Parent-Bias*

Guilherme Lichand† and Juliette Thibaud‡

December 1, 2020

Abstract

This paper uses a lab-in-the-field experiment in Malawi to document two new facts about how parents share resources with their children over time. First, for almost a third of study participants, the further in the future consumption is, the more generous are parents' plans to share it with their children. Second, many participants revise those plans as consumption gets closer, reallocating from children towards themselves – even when consumption is *still in the future*. None of these patterns can be accounted for by present-bias. Instead, both are consistent with a relevant share of parents discounting their future utility of consumption to a greater extent than that of their children. We document that parents characterized by such *asymmetric geometric discounting* display sizable preference reversals every period, a phenomenon we denote *parent-bias*. We find that, despite ambitious plans, those parents actually allocate less to their children in the present than other parents, and that such preferences predict under-investment in children outside the lab just as much as quasi-hyperbolic discounting. Commitment devices designed for present-bias do not mitigate parent-bias. Our findings provide a new explanation for under-investment in children and inform the design of new interventions to address it.

Keywords: Time preferences; Preference reversals; Children's human capital

JEL Classifications: C91, D13, E24

*We are very grateful to Ned Augenblick, Björn Bartling, Patrick Baxter, Lorenzo Casaburi, David Dorn, Ernst Fehr, Andrew Foster, Linn Jaeckle, Michel Maréchal, Nick Netzer, Nathan Nunn, Gautam Rao, Jesse Shapiro, Frank Schilbach, Andrei Shleifer, Dmitry Taubinsky, Neil Thrakal, Séverine Toussaert, Roberto Weber, Jack Willis, David Yagazinawa-Drott and Fabrizio Zilibotti for helpful discussions that greatly benefited this project. Excellent research assistance by Julien Christen, Maite Deambrosi, Matthias Endres, Francis Jere and Faith Millongo, and field supervision by Nicoló Tomaselli. All remaining errors are our own. This research project was supported by the generosity of UNICEF National Committee for Switzerland and Liechtenstein, UZH Excellence Foundation and Maiores Stiftung. Any views and opinions contained in this paper are those of the authors and do not necessarily reflect the views or opinions of UNICEF.

†University of Zurich; guilherme.lichand@econ.uzh.ch

‡University of Zurich; juliette.thibaud@econ.uzh.ch

"The task of anthropologists, and of economists also, is to collect anecdotes that may or may not lend themselves to testable conjecturing. In actual economic life, and in the writings of economists, the phenomena of how people treat the present and the future (and the past, too) offer a rich treasure trove of varied behaviors." – Samuelson (2008, p. 1)

1 Introduction

In the film *The Boy Who Harnessed the Wind*, a brilliant Malawian boy is barred from attending school after his parents fail to pay his school fees despite repeated promises. The lead character builds a windmill that helps his town escape famine, and eventually becomes an engineer. While his fate is exceptional, the starting point of the movie is all too common. Parents, all over the world, frequently fail to follow through on ambitious plans to invest in their children's health and education.[1]

While liquidity is likely to severely constrain those investments among the poor, interventions designed to relieve poverty often have only small effects on investments in children.[2] Instead, quasi-hyperbolic discounting is a leading explanation for gaps between past intentions and present actions. Time inconsistencies, typically in favor of instant gratification (present-bias), come to rationalize broken promises in general, and under-investment in children in particular. Having said that, present-bias does not fully account for deviations from geometric discounting in the way parents plan investments in children, or in the way the latter more often than not renege on those plans. To that effect, when the Malawian father in the movie realizes the family does not have as many resources as initially expected, he does not decrease spending equally across household members: adjustments to the family budget hurt investments in children *disproportionately*. Present-bias cannot rationalize that asymmetry.[3] In this paper, we document that such *asymmetric*

[1] According to the 2018 World Development Report, while primary enrollment has increased substantially worldwide since 1970, secondary enrollment is still only 40% across Sub-Saharan Africa (p. 59), with huge differences between urban and rural areas (p. 62).

[2] Glennerster & Kremer (2012) and Kremer & Holla (2009) document that demand for educational and health investments in children falls very steeply as soon as prices are above zero (even if very low). Microcredit does not systematically increase investments in children (Banerjee, 2013). While conditional cash transfers and earmarked loans increase those investments, there is evidence that such effects might be driven by other mechanisms – from salience to present-bias (Glennerster & Kremer, 2012).

[3] Present-bias predicts only a *general tendency* to cut down on investment plans when their costs become immediate, *not* a tendency to cut down *differentially* across investment plans for parents and children. While asymmetric present-bias could explain this, we discuss in detail below how the phenomena we document are distinct from it.

reversals emerge as part of systematic patterns in how parents share resources with their children over time. We then introduce a new type of time preferences that can rationalize those patterns. We find such preferences to be as correlated with real investments in children's human capital as quasi-hyperbolic discounting, and uncorrelated with present-bias.

To document how parents plan to (and effectively do) share resources with their children in the future, and how those plans evolve over time, we conduct a lab-in-the-field experiment with 1,627 parents in Malawi. Measuring parents' time preferences is challenging, as it requires observing how they plan to share consumption within the household over time, and whether they actually stick to those plans. We follow the literature in moving away from using decisions about monetary payments to make inferences on time preferences; instead, we draw upon real consumption decisions that can be perfectly observed at the time of the experiment.[4] We study parents' decisions concerning a non-fungible, tempting and nutritious good – peanuts –, documenting their plans to split consumption with their children over different time horizons. The experiment takes place during the lean season in Malawi; as such, we can have subjects undertake payoff-relevant decisions in this setting at a reasonably low cost. Enumerators visit the parents three times. During the first visit (round 1), parents choose how they want to share the peanuts with one of their children by the time enumerators revisit them two days later (round 2) and four weeks later (round 3). There is no consumption at round 1. At round 2, parents have the opportunity to set new consumption plans (potentially different from their round 1 decision), allocating consumption in the present and in little less than four weeks. There is no decision at round 3; only consumption. Every allocation set by parents can be drawn to be implemented with positive probability, and peanuts are consumed in front of the enumerators – precluding side transfers. Such design features several advantages. Peanuts shut down concerns with fungibility and arbitrage outside the lab. Moreover, different from the typical experiment to elicit time preferences, our design does not allow subjects to transfer resources over time. This feature allows us to easily capture

[4] As money is fungible, the timing of the payments and of the consumption acquired with it are not necessarily connected (Augenblick & Rabin, 2019). In addition, subjects may have arbitrage opportunities and their choices over monetary payments could simply reflect the interest rates they could access outside of the lab (Augenblick et al., 2015; Cubitt & Read, 2007). What is more, since parents control how money is spent outside the lab, decisions about how to split money between themselves and their children would be non-committal in the context of our experiment. Examples of real consumption in the time-preferences literature include real effort tasks (Augenblick et al., 2015; Augenblick & Rabin, 2019; Barton, 2015), irritating noises (Solnick & Waller, 1980) and squirts of juice (Brown & Camerer, 2009; McClure et al., 2007).

how parents trade off their own and their children's consumption within period, and how those trade-offs vary with the time gap between plans and consumption. Last, allowing parents to revise their former decision enables us to document preference reversals.

We start by documenting two new facts about how a relevant fraction of parents allocate resources over time. The first is that many parents' planned budget shares allocated to children tend to increase with the time gap between plans and actual consumption. At round 2, almost one third of parents allocate a higher share of the total budget to their children in the future – nearly 46% higher than what they get allocated in the present. The second is that many parents revise their plans away from children's future allocation when the decision gets closer to actual consumption (even if still 28 days away). Such preference reversals are almost 1/3 as common as present-bias in our sample, and are quantitatively large: for those parents, the round-3 budget share allocated to children decreases by 37% on average.[5]

We implement many design choices to rule out that such patterns could be rationalized under standard time preferences. First, we can rule out learning about preferences between decision rounds: parents taste some peanuts before each round to minimize the risk of projection bias (Loewenstein et al., 2003) and children do not consume peanuts before the end of round 2. Second, rounds 1 and 2 are only two days apart, minimizing concerns with other shocks that could affect parents' information set when they can revise their plans. Comparing parents' responses to different interest rates at rounds 1 and 2, we can also rule out directly that parents' (expected) marginal utility of consumption at round 2 has changed between rounds.[6,7] We also document that parents do not strategically adjust their child's consumption outside of the experiment between visits.[8]

Most importantly, none of those facts can be accounted for by present-bias. Without shocks to the expected marginal utility of consumption between rounds, present-bias predicts that children's budget shares should be constant over time, and that plans should only be revised once they have immediate consequences. Instead, those stylized facts are consistent with *asymmetric geometric discounting* (AGD). For concreteness, consider a mother who discounts her children's future

[5]They are also consequential: parents have no further opportunity to revise their plans at round 3.

[6]See Supplementary Appendix S3.

[7]In addition, we survey parents about liquidity constraints and hunger at every round and control for these confounding factors.

[8]See Appendix C.3.

consumption to a lesser extent than her own. Those time preferences are a deviation from geometric discounting; she thinks that her future self will (or wishes she would) be more generous towards her children than her present self, which in turn leads to systematic preference reversals: every period, she reallocates budget shares, away from her children's planned consumption and towards her own. For this reason, we call reversals generated by AGD preferences *parent-bias*. Underinvestment in children (relative to her original plans) arises as a direct implication of being relatively less patient about her own future consumption. While preference reversals have been shown to be a necessity when multiple decision-makers with different discount rates allocate a single payoff stream for the group (Jackson & Yariv, 2015), parent-bias follows from a *single decision-maker* – a unitary parent – with *multiple discount rates* deciding on how to allocate a single payoff stream across multiple group members – herself and her child.

Parent-bias has important implications for the following reasons. First, because AGD preferences are likely to be prevalent; after all, time inconsistencies often arise as a conflict between deliberation and affect (Loewenstein, 2018), and children bring about sharp tensions between the two: parents often aspire to provide a better future for their children, but are pressed against present needs. Second, because commitment devices that address present-bias, such as lock-boxes or illiquid accounts, do not mitigate parent-bias: preventing within-household reallocation requires commitment devices that are specifically designed to limit decision-makers ability to change past plans *differentially* across household members.

Using only round-1 decisions, we document that about 30% of subjects exhibit asymmetric geometric discounting, allocating higher budget shares to their child in later rounds. For those subjects, the average share allocated to their children increases from 43% two days later to nearly 65% thirty days later. Allocating a larger share of consumption to the child in later time periods is not necessarily irrational: parents may, for instance, expect their child to have a higher marginal utility of consumption than their own, or a higher probability of survival further in the future. However, such time preferences strongly predict reversals. In just two days, AGD parents reallocate their children's round-3 consumption share towards their own roughly five times as often as other parents, by almost 20% of the mean consumption share they had previously planned to allocate to children in the future. Strikingly, despite ambitious plans, AGD subjects end up allocating *less* to their children in the present (about 4% less than other parents, a statistically significant difference).

Could asymmetric preference reversals be driven instead by asymmetric quasi-

hyperbolic discounting, whereby subjects display present-bias only (or to a greater extent) towards their own consumption?[9] The answer is no: present-bias – even if it affected parents' and children's future consumption to different extents – cannot generate preference reversals when plans are still in the future. Our experiment can capture (asymmetric) present-bias: we have parents also allocate peanuts for their own consumption over time, defining as present-biased those who revise their round-3 consumption plans downwards when we revisit them two days later. We find that present-bias and parent-bias are indeed different phenomena linked to how parents treat the present and the future (in the words of the opening quote by Samuelson): the latter is not systematically correlated with the former. Moreover, we can reject directly that AGD parents in our experiment are systematically less present-biased about their children's future consumption than about their own.[10]

We also rule out several other potential confounders for preference reversals by parents who, at round 1, plan to allocate time-increasing budget shares to their children. Specifically, we show that reversals are *not* an artifact of (1) indifference between different allocations driven by indivisibilities in how peanuts could be split between parents and children; (2) changes across rounds in the salience of fairness with respect to how parents allocate resources across the child participating in the experiment and their other children; (3) measurement error in the extent to which choices reflect preferences; (4) communication between parents of different types between decision rounds; or (5) correlation of AGD preferences with other preference features, such as the strength of their preference for peanuts.

While our experiment is based on consumption allocation trade-offs within periods (as parents cannot transfer resources across periods), most investments in children actually also involve trading off resources across periods. In fact, education or preventive health care are prototypical examples of investments in children's human capital with upfront costs. Do AGD preferences generate under-investment when such trade-offs are present? Using a simple model, we show that, in the investment case, even parents sophisticated about their own bias are prone to preference reversals, and that naive AGD parents massively under-invest relative to otherwise identical symmetric parents.[11] Calibrating the model with parameter estimates based on allocation decisions in the experiment, we show that AGD parents' under-investments in their children lead to large welfare losses: for

[9] Koelle & Wozny (2018) provide evidence that agents are less likely to exhibit present-bias when making decisions on behalf of others.

[10] If anything, the opposite is true; see Supplementary Appendix S3.1.

[11] All our comparative statics between AGD and symmetric parents hold fixed the discount factor that applies to children's future utility of consumption.

investments with a 1-year horizon until returns pay out, naive parents' long-term utility is equivalent to that of symmetric parents with 24.5% lower income. Sophistication only partially mitigates welfare losses (18% lower income-equivalent long-term utility). We also provide evidence that AGD preferences predict investments in children outside the lab to the same extent as present-bias. We survey parents at the end of round 3 about recent investments in the education and health of the child involved in the experiment. The correlation between AGD preferences and actual investments in children is nearly identical to that between the latter and quasi-hyperbolic discounting, holding fixed the extent to which parents discount their children's future utility of consumption. Among parents near the high-end of the patience distribution, AGD preferences (just as present-biased preferences) predict under-investments in children relative to otherwise identical parents; for instance, for parents with a daily discount factor of 0.996, being AGD decreases investments in 6-12 year-old children by approximately a third of the effect of being downgraded from primary school to having no education.

What policy instruments could potentially mitigate the reallocation effects of asymmetric geometric discounting? We start by randomly assigning a subsample of participants to a framing intervention: at the beginning of round 2, they are reminded of their round-1 allocation. As such, each parent's previous allocation to children's consumption is made salient at the time they have the opportunity to revise it. The intervention, however, has no effects: we find that AGD parents assigned to the intervention reallocate away from their children's planned consumption just as much.

Next, we turn to commitment devices.[12] We offer subjects the possibility of committing to their allocation plans set at round 1, varying the price of commitment experimentally. We find no clear pattern linking AGD preferences to demand for commitment in the lab. Strikingly, outside the lab, we find that AGD parents are actually *less likely to commit* to future plans. In a 6-month follow-up, we enter all parents into a lottery, and offer them the opportunity to commit lottery proceeds to tutoring for their child in case they win. We find that AGD parents display a 23% lower willingness to pay for commitment than other parents.

Incidentally, this follow-up experiment allows us to document that the implications of AGD preferences are confined to one's children, different from time-increasing altruism towards others more generally. We do so by eliciting parents'

[12]Demand for commitment designed to address present-bias is often low, even among sophisticated subjects (Ashraf et al., 2006; Augenblick et al., 2015; Kaur et al., 2015; Laibson, 2015).Exceptions are Ariely & Wertenbroch (2002); Beshears et al. (2020); Casaburi & Willis (2018); Schilbach (2019).

willingness to commit lottery proceeds to tutoring someone else's child. While all parents are systematically less likely to allocate lottery proceeds to tutoring another child, AGD parents are not differentially likely to do so.

Our study contributes to an active literature about behavioral biases linked to investments in children's human capital. Present-bias can explain under-investment in children, as naive parents systematically over-estimate the extent to which they will trade off costly investments and later returns in the future.[13] Preference reversals associated with present-bias have been extensively studied (Augenblick et al., 2015; DellaVigna & Malmendier, 2006), particularly within Development Economics (Ashraf et al., 2006; Gine et al., 2016; Tarozzi & Mahajan, 2011). In contrast, parent-bias has been overlooked, even though we find it to be as predictive of real investments in children's human capital as present-bias in our sample.

This paper also relates to a growing literature that studies systematic deviations from geometric discounting in richer ways than a sharp discontinuity between the present and all future periods.[14] In particular, while several papers investigate whether subjects tend to be more generous towards others in the future than in the present, our study is the first to elicit asymmetries in discount rates for parents' and their children's future consumption, and to document the consequences of such asymmetries for within-household allocation.[15] A few papers have posited that discount rates differ across consumption goods (Ubfal, 2016), which could lead to poverty traps (Mullainathan & Shafir, 2013).[16] We extend those consumption models to also account for investments in children, and estimate the welfare implications of sub-optimal investments.

Last, our findings inform a range of new policy instruments that could address parent-bias and, potentially, increase investments in children – from school meals

[13]The theoretical link between present-bias and investments in children is noted in Glennerster & Kremer (2012), but there is limited empirical evidence documenting that link. One notable exception is Ringdal & Sjursen (2017) which shows that increasing the bargaining power of the most patient parent increases investments in children.

[14]See, for instance, theory and evidence that feelings of anticipation can generate rich patterns for future plans and dynamic inconsistencies (e.g. Caplin & Leahy, 2001; Thakral & To, 2020).

[15]In the absence of trade-offs between one's own and other's consumption, subjects tend to be paternalist, aligning choices for others with their own preferences, (Ambuehl et al., 2019; Krawczyk & Wozny, 2017; Uhl, 2011) or to display more patience for others than for themselves (Shapiro, 2010). When there are trade-offs, subjects tend to display time-inconsistent generosity (Koelle & Wozny, 2018). There is related evidence that generosity might interact not only with time preferences, but also with risk preferences (Exley, 2015) and preferences over fairness (Andreoni et al., 2018).

[16]There is a also a literature on different discount rates between different decision-makers within the household and its consequences for time inconsistencies (e.g. Jackson & Yariv, 2015). More generally, for the consequences of intra-household bargaining, see Chiappori (1988) and Baland & Ziparo (2017).

to illiquid savings accounts earmarked to children. At the same time, our results on low demand for commitment outside the lab suggest that it might be challenging to prevent parents from reallocating away from planned investments in children even if those instruments were in place.

The remainder of the paper is organized as follows. Section 2 presents a simple model of investments in children, highlighting the implications of asymmetric geometric discounting for parents' dynamic allocation patterns. Section 3 describes our empirical strategy for testing the model's predictions. Section 4 then documents the two new stylized facts about how parents plan and effectively split resources with their children over time, followed by rigorous tests of model's predictions using our experimental data in Section 5, including analyses of AGD preferences' correlation with real-life investments in children and a calibration of its welfare consequences. Section 6 tests whether framing interventions can mitigate parent-bias, and elicits AGD parents' demand for commitment. Section 7 concludes the paper.

2 A model of parental investments in children

In this section, we present a simple model of parental allocation decisions between themselves and their children over time, characterizing the predictions for within-household allocation trajectories in the presence of asymmetric geometric discounting (AGD). Subsection 2.1 starts with the consumption case, which abstracts from inter-temporal trade-offs and sophistication about preference reversals. Next, we consider the investment case in subsection 2.2.

2.1 The consumption case

We depart from a simple three-period model of investments in children. Each household consists of one child and one parent. In each period $t \in \{1, 2\}$, the parent decides how to allocate income (y, constant over time, for simplicity) between her own consumption (x_t) and that of her child (z_t) at $t \in \{2, 3\}$. To focus on the dynamics of future plans, we abstract from consumption at $t = 1$; the parent only sets future allocations at this period, deciding on how to split consumption between herself and her child at $t = 2$ and $t = 3$. S/he revisits those two decisions at $t = 2$. At $t = 3$, consumption decisions can no longer be changed: the parent and her child consume according to the choice made at $t = 2$. We assume there is no uncertainty and no technology to smooth consumption over time (we relax

the latter in the next subsection).

The parent derives instantaneous utility from consumption $u(x_t)$, and the child, $v(z_t)$, both increasing, strictly concave and continuously differentiable. The parent weights her child's instantaneous utility by an imperfect altruism parameter, $\alpha \geq 0$. S/he discounts her child's future consumption one period ahead by $\delta \in [0,1]$, and her own consumption one period ahead by $\theta\delta$, with $\theta \in [0,1]$.[17] While we abstract from present-bias in this simple formulation, Supplementary Appendix S3.1 augments the model by allowing for quasi-hyperbolic discounting, with potentially different β's applying to the parent's and the child's future utility of consumption.

Formally, the parent's utility maximization problem at $t = 1$ is as follows:

$$\underset{\{z_t,x_t\}_{t=2,3}}{\text{Max}} \; \theta\delta u(x_2^1) + \alpha\delta v(z_2^1) + (\theta\delta)^2 u(x_3^1) + \alpha\delta^2 v(z_3^1) \tag{1}$$

$$\text{s.t.}$$
$$\begin{cases} x_2^1 + z_2^1 \leq y \\ x_3^1 + z_3^1 \leq y, \end{cases}$$

where superscripts indicate that the decision is made at $t = 1$.

At $t = 2$, the parents' utility maximization problem becomes:

$$\underset{\{z_t,x_t\}_{t=2,3}}{\text{Max}} \; u(x_2^2) + \alpha v(z_2^2) + \theta\delta u(x_3^2) + \alpha\delta v(z_3^2) \tag{2}$$

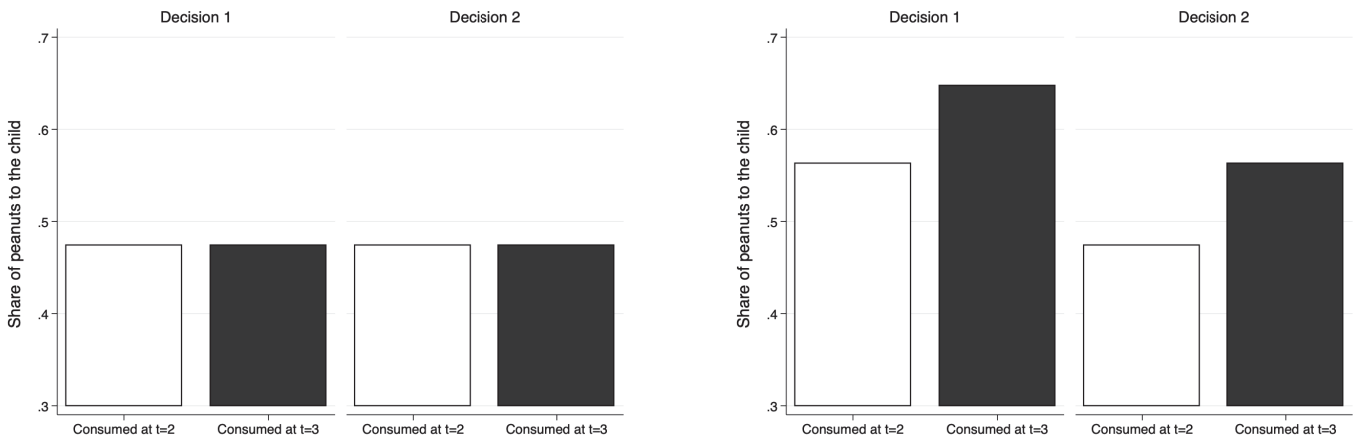
$$\text{s.t.}$$
$$\begin{cases} x_2^2 + z_2^2 \leq y \\ x_3^2 + z_3^2 \leq y, \end{cases}$$

First Order Conditions at $t = 1$ are given by: $\frac{u'(x_2^1)}{v'(z_2^1)} = \frac{\alpha}{\theta}$ and $\frac{u'(x_3^1)}{v'(z_3^1)} = \frac{\alpha}{\theta^2}$. At $t = 2$, they become : $\frac{u'(x_2^2)}{v'(z_2^2)} = \alpha$ and $\frac{u'(x_3^2)}{v'(z_3^2)} = \frac{\alpha}{\theta}$.

If the parent discounts her consumption and that of her child to the same extent, i.e. if $\theta = 1$, then FOCs are identical in both periods, and the share of consumption allocated to the child is constant over time. In other words, the parent will not deviate from plans made at $t = 1$ when revisiting the decision at

[17] While this formulation is related to Ubfal (2016), which estimates heterogeneous discount rates across different goods, and to Banerjee & Mullainathan (2010), in which temptation goods are differentially discounted in the future, in our model, differences in discount rates can arise across consumption of *different subjects*, with consequential implications for parental decisions on behalf of the child. Moreover, our model gives rise to new insights, in particular when it comes to issues of sophistication and commitment to future plans (see Section 2.2) and welfare implications (see Section 5.6).

Figure 1: Budget shares allocated to the child in each period at each decision round



(a) Symmetric parents ($\theta = 1$)

(b) AGD parents ($\theta < 1$)

Notes: Panels (a) and (b) plot the budget share allocated to the child in each period at each decision round. In each panel, the first two bars refer to the decision made at $t = 1$, and the last two bars, to the decision made at $t = 2$. Within each decision round, the white bar represents children's consumption share at $t = 2$ and the black bar, that at $t = 3$. All figures assume that the parent's and the child's instantaneous utility of consumption is logarithmic, that parent's coefficient of imperfect altruism, α, is equal to 0.9 and that $\theta = 0.7$.

$t = 2$.

Conversely, if the parent discounts her own future consumption to a greater extent than that of her child, i.e. if $\theta < 1$, then s/he will plan to allocate a *larger share* of consumption to her child further in the future, an immediate implication of decreasing marginal utility of consumption. Moreover, this will lead to preference reversals at $t = 2$: an AGD parent will deviate from her $t = 1$ plans when given the chance to update her decision at $t = 2$, reallocating consumption away from her child, towards herself. Those predictions are indistinguishable from an alternative formulation in which the coefficient of imperfect altruism is time-increasing, with $\alpha_k^j < \alpha_{k+1}^j$. Empirically, we consider the distinction between parent-bias and time-increasing altruism (not specific to one's children) in subsection 6.3.4.

Figure 1 illustrates the differences in (planned) allocations between symmetric and asymmetric geometric discounters who face otherwise identical utility functions with a numerical example. The left-hand-side panel showcases the decisions of the former: they allocate the same budget share to their children to be consumed at $t = 2$ and $t = 3$, irrespective of when that decision is made. The right-hand-side panel highlights that, in contrast, when AGD parents make allocation decisions in the first time period (the first two bars), they plan to be more generous towards their child in the later consumption period relative to when revising their

decision in the second time period (the last two bars). Allowing parents to be present-biased does not change Figure 1 as long as the same β applies for both the parent's and the child's future utility of consumption; Supplementary Appendix S3.1 discusses how that changes when we allow β's to be different.

Proposition 1 generalizes the model's predictions in the presence of AGD preferences.

Proposition 1: *AGD parents (1) allocate budget shares to their children increasing in the time gap between the decision and consumption, and (2) reallocate away from their children planned consumption (towards their own) at every period.*

Proof: See Appendix A.

2.2 The investment case

In this section, we allow for inter-temporal transfers to study if the insights of the consumption model still hold in the presence of dynamic trade-offs, especially among sophisticated parents. In this model, AGD parents are still subject to preference reversals when it comes to revising their child's planned consumption; we investigate whether anticipating reallocation leads sophisticated AGD parents to compensate by *over-investing* in children in the present, relative to symmetric parents.

To keep the analysis simple, we restrict attention to a two-period model, allowing parents to revise consumption decisions but not investment decisions over time. At $t = 1$, the parent chooses how much out of income y to consume, how much to allocate to her child's consumption at $t = 1$, and how much to invest in her child, I. Investment yields gross return $R = 1 + r$ at $t = 2$. As in the consumption case, the parent makes plans of how to split income $y + RI$ between herself and her child in $t = 2$. In sum, investments in children can be used as a savings vehicle to smooth consumption over time.[18]

The parents' utility maximization problem at $t = 1$ is as follows:

$$\underset{\{z_t,x_t\}_{t=1,2},I}{\text{Max}} \; u(x_1^1) + \alpha v(z_1^1) + \theta\delta u(x_2^1) + \alpha\delta v(z_2^1) \tag{3}$$

s.t.

[18] Although extremely simple, the model can accommodate more complex elements; for instance, if parents can recover only a fraction of investments in children, that can be expressed as a lower interest rate.

$$\begin{cases} x_1^1 + z_1^1 + I \leq y \\ x_2^1 + z_2^1 \leq y + RI \end{cases}$$

An important additional element of the investment model is the parent's belief about her future utility function. We define $\hat{\theta}$ as the parent's belief at $t = 1$ about the value that θ takes at $t = 2$. More precisely, following O'Donoghue & Rabin (1999), the agent thinks that her $t = 2$ utility function is: $\hat{\theta}u(x_2) + \alpha v(z_2)$, with $\hat{\theta} \in [\theta, 1]$. The sophisticated type anticipates correctly that her $t = 2$ utility function entails $\hat{\theta} = 1$. The naive type incorrectly believes that her $t = 2$ utility function entails $\hat{\theta} \in [\theta, 1)$, (fully naive if $\hat{\theta} = \theta$).

Assuming a specific functional form for the parent's and the child's instantaneous utility function (CRRA) to obtain tractable results for comparative statics, Proposition 2 establishes that AGD parents *under-invest* relative to symmetric discounters.

Proposition 2: *AGD parents of all types choose a lower level of investment than a symmetric geometric discounter with otherwise identical preferences. If the coefficient of constant relative risk aversion is larger than 1, then investments in children by AGD parents increase with their degree of sophistication.*

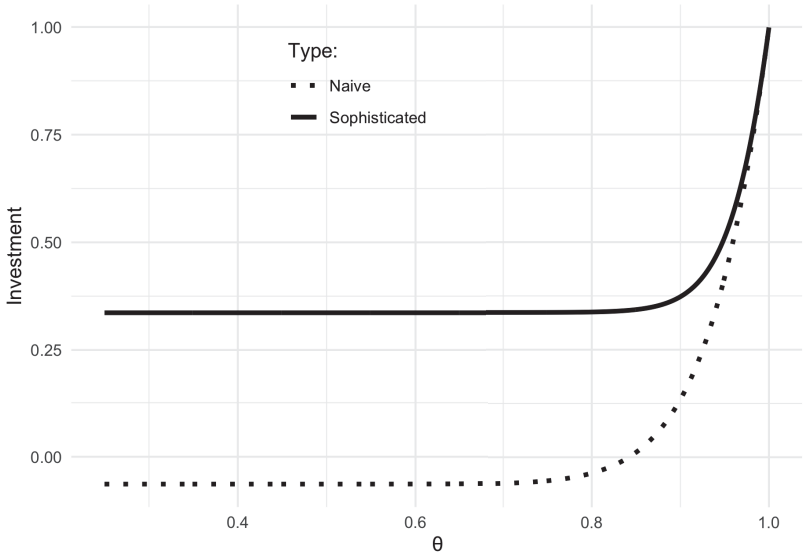
Proof: See Appendix A.

The second part of the proposition holds as long as the coefficient of constant relative risk aversion is larger than 1 (Holden & Quiggin, 2017 finds an average CRRA coefficient of 1.73 for Malawi). In that case, sophisticated AGD parents increase investments in anticipation of future reallocations away from children's planned consumption (although never as much as to completely mitigate the effects of AGD preferences on investments in children).

Figure 2 showcases a numerical example of optimal investment levels for the two extreme types – a naive agent (with $\hat{\theta} = \theta$) and a sophisticated agent (with $\hat{\theta} = 1$). It makes it clear that, the lower $\hat{\theta}$, the larger the gap in investments between naive and sophisticated agents.

The results in this section also make it clear that instruments designed to mitigate present-bias (such as illiquid savings) only imperfectly address parent-bias: while sophisticated AGD parents could use investments in children as a way to mitigate the consequences of within-household reallocation by ensuring a larger resource pool in the future, that does not preclude deviations away from children's planned consumption (as the ratio of marginal utilities in each decision period remain the same as in Section 2.1).

Figure 2: Optimal levels of investments in children



Notes: Optimal investment as a function of θ, derived from the model presented in Section 2.2 with the second consumption period taking place 30 days after the investment decision, and the following parameters: $\alpha = 0.9$, $\gamma = 1.73$, $R = 1.05$, $y = 6.55$ and $\delta = 0.99$, chosen such that the optimal investment of a symmetric parent is equal to 1. As in Section 2.2, the naive agent believes that $\hat{\theta} = \theta$, and the sophisticated agent, that $\hat{\theta} = 1$.

In contrast, Proposition 3 establishes that if investments in children paid out directly as future consumption rather than non-earmarked resources – which, in practice, could be achieved with instruments such as school meal plans –, AGD parents would have (weakly) lesser scope for reallocation (strictly if a corner solution is reached at $t = 2$). This result comes to show that parent-bias requires specific commitment devices to decrease the scope for within-household reallocation in the future.

Proposition 3: *Commitment devices that pre-set future allocations to the child weakly increase her future consumption relative to commitment devices with identical gross returns that merely ensure a larger resource pool in the future.*

Proof: See Appendix A.

3 Empirical strategy

This section describes the design of our experiments, data collection and estimation. Subsection 3.1 discusses how we elicit parents' planned budget allocations between themselves and their child over different horizons, and how we document

preference reversals, followed by a discussion of identification concerns and the design choices we implement to address them in subsection 3.2. Last, subsection 3.3 introduces how we evaluate interventions with the potential to mitigate parent-bias.

All details of the experimental design and a pre-analysis plan were pre-registered at the AEA RCT Registry on November 06, 2018 (AEARCTR-0003535).[19] An additional pre-analysis plan was registered before the follow-up wave.[20] Our baseline experiment was conducted between November 2018 and January 2019, and the follow-up experiment, between June and September 2019. Both experiments followed the same sample of households.

3.1 Documenting AGD and parent-bias

We design a lab-in-the-field experiment whose structure closely matches that of the consumption model (Section 2.1). We visit participants three times: at round 1 ($t = 1$ in the model), round 2 (two days later; $t = 2$) and round 3 (a month later; $t = 3$). At round 1, respondents are asked to make consumption plans for rounds 2 and 3 (Section 3.3 discusses how we offer some participants the opportunity to commit to those plans). At round 2, they make those consumption decisions again (Section 3.3 discusses how we frame that second decision relative to round-1 allocation in different ways). At the end of round 2, one allocation (that set at rounds 1 or 2) is randomly chosen to be implemented; this ensures that both decisions are consequential. Following Augenblick et al. (2015), in the absence of commitment, the round-2 decision is implemented with 90% probability. At round 3, respondents do not make consumption plans; they only consume according to the allocation drawn at round 2.

Our sample consists of 1,627 households across 80 villages in Malawi's Salima district. Households were eligible to be enrolled in the study if both parents lived at home, if they had at least one child aged between 3 and 12 years old, and if no one in the household was allergic to peanuts. If participating households had multiple children in that age range, we randomly selected one to participate in the experiment.

We use peanuts as the experimental currency that participants allocate between themselves and their children over time. Peanuts are a familiar and tempt-

[19] Pre-analysis plan available in full at `https://www.socialscienceregistry.org/trials/3535`. See Supplementary Appendix S1.4 for a detailed discussion about deviations from pre-registration.

[20] `https://www.socialscienceregistry.org/trials/4386`.

ing good, consumed and enjoyed by both parents and children in Malawi: 88% of adults and 97% of children in our sample report enjoying the peanuts we distributed. This experimental currency is payoff-relevant since Malawi is a poor country and our experiments take place during the lean season.

Participants are asked to split the consumption of five packages of peanuts between themselves and their child to be consumed at rounds 2 and 3. Round-1 allocations are set by splitting five tokens between two plates labeled "My child in two days" and "Myself in two days", and five tokens between two plates labeled "My child in a month" and "Myself in a month". Round-2 allocations are set by splitting five tokens between two plates labeled "My child today" and "Myself today", and five tokens between two plates labeled "My child in 28 days" and "Myself in 28 days". At each round, respondents can only choose integer allocations. Each parent makes decisions by herself, in the absence of children.[21]

This simple experimental design allows us to capture AGD preferences and to test the main predictions of the model. First, comparing round-1 allocations set for rounds 2 and 3 allows us to define AGD parents according to the model's prediction of time-increasing budget shares allocated to children. Second, comparing round-2 allocations for rounds 2 and 3 allows us to test the model's prediction that AGD parents allocate time-increasing budget shares to their children (to a greater extent than consistent parents). Third, comparing round-1 and round-2 allocation decisions for round 3 allows us to test the model's prediction that AGD parents display preference reversals for future consumption, reallocating away from their children's planned consumption as the time gap between decisions and actual consumption decreases (to a greater extent than consistent parents).

The comparison to consistent parents remarked in parentheses in the previous paragraph is needed because our definition of AGD preferences entails *measurement error*: there are other (rational) reasons for why parents might set time-increasing budget shares allocated to children at round 1. As such, the experiment allows us to document statistical relationships between the distribution of AGD preferences and those of the behaviors associated with those preferences as predicted by theory. We discuss measurement error in detail in Section 5.4.4.

3.1.1 Definitions

To fix ideas, this subsection rigorously defines how we capture AGD preferences and parent-bias in the data. We define a participant as asymmetric geometric

[21] Supplementary Appendix S4.2 discusses results for a different sample of parents who made decisions in the presence of their children.

discounter based solely on their round-1 decision: if s/he allocates a larger budget share to her child to be consumed at round 3 than at round 2. Formally, let $s^k_{j,i}$ be the share of peanuts allocated to the child's consumption at $t = j$ by parent i when the choice is made at $t = k$. Asymmetric geometric discounting is then defined as: $\mathbb{1}\{\hat{\theta}_i < 1\} \Leftrightarrow s^1_{2,i} < s^1_{3,i}$. We define a participant as symmetric (or consistent) if s/he allocates constant budget shares to her child in rounds 1 and 2.[22]

In turn, we define parent-bias based on whether a participant changes decisions between rounds: if s/he decreases the budget share allocated to their child's round-3 consumption when deciding at round 2 relative to when deciding at round 1. Formally, let $s^k_{j,i}$ be the share of peanuts allocated to the child's consumption at $t = j$ by parent i when the choice is made at $t = k$. Parent-bias is then defined as: $s^2_{3,i} < s^1_{3,i}$.

3.1.2 Sample

Our sample comprises mostly women: only 8.9% of participants are men, as fathers were often away from home working elsewhere during daytime. About 15% of households in our sample are Muslim, the rest identify as Christians. On average, participating households have 2.2 children between 3 and 12 years old. The average age of children taking part in our experiment is 7 years old, with equal participation from boys and girls. Respondents in our sample are poor: on average, respondents state that they would be able to mobilize around three dollars within a week in case of an emergency.

Sample size varies across different specifications of our analyses for the following reasons. First, because we restrict attention to the control group of the framing experiment (see Section 3.3) whenever testing the model's predictions or implications of AGD preferences in the baseline data. As such, whenever we use the follow-up wave (such as when estimating the correlation between AGD preferences and demand for commitment outside the lab), the sample size is larger because there was no framing experiment in that wave. Second, because of small attrition between decision rounds at the baseline experiment, and between baseline and follow-up.[23] Last, when estimating whether AGD preferences are predictive of investments in children, we condition on the values of patience parameter $\hat{\delta}_i$

[22] For completeness, we define a participant as 'child-biased' if s/he allocates a lower budget share to their child to be consumed at round 3 than at round 2.

[23] Appendix C.2 documents attrition and shows that it does not vary systematically with treatment assignment across the different experiments.

calibrated from participants' allocations, which also affects sample size as the calibration algorithm failed to converge for a small number of observations (see Supplementary Appendix S2).

3.2 Identification concerns and design choices

We implement a series of design choices to ensure that our experiment captures time preferences as intended, and to rule out alternative explanations for the stylized facts that we document. Supplementary Appendix S1.1 discusses how we design our experiment to address concerns with fungibility of peanuts with consumption outside the lab, with experimenter demand biases, and with preference reversals being driven by projection bias, shocks to the (expected) marginal utility of consumption, asymmetric quasi-hyperbolic discounting, indifference between allocations, measurement error or communication between rounds. Section 5.4 presents robustness checks that take advantage of those design choices to rule such identification concerns.

3.3 Mitigating parent-bias

Next, we turn to interventions with the potential to mitigate parent-bias. To evaluate the causal effects of those interventions on preference reversals driven by AGD preferences, our experiment cross-randomizes respondents to a framing intervention and to different offers of commitment to their round-1 decisions.

3.3.1 Framing allocation decisions

Thaler (1999) suggests that earmarking funds for specific uses could help individuals resist the temptation to use them for different purposes. In the context of investment decisions by the poor, earmarking lock-boxes to facilitate savings for health care has been evaluated by Dupas & Robinson (2013), with mixed results (effective for emergency spending, but ineffective for preventive health care). We evaluate whether labeling budget shares as *previously allocated to children* prevents AGD parents from parent-biased reallocations in the context of our experiment.

We randomly assign participants to one out of three conditions. In the *control* condition, participants make their round-2 decisions starting from empty plates, just as all participants do at round 1. In the *labeling* condition, participants' decision at round 2 starts from their round-1 allocation: enumerators set up the initial distribution of peanuts across plates so as to match the allocation set by each

participant at the previous round. Last, in the *anchoring* condition, participants start from a random allocation of peanuts across plates.

Across all experimental arms, participants are free to change allocations as they please, regardless of initial conditions. The labeling condition allows us to test whether salience or mental accounting could mitigate parent-bias. If labeling ultimately affects AGD parents' allocations at round 2, the anchoring condition could help understand whether its effect is driven by framing effects in general, or if there is something special about making past promises to children more salient.

3.3.2 Commitment to future plans

Next, we describe how we offer participants the opportunity to commit to their planned allocations. At round 1, after making allocation decisions across both experimental scenarios, all participants are offered the possibility to commit to their round-1 decision by taking up a *probabilistic commitment device*, in the spirit of Augenblick et al. (2015). This device decreases the likelihood that the allocation set by parents at round 2 is implemented. Without commitment, the round-1 allocation is implemented with a 10% probability; with commitment, that probability increases to 90%.[24] We elicit demand for commitment against present-bias and parent-biased separately, allowing participants' take-up decision to vary across experimental scenarios.

We randomize the price of commitment: committing to round-1 decisions require participants to forego packages of peanuts from their own round-3 allocation (0.5, 1 or 1.5, randomly drawn with equal probabilities, and the same in the within-household and inter-temporal scenarios).[25,26] Participants were asked a series of questions to ensure that they understood how the commitment device worked before being asked whether they wanted to commit to their round-1 choice.

Importantly, round-1 allocation decisions are made before the parents are offered the possibility to take up commitment and explained how commitment is billed and parents are not offered the possibility to revise that decision at this stage. For this reason, we do *not* deduct the price of commitment from parents' consumption at round 3 for those who take it up, to avoid artificially making many more subjects look like AGD when they actually did not allocate time-increasing budget shares to children at that point in time.

[24] That design ensure that even those who take up commitment make allocation decisions in both round, and that all allocation decisions are consequential.

[25] Figure S1.2 displays the visual aid the enumerators showed the respondents.

[26] We avoided having subjects pay for commitment early rather than later to avoid low take-up driven by impatience (Casaburi & Willis, 2018).

Appendix C.1.1 shows that learning about costly commitment between decision rounds is not consequential for our analyses. Table C.4 documents that the price of commitment is not systematically associated with parent-biased reversals among those who take it up, and that the association between AGD preferences and the model's predictions is unaffected by allowing commitment price to affect allocations differentially for different planning horizons or at different decision rounds.

Table 1 summarizes the randomization process and the different treatment arms.

Table 1: Distribution of participants per treatment cell in the baseline experiment

Price of commitment \ Framing	Control	Labeling	Anchoring	Total
0.5	263	127	125	*515*
1	254	137	134	*525*
1.5	300	141	146	*587*
Total	*817*	*405*	*405*	*1,627*

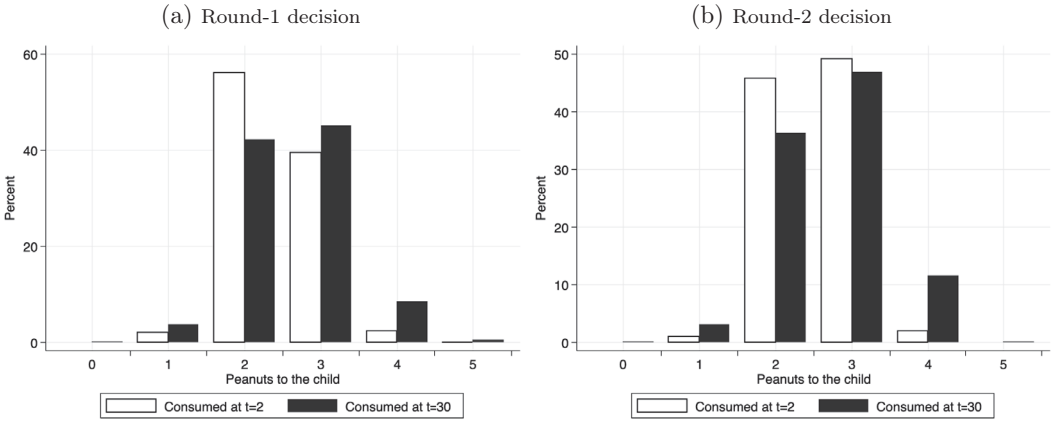
4 Parents' dynamic allocations

This section documents the two new stylized facts about how parents systematically plan to and effectively allocate resources between themselves and their children over time, in subsections 4.1 and 4.2.

4.1 Stylized fact #1: For many parents, budget shares allocated to children *increase with the time gap* between the decision and actual consumption

Figure 3 presents the distribution of parents' round-1 allocations to their child's consumption at rounds 2 and 3. The distribution of children's planned consumption at round 3 is shifted to the right relative to that at round 2. In other words, parents tend to be more generous towards their children further in the future, the larger the time gap between the decision and actual consumption.

Figure 3: Parent's decisions: number of peanuts allocated to the child



Notes: Panel (a) illustrates the distribution of round-1 allocation decisions; Panel (b) focuses on round-2 allocation decisions. On the x-axis is the number of peanuts respondents allocated to be consumed by their children at round 2 (white bars) or round 3 (black bars). The height of each bar represents the percentage of parents who chose each allocation at each round.

At round 2, we observe very similar patterns. Figure 3 shows that, given the chance to revise their round-1 decisions, it is still the case that the distribution of consumption plans allocated to children at round 3 is shifted to the right relative to that set at round 2. 29.2% of parents allocate increasing shares of consumption to their children over time at that point, after which they have no further possibility of revising consumption plans. Among those parents, the average round-3 allocation set to children is 45% larger than the average allocation set for immediate consumption.

Allocating time-increasing budget shares to children is uncorrelated with liquidity constraints or hunger, minimizing concerns that this pattern is merely driven by fungibility with consumption outside of the experiment or projection bias.[27]

Having said that, there are other, rational reasons for why parents might set time-increasing consumption patterns for their children relative to their own: in particular, they might expect different paths for their marginal utility of future consumption, e.g. because they might attribute different survival probabilities to themselves and their children. Alternatively, parents might anticipate that, upon enjoying the consumption of peanuts at round 2, their children will demand a larger amount of peanuts in the next round (akin to habit formation).

In contrast, preference reversals cannot be accommodated by such rational expectations, especially when it comes to reallocation away from consumption

[27]The correlation coefficient between time-increasing budget shares allocated to the child and self-reported liquidity constraints is -0.012 ($p = 0.64$), and that between the former and hunger is -0.001 for that reported by parents ($p = 0.96$) and 0.02 for that reported for children ($p = 0.46$).

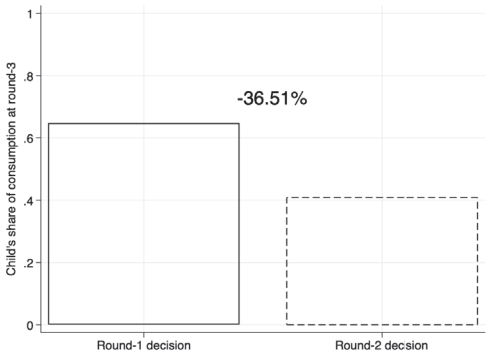
plans still in the future, only two days after those plans were set. The next subsection turns to this stylized fact.

4.2 Stylized fact #2: Many parents revise planned allocations to children downwards *even before consumption time*

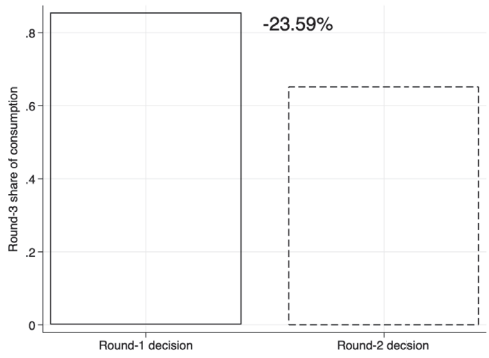
For this analysis, we restrict attention to parents in the control group of the framing intervention, to capture the extent of preference reversals in the absence of interventions with the potential to mitigate them. The probabilistic commitment device does not affect our analysis because round-2 decisions can still be implemented with positive probability (see Section 3.3.2). Figure 4 presents average round-1 and round-2 budget shares allocated to children to be consumed at round 3. 14.3% of parents reallocate away from their children's planned future consumption when given the opportunity to revise their decision, only two days later.

Figure 4: Reallocations

(a) Children's average consumption share at $t = 3$ among parents who reallocate away from children
Parent-bias



(b) Average consumption share at $t = 3$ among parents who reallocate away from the future
Present-bias



Notes: Panel (a) shows the average pattern for children's consumption share at $t = 3$ set by respondents who revise their original planned allocation to their children at that round towards their own consumption (14.3% of our sample). Panel (b) shows the average pattern for one's own consumption at $t = 3$ set by respondents who revise their original planned allocation to their consumption at that round towards the earlier round (43.6% of our sample). In both panels, we restrict attention to the sub-sample is the control group of the framing experiment.

The extent of reallocation is large. The left-hand-side panel of Figure 4 documents a 36.5% reduction in the share of consumption allocated to the child between decision rounds. Reallocating away from children's future consumption is uncorrelated with changes in liquidity constraints or hunger between rounds.[28]

[28]The correlation coefficient between reallocations away from children's future consumption

It is useful to compare those reallocations to those induced by present-bias, presented in the right-hand-side panel of Figure 4, captured through the alternative experimental scenario in which parents set inter-temporal allocations for their own consumption. In that scenario (also restricting attention to parents in the control group of the framing intervention in the main experiment, to ensure comparability), 43.6% of the respondents reallocate away from their future consumption. While present-bias is about three times as prevalent as reallocating away from children's future consumption, the magnitude of present-biased reallocations is *smaller*: for those parents, round-3 consumption decreases by 23.6% on average.

5 Testing the model's predictions

Our simple model yields testable predictions connecting AGD preferences to the two stylized facts from the previous section. First, AGD parents' budget shares allocated to children should increase with the time gap between the decision and actual consumption (to a greater extent than symmetric parents). Second, AGD parents should revise future allocations to children downwards as the planning horizon gets shorter (to a greater extent than symmetric parents). Subsections 5.1 and 5.2 test these hypotheses, followed by a summary of results' heterogeneity in subsection 5.3 and robustness tests in subsection 5.4. Next, we evaluate whether AGD preferences are predictive of real-life investments in children in subsection 5.5. Last, subsection 5.6 computes welfare losses associated with AGD preferences by calibrating the investment model with parameter estimates based on parents' allocation decisions in the experiment.

5.1 Prediction #1: AGD parents' budget shares allocated to children increase with the time gap between the decision and actual consumption

We categorize about 30% of parents in our sample as AGD, based on setting time-increasing budget shares to their children at round 1. The model predicts that AGD parents should allocate time-increasing budget shares to their children also at round 2 (to a greater extent than symmetric parents).

and changes to self-reported liquidity constraints between rounds is -0.01 ($p = 0.84$), and that between the former and changes to hunger between rounds is 0.02 for that reported by parents ($p = 0.32$) and -0.02 for that reported for children ($p = 0.37$).

We test this hypothesis formally with the following regression:

$$s^2_{j,i} = \alpha + \gamma_0 (j-2) + \gamma_1 \left(\mathbb{1}\{\hat{\theta}_i < 1\} \right) + \gamma_2 \left(\mathbb{1}\{\hat{\theta}_i < 1\} \times (j-2) \right) + \lambda X_i + \varepsilon_{ij}, \tag{4}$$

where $s^2_{j,i}$ is the share of peanuts parent i allocates to their child at $t = 2$ to be consumed at $t = j \in \{2, 30\}$; $(j-2)$ is the number of days between the decision and consumption; $\mathbb{1}\{\hat{\theta}_i < 1\}$ equals 1 if parent i sets time-increasing budget shares to her child at round 1, and 0 otherwise; X_i a vector of individual characteristics; and ε_{ij} is an error term. We are interested in testing $\gamma_2 \geq 0$.

Table 2: Testing the model's predictions

Panel A: Prediction #1			Panel B: Prediction #2			Panel C: Present-bias			
	(1)	(2)		(3)	(4)		(5)	(6)	(7)
	$\mathbb{1}\{s^2_{30} > s^2_2\}$	s^2_j		$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	s^k_{30}		$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	s^k_{30}	s^k_{30r}
$j-2$		0.000026	$\mathbb{1}\{k=2\}$		0.0530***			0.0181***	0.0969***
		(0.0002)			(0.0059)			(0.0066)	(0.0070)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.403***	-0.0196**	$\mathbb{1}\{\hat{\theta} < 1\}$	0.270***	0.186***	$\mathbb{1}\{\hat{\beta} < 1\}$	0.0237	-0.0054	0.106***
	(0.0357)	(0.0087)		(0.0323)	(0.0081)		(0.0254)	(0.0101)	(0.0126)
$\mathbb{1}\{\hat{\theta} < 1\}$		0.0031***	$\mathbb{1}\{\hat{\theta} < 1\}$		-0.118***	$\mathbb{1}\{\hat{\beta} < 1\}$		0.001	-0.299***
$\times(j-2)$		(0.0004)	$\times\mathbb{1}\{k=2\}$		(0.0116)	$\times\mathbb{1}\{k=2\}$		(0.0112)	(0.0119)
Control	Yes	Yes		Yes	Yes		Yes	Yes	Yes
AGD parents' mean	0.581	0.536		0.328	0.612		0.328	0.612	0.762
Symmetric parents' mean	0.176	0.513		0.063	0.487		0.063	0.487	0.747
N	795	1590		795	1608		795	1590	4770
Respondents	795	795		795	813		795	795	795
Sample	Control	Control		Control	Control		Control	Control	Control

Notes: Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$; and Panel C studies the impact of present-bias on preference reversals. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In columns (3) and (5), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In columns (4) and (6), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. In column (7), the outcome variable is the share of peanuts allocated to herself at $t = 30$ in the inter-temporal scenario when the decision was made at $t = k$ under interest rate r; participants' allocations for each decision round k and under each interest rate r are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Panel A of Table 2 presents the results, restricting attention to the control group of the framing experiment. Consistent with the model's prediction, AGD parents set time-increasing budget shares to their children at round 2 to a much

greater extent: 58.1% of them do so, compared to only 17.6% of symmetric parents. Column (1) shows that such large and statistically significant difference is nearly unchanged after controlling for individual characteristics. Column (2) tests prediction #1 directly. On the one hand, symmetric parents do not systematically set time-increasing budget shares at round 2 (the coefficient of $j-2$ is not statistically significant and very close to zero); on the other hand, children of AGD parents are allocated a substantially higher share of peanuts 28 days later (8.7 p.p. higher on average, significant at the 1% level and equivalent to 17% of the mean consumption share symmetric parents allocate to children at round 2).

Strikingly, panel A of Table 2 also shows that, despite such ambitious plans, AGD parents actually allocate systematically *less* to their children in the present: at round 2, they set a 1.96 p.p. lower share of peanuts to their children's consumption relative to symmetric parents (statistically significant at the 5% level).

5.2 Prediction #2: AGD parents decrease budget shares allocated to children in the future as it gets closer to consumption time

The model also predicts that, between rounds 1 and 2, AGD parents should decrease budget shares allocated to their children at round 3 (to a greater extent than symmetric parents).

We test this hypothesis formally with the following regression:

$$s^k_{30,i} = \alpha + \gamma_0 \left(\mathbb{1}\{k=2\}\right) + \gamma_1 \left(\mathbb{1}\{\hat{\theta}_i < 1\}\right) + \gamma_2 \left(\mathbb{1}\{\hat{\theta}_i < 1\} \times \mathbb{1}\{k=2\}\right) + \lambda X_{ik} + \varepsilon_{ik}, \tag{5}$$

where $s^k_{30,i}$ is the share of peanuts parent i allocates to their child at $t = k \in \{0, 2\}$ to be consumed at $t = 30$; $\mathbb{1}\{k=2\}$ equals 1 for allocation decisions undertaken at round 2, and 0 otherwise; $\mathbb{1}\{\hat{\theta}_i < 1\}$ equals 1 if parent i sets time-increasing budget shares to her child at round 1, and 0 otherwise; X_{ik} a vector of (time-varying) individual characteristics; and ε_{ik} is an error term. We are interested in testing $\gamma_2 \leq 0$.

Panel B of Table 2 presents the results, restricting attention to the control group of the framing experiment. AGD parents are 5 times as likely as symmetric parents to reallocate away from their children's consumption at round 2 relative to plans set just two days before. Column (3) shows that such large and statistically significant difference is nearly unchanged after controlling for individual

characteristics. Column (4) tests prediction #2 directly. At round 2, AGD parents reallocate 11.8 p.p. more away from their children's future consumption plans set at round 1 relative to symmetric parents. This is a large effect size, 24% of the average budget share allocated by symmetric parents to be consumed by children at round 3 (statistically significant at the 1% level).

Last, Panel C of Table 2 shows that present-bias does *not* predict parent-bias in our sample. If anything, parents with quasi-hyperbolic discounting preferences tend to *increase* their children's future allocation between decision rounds, relative to symmetric respondents (column 5; not statistically significant). In the specification that specifically tests prediction #2 (column 6), the coefficient of present-bias is nearly zero. That is the case even though quasi-hyperbolic discounting is associated with sizable reallocation away from participants' own future consumption at round 2 (column 7), when the latter is traded off against current consumption.

5.3 Heterogeneity

Appendix B investigates which household and individual characteristics predict AGD preferences, and replicates the analyses of the two previous subsections splitting the sample by the child's gender, age and birth order. We find that no characteristic systematically predicts AGD preferences, and that the impacts of those preferences on dynamic allocation patterns, including parent-bias, do not vary systematically with the child's characteristics.

AGD preferences do seem to matter differentially when it comes to real-life investments, but slicing the sample by characteristics makes it challenging to precisely detect differences in coefficients of the interaction between AGD and the discount factor. Among 6-12 year-old children, investments in health and education seem to decrease much faster at the high-end of the patience distribution among AGD parents of girls and first-born children, although differences in coefficients are not statistically significant at the 10% level.

5.4 Robustness checks

As discussed in Section 3.2, we implement design choices in our experiment to minimize concerns with issues such as fungibility, projection bias and experimenter demand bias. Moreover, Section 4 discusses correlational evidence that liquidity constraints and projection bias (or changes in those) are uncorrelated with the prevalence of the behaviors characterized by the two facts across parents in our sample. In this section, we summarize a series of robustness tests to rule

out additional identification concerns. Subsection 5.4.1 presents direct tests of the hypothesis that our experimental currency was not fungible with children's consumption outside the lab. Subsection 5.4.2 documents direct tests of the hypothesis that parent-bias is not driven by shocks between decision rounds. Next, we show that our findings for model's prediction #2 are not driven by indifference between integer allocations (subsection 5.4.3) neither by measurement error or communication across subjects between rounds (subsection 5.4.4). Subsection 5.4.5 showcases that parent-bias is also not driven by changes in the salience of fairness towards siblings of the child participating in the experiment between decision rounds. Last, subsection 5.4.6 documents that the predictive power of AGD preferences for parent-biased reversals is not an artifact of other preference features that could be associated with setting time-increasing budget shares at round 1, controlling flexibly for the consumption shares set to children at that round when testing model's prediction #2.

5.4.1 Parents consider adjusting consumption outside of the experiment, but fail to follow through

Appendix C.3 analyzes the correlation between the time elapsed since children's last meal (reported by parents before consumption decisions are implemented) and budget shares allocated to children within the experiment. Interestingly, parents seem to systematically adjust children's consumption outside of the experiment in anticipation of peanuts set to be consumed by them at round 2: the number of hours since children last ate correlates positively and significantly with their consumption share assigned by parents two days before.

Having said that, parents do not follow through on those plans: there is no systematic correlation between the time since children's last meal and their actual consumption share allocated by parents at round 2 – when allocation decisions can no longer be revised. Moreover, when we ask parents about whether children participating in the experiment are hungry (again, before consumption decisions are implemented), even though nearly half of them responds affirmatively, their answers do not systematically correlate with round-2 allocation plans set to children either in the present or two days before.

5.4.2 Parent-bias is not driven by shocks to the (expected) marginal utility of consumption between rounds

To rule out that shocks between decision rounds induce preference reversals, we estimate whether parents' round-2 consumption in the inter-temporal experimental scenario responds differentially to interest rates across decision rounds. In an interior solution, parents equalize the ratio of marginal utilities of consumption within each round to the ratio of gross interest rates under which parents trade off consumption over time; as such, shocks to parents' (expected) marginal utility of consumption should translate into changes in their responses to interest rates. Supplementary Appendix S3 shows that is not the case: the slope of parents' round-2 consumption with respect to interest rates does not systematically change between rounds 1 and 2.

Because this test cannot directly rule out shocks to *children's* (expected) marginal utility of consumption, Appendix C.1.1 provides additional tests for the hypothesis of no systematic shocks between decision rounds. First, we control for village fixed-effects, as negative shocks (such as droughts or floods) often affect many households in a village. This does not affect our estimates in testing model's predictions for AGD preferences, nor their significance levels. Second, we cluster standard errors at the village level, allowing shocks to be arbitrarily correlated at that level. If anything, this increases the precision of our estimates. Last, we control for the day of the month in which the enumerator's visit took place, in order to avoid any payday effects that could induce systematic differences between decision rounds, for instance if distance to payday increases on average between rounds for those who set time-increasing consumption shares to their children at round 1. For prediction #1, we interact the number of days between the decision and consumption $(j-2)$ with an indicator of when the visit took place, whether in the first or second half of the month. For prediction # 2, we interact this indicator of the visit date with an indicator of allocation decisions undertaken at round 2 $(\mathbb{1}\{k=2\})$.[29] Our estimates and their precision levels in each case are unaffected.

5.4.3 Parent-bias is not driven by indifference between integer allocations

In our baseline experiment, participants have to make integer allocations decisions over an odd number of packages within each period. This design choice aims at

[29] We restrict our sample to observations for whom visit dates were accurately recorded by enumerators, which leads to approximately 80 observations being dropped across different specifications.

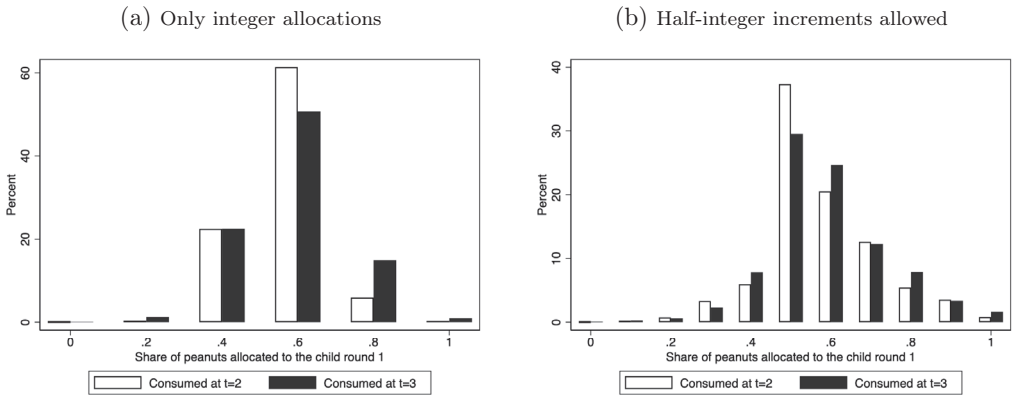
mitigating experimenter demand bias by ruling out the possibility of an egalitarian split focal point. Having said that, this feature could have brought about a different concern: unable to implement even splits within each round, some participants might have tried to set up even splits *on average* – i.e. (2,3) and (3,2) allocations to be consumed at rounds 2 and 3 by themselves and by their child, respectively, or, similarly, (3,2) and (2,3). In the latter case, we would classify those parents as AGD preferences when, in truth, they are merely trying to enforce equal splits. What is more, those parents might reverse future consumption plans for their children between rounds 1 and 2 not because of parent-bias but, rather, because they are indifferent between (2,3) and (3,2) allocations at round 3.

Figure 3 shows that, consistent with the desire to be egalitarian across rounds, the majority of parents choose [(2,3);(3,2)] or [(3,2);(2,3)] allocations – where the first and second parentheses indicate allocations to be consumed at rounds 2 and 3 by themselves (first argument) and by their child (second argument), respectively. Having said that, even among such egalitarian parents, 33.7% choose [(2,3);(3,2)] and 66.3% choose [(3,2);(2,3)]; as such, the majority is more generous towards their child later in the future.

To deal with that concern directly, we re-run round 1 of the experiment *twice* in a follow-up wave, conducted six months after the first one: one version exactly as in the baseline, and another version in which participants split consumption between themselves and their children allowing for half-package increments – including the possibility of splitting peanuts equally with the child within each round. Figure 5 shows that allowing for non-integer allocations (including even splits in the parents' budget set within each round) at the follow-up wave still yields the same stylized fact.

Figure 5's left-hand-side panel displays round-1 allocations at the follow-up wave when parents were only allowed to choose integer allocations, while the right-hand-side panel displays allocations when half-package increments were allowed (including even splits within each period). In the RHS panel, 20.3% of respondents still allocate increasing budget shares to their children (compared to 23.2% in the LHS panel). This rules out that time-increasing budget shares allocated to children are an artifact of parents trying to implement even splits *on average*.

Figure 5: Share of peanuts allocated to the child
Follow-up data collection, $t = 1$ decision



Notes: Distribution of parents' allocation decisions at round 1 in the follow-up wave of data collection. On Panel (a), parents were only allowed to change allocations by integer increments; on Panel (b), parents were allowed to change allocations by half-package increments.

5.4.4 Parent-bias is not driven by measurement error or communication between rounds

Replicating the experiment at the follow-up wave also allows us to deal with two other critical identification concerns. The first is measurement error. As discussed, choices might express preferences with error. If such error is correlated over time, then it would generate spurious correlation between setting time-increasing consumption shares to children across both rounds (model's prediction #1). Moreover, unless measurement error is perfectly correlated over time, it would also generate a spurious correlation between setting a high budget share to be consumed by children at the very last round (naturally more common among AGD parents) and downwards revisions at round 2 (model's prediction #2).

The second additional identification concern is communication between rounds. If parents who set different dynamic patterns for children's consumption at round 1 learn about each other's plans before round 2 and feel pressured to change their plans in the next round (e.g. because of social expectations), that would make downwards revisions by AGD parents relative to symmetric parents prevalent at round 2 (model's prediction #2). Although very different in nature, that issue ends up looking exactly like measurement error, since it would tend to make allocation plans for different sets of parents to move in different directions across rounds.

We address concerns with measurement error and communication between different sets of parents between decision rounds by restricting our sample to parents who were consistently AGD or consistently symmetric across the baseline

and follow-up experiments. The rationale is that parents who set time-increasing consumption shares to their children at round 1 are less likely to be subject to measurement error or to have changed their round-2 allocations due to conformity pressures in the baseline experiment. In doing so, we drop 338 participants from our sample. Table 3 tests model's predictions for the remaining sample.

Results are very close to those in Table 2. In this sub-sample, the share of AGD parents who set time-increasing shares to children at round 2 becomes even larger, and their allocation increases to a greater extent with the planning horizon (prediction #1). Although effect sizes marginally decrease for prediction #2, AGD parents are still 20 p.p. more likely than symmetric parents to reallocate away from their children's future consumption relative to their original plans (27 p.p. in the full sample), and decrease those plans by 11.1 p.p. on average (11.8 p.p. in the full sample). All the estimates remain very precisely estimated, significant at the 1% level.

5.4.5 Parent-bias is not driven by changes in the salience of fairness towards other children between rounds

Given more time to consider allocation decisions, it could be the case that certain parents realize that they set too high consumption shares to the child participating in the experiment, especially when confronted with his/her siblings between decision rounds, inducing preference reversals. Since preferences are not randomly distributed, if family structure is systematically correlated with AGD preferences, that could lead to a spurious correlation between those preferences and present-biased reversals.

Although Appendix B showcases that individual or household characteristics do not systematically predict AGD preferences, we can also test for this hypothesis directly by restricting attention to participants with just one child. For those subjects, there can be no influence of siblings in subsequent allocation decisions.

Table 4 shows that restricting the analysis to that sub-sample does not weaken the correlation between AGD and time-increasing shares nor the correlation between AGD and parent-biased reversals.

5.4.6 AGD preferences are not an artifact of other preference features

More broadly, we consider whether the correlation between AGD preferences and parent-bias could be driven by the fact that the former are also systematically different when it comes to other preference parameters. In particular, it could be

Table 3: Robustness to dropping subjects who switch from AGD to symmetric (or the other way around) between waves

Panel A: Prediction #1			Panel B: Prediction #2		
	(1)	(2)		(3)	(4)
	$\mathbb{1}\{s_{30}^2 > s_2^2\}$	s_j^2		$\mathbb{1}\{s_{30}^0 > s_{30}^2\}$	s_{30}^k
$j-2$		0.0000356	$\mathbb{1}\{k=2\}$		0.0532***
		(0.000221)			(0.00646)
$\mathbb{1}\{\hat{\theta}<1\}$	0.533***	-0.0397**	$\mathbb{1}\{\hat{\theta}<1\}$	0.198***	0.195***
	(0.0551)	(0.0156)		(0.0368)	(0.0132)
$\mathbb{1}\{\hat{\theta}<1\}$		0.00443***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.111***
$\times(j-2)$		(0.000631)	$\times\mathbb{1}\{k=2\}$		(0.0193)
Control	Yes	Yes		Yes	Yes
N	457	914		457	922
Respondents	457	457		457	465
Sample	Control	Control		Control	Control

Notes: This table reports the estimated impact of AGD on allocation decisions by restricting the sample to participants who either (1) set time-increasing consumption shares to children at round 1 in both the baseline and the follow-up waves; or (2) do not set time-increasing consumption shares to children at round 1 in either the baseline or the follow-up waves. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. In addition, the sample is further restricted to respondents whom where classified as AGD or not in the same way in the original survey and in the six-month follow-up survey in the task where parents were allowed to split peanuts equally. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table 4: Robustness to dropping households with multiple children

Panel A: Prediction #1			Panel B: Prediction #2		
	(1)	(2)		(3)	(4)
	$\mathbb{1}\{s_{30}^2 > s_2^2\}$	s_j^2		$\mathbb{1}\{s_{30}^0 > s_{30}^2\}$	s_{30}^k
$j-2$		-0.000685*	$\mathbb{1}\{k=2\}$		0.0397***
		(0.000386)			(0.0104)
$\mathbb{1}\{\hat{\theta}<1\}$	0.452***	-0.0114	$\mathbb{1}\{\hat{\theta}<1\}$	0.300***	0.189***
	(0.0663)	(0.0160)		(0.0624)	(0.0145)
$\mathbb{1}\{\hat{\theta}<1\}$		0.00379***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.0963***
$\times(j-2)$		(0.000761)	$\times\mathbb{1}\{k=2\}$		(0.0201)
Control	Yes	Yes		Yes	Yes
AGD parents' mean	0.594	0.542		0.362	0.614
Symmetric parents' mean	0.136	0.507		0.075	0.479
N	215	430		215	436
Sample	Control and only child's households				

Notes: This table reports the estimated impact of AGD on allocation decisions by restricting the sample to children without siblings. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. In addition, the sample is further restricted to households where there is only one child. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

the case that AGD parents actually care less about peanuts than other parents, leading them to take their decisions within the experiment less seriously. If that were the case, AGD preferences would not only exhibit reversals to a greater extent than other parents, but higher variance more generally.

Table 5 captures participants' preferences by controlling flexibly for their round-1 allocation decisions when testing for the relationship between AGD preferences and parent-bias, in columns (1) and (2), and by testing directly for whether AGD preferences lead to higher variance of allocations, in columns (3) and (4). Columns (1) and (2) feature the change between rounds in allocations set to be consumed by children at round 3 as the outcome, controlling for the round-1 consumption share allocated to the child two days ahead in column (1), and for its quadratic and cubic terms in column (2). Columns (3) and (4) feature the square of the change between rounds in allocations set to be consumed by children at round 3 as the outcome, as a measure of variance.

In Table 5, controlling flexibly for round-1 decisions does not change results: if anything, that actually increases the estimated effect of AGD preferences on parent-biased reallocations relative to symmetric parents. We also do not find evidence that AGD parents exhibit systematically higher variance of allocations than symmetric parents.

5.5 Investments in children outside the lab

Even though Section 2.2 establishes that the patterns for how AGD parents allocate consumption over-time between themselves and their children in the absence of inter-temporal trade-offs should also translate to the investment case, the artificial setup of our experiment has limits in what it allows us to say about the connection between AGD preferences and investments in children outside the lab. For this reason, we survey parents about real-life investments in children's health and education at the end of round 3 of the experiment, to study whether AGD parents actually invest less in their children than symmetric geometric discounters.

We analyze the correlation between AGD preferences and self-reported investments in children's education and health. We pre-registered that we would analyze separately investments in children between 3 and 5 years old and those between 6 and 12 years old. Since we elicit multiple outcomes to capture parental investments in children, we control for family-wise error rate in the context of multiple hypotheses testing by building a summary measure of investments in children Kling et al. (2007). Each summary index measure is the equally weighted average

Table 5: Robustness to correlation between AGD preferences and other preference features

	(1) Δs_{30}	(2) Δs_{30}	(3) $\Delta^2 s_{30}$	(4) $\Delta^2 s_{30}$
$\mathbb{1}\{\hat{\theta} < 1\}$	-0.137***	-0.139***	0.00519	0.00613
	(0.0118)	(0.0118)	(0.00434)	(0.00455)
s_2^0	-0.259***	-0.469		0.0121
	(0.0504)	(1.219)		(0.0227)
$(s_2^0)^2$		-0.638		
		(2.325)		
$(s_2^0)^3$		1.066		
		(1.470)		
Control variables	Yes	Yes	Yes	Yes
Mean	0.0186	0.0186	0.0235	0.0235
N	795	795	795	795

Notes: Columns (1) and (2) estimate whether AGD preferences still predict parent-biased reversals when controlling flexibly for other preference features, proxied by the round-1 consumption share set to children two days ahead. Columns (3) and (4) test directly whether AGD parents have a weaker preference for peanuts by estimating whether AGD preferences predict the variance of round-3 allocations (since more prevalent indifferences would induce higher variance). In columns (1) and (2), the outcome variable is the change in the child's planned consumption share at $t = 3$ between decision rounds; in columns (3) and (4), it is that variable squared (a measure of round-3 allocations' variance). Columns (4) also controls for the round-1 consumption share set to children two days ahead. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

of its standardized components.[30] We normalize each component by the mean and standard deviations of symmetric parents.

In the analyses, we control for individual characteristics, including the extent to which parents discount their own future consumption elicited in the inter-temporal experimental scenario.[31] We allow the effects of AGD preferences to vary with parents' discount rate, $\hat{\delta}_i$. The reason is that, while there is only one way to be present-biased ($\beta < 1$), there are multiple ways to be parent-biased: without holding δ constant, there is no constraint on the extent to which an AGD parent is more or less patient than a symmetric parent with respect to their children's future consumption or their own.

We formally test whether AGD preferences are predictive of real-life investments in children with the following regression:

$$Y_i = \gamma_0 + \gamma_1 \left(\mathbb{1}\{\hat{\theta}_i < 1\} \right) + \gamma_2 \hat{\delta}_i + \gamma_3 \left(\mathbb{1}\{\hat{\theta}_i < 1\} \times \hat{\delta}_i \right) + \lambda X_i + \varepsilon_i, \tag{6}$$

where Y_i is a summary measure of investments in parent i' child; $\mathbb{1}\{\hat{\theta}_i < 1\}$ equals 1 if parent i sets time-increasing budget shares to her child at round 1, and 0 otherwise; $\hat{\delta}_i$ is parent i's discount rate inferred from the experiment; X_i is a vector of individual characteristics; and ε_i is the error term. We are interested in testing $\gamma_1 \leq 0$ and $\gamma_3 \leq 0$.

Table 6 shows the results, restricting attention to 3-5 year-olds in the two columns reported under (1) and 6-12 year-olds in the two columns reported under (2).

Both columns showcase similar patterns: the correlation between AGD preferences with real-life investments in children is *nearly the same* as that between the latter and quasi-hyperbolic discounting. Among impatient parents, AGD and quasi-hyperbolic preferences are actually associated with higher investments in children. As parents' patience increases, however, the correlation between each type of time preferences and investments decreases, and eventually becomes *negative* at the high-end of the patience distribution. As a benchmark, for parents with $\hat{\delta} = 0.996$, the effect of AGD preferences on investments in 6-12 year-old children is sizable, roughly 1/3 of the effect of being downgraded from primary school to having no education.

In sum, parent-bias seems to matter for real-life investments in children's hu-

[30] Supplementary Appendix S1.3 lists all the components of each index.

[31] Supplementary Appendix S2 presents all details on how we compute preference parameters for each participant based on their allocation decisions in the experiment.

Table 6: AGD preferences and real-life investments in children

	(1)		(2)	
	Index of investments, 3-5 years old		Index of investments, 6-12 years old	
$\mathbb{1}(\hat{\theta} < 1)$	0.0261	5.7702	0.0132	7.5130***
	(0.0309)	(4.3139)	(0.0213)	(2.8821)
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$		-5.7884		-7.5638***
		(4.3454)		(2.9051)
$\hat{\delta}$		2.8095		3.3674**
		(2.3921)		(1.6742)
$\mathbb{1}(\hat{\beta} < 1)$	0.0023	2.2392	0.0020	1.7086
	(0.0285)	(3.3076)	(0.0201)	(2.4799)
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$		-2.2611		-1.7256
		(3.3451)		(2.5061)
Control variables	No	No	No	No
N	851	851	1,485	1,485
$\mathbb{1}(\hat{\theta} < 1)$=$\mathbb{1}(\hat{\beta} < 1)$ (p-value)	0.5749	0.5444	0.7004	0.1457
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$=$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$ (p-value)		0.5486		0.1473
Correlation with education		0.0443		0.0650

Notes: Across all columns, the outcome variable is a summary index variable of investments in children, whose components are described in Supplementary Appendix S1.3; each component is normalized with respect to its mean and standard deviation among symmetric parents. Columns (1) and (2) restrict attention to investments in children 5 years old and younger; columns (3) and (4), to investments in children 6 years old and older. $\hat{\delta}$ is the discount factor each parent attaches to their child's future utility of consumption, calibrated from their allocation decisions in the experiment (see Supplementary Appendix S2). As a benchmark, we also report the correlation coefficient of each index with participants' education level, restricting the sample to caregivers who have no education or attended primary school only (87% of the sample). We do not include other controls because present-bias correlates with some household and individual characteristics. Standard errors in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

man capital just as much as present-bias, especially among parents who value the future to a greater extent.

5.6 Welfare consequences of AGD preferences

Last, we estimate welfare losses due to AGD preferences, calibrating our simple investment model (Section 2.2) with the parameter estimates based on parents' allocation decisions in the main experiment.

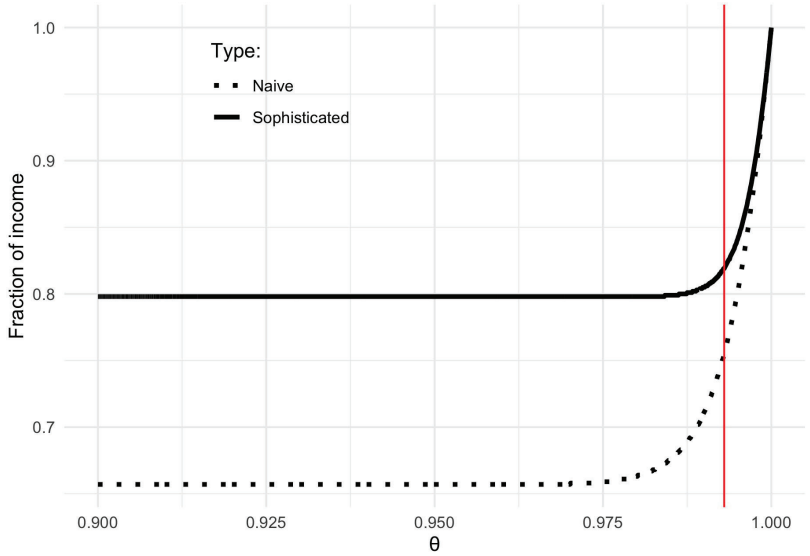
Welfare comparisons for subjects with time-inconsistent preferences are challenging, since welfare analyses traditionally assumes stable preferences (Bernheim & Taubinsky, 2019). Analogously to O'Donoghue & Rabin (1999), we use the symmetric geometric discounter agent as the normative standard ($\theta = 1$). We further assume that this agent's decision at $t = 1$ perfectly overlaps with the child's preferences at that point. This allows us to derive parent's long-run utility, and estimate welfare losses as the monetary compensation that a symmetric agent would require to achieve the same long-run utility as an AGD parent (as a % of their income).

As in Section 2.2, we assume that the parent's and the child's instantaneous utility function is CRRA, with coefficient of relative risk aversion γ. We calibrate welfare calculations with the median values of $\hat{\alpha}$, and $\hat{\delta}$ among AGD parents in

our sample (inferred from allocation decisions with the help of this functional form assumption), and setting $\gamma = 1.73$ (the average value of the CRRA coefficient in Holden & Quiggin, 2017 within a sample of Malawian subjects). All details of the calibration procedure are presented in Supplementary Appendix S2.

Because of model's prediction #1, our simulations yield welfare losses proportional to the time gap between decisions and investment payout (what we call the *planning horizon*). When the planning horizon is short (e.g. a month), welfare losses are rather small: at the median value of $\hat{\theta}$ in our sample, they are less than 1% of the income-equivalent long-run utility of the symmetric agent even for the naive type. When the planning horizon is long, however, welfare losses are substantial. Figure 6 plots the income that a symmetric agent would require to obtain the same long-term utility as naive/sophisticated AGD parent with income $y = 1$ when the planning horizon is a year.

Figure 6: Welfare calibration



Notes: Income that a symmetric agent would require to obtain the same long-run utility as naive and sophisticated AGD agents with income $y = 1$ and asymmetric geometric discounting parameter θ, computed from the model in Section 2.2 with the second consumption period taking place 365 days after the investment decision, and the following parameters, given by the median values in our sample calibrated from parents' allocation decisions in the experiment (see Supplementary Appendix S2): $\alpha = 0.9058$, $\gamma = 1.73$, and $\delta = 0.9939$. As in Section 2.2, the naive agent believes that $\hat{\theta} = \theta$, and the sophisticated agent, that $\hat{\theta} = 1$. The red vertical line marks the median value of θ among AGD parents in our sample. We use $1 + r = 1.000589$ to match Malawi's 24% real yearly interest rate in 2018.[a]

[a]Source: World Bank; see https://data.worldbank.org/indicator/FR.INR.RINR?locations=MW.

Figure 6 shows that naive agents with $\theta = 0.993$ (the median value of $\hat{\theta}$ among AGD parents in our sample) reach a long-run utility equivalent to that of a sym-

metric agent with 24.5% lower income. Sophistication only partially mitigates welfare losses: sophisticated AGD agents' long-run utility is equivalent to that of a symmetric agent with 18% lower income in that case.

6 Testing interventions to mitigate parent-bias

In this section, we analyze interventions designed to mitigate parent-bias, with the help of additional experiments. We start by documenting in subsection 6.1 that participants' characteristics are balanced across Figure 1's treatment cells, and that attrition between rounds is not systematically correlated with treatment assignment. Next, subsection 6.2 investigates whether a framing intervention – reminding parents of their initial decision – decreases reallocation away from children's consumption among AGD parents. Last, subsection 6.3 investigates AGD parents' demand for commitment, in the lab and when it comes to real-life decisions.

Supplementary Appendix S4 presents all details of an alternative intervention which we pre-registered as a potential alternative to mitigate present-bias: child participation. We document the effects of having the child be present at parent's round-2 decision, and elicit parents' willingness to pay for it. We relegate those results to supplementary materials because, as discussed in this Supplementary Appendix, including children as part of the decision problem yields ambiguous theoretical predictions for dynamic allocation decisions.

6.1 Balance and selective attrition tests

Table 7 shows that our sample is balanced across the different treatment arms of the framing experiment and across the different prices of commitment in the baseline experiment. Appendix C.2 discusses selective attrition. Roughly 97% of participants were surveyed in all three rounds of the baseline experiment, and we were able to track 91% of participants in the 6-month follow-up. Attrition across rounds within the baseline experiment, or across the baseline and follow-up waves, is uncorrelated with treatment assignment or with AGD preferences.

6.2 Framing allocation decisions

Does reminding AGD parents about their previous decision mitigate parent-bias? To study that question, we randomly assign the starting point of participants'

Table 7: Summary statistics and balance across treatment arms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Framing experiment				Panel B: Commitment prices			
	Control	Labeling	Anchoring	p-value	0.5	1	1.5	p-value
Mother	0.924	0.901	0.894	0.16	0.903	0.916	0.913	0.73
	(0.265)	(0.299)	(0.308)		(0.296)	(0.277)	(0.282)	
Islam	0.191	0.168	0.136	0.05	0.165	0.185	0.165	0.62
	(0.393)	(0.374)	(0.343)		(0.372)	(0.388)	(0.372)	
Number of children	2.193	2.163	2.281	0.21	2.268	2.139	2.216	0.12
	(0.992)	(0.979)	(1.076)		(1.043)	(0.983)	(1.005)	
Credit constraint, round 1	3.626	3.367	3.037	0.18	3.703	3.358	3.212	0.29
	(5.711)	(5.318)	(4.211)		(5.897)	(5.260)	(4.685)	
Female child	0.531	0.521	0.514	0.84	0.567	0.503	0.506	0.06
	(0.499)	(0.500)	(0.500)		(0.496)	(0.500)	(0.500)	
Child age	7.034	7.202	7.030	0.59	7.052	7.173	7.007	0.62
	(2.891)	(2.966)	(2.803)		(2.879)	(2.954)	(2.836)	
Spending on preventative healthcare, USD	0.195	0.168	0.289	0.38	0.249	0.180	0.209	0.70
	(1.088)	(0.977)	(1.899)		(1.111)	(1.045)	(1.657)	
Index of investments in health	-0.022	0.003	0.013	0.28	0.006	-0.014	-0.012	0.63
	(0.378)	(0.404)	(0.370)		(0.408)	(0.362)	(0.378)	
Index of investments in education	-0.041	0.035	-0.033	0.27	0.002	-0.002	-0.056	0.38
	(0.661)	(0.990)	(0.562)		(0.934)	(0.654)	(0.593)	
s_2^0	0.486	0.482	0.484	0.90	0.486	0.486	0.482	0.76
	(0.112)	(0.121)	(0.118)		(0.116)	(0.120)	(0.112)	
Observations	817	405	405		515	525	587	
F-test (p-value)				0.118				0.475

Notes: Columns (1)-(3) and (5)-(7) report the mean in each treatment group of each of the variables listed. P-values in columns (4) and (8) are from an ANOVA test of equality of means for each variable across treatment groups. In the last row, we report p-values of F-tests from multinomial logistic regression, with a categorical variable containing treatment assignment as dependent variable. Standard deviations in parentheses.

allocation decisions at round 2, as discussed in Section 3.3.1. We are interested in the effects of highlighting parent's allocation plans set at the previous round, which we call *labeling*, whereby participant's decision at round 2 starts with enumerators sorting peanut packages across plates to match their round-1 allocations. Because plates were explicitly labeled as 'My own consumption' and 'My child's consumption' at each decision round, we believe such framing intervention should prime parents explicitly about past promises made to children (or to themselves, on their behalf). In contrast, participants in the *control* condition start with empty plates. The framing experiment also entailed a third condition, which we call *anchoring*, whereby participant's decision at round 2 starts with enumerators randomly distributing peanut packages across plates, to help us understand the mechanism behind labeling effects (if any).

We estimate the effects of the framing intervention through the following regression:

$$\Delta s_{30,i} = \alpha + \gamma_0 \mathbf{T_i} + \gamma_1 \mathbb{1}\{\hat{\theta}_i < 1\} + \gamma_2 \mathbb{1}\{\hat{\theta}_i < 1\} \times \mathbf{T_i} + \lambda X_{ki} + \varepsilon_i \tag{7}$$

where $\Delta s_{30,i}$ is the difference in the share of peanuts parent i allocates to their child to be consumed at $t = 30$ when deciding at $t = 2$ relative to $t = 0$; $\mathbf{T_i}$

equals 1 if parent i is assigned to condition **T**, and 0 otherwise; $\mathbb{1}\{\hat{\theta}_i < 1\}$ equals 1 if parent i sets time-increasing budget shares to her child at round 1, and 0 otherwise; X_i is a vector of individual characteristics; and ε_i is the error term. We are interested in testing $\gamma_2 \leq 0$.

Table 8 presents the results.

Table 8: Effects of framing children's consumption on parent-bias

	(1) Δs_{30}		(2) Δs_{30}
Labeling	0.0111	*Anchoring*	0.00202
	(0.0107)		(0.0134)
$\mathbb{1}\{\hat{\theta} < 1\}$	-0.117***		-0.130***
	(0.0115)		(0.0150)
Labeling	-0.0123	*Anchoring*	-0.0118
$\times\mathbb{1}\{\hat{\theta} < 1\}$	(0.0188)	$\times\mathbb{1}\{\hat{\theta} < 1\}$	(0.0256)
Control variables	Yes		Yes
Mean Δs_3^1 in Control group for $\mathbb{1}\{\hat{\theta} < 1\}$	-0.062		-0.062
N	1194		796
Sample	*Labeling* × *Control*		*Labeling* × *Anchoring*

Notes: Across all columns, the outcome variable is the change in the child's planned consumption share at $t = 3$ between decision rounds. Column (1) estimates the effect of labeling relative to the control group, and column (2), the effect of anchoring relative to that of labeling. Labeling equals 1 if, at the beginning of round 2, the starting point for participant's allocation decision is their round-1 allocation, and 0 otherwise. Anchoring equals 1 if, at the beginning of round 2, the starting point for participant's allocation decision a random allocation, and 0 otherwise. In the control group, participants start from empty plates. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

While labeling has a positive although small and insignificant effect on children's consumption share, it does not mitigate parent-bias: if anything, it even *magnifies* the extent of reallocation away from children's planned consumption among AGD parents (column 1, also statistically insignificant). Column (2) shows that labeling effects are undistinguishable from mere anchoring effects in our experiment, which also fail to prevent parent-biased reallocations among AGD parents.

6.3 Commitment to future plans

This section studies AGD parents' demand for commitment. We start by documenting the extent to which participants are sophisticated about parent-bias in subsection 6.3.1. Next, subsections 6.3.2 and 6.3.3 present results on demand for commitment, testing whether AGD preferences predict differences in take-up of

different types of commitment devices, in the lab and when it comes to real-life decisions, respectively.

6.3.1 Sophistication

As underscored by Section 2.2, sophistication matters for parental investments in children and, naturally, should influence their demand for commitment. To what extent are parent-biased individuals sophisticated?

We measure sophistication through participants' predictions of future behavior. Following Augenblick & Rabin (2019) and Toussaert (2018), we elicit parents' beliefs about the future behavior of other respondents, incentivizing their responses.[32]

Supplementary Appendix S3.3 summarizes our results. We find that 1/3 of parent-biased participants are sophisticated, quite similarly to the prevalence of sophistication among present-biased participants (36.5%; a small and statistically insignificant difference). Incidentally, Supplementary Appendix S3.3 documents that sophistication is not a unidimensional feature across biases: less than 15% of participants are sophisticated about both present-bias and parent-bias, while roughly 34.5% are sophisticated about one but not the other.

6.3.2 Demand for commitment in the lab

Demand for commitment in the lab is extremely high, surpassing 90% when it comes to the experimental scenario involving within-household allocations, presumably due to delayed payment, such as in Casaburi & Willis (2018). While AGD parents are slightly more likely to take-up commitment than other parents, that difference is driven by substantially higher demand at its lowest price – as commitment price increases, however, demand by AGD parents falls significantly more steeply than among other parents. All in all, there is no clear pattern linking AGD preferences and demand for commitment in the lab. Appendix C.4 compiles those findings.

[32] Concretely, enumerators state: "We are asking many other households to make those same decisions. Do you think that, two days from now, most other people will choose to give less, more or the same amount of peanuts to the child than they did today?" If respondents have accurately guessed the behavior of the majority of other respondents, they get two extra packages of peanuts at the end of round 3, outside of the experiment. Respondents learn whether their guesses were correct at the end of round 3.

6.3.3 Demand for commitment outside the lab

The extent to which AGD parents demand commitment in the lab could be a misleading indication of their demand for commitment outside of it, especially in the presence of experimenter demand bias. For this reason, in the follow-up wave we elicit parents' willingness to commit resources to a real investment in their children's education.

To do that, we enter study participants into a lottery for a chance to win 2,000 kwachas (about 2 dollars at the time of the experiment), which they would receive approximately two months after the survey in case they won. Before learning the outcome of the lottery, parents could choose to either receive the prize in cash or to commit those proceeds to one week of tutoring for their child (1 hour a day, for a week), delivered by a local NGO that offers those services regularly outside of the experiment. To elicit parents' willingness to pay for commitment, the flexible option comes with extra cash (a bonus), and we ask participants to choose between flexibility or commitment for different bonus amounts. At the end of the survey, one bonus is randomly picked and the participants' decision for that amount is implemented – a Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) that ensures that respondents' answers are incentive-compatible.[33] Since attaching a bonus to the flexible option may induce experimenter demand bias, we also ask participants to chose between the flexible or commitment option with extra cash attached to commitment, following Carrera et al. (2019).[34] For parents with no school-age children, we ask this question hypothetically. We analyze parent's interest in taking up commitment and their willingness to pay for it.

The take-up of commitment to tutoring is high: 84% of parents express interest in receiving tutoring for their child rather than 2,000 kwachas in cash cards in case they are drawn as lottery winners. What is more, 32% of those are actually willing to pay for commitment. Table 9 documents whether AGD parents are differentially likely to do so.

[33] To ensure that the participants understand this experiment, they first do a practice run to measure their willingness to pay for a bar of soap.

[34] This has the additional advantage of allowing us to assess how well understood this part of the experiment was by identifying respondents who exhibit a positive WTP for both the flexible option and the commitment. This is the case for only 18% of respondents who had chosen to commit lottery proceedings to tutoring; we exclude those participants from our regression analyses.

Table 9: AGD preferences and demand for commitment outside the lab

	(1)	(2)	(3)	(4)
	Own child		Other child	
	Intends to commit	WTP	Intends to commit	WTP
$\mathbb{1}\{\hat{\theta} < 1\}$	-0.0540***	-56.5797***	-0.0017	-18.26
	(0.0181)	(20.6995)	(0.0246)	(23.8691)
Control variables	Yes	Yes	Yes	Yes
AGD parents' mean	0.803	204.34	0.611	242.33
Symmetric parents' mean	0.855	249.31	0.623	266.84
N	2169	834	2171	600

Notes: In columns (1) and (3), the outcome variable equals to 1 if the parent decided to commit lottery proceedings to tutoring for their child, and 0 otherwise. In columns (2) and (4), the outcome variable is parents' willingness to pay for commitment. Columns (2) and (4) restrict attention to the sub-sample of participants who express interest in commitment, who do not express willingness to pay for both flexibility and commitment, and for whom there is at least one non-negative price at which they choose to commit. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, the order in which the inter-temporal and main scenarios were presented to the respondent, and the level of education of the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table 9 shows that, strikingly, AGD parents are actually *less* likely to commit lottery proceeds to tutoring than other parents. Column (1) shows that those are 5.4 p.p. less likely to choose tutoring over cash (significant at the 1% level, and robust to excluding parents with children younger than 6 years old, for whom that question was asked hypothetically). Next, column (2) documents that AGD preferences negatively correlate with willingness to pay, a large and precisely estimated coefficient (22% of the mean willingness to pay of symmetric parents, significant at the 1% level).

We also offer parents the opportunity to commit lottery proceeds to a savings account under the child's name. We actually follow through on this decision for lottery winners who opt into the account, having enumerators accompany them to a local bank. To disentangle demand for commitment against parent-bias from demand for a savings account more generally, we also elicit parents' willingness to pay for a bank account earmarked to their child relative to an account under their own name. Appendix C.5 compiles our findings. In that case, despite expressing significantly higher interest in commitment than symmetric parents, AGD parents are also less willing to pay for it relative to receiving cash (a small and statistically insignificant effect). To rule out that demand for commitment against parent-bias conflates demand for commitment against present-bias, we ask parents who express interest in committing lottery proceeds to a bank account whether they would rather have that account earmarked to their child or themselves. We find that a very low share of parents actually prefer an account earmarked to their child, and that AGD parents are no more likely to do so. What is more, AGD

parents' willingness to pay for an account under their child's name relative to an account under their own name is 25% lower than that of symmetric parents (this effect is, however, only imprecisely estimated, as the share of participants who express interest in committing lottery proceeds to the child in that case is rather small).

All in all, results indicate that AGD parents do not demand commitment to investments in children to a greater extent, either in the lab or outside of it – where, if anything, they are less willing to pay for it.

6.3.4 AGD preferences vs. Time-increasing altruism

Incidentally, this elicitation procedure allows us to document that the implications of AGD preferences are confined to decisions about one's children – rather than about anyone else as a result of time-increasing altruism towards others in general. To do so, we re-elicit parents' demand for commitment when it comes to allocating lottery proceeds to tutoring *someone else's child*. Since we actually implement decisions for lottery winners, we take a few measures to prevent ceiling effects potentially driven by reciprocity motives (Fark & Fischbacher, 2006): participants are informed that the other child would be randomly chosen among other respondents' children, that they would not be informed of the child's identity, and that the child's family would not be informed of the sender's identity.

Columns (3) and (4) of Table 9 present the results. While a lower share of respondents are willing to commit lottery proceeds to tutoring someone else's child (62% of the symmetric respondents, relative to 85.5% in the previous subsection), column (3) documents that AGD parents are no less likely to do so. Along those lines, column (4) shows that AGD parents are also no less willing to pay for commitment in that case.

We can rule out that the coefficients in columns (1) and (3) are statistically identical (at the 10% level). The contrast between the findings in this and the previous subsection suggests that AGD preferences are unlikely to merely express time-increasing altruism towards others in general (or towards any other children). Rather, they likely apply specifically to how parents plan and effectively share resources with their children over time.

7 Concluding remarks

This paper was born from simple observation in the field. As we conducted focus groups with Malawian parents for a different project, in 2017, to gauge whether they understood questions about allocation decisions over different time horizons, we noticed that parents indicated they would like to allocate more and more resources to their children in the future. That observation led us to try to investigate whether that phenomenon was actually prevalent, whether it would lead to systematic preference reversals, and whether it would matter for investments in children. It was only later that we rationalized that behavior with a new type of time preferences. Even though it feels natural to structure the paper the other way around, we still decided to frame it around the stylized facts at the origin of this research, which feed the intellectual curiosity of economists (and anthropologists alike, as Samuelson put it).

This paper is the first to document asymmetric geometric discounting and its consequences for parental investments in children. Those preferences are prevalent and lead to sizable reallocations away from past plans to set resources to children. They also correlate with real-life investments in children just as much as quasi-hyperbolic discounting, and lead to large welfare consequences for decisions with long planning horizons – such as whether or not to pay for school fees in the following year, as in *The Boy Who Harnessed the Wind*. While AGD preferences might be prevalent everywhere, their implications are expected to be particularly dramatic in developing countries, where many investments in children's human capital (from immunization to cooking with clean water) are not institutionalized. In those settings, parents have to *often* and *actively* decide to follow through on past plans to invest in their children, making their time preferences much more consequential.

Results also suggest that parent-bias is likely hard to mitigate. First, in the absence of commitment, we found that reminding parents of their past decisions did not decrease parent-biased reallocations. Second, parent-bias not only requires different commitment devices than present-bias (since lock-boxes do not prevent within-household reallocation), but also, even when offered opportunities to commit to investments in children, AGD parents were no more likely – and, in some cases, even less likely – to take them up.

Our findings open the door to several additional research questions. For instance, do parents discount the future consumption of different children in the household to different extents? If so, does that lead them to systematically reallo-

cate resources over time despite ambitious plans to equalize inputs across siblings (Berry et al., 2020)? Do different decision-makers in the household vary in the extent to which they display AGD preferences? If so, how do they interact in determining household investments in children? What are effective and *attractive* commitment devices to mitigate parent-bias? We hope that those conjectures might lead to new observations in the field that, ultimately, "*lend themselves to testable conjecturing*".

References

Ambuehl, S. B., Bernheim, D., & Ockenfelss, A. (2019). Projective paternalism. *mimeo*.

Andreoni, J., Aydin, D., Barton, B., Bernheim, B. D., & Naecker, J. (2018). When fair isn't fair: Understanding choice reversals involving social preferences. *NBER Working Paper*(25257).

Ariely, D., & Wertenbroch, K. (2002). Procastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, *13*(3).

Ashraf, N., Karlan, D., & Yin, W. (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *Quarterly Journal of Economics*, *121*(1), 635–672.

Augenblick, N., Niederle, M., & Sprenger, C. (2015). Working over time: Dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics*, *130*(3), 1067–1115.

Augenblick, N., & Rabin, M. (2019). An experiment on time preference and misprediction in unpleasant tasks. *Review of Economic Studies*, *86*(3), 941–975.

Baland, J.-M., & Ziparo, R. (2017). Intra-household bargaining in poor countries.

Banerjee, A. (2013). Microcredit under the microscope: What have we learned in the past two decades, and what do we need to know? *Annual Review of Economics*, *5*, 487–519.

Banerjee, A., & Mullainathan, S. (2010). The shape of temptation: Implications for the economic lives of the poor. *NBER Working Papers*, *15973*.

Barton, B. (2015). Interpersonal time inconsistency and commitment. *mimeo*.

Becker, G., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, *9*, 226–236.

Bernheim, D., & Taubinsky, D. (2019). Behavioral public economics. In D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics* (Vol. 1, p. 381-516). Elsevier, New York.

Berry, J., Dizon-Ross, R., & Jagnani, M. (2020). *Not playing favorites: An experiment on parental fairness preferences* (Tech. Rep.). NBER Working Paper No. w26732, February 2020.

Beshears, J., Choi, J. J., Harris, C., Laibson, D., Madrian, B. C., & Sakong, J. (2020). Which early withdrawal penalty attracts the most deposits to a commitment savings account? *Journal of Public Economics*, *183*(104144).

Brown, Z. E. C., Alexander L., & Camerer, C. F. (2009). Learning and visceral temptation in dynamic saving experiments. *Journal of Neuroscience*, *124*(1), 197—-231.

Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics*, *116*, 55–79.

Carrera, M., Royer, H., Stehr, M., Sydnor, J., & Taubinsky, D. (2019). How are preferences for commitment revealed? theory, evidence, and policy implications. *mimeo*.

Casaburi, L., & Willis, J. (2018). Time versus state in insurance: Experimental evidence from contract farming in kenya. *American Economic Review*, *108*(2), 3778–3813.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica*, *56*, 63–90.

Cubitt, R. P., & Read, D. (2007). Can intertemporal choice experiments elicit preferences for consumption? *Experimental Economics*, *10*(4), 369—389.

DellaVigna, S., & Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review*, *96*, 694—719.

Dupas, P., & Robinson, J. (2013). Why don't the poor save more? evidence from health savings experiments. *American Economic Review*, *103*(4), 1138—1171.

Exley, C. (2015). Excusing selfishness in charitable giving: The role of risk. *Review of Economics Studies*, *83*(2), 587-628.

Fark, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, *54*(2), 293-315.

Gine, X., Goldberg, J., Silverman, D., & Yang, D. (2016). Revising commitments: Field evidence on the adjustment of prior choice. *The Economic Journal*, *128*(February), 159–188.

Glennerster, R., & Kremer, M. (2012). Improving health in developing countries: Evidence from randomized evaluations. In *Handbook of health economics* (chap. 2).

Holden, S. T., & Quiggin, J. (2017). Climate risk and state-contingent technology adoption: shocks, drought tolerance and preferences. *European Review of Agricultural Economics*, *44*(2), 285–308.

Jackson, M. O., & Yariv, L. (2015). Collective dynamic choice: The necessity of time inconsistency. *American Economic Journal: Microeconomics*, *7*(4), 150—-178.

Kaur, S., Kremer, M., & Mullainathan, S. (2015). Self-control at work. *Journal of Political Economy*, *123*(6), 1227–1277.

Kling, J. R., Liebman, J. B., & Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, *75*(1), 83—119.

Koelle, M., & Wozny, L. P. (2018). Present-biased generosity: Time inconsistency across individual and social contexts. *CeDEx Discussion Paper Series*, *2018-02*.

Krawczyk, M., & Wozny, L. P. (2017). An experiment on temptation and attitude towards paternalism. *mimeo*.

Kremer, M., & Holla, A. (2009). Improving education in the developing world: What have we learned from randomized evaluations? *Annual Review of Economics*, *One*, 513–542.

Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review Papers and Proceedings*, *105*(5), 267–272.

Loewenstein, G. (2018). Self-control and its discontents. a commentary on "beyond willpower: Strategies for reducing failures of self-control" by angela duckworth, david laibson and katie milkman. *Psychological Science in the Public Interest*, *19*(3), 95–101.

Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics*, *118*(4), 1209–1248.

McClure, S., Laibson, D., Loewenstein, G., & Cohen, J. (2007). Time discounting for primary rewards. *Journal of Neuroscience*, *27*(21), 5796—5804.

Mullainathan, S., & Shafir, E. (2013). *Scarcity: Why having too little means so much*. New York, NY: Time Books, Henry Holt and Company LLC.

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review*, *89*(1), 103–124.

Ringdal, C., & Sjursen, I. H. (2017). Household bargaining and spending on children: Experimental evidence from tanzania. *WIDER Working Paper*, *2017/128*.

Schilbach, F. (2019). Alcohol and self-control: A field experiment in india. *American Economic Review*, *109*(4), 1290–1322.

Shapiro, J. (2010). Discounting for you, me and we: Time preference in groups and pairs. *mimeo*.

Solnick, C. H. K. D. A. E., Jay V., & Waller, M. B. (1980). An experimental analysis of impulsivity and impulse control in humans. *Experimental Economics*, *11*, 61—77.

Tarozzi, A., & Mahajan, A. (2011). Time inconsistency, expectations and technology adoption: The case of insecticide treated nets. *ERID Working Paper*, *No. 105*.

Thakral, N., & To, L. (2020). *Anticipation and consumption* (Tech. Rep.). Brown University.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*(3), 183—206.

Toussaert, S. (2018). Eliciting temptation and self-control through menu choices: A lab experiment. *Econometrica*, *86*(3), 859–889.

Ubfal, D. (2016). How general are time preferences? eliciting good-specific discount rates. *Journal of Development Economics*, *118*, 150-170.

Uhl, M. (2011). Do self-committers mind commitment by others? an experiment on weak paternalism. *Rationality, Markets and Morals*, *2*, 13–34.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, *12*, 187–191.

Appendices

A [Online Appendix] Proofs

A.1 Proof of Proposition 1

Generalizing the model's FOCs to decisions made at $t = j$ about consumption at $t = k$ yields:

$$\frac{u'(x_k^j)}{v'(z_k^j)} = \frac{\alpha}{\theta^{k-j}} \tag{A.1}$$

Due to decreasing marginal utility of consumption, if $\theta < 1$, then the child's budget share increases in $k - j$ (part 1) and decreases in j (part 2). ■

A.2 Proof of Proposition 2

For this proof, we assume that the instantaneous utility of consumption of both the parent and the child exhibits constant relative risk aversion (CRRA) γ:

$$v(c) = u(c) = \begin{cases} \frac{1}{1-\gamma} c^{1-\gamma}, & \text{if } \gamma \neq 1 \\ ln(c), & \text{otherwise} \end{cases} \tag{A.2}$$

The parent anticipates that her optimal allocation at $t = 2$ will be:

- $z_2^* = \alpha^{\frac{1}{\gamma}} \frac{(y+IR)}{\hat{\theta}^{\frac{1}{\gamma}} + \alpha^{\frac{1}{\gamma}}}$
- $x_2^* = \hat{\theta}^{\frac{1}{\gamma}} \frac{y+IR}{\hat{\theta}^{\frac{1}{\gamma}} + \alpha^{\frac{1}{\gamma}}}$

Let $C = \left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}} + \alpha^{\frac{1}{\gamma}}} \right)^{1-\gamma}$ and $B = \left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}} + \alpha^{\frac{1}{\gamma}}} \right)^{1-\gamma}$. Plugging the anticipated values of z_2 and x_2 into the $t = 1$ utility function yields: $\frac{(y-I-z_1)^{1-\gamma}}{1-\gamma} + \alpha \frac{z_1^{1-\gamma}}{1-\gamma} + (\theta\delta B + \alpha\delta C)(\frac{y+IR}{1-\gamma})^{1-\gamma}$.

Next, the first order condition with respect to z_1 yields: $z_1^* = \frac{(y-I)\alpha^{\frac{1}{\gamma}}}{1+\alpha^{\frac{1}{\gamma}}}$.

That with respect to I yields: $(y - I^* - z_1^*)^{-\gamma} = (\theta\delta BR + \alpha\delta CR)(y + I^*R)^{-\gamma}$

$$\Leftrightarrow y - I^* - z_1^* = (\theta\delta BR + \alpha\delta CR)^{-\frac{1}{\gamma}}(y + I^*R)$$

$$\Leftrightarrow I^* = \frac{y\left((BR\delta\theta+CR\alpha\delta)^{\frac{1}{\gamma}}-1-\alpha^{\frac{1}{\gamma}}\right)}{(BR\delta\theta+CR\alpha\delta)^{\frac{1}{\gamma}}+R(1+\alpha^{\frac{1}{\gamma}})}$$

$$\Leftrightarrow I^* = \frac{y\left(\left(\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\delta\theta+\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\alpha\delta\right)^{\frac{1}{\gamma}}-1-\alpha^{\frac{1}{\gamma}}\right)}{\left(\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\delta\theta+\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\alpha\delta\right)^{\frac{1}{\gamma}}+R\left(1+\alpha^{\frac{1}{\gamma}}\right)}$$

Hence, parents' optimal investment at $t = 1$ depends on the parent's degree of sophistication (given by $\hat{\theta}$).

Next, define: $X(\hat{\theta}) = \left(\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\delta\theta+\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{1-\gamma}R\alpha\delta\right)$

This leads to $I^*(X(\hat{\theta})) = \frac{y\left(X(\hat{\theta})^{\frac{1}{\gamma}}-1-\alpha^{\frac{1}{\gamma}}\right)}{X(\hat{\theta})^{\frac{1}{\gamma}}+R(1+\alpha^{\frac{1}{\gamma}})}$, with $\frac{\partial I^*}{\partial\hat{\theta}} = \frac{\partial I^*}{\partial X}\frac{\partial X}{\partial\hat{\theta}}$

The first term is $\frac{\partial I^*}{\partial X} = \frac{(R+1)\,X^{\frac{1}{\gamma}-1}y\left(\alpha^{\frac{1}{\gamma}}+1\right)}{\left(R\,\alpha^{\frac{1}{\gamma}}+X^{\frac{1}{\gamma}}+R\right)^2\gamma} > 0$

As for the second term:

$$\frac{\partial X}{\partial\hat{\theta}} = \frac{R\,\alpha^{\frac{1}{\gamma}}\,(\gamma-1)\delta\,\hat{\theta}^{\frac{1}{\gamma}-1}\left(\alpha\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}-\theta\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}\right)}{\gamma\left(\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}\right)^2\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}} = \frac{R\delta(1-\gamma)\alpha^{\frac{1+\gamma}{\gamma}}\hat{\theta}^{\frac{1}{\gamma-1}}(\theta-\hat{\theta})}{\left(\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}\right)^{\gamma}\gamma\left(\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}\right)^2\left(\frac{\alpha^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}\left(\frac{\hat{\theta}^{\frac{1}{\gamma}}}{\hat{\theta}^{\frac{1}{\gamma}}+\alpha^{\frac{1}{\gamma}}}\right)^{\gamma}}$$

Since $\hat{\theta} \in [\theta, 1]$, $\frac{\partial X}{\partial\hat{\theta}} \geq 0$ if $\gamma > 1$, and < 0 otherwise.

Putting the two terms together, $\frac{\partial I^*}{\partial\hat{\theta}} = \frac{\partial I^*}{\partial X}\frac{\partial X}{\partial\hat{\theta}} \geq 0$ if $\gamma > 1$, and < 0 otherwise. ■

A.3 Proof of Proposition 3

Denote as $\{z_t^*, x_t^*\}_{t=1,2}, I^*$ the solution to the parent's problem when investment pays out in cash in period $t = 2$. Now, let us modify the parents' utility maximization problem at $t = 1$ is as follows:

$$\underset{\{z_t,x_t\}_{t=1,2},I}{\text{Max}}\; u(x_1^1) + \alpha v(z_1^1) + \theta\delta u(x_2^1) + \alpha\delta v(z_2^1 + RI) \tag{A.3}$$

$$\text{s.t.}$$
$$\begin{cases} x_1^1 + z_1^1 + I \leq y \\ x_2^1 + z_2^1 \leq y \end{cases}$$

In the modified problem, investment I pays out directly as child's consumption in $t = 2$, under the same gross interest rate $R = 1 + r$. Denote the solution to this new problem $\{\hat{z}_t, \hat{x}_t\}_{t=1,2}, \hat{I}$.

In an interior solution, the ratio of marginal utilities at $t = 2$ is the same in both

problems:

$$\frac{u'(x_2^*)}{v'(z_2^*)} = \frac{u'(\hat{x}_2)}{v'(\hat{z}_2)} = \alpha \tag{A.4}$$

By that argument, if $R\hat{I} \leq z_2^*$, then we are at the interior solution and $\hat{z}_2^1 - \hat{z}_2^2 = z_2^{*,1} - z_2^{*,2}$.
In contrast, if $R\hat{I} > z_2^*$, then $\hat{z}_2 = R\hat{I}$, a corner solution, and it must be that $\hat{z}_2^1 - \hat{z}_2^2 < z_2^{*,1} - z_2^{*,2}$. ■

B [Online Appendix] Heterogeneity

Table B.1 documents that AGD preferences are *not* predicted by household or individual characteristics that we observe, neither at the baseline or at the follow-up wave. Next, we estimate heterogeneous effects of AGD preferences on parents' dynamic allocations (subsection B.1) and real-life investments in children (subsection B.2) by children's gender, age and birth order. These tests show that the effects of AGD preferences do not vary systematically with the child's characteristics, although some large differences suggest that we have limited statistical power to precisely detect them.

Table B.1: AGD and individual characteristics

	(1)	(2)
	AGD indicator	
Income	0.000	0.000
	(0.000)	(0.000)
First born	0.042	
	(0.054)	
Birth order	0.023	-0.009
	(0.028)	(0.015)
Number of children in household	0.010	0.002
	(0.012)	(0.018)
Child age	-0.008	-0.004
	(0.006)	(0.006)
Father	-0.097	0.021
	(0.061)	(0.047)
Credit constraint, round 1	-0.002	-0.001
	(0.003)	(0.002)
Education		-0.002
		(0.016)
Observations	813	1153
Sample	Control	Follow-up

Notes: This table estimates whether AGD preferences are predicted by household or individual characteristics. In column (1), the outcome variable is an indicator of AGD preferences determined in the baseline experiment. The sample in this column is restricted to the respondents assigned to the control group in the baseline experiment. In column (2), the outcome variable is an indicator of AGD preferences determined in the follow-up data collection. This column excludes from the follow-up sample the respondents who were not part of the main sample at baseline. * $p < 0.1$, ** $p < .05$, *** $p < .01$

B.1 Heterogeneous effects on dynamic allocation decisions

In each table, Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with the time gap between decisions and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) is revised downwards as consumption gets closer (although still in the future). Table B.2 estimates heterogeneous effects by the child's gender. Table B.3 estimates heterogeneous effects by the child's birth order, interacting the indicator of AGD preferences with an indicator of whether the child is firstborn. Table B.4 estimates heterogeneous effects by the child's age, interacting the indicator of AGD preferences with an indicator of whether the child is above or below the sample median age (8 years old). Last, Table B.5 estimates heterogeneous effects by the child's gender and birth order.

Table B.2: Testing the model's predictions: Heterogeneity by the child's gender

	Panel A: Prediction #1		Panel B: Prediction #2	
	(1)	(2)	(3)	(4)
	$\mathbb{1}\{s^2_{30} > s^2_2\}$	$s^2_{30} - s^2_2$	$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	$s^2_{30} - s^0_{30}$
$\mathbb{1}\{\hat{\theta} < 1\}$	0.373***	0.0735***	0.267***	-0.124***
	(0.0520)	(0.0170)	(0.0462)	(0.0158)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0569	0.0249	0.00492	0.0133
$\times Trait =$ Girl	(0.0715)	(0.0227)	(0.0640)	(0.0228)
Control	Yes	Yes	Yes	Yes
AGD parents' mean	0.581	0.087	0.328	-0.062
Symmetric parents' mean	0.176	0.001	0.063	0.054
N	795	795	795	795
Sample	Control	Control	Control	Control

Notes: Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In each column, we interact the indicator of AGD preferences with an indicator of the child's gender. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the difference between the share of peanuts set to be consumed by the child at $t = 30$ and $t = 2$ for decisions made in round-2. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the difference between the decision made at $t = 2$ and $t = 0$ for the share of peanuts allocated to the child at $t = 30$. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table B.3: Testing the model's predictions: Heterogeneity by the child's birth order

	Panel A: Prediction #1		Panel B: Prediction #2	
	(1)	(2)	(3)	(4)
	$\mathbb{1}\{s^2_{30} > s^2_2\}$	$s^2_{30} - s^2_2$	$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	$s^2_{30} - s^0_{30}$
$\mathbb{1}\{\hat{\theta} < 1\}$	0.393***	0.0829***	0.258***	-0.124***
	(0.0435)	(0.0133)	(0.0386)	(0.0142)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0323	0.0124	0.0321	0.0227
$\times Trait$ = First born	(0.0771)	(0.0259)	(0.0706)	(0.0249)
Control	Yes	Yes	Yes	Yes
AGD parents' mean	0.581	0.087	0.328	-0.062
Symmetric parents' mean	0.176	0.001	0.063	0.054
N	795	795	795	795
Sample	Control	Control	Control	Control

Notes: Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In each column, we interact the indicator of AGD preferences with an indicator of whether the child is firstborn. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the difference between the share of peanuts set to be consumed by the child at $t = 30$ and $t = 2$ for decisions made in round-2. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the difference between the decision made at $t = 2$ and $t = 0$ for the share of peanuts allocated to the child at $t = 30$. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table B.4: Testing the model's predictions: Heterogeneity by the child's age

	Panel A: Prediction #1		Panel B: Prediction #2	
	(1)	(2)	(3)	(4)
	$\mathbb{1}\{s^2_{30} > s^2_2\}$	$s^2_{30} - s^2_2$	$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	$s^2_{30} - s^0_{30}$
$\mathbb{1}\{\hat{\theta} < 1\}$	0.386***	0.0853***	0.270***	-0.117***
	(0.0459)	(0.0146)	(0.0400)	(0.0150)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0352	0.00160	0.00291	-0.00152
$\times Trait$ = Over 8 years old	(0.0737)	(0.0232)	(0.0668)	(0.0233)
Control	Yes	Yes	Yes	Yes
AGD parents' mean	0.581	0.087	0.328	-0.062
Symmetric parents' mean	0.176	0.001	0.063	0.054
N	795	795	795	795
Sample	Control	Control	Control	Control

Notes: Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In each column, we interact the indicator of AGD preferences with an indicator of whether the child is above or below the sample median age (8 years old). In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the difference between the share of peanuts set to be consumed by the child at $t = 30$ and $t = 2$ for decisions made in round-2. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the difference between the decision made at $t = 2$ and $t = 0$ for the share of peanuts allocated to the child at $t = 30$. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table B.5: Testing the model's predictions: Heterogeneity by the child's birth order and gender

	Panel A: Prediction #1		Panel B: Prediction #2	
	(1)	(2)	(3)	(4)
	$\mathbb{1}\{s^2_{30} > s^2_2\}$	$s^2_{30} - s^2_2$	$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	$s^2_{30} - s^0_{30}$
$\mathbb{1}\{\hat{\theta} < 1\}$	0.369***	0.0723***	0.248***	-0.126***
	(0.0627)	(0.0196)	(0.0545)	(0.0201)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0109	0.00324	0.0633	0.00703
$\times Trait =$ First born	(0.114)	(0.0392)	(0.103)	(0.0323)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0363	0.0156	-0.0594	0.0288
$\times Trait =$ First born $\times$ Girl	(0.155)	(0.0518)	(0.143)	(0.0491)
$\mathbb{1}\{\hat{\theta} < 1\}$	0.0452	0.0203	0.0192	0.00504
$\times Trait =$ Girl	(0.0870)	(0.0263)	(0.0768)	(0.0280)
Control	Yes	Yes	Yes	Yes
AGD parents' mean	0.581	0.087	0.328	-0.062
Symmetric parents' mean	0.176	0.001	0.063	0.054
N	795	795	795	795
Sample	Control	Control	Control	Control

Notes: Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In each column, we interact the indicator of AGD preferences with an indicator of the child's gender, with an indicator of whether the child is firstborn or not, and with both indicators interacted together. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the difference between the share of peanuts set to be consumed by the child at $t = 30$ and $t = 2$ for decisions made in round-2. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the difference between the decision made at $t = 2$ and $t = 0$ for the share of peanuts allocated to the child at $t = 30$. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children in the household, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

B.2 Real-life investments in children

Each table documents the correlation between AGD preferences and investments in children, also estimating the latter's correlation with present-bias as a benchmark. In each table, column (1) restricts attention to 3-5 year-old children, and column (2), to 6-12 year-old. Table B.6 estimates heterogeneous effects by the child's gender. Table B.7 estimates heterogeneous effects by the child's birth order, splitting

the sample by first-born children and all others. We estimate seemingly unrelated regressions (SUR) to test for equality of coefficients across different regressions.

Table B.6: AGD and investments in children: Heterogeneity by the child's gender

	(1) Index of investments, 3-5 years old		(2) Index of investments, 6-12 years old	
	(1) Boys	(2) Girls	(3) Boys	(4) Girls
$\mathbb{1}(\hat{\theta} < 1)$	4.8368	8.9379	4.5465	9.8061**
	(6.5396)	(5.7858)	(3.9616)	(4.2463)
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$	-4.8334	-8.9846	-4.5646	-9.8781**
	(6.5810)	(5.8322)	(3.9937)	(4.2797)
$\hat{\delta}$	4.2531	1.0103	0.8702	5.5154**
	(3.3926)	(3.3450)	(2.3090)	(2.4538)
$\mathbb{1}(\hat{\beta} < 1)$	3.9584	-0.3479	-1.3478	4.4773
	(4.7593)	(4.5811)	(3.4552)	(3.6027)
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$	-4.0871	0.4364	1.4092	-4.5672
	(4.8119)	(4.6343)	(3.4926)	(3.6401)
Tests of equality of coefficients for boys vs girls (p-value)				
$\mathbb{1}(\hat{\theta} < 1)$		0.6711		0.3808
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$		0.6691		0.3797
$\mathbb{1}(\hat{\beta} < 1)$		0.5473		0.2480
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$		0.5317		0.2412
Control variables	No	No	No	No
N	389	462	759	726
$\mathbb{1}(\hat{\theta} < 1)$=$\mathbb{1}(\hat{\beta} < 1)$ (p-value)	0.9183	0.2469	0.2663	0.3805
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$=$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$ (p-value)	0.9311	0.2450	0.2645	0.3867

Notes: Across all columns, the outcome variable is a summary index variable of investments in children, whose components are described in Supplementary Appendix S1.3; each component is normalized with respect to its mean and standard deviation among symmetric parents. Columns (1) and (2) restrict attention to investments in children 5 years old and younger; columns (3) and (4), to investments in children 6 years old and older. Columns (1) and (3) display results for the sub-sample of boys in our data, while columns (2) and (4) are for the sub-sample of girls. $\hat{\delta}$ is the discount factor each parent attaches to their child's future utility of consumption, calibrated from their allocation decisions in the experiment (see Supplementary Appendix S2). We do not include other controls because present-bias correlates with some household and individual characteristics. Standard errors in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table B.7: AGD and investments in children: Heterogeneity by the child's birth order

	(1)		(2)	
	Index of investments, 3-5 years old		Index of investments, 6-12 years old	
	$\neq$ first born	= first born	$\neq$ first born	= first born
$\mathbb{1}(\hat{\theta} < 1)$	5.5269	5.7170	6.6338**	8.2623
	(5.4347)	(7.1106)	(3.3109)	(6.0162)
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$	-5.5695	-5.6768	-6.6916**	-8.2809
	(5.4757)	(7.1607)	(3.3358)	(6.0722)
$\hat{\delta}$	3.0049	2.2982	3.2724*	3.7161
	(3.1400)	(3.6965)	(1.8548)	(3.9289)
$\mathbb{1}(\hat{\beta} < 1)$	-0.3712	7.7877	0.4108	6.7298
	(4.1922)	(5.4193)	(2.7974)	(5.4286)
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$	0.3517	-7.8369	-0.3976	-6.8675
	(4.2401)	(5.4805)	(2.8266)	(5.4889)
Tests of equality of coefficients for first born vs not first born (p-value)				
$\mathbb{1}(\hat{\theta} < 1)$		0.9848		0.8272
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$		0.9915		0.8327
$\mathbb{1}(\hat{\beta} < 1)$		0.2846		0.3202
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$		0.2878		0.3144
Control variables	No	No	No	No
N	486	362	1,166	315
$\mathbb{1}(\hat{\theta} < 1)$=$\mathbb{1}(\hat{\beta} < 1)$ (p-value)	0.4127	0.8339	0.1704	0.8576
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$=$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$ (p-value)	0.4153	0.8285	0.1696	0.8699

Notes: Across all columns, the outcome variable is a summary index variable of investments in children, whose components are described in Supplementary Appendix S1.3; each component is normalized with respect to its mean and standard deviation among symmetric parents. Columns (1) and (2) restrict attention to investments in children 5 years old and younger; columns (3) and (4), to investments in children 6 years old and older. Columns (1) and (3) display results for the sub-sample of children who are *not* firstborn in our data, while columns (2) and (4) are for the sub-sample of children who are firstborn. $\hat{\delta}$ is the discount factor each parent attaches to their child's future utility of consumption, calibrated from their allocation decisions in the experiment (see Supplementary Appendix S2). We do not include other controls because present-bias correlates with some household and individual characteristics. Standard errors in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

C [Online Appendix] Additional results

This Appendix compiles additional results. Section C.1 documents additional robustness checks for our results on testing the model's predictions and real-life investments in children. Specifically, we show that those results are robust to clustering standard errors at the village level to allow for spatially correlated shocks, to absorbing village fixed-effects, and for controlling flexibly for payday effects and for the price of commitment, to rule out identification concerns discussed in the main text. Section C.2 documents selective attrition tests. Section C.3 documents whether parents' decisions during the experiment affects children's consumption outside of the experiment. Section C.4 presents results on whether AGD parents are more likely to demand commitment to round-1 decisions in the experiment. Last, Section C.5 presents additional results on whether AGD parents are more likely to demand commitment outside the lab.

C.1 Additional robustness checks

C.1.1 Dynamic allocation decisions

In each table, Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with the time gap between decisions and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) is revised downwards when as consumption gets closer (although still in the future). We control for village fixed-effects in Table C.1, cluster standard errors at the village level in Table C.2, and allow the slope of the allocation schedule as well as reversals to vary with the date of the visit in Table C.3 and with the price of commitment in Table C.4.

Table C.1: Testing the model's predictions: Village-level shocks

Panel A: Prediction #1			Panel B: Prediction #2		
	(1)	(2)		(3)	(4)
	$\mathbb{1}\{s_{30}^2 > s_2^2\}$	s_j^2		$\mathbb{1}\{s_{30}^0 > s_{30}^2\}$	s_{30}^k
$j-2$		0.000026	$\mathbb{1}\{k=2\}$		0.0534***
		(0.0002)			(0.0061)
$\mathbb{1}\{\hat{\theta}<1\}$	0.390***	-0.0197**	$\mathbb{1}\{\hat{\theta}<1\}$	0.274***	0.187***
	(0.0370)	(0.0086)		(0.0330)	(0.0081)
$\mathbb{1}\{\hat{\theta}<1\}$		0.0031***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.118***
$\times(j-2)$		(0.0004)	$\times\mathbb{1}\{k=2\}$		(0.0118)
Village fixed effects	Yes	Yes		Yes	Yes
Control	Yes	Yes		Yes	Yes
AGD parents' mean	0.581	0.536		0.328	0.612
Symmetric parents' mean	0.176	0.513		0.063	0.487
N	795	1590		795	1608
Respondents	795	795		795	813
Sample	Control	Control		Control	Control

Notes: This table reports the estimated impact of AGD on allocation decisions by controlling for village fixed effects. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. We further control for village fixed effects. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.2: Testing the model's predictions: Higher-level clustering

Panel A: Prediction #1			Panel B: Prediction #2		
	(1)	(2)		(3)	(4)
	$\mathbb{1}\{s_{30}^2 > s_2^2\}$	s_j^2		$\mathbb{1}\{s_{30}^0 > s_{30}^2\}$	s_{30}^k
$j-2$		0.000026	$\mathbb{1}\{k=2\}$		0.053***
		(0.0002)			(0.0055)
$\mathbb{1}\{\hat{\theta}<1\}$	0.403***	-0.0196***	$\mathbb{1}\{\hat{\theta}<1\}$	0.270***	0.186***
	(0.0362)	(0.0074)		(0.0340)	(0.0076)
$\mathbb{1}\{\hat{\theta}<1\}$		0.0031***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.118***
$\times(j-2)$		(0.0004)	$\times\mathbb{1}\{k=2\}$		(0.0120)
Control	Yes	Yes		Yes	Yes
AGD parents' mean	0.581	0.536		0.328	0.612
Symmetric parents' mean	0.176	0.513		0.063	0.487
N	795	1590		795	1608
Respondents	795	795		795	813
Sample	Control	Control		Control	Control

Notes: This table reports the estimated impact of AGD on allocation decisions by clustering the standard errors at the village level. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. We further control for village fixed effects. Standard errors clustered at the village level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.3: Testing the model's predictions: Controlling for visit days

Panel A: Prediction #1			Panel B: Prediction #2		
	(1)	(2)		(3)	(4)
	$\mathbb{1}\{s_{30}^2 > s_2^2\}$	s_j^2		$\mathbb{1}\{s_{30}^0 > s_{30}^2\}$	s_{30}^k
$j-2$		0.0000164	$\mathbb{1}\{k=2\}$		0.047***
		(0.0003)			(0.0086)
$\mathbb{1}\{\hat{\theta}<1\}$	0.403***	-0.0209**	$\mathbb{1}\{\hat{\theta}<1\}$	0.270***	0.186***
	(0.0357)	(0.0089)		(0.0323)	(0.0084)
$\mathbb{1}\{\hat{\theta}<1\}$		0.0032***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.117***
$\times(j-2)$		(0.0004)	$\times\mathbb{1}\{k=2\}$		(0.0120)
Control	Yes	Yes		Yes	Yes
AGD parents' mean	0.581	0.536		0.328	0.612
Symmetric parents' mean	0.176	0.513		0.063	0.487
N	795	1508		795	1523
Respondents	795	754		795	769
Sample	Control	Control		Control	Control

Notes: This table reports the estimated impact of AGD on allocation decisions by controlling for a dummy variable indicating whether the visit by enumerators took place in the first-half (dummy=0) or second-half (dummy=1) of the month. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (1), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (2), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (3), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (4), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. It is also further restricted to households for which the day of the visit was correctly recorded in the data. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. We further control for a dummy variable indicating whether the visit by enumerators took place in the first-half (dummy=0) or second-half (dummy=1) of the month. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.4: Robustness to learning about costly commitment between decision rounds

		Panel A: Prediction #1			Panel B: Prediction #2		
	(1)		(2)	(3)		(4)	(5)
	Δs_{30}		$\mathbb{1}\{s^2_{30} > s^2_2\}$	s^2_j		$\mathbb{1}\{s^0_{30} > s^2_{30}\}$	s^k_{30}
		$j-2$		0.0000257	$\mathbb{1}\{k=2\}$		0.053***
				(0.0002)			(0.0059)
		$\mathbb{1}\{\hat{\theta}<1\}$	0.360***	-0.0198**	$\mathbb{1}\{\hat{\theta}<1\}$	0.225***	0.185***
			(0.0934)	(0.0087)		(0.0844)	(0.0081)
		$\mathbb{1}\{\hat{\theta}<1\}$		0.0029***	$\mathbb{1}\{\hat{\theta}<1\}$		-0.130***
		$\times(j-2)$		(0.0009)	$\times\mathbb{1}\{k=2\}$		(0.0249)
Commitment price	0.0110	Commitment price	-0.0404	-0.0054	Commitment price	-0.0480*	-0.0131
	(0.0134)		(0.0381)	(0.0095)		(0.0261)	(0.0091)
		$\mathbb{1}\{\hat{\theta}<1\}\times$ Price	0.0413		$\mathbb{1}\{\hat{\theta}<1\}\times$ Price	0.0435	
			(0.0845)			(0.0763)	
		$\mathbb{1}\{\hat{\theta}<1\}\times$ Price		0.0001	$\mathbb{1}\{\hat{\theta}<1\}\times$ Price		0.0117
		$\times(j-2)$		(0.0008)	$\times\mathbb{1}\{k=2\}$		(0.0229)
Control	No		Yes	Yes		Yes	Yes
AGD parents' mean	-0.062		0.581	0.536		0.328	0.612
Symmetric parents' mean	0.054		0.176	0.513		0.063	0.487
N	719		795	1590		795	1608
Respondents	719		795	795		795	813
Sample	Control		Control	Control		Control	Control

Notes: This table reports the estimated impact of the price of commitment on: the change in the child's planned consumption share at t=3 between decision rounds in column (1); the model's predictions in Panels A and B. The price of commitment takes values 0.5, 1, and 1.5. In column (1), the outcome variable is the change in the child's planned consumption share at t=3 between decision rounds. Panel A tests model's prediction #1 that AGD parents' allocations set to children increase with time horizon between the decision and consumption; and Panel B tests model's prediction #2 that AGD parents' allocation set to be consumed by children at the last round ($t = 30$) decreases when the decision is made at $k = 2$ relative to when it is made at $k = 1$. In column (2), the outcome variable equals 1 if the parent set time-increasing budget shares to her child at round 2, and 0 otherwise. In column (3), the outcome variable is the share of peanuts set to be consumed by the child at $t = j$ for decisions made in round-2 ($k = 2$); participants' allocations for each consumption horizon j are stacked for the analysis. In column (4), the outcome variable equals 1 if the parent set a lower consumption share to be consumed by her child at $t = 30$ when the decision was made at $t = 2$ relative to when it was made at $t = 0$. In column (5), the outcome variable is the share of peanuts allocated to the child at $t = 30$ when the decision is made at $t = k$; participants' allocations for each decision round k are stacked for the analysis. Across all columns, the sample is restricted to the control group of the framing experiment. Control variables in Panels A and B include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

C.1.2 Real-life investments in children

Each table documents the correlation between AGD preferences and investments in children, also estimating the latter's correlation with present-bias as a benchmark. In each table, column (1) restricts attention to 3-5 year-old children, and column (2), to 6-12 year-old. We control for village fixed-effects in Table C.5 and cluster standard errors at the village level in Table C.6.

Table C.5: AGD and investments in children: Village-level shocks

	Index of investments, 3-5 years old		Index of investments, 6-12 years old	
	(1)	(2)	(3)	(4)
$\mathbb{1}(\hat{\theta} < 1)$	0.0246	2.9500	0.0142	6.9119**
	(0.0315)	(4.3871)	(0.0213)	(2.8776)
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$		-2.9359		-6.9651**
		(4.4192)		(2.9007)
$\hat{\delta}$		1.7319		4.5209***
		(2.4642)		(1.6956)
$\mathbb{1}(\hat{\beta} < 1)$	-0.0174	4.2145	-0.0019	2.6935
	(0.0290)	(3.3900)	(0.0204)	(2.4969)
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$		-4.2756		-2.7290
		(3.4283)		(2.5235)
Village fixed effects	Yes	Yes	Yes	Yes
Control variables	No	No	No	No
N	851	851	1,485	1,485

Notes: Across all columns, the outcome variable is a summary index variable of investments in children, whose components are described in Supplementary Appendix S1.3; each component is normalized with respect to its mean and standard deviation among symmetric parents. Columns (1) and (2) restrict attention to investments in children 5 years old and younger; columns (3) and (4), to investments in children 6 years old and older. $\hat{\delta}$ is the discount factor each parent attaches to their child's future utility of consumption, calibrated from their allocation decisions in the experiment (see Supplementary Appendix S2). We control for village fixed effects in all columns. We do not include other controls because present-bias correlates with some household and individual characteristics. Standard errors in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.6: AGD and investments in children: Higher-level clustering

	Index of investments, 3-5 years old		Index of investments, 6-12 years old	
	(1)	(2)	(3)	(4)
$\mathbb{1}(\hat{\theta} < 1)$	0.0261	5.7702	0.0132	7.5130**
	(0.0310)	(4.5665)	(0.0215)	(2.9872)
$\hat{\delta} \times \mathbb{1}(\hat{\theta} < 1)$		-5.7884		-7.5638**
		(4.5959)		(3.0122)
$\hat{\delta}$		2.8095		3.3674*
		(2.2176)		(1.7667)
$\mathbb{1}(\hat{\beta} < 1)$	0.0023	2.2392	0.0020	1.7086
	(0.0308)	(3.6830)	(0.0211)	(2.2512)
$\hat{\delta} \times \mathbb{1}(\hat{\beta} < 1)$		-2.2611		-1.7256
		(3.7237)		(2.2780)
Control variables	No	No	No	No
N	851	851	1,485	1,485

Notes: Across all columns, the outcome variable is a summary index variable of investments in children, whose components are described in Supplementary Appendix S1.3; each component is normalized with respect to its mean and standard deviation among symmetric parents. Columns (1) and (2) restrict attention to investments in children 5 years old and younger; columns (3) and (4), to investments in children 6 years old and older. $\hat{\delta}$ is the discount factor each parent attaches to their child's future utility of consumption, calibrated from their allocation decisions in the experiment (see Supplementary Appendix S2). We do not include other controls because present-bias correlates with some household and individual characteristics. Standard errors clustered at the village level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

C.2 Selective attrition tests

This Appendix tests for selective attrition between decision rounds within the baseline wave, and across the baseline and follow-up waves. Tables C.7 and C.8 show that attrition is not systematically affected by treatment assignment at the framing experiment or by commitment price, neither at baseline or at the follow-up wave. Tables C.9 and C.10 show that AGD preferences do not systematically correlate with attrition within each treatment arm.

Table C.7: Test of differential attrition across rounds (baseline)

	(1) Commitment	(2) Framing	(3) Baseline variables
Probabilistic, 1	-0.00714		
	(0.0103)		
Probabilistic, 1.5	-0.00201		
	(0.0100)		
Labeling		0.0158	
		(0.0101)	
Anchoring		0.00342	
		(0.0101)	
AGD indicator			-0.00253
			(0.0094)
Mother			0.0368**
			(0.0145)
Islam			-0.000589
			(0.0110)
Number of children			0.00656
			(0.00413)
Credit constraint, round 1			0.00123
			(0.000778)
Female child			0.00135
			(0.00823)
Child age			-0.000284
			(0.00145)
s_2^0			-0.0563
			(0.0373)
Constant	0.975***	0.967***	0.949***
	(0.00731)	(0.00580)	(0.0276)
N	1627	1627	1627
Mean	0.972	0.972	0.972
F-test (p-value)	0.773	0.289	0.075

Notes: This table displays the results of three OLS regressions. The outcome variable is a dummy variable equal to 1 if the respondent was observed in all three waves of baseline data collection and 0 otherwise. Column (1) shows how attrition varies according to the type and price of commitment offered to the respondent. The omitted category is having been assigned to the probabilistic commitment device for a price of 0.5 packets of peanuts. Column (2) shows how attrition varies according to the framing of the second visit. The omitted categoy is the control group. Column (3) shows the correlation between attrition and baseline characteristics. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.8: Test of differential attrition across survey waves

	(1) Commitment	(2) Framing	(3) Baseline variables
Probabilistic, 1	-0.00196		
	(0.0179)		
Probabilistic, 1.5	0.0170		
	(0.0174)		
Labeling		0.0177	
		(0.0175)	
Anchoring		0.0127	
		(0.0175)	
AGD indicator			-0.0152
			(0.0163)
Mother			-0.0476*
			(0.0253)
Islam			-0.0174
			(0.0191)
Number of children			0.00755
			(0.00719)
Credit constraint, round 1			-0.00258*
			(0.00136)
Female child			0.00928
			(0.0143)
Child age			0.00187
			(0.00252)
s_2^0			0.00876
			(0.0650)
Constant	0.903***	0.901***	0.890***
	(0.0127)	(0.0101)	(0.0414)
N	1627	1627	1627
Mean	0.908	0.908	0.908
F-test (p-value)	0.479	0.553	0.207

Notes: This table displays the results of three OLS regressions. The outcome variable is a dummy variable equal to 1 if the respondent was observed in the baseline data collection and in the follow-up and 0 otherwise. Column (1) shows how attrition varies according to the type and price of commitment offered to the respondent. The omitted category is having been assigned to the probabilistic commitment device for a price of 0.5 packets of peanuts. Column (2) shows how attrition varies according to the framing of the second visit. The omitted categoy is the control group. Column (3) shows the correlation between attrition and baseline characteristics. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.9: Test of differential attrition per treatment arm (across rounds, baseline)

	(1)	(2)	(3)	(4)	(5)	(6)
	Control	Labeling	Anchoring		Commitment	
				0.5	1	1.5
AGD indicator	0.000701	-0.0166	0.00487	-0.00782	0.0218	-0.0172
	(0.0143)	(0.0146)	(0.0198)	(0.0157)	(0.0177)	(0.0154)
Observations	817	405	405	515	525	587

Notes: This table displays the results of OLS regressions. The outcome variable is a dummy variable equal to 1 if the respondent was observed in all three waves of baseline data collection and 0 otherwise. The outcome variable is regressed on the same baseline variables as those reported in column (3) in Table C.7, separately for each probabilistic commitment device (in columns (1) to (3)) and price (in columns (4) to (6)). We only report the coefficient for the indicator of AGD preferences. Standard errors in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.10: Test of differential attrition per treatment arm (across survey waves)

	(1)	(2)	(3)	(4)	(5)	(6)
	Control	Labeling	Anchoring		Commitment	
				0.5	1	1.5
AGD indicator	-0.0502*	0.0317	0.0107	-0.0553	0.0222	-0.0120
	(0.0240)	(0.0308)	(0.0323)	(0.0296)	(0.0303)	(0.0256)
Observations	817	405	405	515	525	587

Notes: This table displays the results of OLS regressions. The outcome variable is a dummy variable equal to 1 if the respondent was observed in the baseline data collection and in the follow-up and 0 otherwise. The outcome variable is regressed on the same baseline variables as those reported in column (3) in Tables C.7 and C.8, separately for each probabilistic commitment device (columns (1) to (3)) and price (columns (4) to (6)). We only report the coefficient for the AGD indicator. Standard errors in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.3 Fungibility of peanuts and consumption outside the lab

Table C.11 regresses measures of children's consumption outside of the experiment at the beginning of round 2 (the number of hours since the child last ate and whether the child is hungry) on parent's allocation decisions on that date and two days before. The correlation between children's consumption outside of the experiment and parents' round-1 allocations to children express parents' plans to adjust children's consumption outside of the experiment; that between the former and parents' round-2 allocations express whether they follow through on those plans. The table shows that while parents systematically plan to adjust their child's consumption in response to the share of peanuts they are bound to receive in the experiment, they do not follow through on those plans. Moreover, child's hunger is unrelated to consumption shares allocated in the experiment at either decision round.

Table C.11: Substitution between allocation decisions and consumption outside of the experiment (beginning of round 2)

	(1)	(2)	(3)	(4)
	Number of hours since the child last ate		Child is hungry	
s^k_{ji}	3.467**	1.272	0.0827	0.0304
	(1.747)	(1.671)	(0.129)	(0.126)
$s^k_{ji} \times \mathbb{1}\{j = 2\}$	-0.886	0.407	0.0213	0.0318
	(2.066)	(1.955)	(0.167)	(0.162)
$\mathbb{1}\{j = 2\}$	0.550	-0.157	-0.0066	-0.0141
	(1.023)	(1.004)	(0.0828)	(0.0835)
Control variables	Yes	Yes	Yes	Yes
k	1	2	1	2
Mean	5.30	5.30	0.48	0.48
N	1,590	1,590	1,590	1,590
Respondents	795	795	795	795

Notes: This table reports the results from two sets of Ordinary Least Squares (OLS) regressions which assess whether parents adjust the child's outside consumption in anticipation of round 2. In colums (1-2) the outcome variable is the number of hours since the child last ate by the start of round 2. In colums (3-4) the outcome variable is a dummy variable equal to 1 if the respondent reports that the child is hungry at the beginning of round 2, 0 otherwise. In columns (1) and (3), the independent variables are the shares of peanuts that the parents allocated to the child in their first round decision. In columns (2) and (4), the independent variables are the shares of peanuts that the parents allocated to the child in their second round decision. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. Standard errors clustered at the household level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

C.4 Demand for commitment in the lab

We test whether AGD parents demand commitment to past plans to a greater extent than symmetric parents with the following regression:

$$Y_i = \alpha + \gamma_0 \mathbb{1}\{\hat{\theta}_i < 1\} + \gamma_1 Price_i + \gamma_2 \left(Price_i \times \mathbb{1}\{\hat{\theta}_i < 1\} \right) + \lambda X_i + \varepsilon_i, \tag{C.1}$$

where Y_i equals 1 if the parent i takes up commitment at $Price_i$ (randomly assigned), and 0 otherwise. We are interested in testing $\gamma_0 \geq 0$ and $\gamma_2 \geq 0$.

Table C.12 shows that while AGD parents are more likely than symmetric parents to demand commitment at low prices, their demand falls more steeply with prices than that of other parents. All in all, commitment in the lab is very high, and there is no systematic relationship between AGD preferences and demand for commitment in the experiment. As a benchmark, we also offered participants the possibility of committing to their round-1 decisions in their intertemporal scenario. Table C.12 also shows very high commitment on average, but no systematic relationship between present-bias and demand for commitment.

Table C.12: Demand for commitment in the lab

	(1)	(2)	(3)	(4)	(5)
	Takes up the probabilistic commitment device				
	Main Scenario			Inter-temporal Scenario	
$\mathbb{1}\{\hat{\theta} < 1\}$	0.103**	0.0148			
	(0.0430)	(0.0161)			
$\mathbb{1}\{\hat{\theta} < 1\} \times Price$	-0.0867**				
	(0.0391)				
$\mathbb{1}\{\hat{\beta} < 1\}$				0.0334	0.0190
				(0.0404)	(0.0150)
$\mathbb{1}\{\hat{\beta} < 1\} \times Price$				-0.0139	
				(0.0366)	
Belief Parent-Bias			-0.0136		
			(0.0411)		
Price × Belief Parent-Bias			0.0396		
			(0.0372)		
Price	-0.0194		-0.0616***	-0.0358	
	(0.0215)		(0.0228)	(0.0239)	
Mean for $\mathbb{1}\{\hat{\theta} < 1\}$	91.1 %	91.1%			
Mean for $\mathbb{1}\{\hat{\beta} < 1\}$				91.3%	91.3%
Price fixed effects	No	Yes	No	No	Yes
Control variables	Yes	Yes	Yes	Yes	Yes
N	1622	1622	1622	1591	1591

Notes: This table reports the estimated impact of AGD (col. (1-2)), beliefs over reversals (col. (3)) and present-bias (col. (4-5)) on the demand for the probabilistic commitment device in the main Scenario and the inter-temporal Scenario respectively. The outcome variable is a dummy equal to 1 if the respondent takes up the probabilistic commitment device, zero otherwise. Respondents choose whether to take up the commitment device after making allocation decisions during the first visit. Taking up the commitment device reduces the probability that round 2 decision will be implemented. The commitment device comes at a random price (0.5, 1 or 1.5 packages of peanuts deducted from the parents' allocation at $t = 3$.). Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, an indicator variable of whether the household is Muslim or not, and the order in which the inter-temporal and main scenarios were presented to the respondent. The sample is restricted to parents having been offered the probabilistic commitment device. Belief Parent-Bias is an indicator variable equal to 1 if respondents predict that the majority of respondents would allocate more peanuts to themselves in the second round, in the incentivized task, and 0 otherwise. Standard errors clustered at the household level in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$

C.5 Demand for commitment outside the lab

This Appendix documents participants' willingness to commit lottery proceedings to a savings account, varying whether that account was earmarked to their child or not. Table C.13 presents the results. While we find that AGD parents are significantly more likely to express interest in commitment in that case, they are no more willing to pay for it than symmetric parents. What is more, their demand actually seems to be more closely connected to a savings account than to commitment against parent-bias, as they are actually much less willing to pay for an account earmarked to their children relative to one under their own name (the effect size is large, but imprecisely estimated as less that 12% of respondents express interest in opening a bank account under their child's name).

Table C.13: AGD preferences and demand for a savings account (earmarked to their child or not)

	(1)	(2)	(3)	(4)
	Interest in committing lottery prize to child's savings account	WTP	Interest in child's savings account (rather than own account)	WTP
$\mathbb{1}\{\theta < 1\}$	0.027*	-55.71	0.016	-133.95
	(0.015)	(53.11)	(0.018)	(174.73)
AGD parents' mean	0.935	3607.88	0.138	536.67
Symmetric parents' mean	0.913	3638.37	0.118	639.87
Control variables	Yes	Yes	Yes	Yes
N	1989	1826	1887	237

Notes: This table looks at the impact of AGD on parents' willingness to pay for a savings account. In column (1), the outcome variable is a dummy equal to 1 if the parent chooses to open a bank account in its child's name instead of receiveing 2,000 kwachas in cash if they earned it in the lottery, even when the savings account's option is free. In column (2), the outcome variable is the willingness to pay for this savings account, restricting the sample to parents for which there exists a non-negative price at which they would take up the savings account. In column (3), the outcome variable is a dummy equal to 1 if the parent chooses to open a savings account in the name of its child rather than its own. In column (4), the outcome variable is the willingness to pay for this savings account, among parents for which there exists a non-negative price at which they would take up the savings account in their child's name. Control variables include the gender and age of the respondent, the gender and age of the child, a measure of credit constraints, the number of children, the religion of the household, the level of education of the respondent. Standard errors in parentheses. * $p < 0.1$, ** $p < .05$, *** $p < .01$