 range).

Figure 2: Consistent with inequality aversion over inputs, the share of baseline educational expenditures on Child L has a spike at 50%

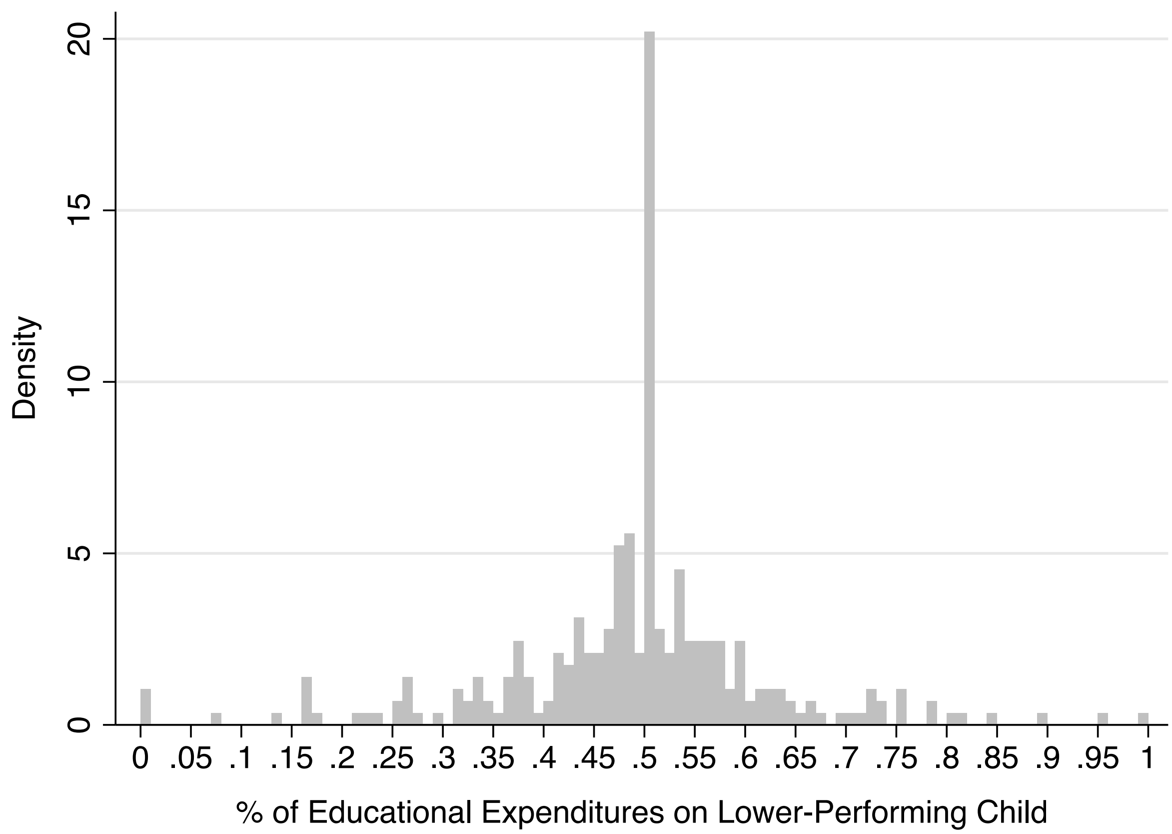

Notes: Figure shows the distribution of the percent of educational expenditures on Child L, as a share of the total educational expenditures on both sampled children, in the study sample. Some households had more than 2 children. In this case, we calculate percent spending using expenditure data for the 2 children in our study sample, as we did not gather expenditure data for the other children in the household.

Figure 3: Experimental ticket allocations, pooled across scenarios 1-5, show a meaningful spike at 50%

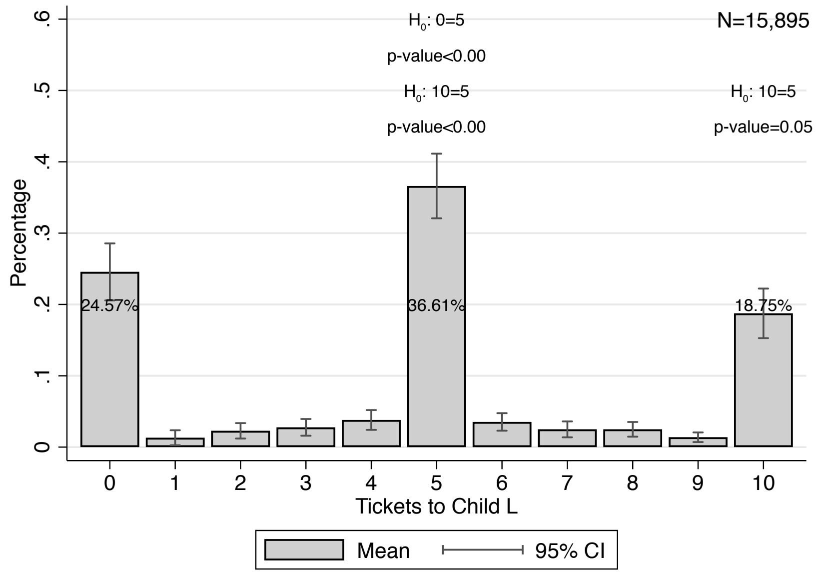

Notes: This figure presents the allocation of lottery tickets across children, pooled across scenarios 1-5 and 95% confidence intervals.

Figure 4: A substantial share of parents choose equal inputs in each scenario

Ticket allocations, by scenario

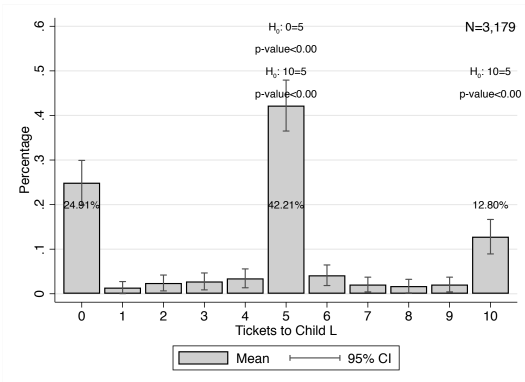

(a) Scenario 1 (*Base Case*)

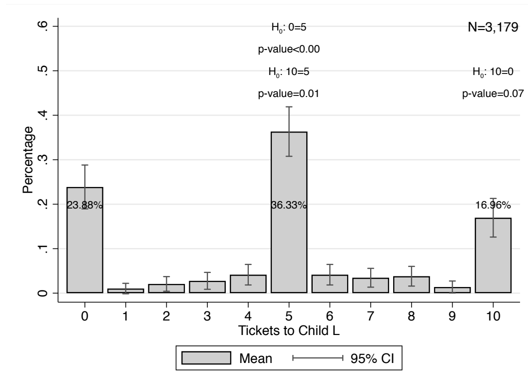

(b) Scenario 2 (*Lump Sum to Child L*)

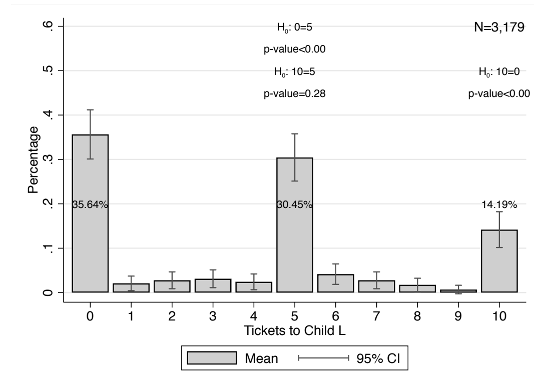

(c) Scenario 3 (*Higher Returns to Child H*)

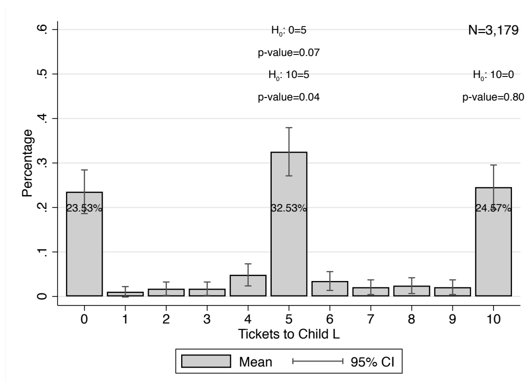

(d) Scenario 4 (*Higher Returns to Child L & Lump Sum to Child H*)

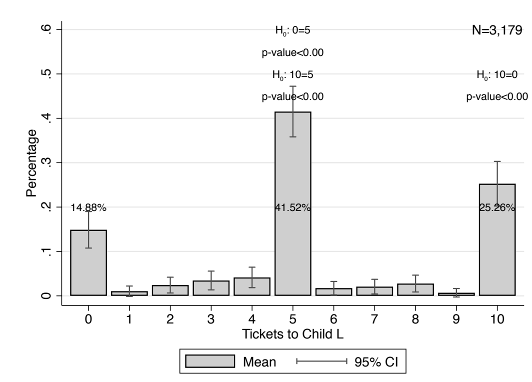

(e) Scenario 5 (*Higher Returns to Child L*)

Notes: This figure presents the allocation of lottery tickets across children and 95% confidence intervals for each scenario.

Figure 5: Inconsistent with IAO, when the lump sum transfer to Child L increases, parents do not reallocate inputs to the higher-performing child.

Relative to Scenario 1 (Base Case), Scenario 2 (Lump Sum to Child L) increases the lump sum relatively more for Child L.

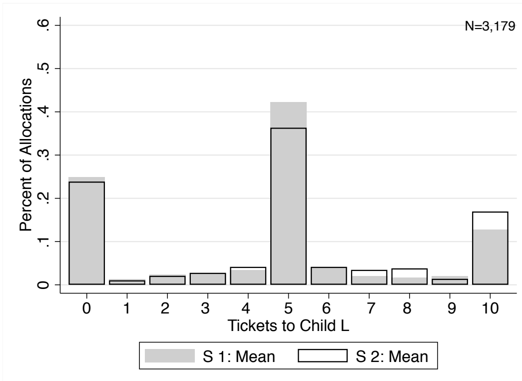

(a) Raw Ticket Allocations

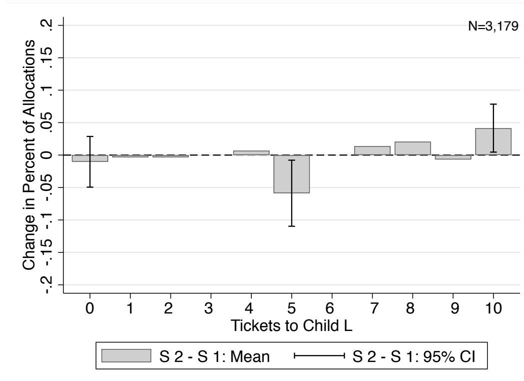

(b) Change in Ticket Allocations

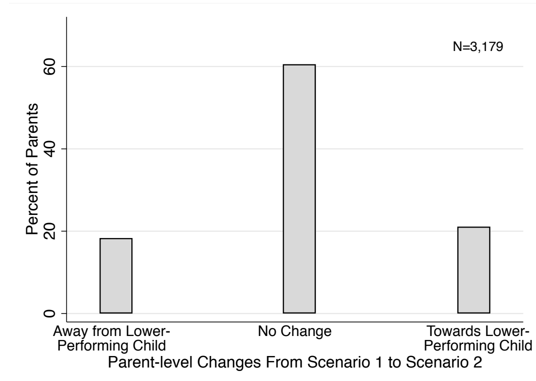

(c) Individual Parent-Level Changes

Relative to Scenario 4 (Higher Returns to Child L & Lump Sum to Child H), Scenario 5 (Higher Returns to Child L) increases the lump sum relatively more for Child L.

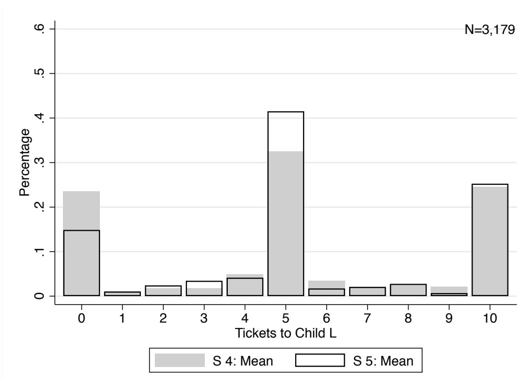

(d) Raw Ticket Allocations

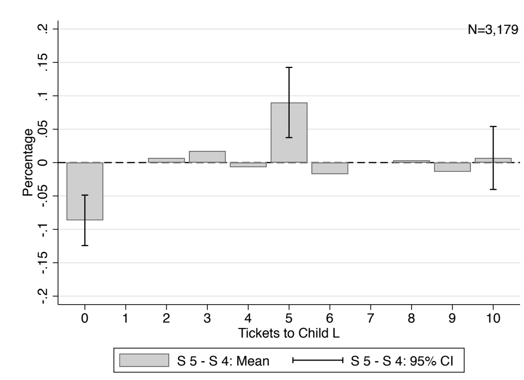

(e) Change in Ticket Allocations

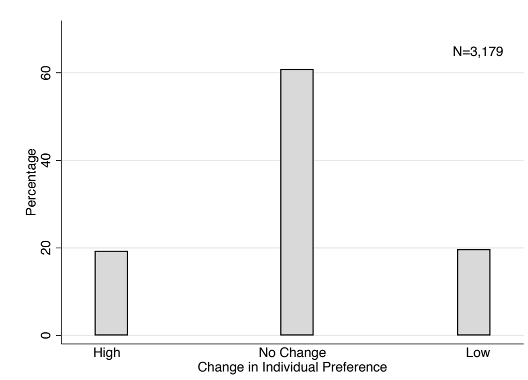

(f) Individual Parent-Level Changes

Notes: Panels (a), (b), and (c) present the distribution of ticket allocation by parents' preferences across Scenarios 1 (Base Case) and 2 (Lump Sum to Child L), while panels (d), (e), and (f) present the distribution of ticket allocation by parents' preferences across Scenarios 4 (Higher Returns to Child L & Lump Sum to Child H) and 5 (Higher Returns to Child L).

Figure 6: There are meaningful differences across parents' choice set in expected total payments and expected outcomes inequality

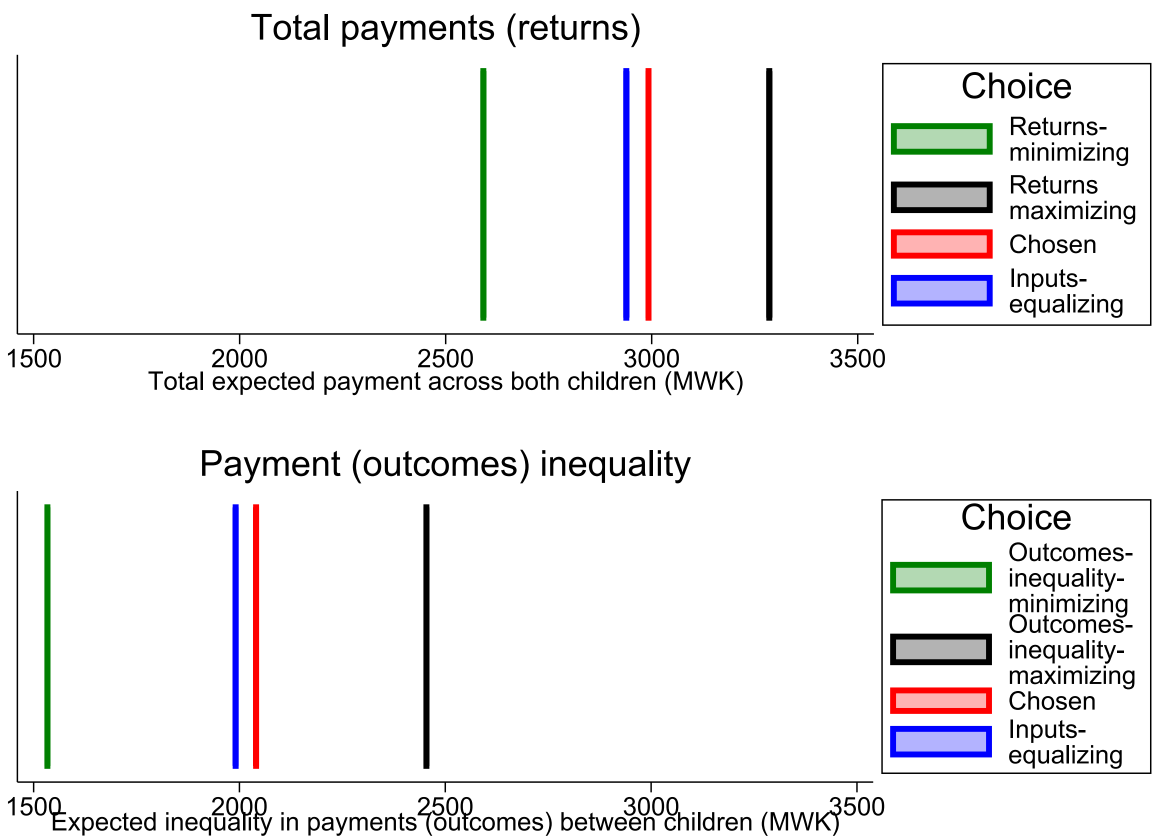

Notes: Both panels represent averages across all parents and scenarios. The first panel shows the average total expected payments (returns) associated with four different potential ticket allocations parents could have chosen: the returns-maximizing choice (i.e., the choice that maximized total expected payments), the returns-minimizing choice, the "IAI" allocation that minimized the gap in inputs/tickets, and the allocation parents actually chose. The second panel shows the analogous statistics but for outcomes inequality; in particular, it shows the average expected *gap* in payments across children associated with: the IAO choice (i.e., the choice that minimized the expected gap in payments), the choice that maximized the expected gap in payments, the IAI choice (input-inequality-minimizing choice), and the allocation parents actually chose. Appendix Figure A.6 shows these averages separately for Scenarios 1-5.

Figure 7: When the per-point reward for a child increases, parents shift inputs toward that child, consistent with returns maximization over inequality aversion in outcomes.

Compared with Scenario 3 (Higher Returns to Child H), Scenario 5 (Higher Returns to Child L) has a higher per-point reward for Child L relative to Child H.

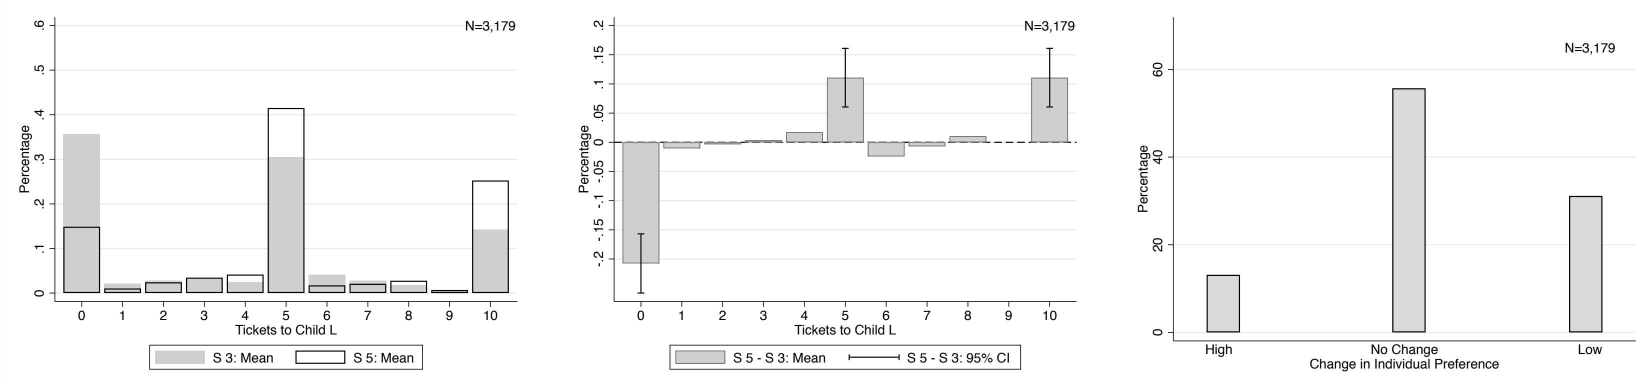

(a) Raw Ticket Allocations

(b) Change in Ticket Allocations

(c) Individual Parent-Level Changes

Compared with Scenario 3 (Higher Returns to Child H), Scenario 1 (Base Case) has a higher per-point reward for Child L relative to Child H.

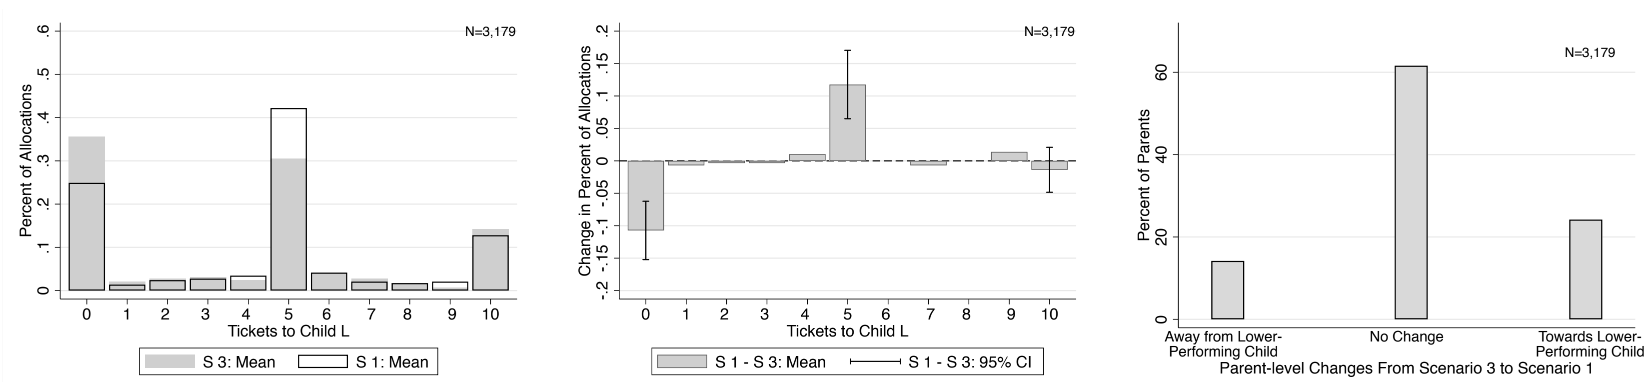

(d) Raw Ticket Allocations

(e) Change in Ticket Allocations

(f) Individual Parent-Level Changes

Notes: Panels (a), (b), and (c) present the distribution of ticket allocations to Child L for Scenarios 3 (Higher Returns to Child H) and 5 (Higher Returns to Child L), while panels (d), (e), and (f) present the distribution of ticket allocation to Child L for Scenarios 1 (Base Case) and 3 (Higher Returns to Child H).

Figure 8: Tickets allocated to Child L across scenarios

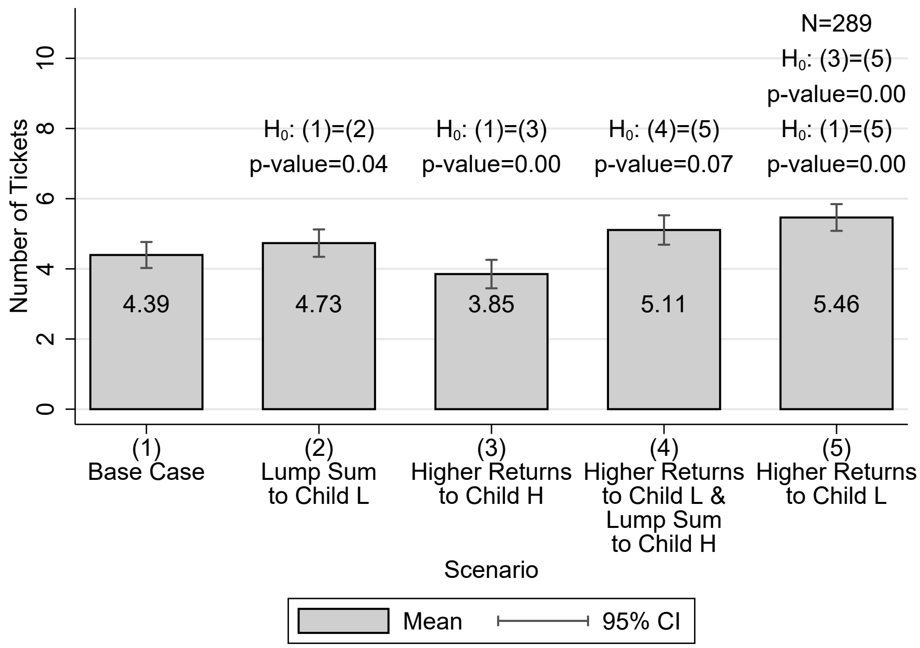

Summary of Scenarios: Scenario 1 (Base Case) pays each child 10 MWK for each point above the threshold. We now summarize the other scenarios relative to the Base Case. Scenario 2 adds a 1,000 MKW lump sum transfer to Child L. Scenario 3 has higher returns to Child H, with Child H receiving 100 MWK for each point above the threshold. Scenario 4 has higher returns for Child L and a lump sum transfer to Child H, with Child L receiving 100 MWK for each point above the threshold and adding a 6,000 MWK (8.26 USD) lump sum transfer to Child H. Scenario 5 has higher returns to child L, with Child L receiving 100 MWK for each point above the threshold.

Notes: This figure presents lottery tickets allocated (out of 10) to Child L (and 95% confidence intervals). We present p-values for tests that ticket allocation to Child L is the same for (i) scenario 1 and scenario 2, (ii) scenario 1 and scenario 3, (iii) scenario 4 and scenario 5, (iv) scenario 1 and scenario 5, and (v) scenario 3 and scenario 5.

Figure 9: Ticket allocations, by whether *inequality aversion in outcomes* (IAO) and *returns maximization* (RM) have the same or opposite predictions

Versions where x-axis represents tickets to Child L

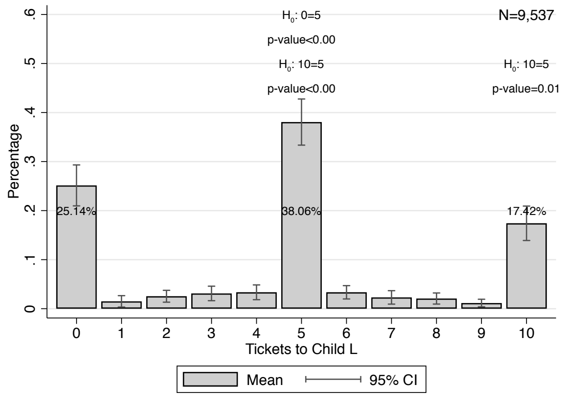

(a) Scenarios where IAO and RM have opposite predictions (1 - *Base Case*, 3 - *Higher Returns to Child H*, and 5 - *Higher Returns to Child L*)

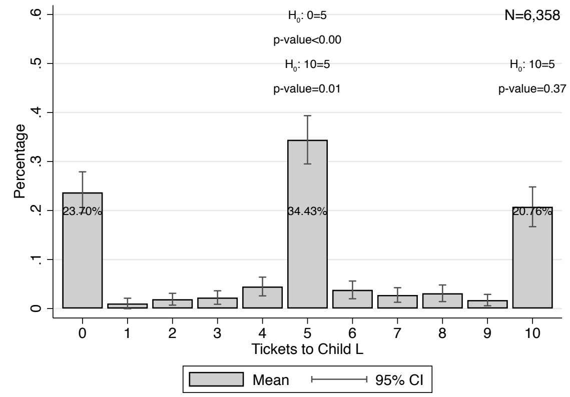

(b) Scenarios where IAO and RM have the same prediction (2 - *Lump Sum to Child L* and 4 - *Higher Returns to Child L & Lump Sum to Child H*)

Versions where x-axis represents alignment with predictions

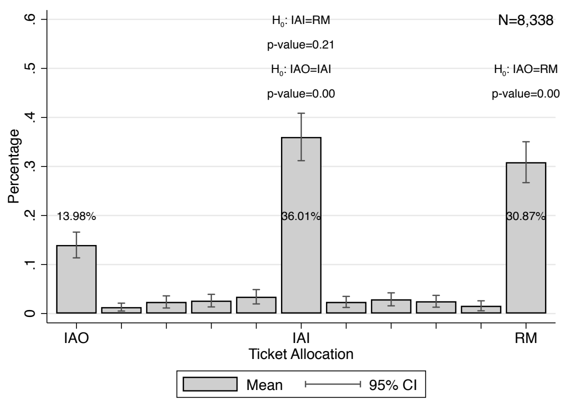

(c) Scenarios where IAO and RM have opposite predictions (1 - *Base Case*, 3 - *Higher Returns to Child H*, and 5 - *Higher Returns to Child L*)

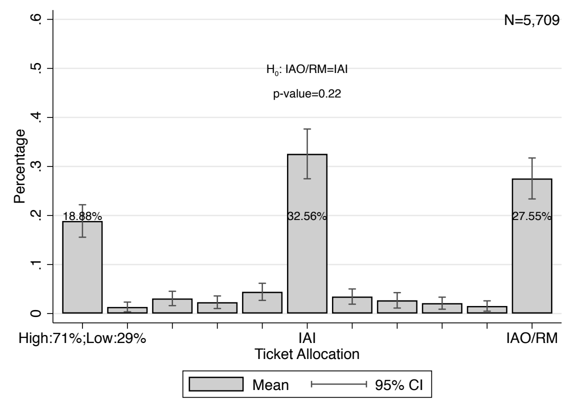

(d) Scenarios where IAO and RM have the same prediction (2 - *Lump Sum to Child L* and 4 - *Higher Returns to Child L & Lump Sum to Child H*)

Notes: These figures presents the allocation of lottery tickets across children. The left figures show all scenarios where inequality aversion over outcomes (IAO) and returns maximization (RM) have opposite predictions. The right figures show all scenarios where inequality aversion over outcomes (IAO) and returns maximization (RM) have the same prediction. The y-axis shows the percent of scenarios in which each allocation was chosen. For panels (a) and (b), the x-axis is the number of tickets given to Child L. For panels (c) and (d), the x-axis shows alignment with the predictions. For panel (c), the x-axis shows the 11 potential allocations (from giving all 10 tickets to one child to all tickets to the other), ordered from the allocation that gives all tickets to the child who would be the choice if the parent only cared about IAO, to the allocation that gives all tickets to the returns-maximizing child. The middle allocation (5 to each child) is labeled IAI as that is the allocation which parents would choose if they only placed weight on IAI. For panel (d), the x-axis is now ordered from the allocation that gives all tickets to the child whom neither IAO nor RM would suggest the parent should give the tickets to, to the allocation that gives all tickets to the child whom both IAO and RM would predict the parent should give all tickets to. For subfigures (c) and (d), we omit scenarios where there were not clean predictions for RM; thus the sample is slightly different than for (a) and (b) which include all scenarios.

Figure 10: Foregone Variable Earnings, Benchmarked Against Maximum Possible Foregone Earnings

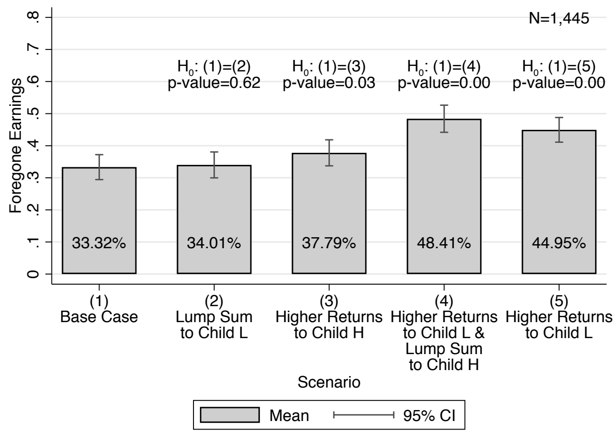

Notes: This figure presents variable income foregone by parents (and 95% confidence intervals) as a percentage of maximum possible foregone earnings. We present p-values for tests that earnings foregone is the same for (i) scenario 1 (Base Case) and scenario 2 (Lump Sum to Child L), (ii) scenario 1 (Base Case) and scenario 3 (Higher Returns to Child H), (iii) scenario 1 (Base Case) and scenario 4 (Higher Returns to Child L & Lump Sum to Child H), and (iv) scenario 1 (Base Case) and scenario 5 (Higher Returns to Child L).

Figure 11: Foregone Earnings

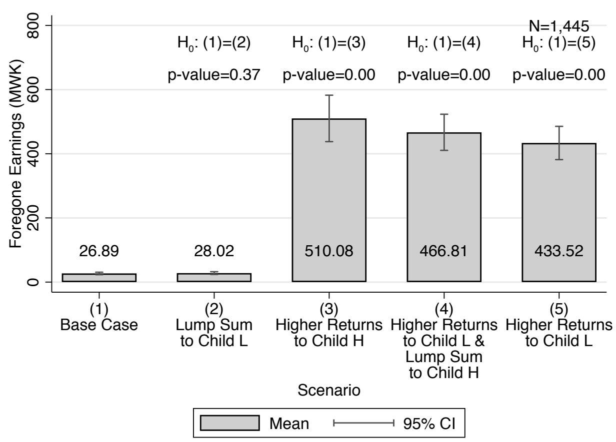

Notes: This figure presents income foregone by parents (and 95% confidence intervals). Foregone earnings are calculated as the difference between expected income if parents were maximizing returns and expected income based on their preferred lottery ticket allocation. We present p-values for tests that earnings foregone is the same for (i) scenario 1 (Base Case) and scenario 2 (Lump Sum to Child L), (ii) scenario 1 (Base Case) and scenario 3 (Higher Returns to Child H), (iii) scenario 1 (Base Case) and scenario 4 (Higher Returns to Child L & Lump Sum to Child H), and (iv) scenario 1 (Base Case) and scenario 5 (Higher Returns to Child L).

Figure 12: Evidence of equalizing in a setting where parents cannot ex-post equalize

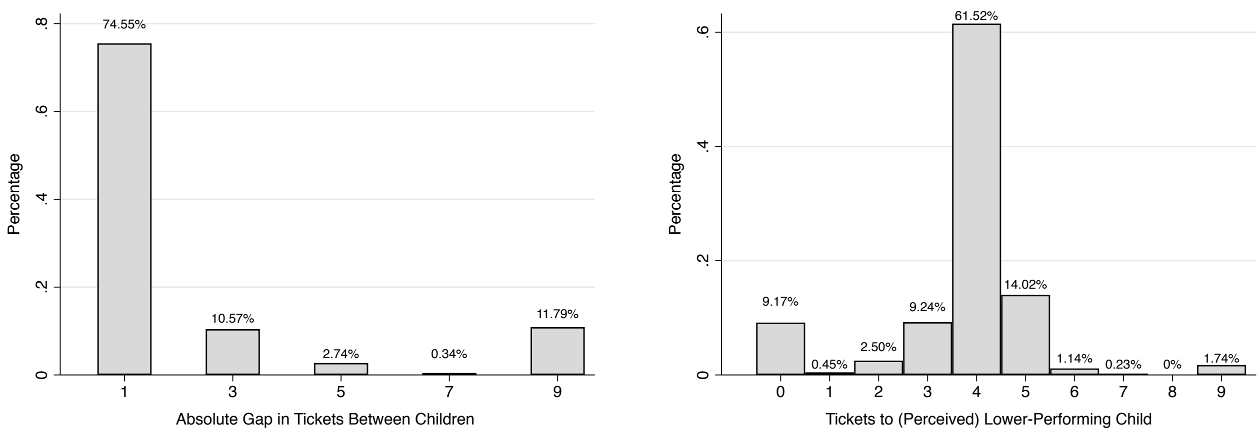

(a) Absolute Gap in Tickets (b) Tickets Allocated to Lower-Performing Child

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."

Figure 13: Fraction of Equalizers by Bin of Foregone Earnings

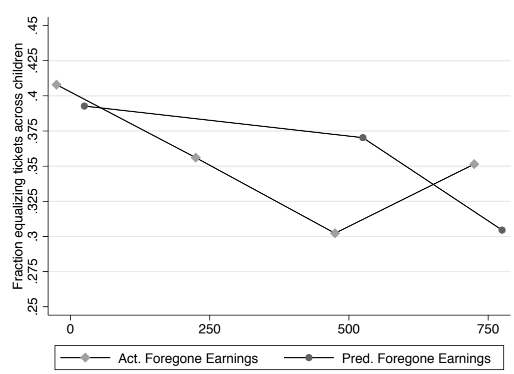

Notes: This figure presents the fraction of parents who equalize tickets by bin of potential foregone earnings. For each household and scenario, potential foregone earnings are calculated as the difference between expected earnings from the returns-maximizing choice and expected earnings from the input-equalizing choice. The lighter diamonds represent actual foregone earnings, and the darker circles represent foregone earnings predicted by average foregone earnings for that scenario.

Figure 14: Mixed logit estimates of willingness to pay for different ticket allocations

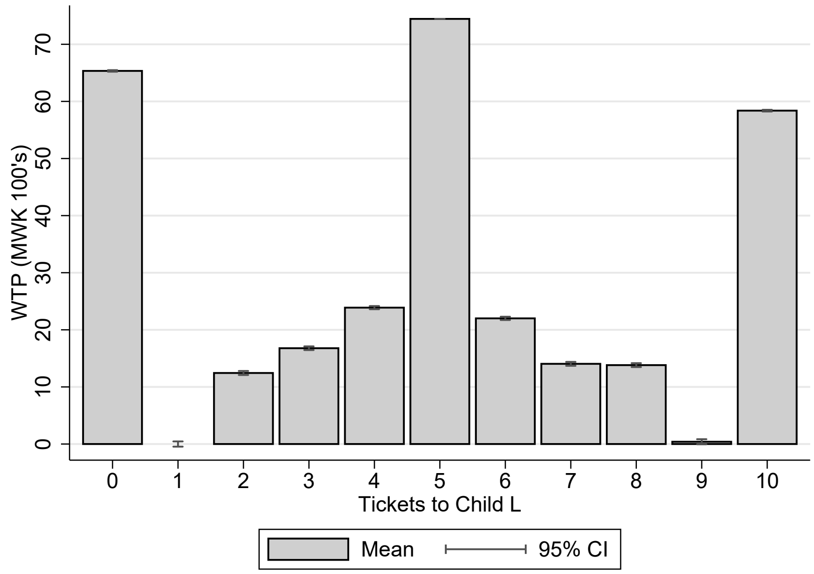

Notes: Figure present estimates of parents' preferences, represented in willingness-to-pay units, for different ticket allocations. Estimates are created by estimating a mixed logit specification based on (5) but where we allow for ticket-allocation fixed effects for all potential choices. WTP is calculated by dividing each coefficient by the coefficient on the *Household earnings* term. WTP numeric estimates are all relative to the least-preferred option (the option giving exactly 1 ticket to Child L). Confidence intervals are for tests for equality for each option vs. choice 5 (equal allocation).

Tables

Table 1: Summary Statistics

Total Households		289	
A. Pre-Experiment Survey Question:			
Compared to child L, how much do you spend on the education of Child H?	**Equal**	**Less**	**More**
Percentage of households (%)	85.12%	3.46%	11.42%
B. Pre-Experiment 'Falsification' Test:			
Choose value of 10 lottery tickets; 50 MWK vs. 100 MWK?	**50 MWK**	**100 MWK**	
Mean ticket allocation (out of 10)	0.12 Tickets	9.88 Tickets	
C. Respondent Characteristics:			
Female (%)		0.94	
		(0.24)	
Education > Class 8 (% of parents)		0.16	
		(0.36)	
Thought tutoring had positive returns for at least one child (% of parents)		0.99	
		(0.08)	
Thought tutoring had positive returns for both children (% of parents)		0.93	
		(0.26)	
D. Child Characteristics:	**Child L**	**Child H**	
Grade			
Class 5 (% of children)	0.41	0.38	
Class 6 (% of children)	0.32	0.36	
Class 7 (% of children)	0.27	0.26	
Female (%)	0.55	0.56	
Parents beliefs without tutoring			
Mean (out of 100)	53.14	64.43	
Parents beliefs with tutoring			
Mean (out of 100)	64.43	82.43	
Parent thought child had positive returns to tutoring			
(% of children)	0.96	0.96	
Parent thought child had strictly higher returns to tutoring than sibling			
(%of children)	0.14	0.66	
Who received tutoring?			
(% received tutoring)	0.43	0.57	
Math test score			
Mean (out of 100)	41.92	44.14	
Average annual household education expenditure			
Mean (in MWK)	8412.06	8372.10	

Notes: This table presents baseline summary statistics for the sample. Standard deviations are in parentheses. The parent believe tutoring had "positive returns" means that the parent perceived that the child's test score would increase as a result of tutoring.

Table 2: Stability of Preferences

Total Households	289
Stability of Preferences Across Scenarios:	
IAI all scenarios (% of parents)	0.19
RM all scenarios (% of parents)	0.06
IAO all scenarios % of parents)	0.00

Notes: This table presents the proportion of parents who only preferred to equalize inputs, maximize returns, or equalize outcomes for all scenarios.

Table 3: Foregone Earnings and Input Equalization

	Dependent Variable: Split Allocation 5/5	
	(1)	(2)
	OLS	IV
	β / SE	β / SE
Foregone Earnings from Splitting ('00)	-0.0050	-0.0103***
	(0.0035)	(0.0031)
Constant	0.3834***	0.4017***
	(0.0263)	(0.0260)
Observations	1445	1445
R^2	0.002	.

Notes: This table presents regressions of measures of input-equalizing choices on the difference between earnings under the returns-maximizing allocation of tickets and the equalizing allocation. The dependent variable is an indicator for whether the parent split the tickets equally between children. Column (2) instruments foregone earnings for the household and card with average earnings for that card across all households. Standard errors are clustered at the household level.

Table 4: Mixed Logit Estimates of Parental Preferences for Investment

	(1) Mixed Logit β / SE	(2) Mixed Logit β / SE	(3) OLS β / SE	(4) IV β / SE
Household earnings (MWK100)	0.0779***	0.1659***	0.0037***	0.0050***
	(0.0238)	(0.0312)	(0.0009)	(0.0010)
Absolute difference in inputs	-0.2574***			
	(0.0380)			
Inputs not equally split (0/1)		-2.7257***	-0.3027***	-0.3027***
		(0.1904)	(0.0269)	(0.0256)
Absolute difference in earnings (MWK100)	0.0252	0.0313	0.0008	0.0011
	(0.0184)	(0.0206)	(0.0006)	(0.0007)
SD Household Earnings	0.29	0.43		
SD Absolute difference in inputs	1.13			
SD Inputs not equally split		3.79		
SD Absolute difference in earnings	0.11	0.20		
WTP for 1 unit lower input inequality (MWK100)	3.3			
WTP for equal inputs (MWK100)		16.4		
Observations	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. The dependent variable is a dummy for whether the parent chose that allocation for that scenario. Columns 1 and 2 are estimated using a mixed logit model, and the WTP's are calculated by dividing the "Absolute difference in inputs" (col. 1) and "Inputs not equally split" (col. 2) coefficients by the "Household earnings" coefficient. Column 3 is estimated using OLS. Column 4 is estimated using 2SLS, with dummies for the scenario × ticket allocation as instruments for the first two regressors. Standard errors in parentheses, clustered at the household level.

Appendices

Appendix A: Tables and Figures

Figures

Appendix Figure A.1: Sequence of Events

Household survey with the primary caregiver

↓

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

↓

"Training": practice scenarios; explain what to do to maximize returns, equalize outcomes, or equalize inputs; practice money lottery

↓

Experiment: ticket allocation across 5 (real) scenarios; randomly choose scenario; lottery

↓

Winner of the lottery gets an hour of tutoring; both children take the math test

↓

Post-experiment survey with primary caregiver

↓

Reward money handed to children in individual envelopes

Appendix Figure A.2: Cross-Country Data from the Young Lives Survey (YLS) Also Shows Excess Density in Spending at 50%

Percent of Educational Expenditures on the "YLS child" in YLS Households with 2 School-Age Children

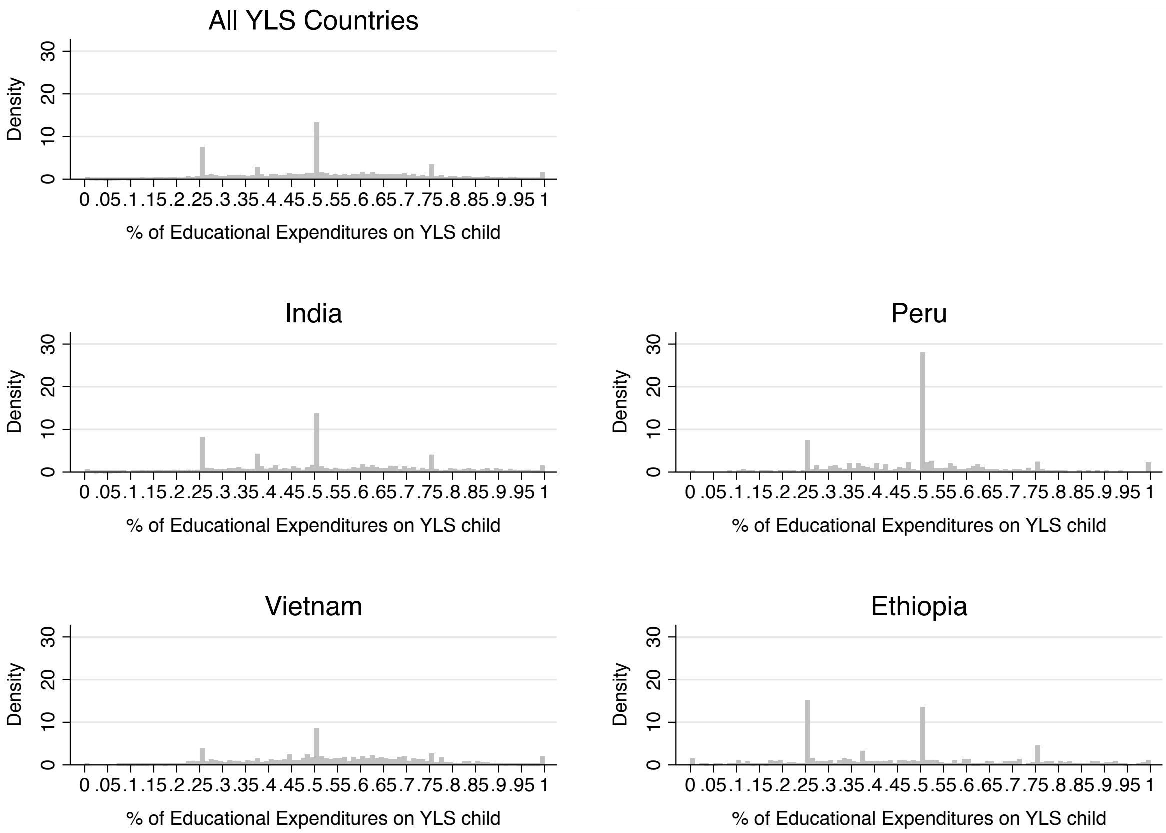

Notes: Figure shows distribution of percent of educational expenditures on the "YLS child" (i.e., the child who was the chosen respondent for the survey) in households with 2 school-age children, using data from the Young Lives Survey. The survey asks for total household expenditure in multiple subcategories of education spending, and also asks what share of that spending is spent on the YLS child. The response options (none, less than half, roughly half, more than half but not all, all) were coded as 0, 0.25, 0.5, 0.75, and 1, respectively. We then calculate total spent on YLS child across all education subcategories and divide by the total household spending on education. Because of the coding method, we naturally see spikes at 0.25, 0.5, and 0.75. However, among 2-children households, the highest density is seen at 0.5, which may be suggestive of an equal division of spending between the 2 children.

Appendix Figure A.3: Heterogeneity in allocations for Scenarios 2-5, by whether parents allocated more tickets to Child H in Scenario 1

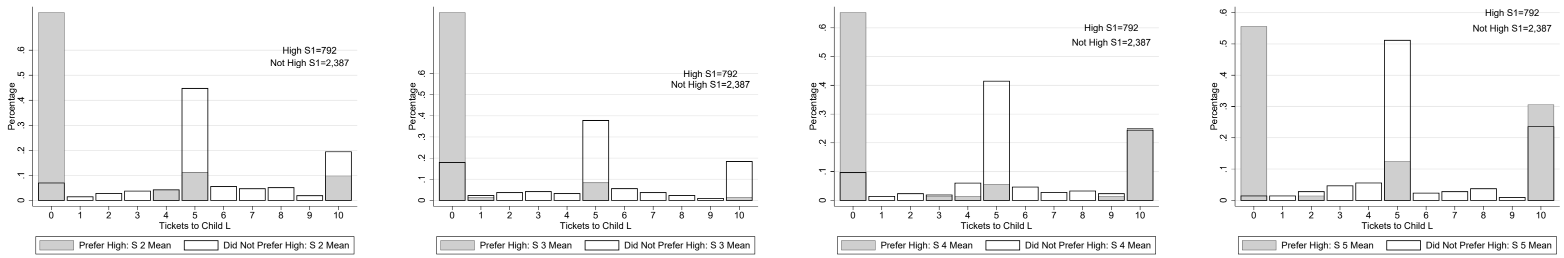

(a) Scenario 2 *(Lump Sum to Child L)* (b) Scenario 3 *(Higher Returns to Child H)* (c) Scenario 4 *(Higher Returns to Child L, Lump Sum to Child H)* (d) Scenario 5 *(Higher Returns to Child L)*

Notes: This figure presents the allocation of lottery tickets across children and 95% confidence intervals for each scenario.

Appendix Figure A.4: Heterogeneity in allocations for Scenarios 2-5, by whether parents allocated more tickets to Child L in Scenario 1

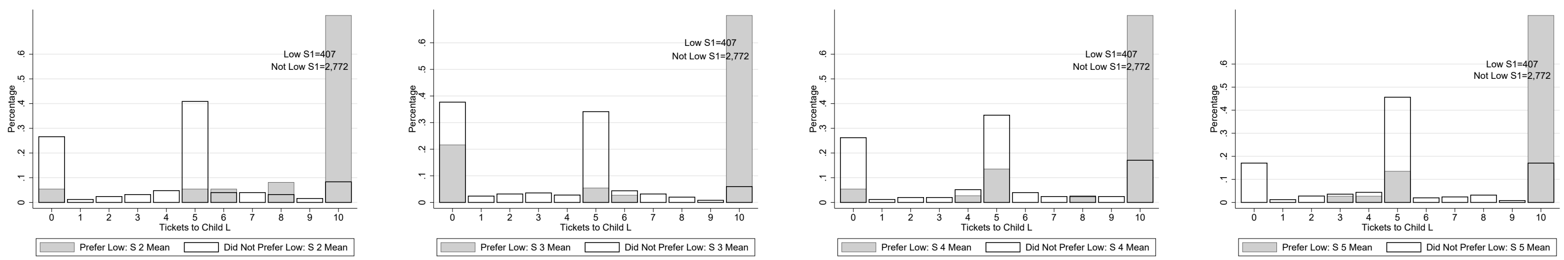

(a) Scenario 2 *(Lump Sum to Child L)* (b) Scenario 3 *(Higher Returns to Child H)* (c) Scenario 4 *(Higher Returns to Child L, Lump Sum to Child H)* (d) Scenario 5 *(Higher Returns to Child L)*

Notes: This figure presents allocation of lottery tickets across children for Scenarios 2-5, and 95% confidence intervals for each scenario, separately by how parents allocated their tickets in Scenario 1 (Base Case).

Appendix Figure A.5: Heterogeneity in ticket allocations by parents' beliefs uncertainty

Heterogeneity by parents' baseline measure of uncertainty

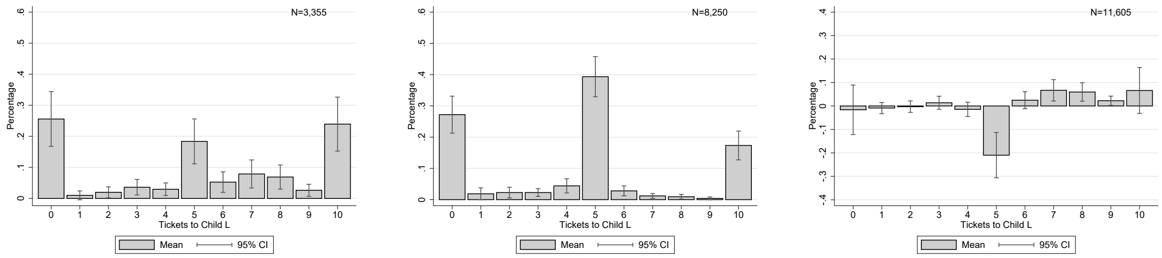

(a) Parents who are uncertain about their beliefs

(b) Parents who are certain about their beliefs

(c) Difference in means (uncertain - certain)

Heterogeneity by whether parents changed beliefs between baseline survey and experiment

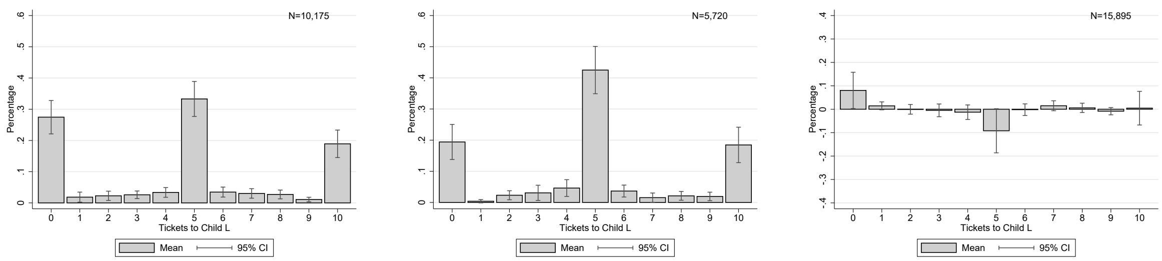

(d) Parents who changed their beliefs

(e) Parents who did not change their beliefs

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

Heterogeneity by difference between parents' beliefs and actual scores

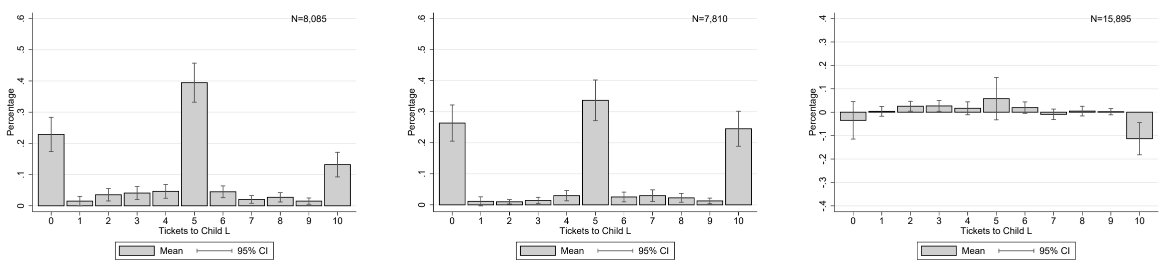

(g) Difference between beliefs and actual scores $\geq$ p50

(h) Difference between beliefs and actual scores $<$ p50

(i) Difference in means ($\geq$ p50 - $<$ p50)

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.

Appendix Figure A.6: There are meaningful differences in expected payments and expected outcomes inequality across the choice set

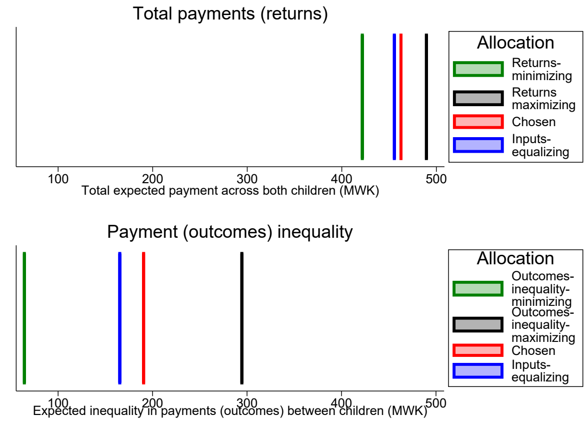

(a) Scenario 1 *(Base Case)*

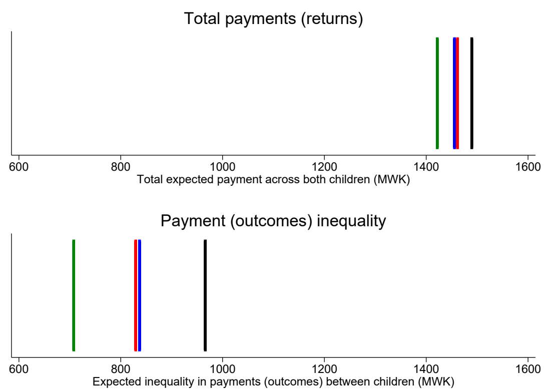

(b) Scenario 2 *(Lump Sum to Child L)*

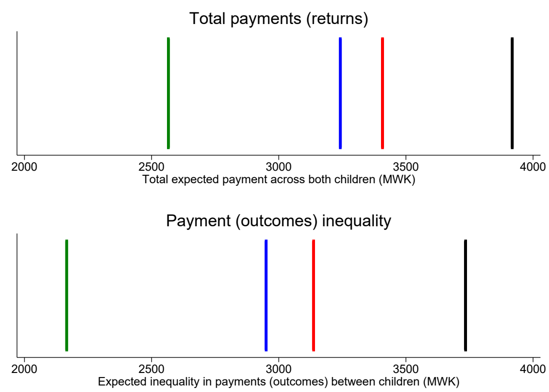

(c) Scenario 3 *(Higher Returns to Child H)*

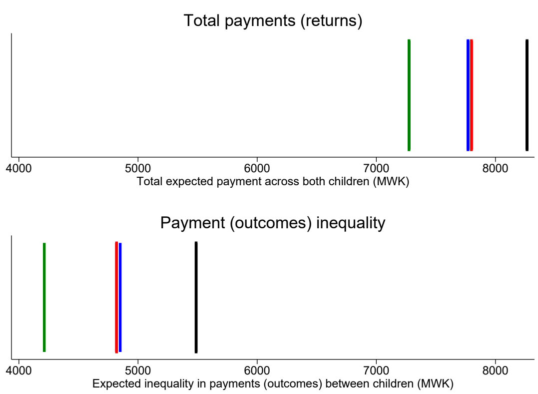

(d) Scenario 4 *(Higher Returns to Child L & Lump Sum to Child H)*

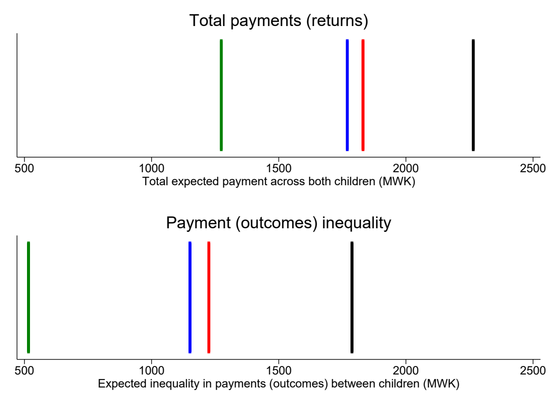

(e) Scenario 5 *(Higher Returns to Child L)*

Notes: Each figure shows the statistics from Figure 6, but now calculated separately at the scenario level. The legend from subfigure (a) applies to all subfigures.

Tables

Appendix Table A.1: No evidence of IAO even among parents who can perfectly equalize outcomes

Scenario 1 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. 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. The P-value reported in column 3 tests for a difference in means between columns 1 and 2.

Appendix Figure A.7: Cross-Country Data from the Young Lives Survey (YLS) Also Shows Excess Density in Spending at 50%

Percent of Educational Expenditures on the "YLS child" in YLS Households with 3 School-Age Children

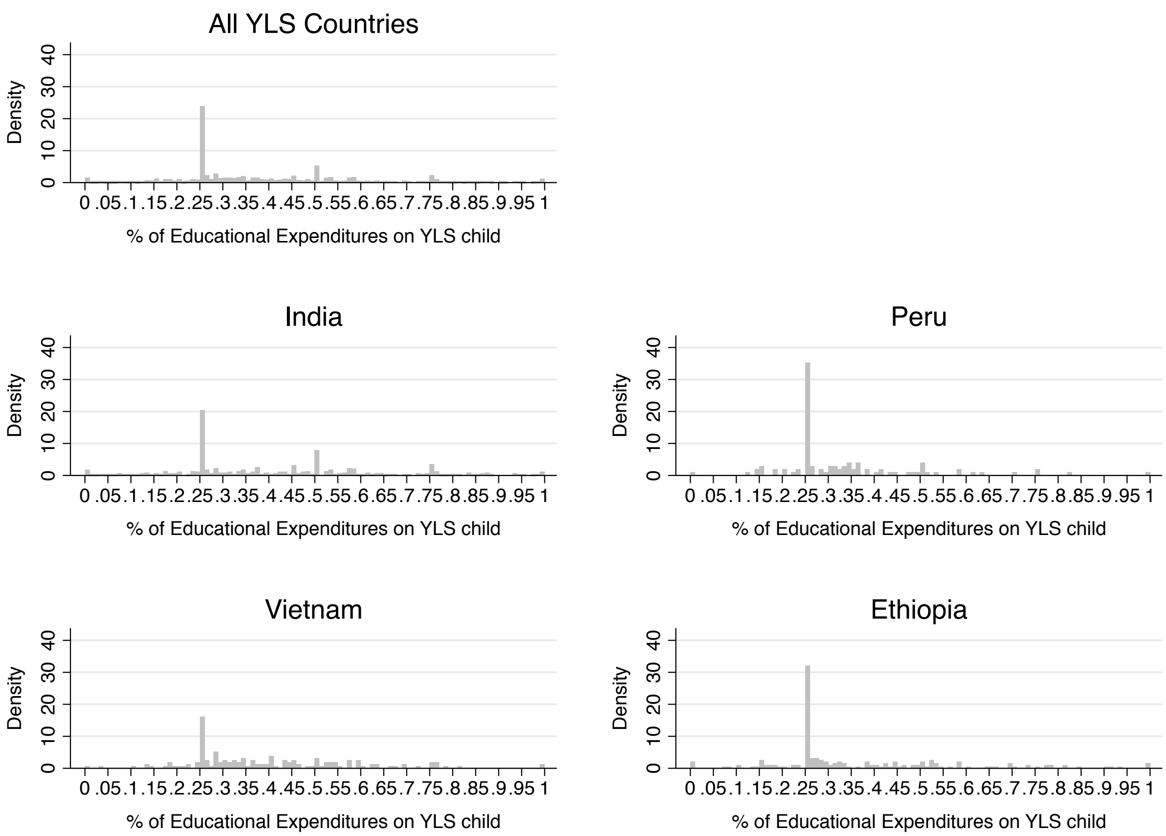

Notes: Figure shows distribution of percent of educational expenditures on YLS child in households with 3 school-age children, using data from the Young Lives Survey. The survey asks for total household expenditure in multiple subcategories of education spending, and also asks what share of that spending is spent on the YLS child. The response options (none, less than half, roughly half, more than half but not all, all) were coded as 0, 0.25, 0.5, 0.75, and 1, respectively. We then calculate total spent on YLS child across all education subcategories and divide by the total household spending on education. Because of the coding method, we naturally see spikes at 0.25, 0.5, and 0.75. However, among 3-children households, the highest density is seen at 0.25, which may be suggestive of a roughly equal division of spending between the 3 children.

Appendix Table A.2: Heterogeneity Results

	(1) Child A Tickets β / SE	(2) IAI β / SE	(3) RM β / SE	(4) IAO β / SE
Caregiver > Class 8	0.63*	0.12*	0.01	-0.03
	(0.38)	(0.07)	(0.05)	(0.03)
Constant	4.61***	0.35***	0.27***	0.13***
	(0.17)	(0.02)	(0.02)	(0.01)
Observations	1445	1445	867	867
R^2	0.004	0.008	0.000	0.001

Notes: 'Caregiver > Class 8' is a binary variable that takes the value 1 if the primary caregiver has at least completed Grade 8; 0 otherwise. Standard errors in parentheses, clustered at the household level.

Appendix B: Sample Script for Scenario 1

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 1

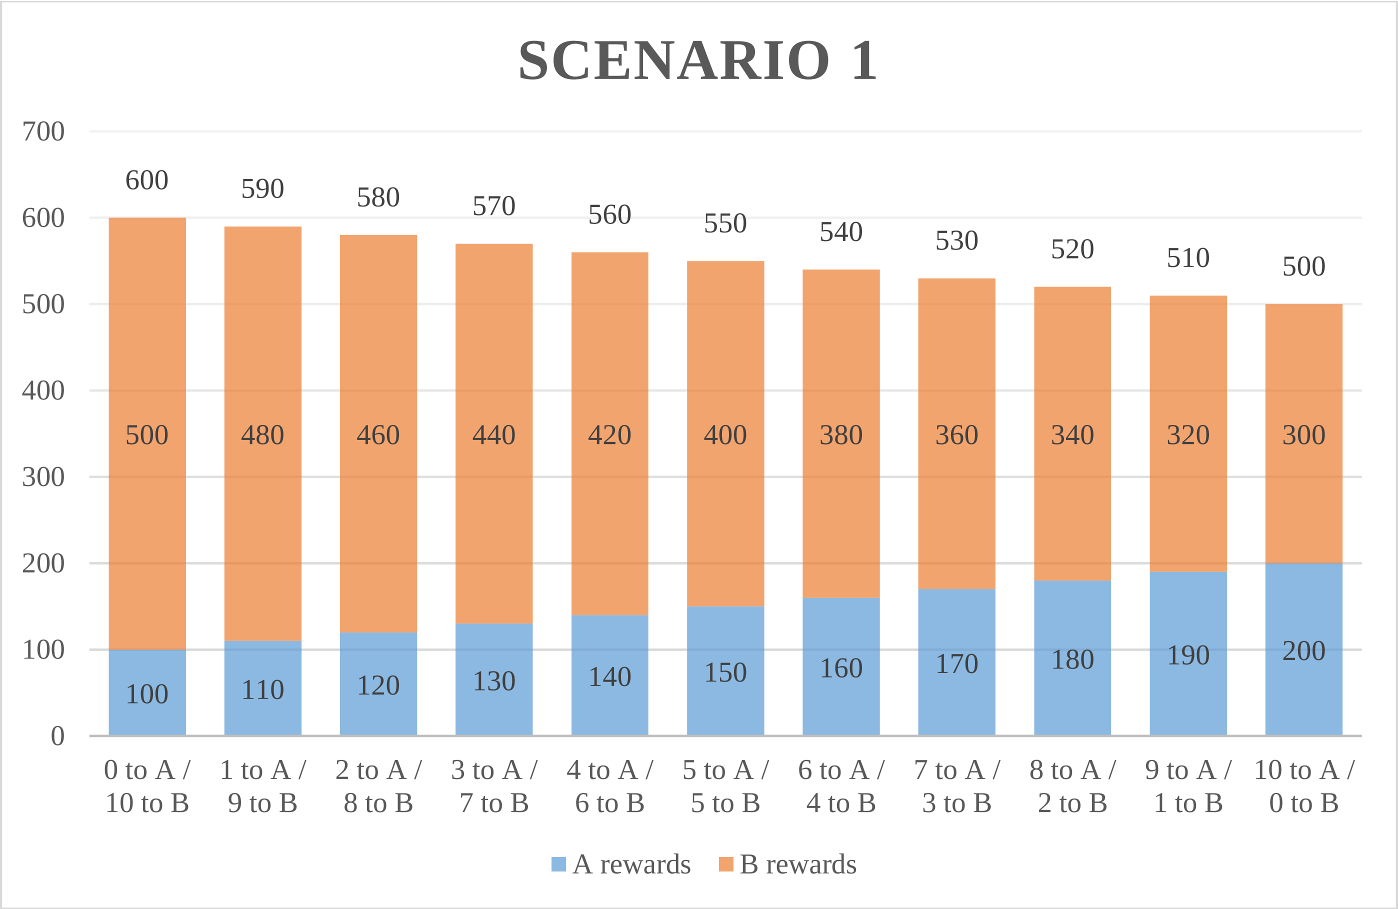

Scenario 1: Visual Aid 2

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