The Family Care Contract: Childcare, Long-Term Care, and Intergenerational Insurance

Neha Bairoliya* Karen Kopecky† Bettina Siflinger‡
Neha.Bairoliya@usc.edu karen.kopecky@clev.frb.org B.M.Siflinger@uvt.nl

October 4, 2026

Abstract

Families exchange childcare and informal long-term care over the life cycle, thus linking the well-being of parents and adult children. Using the Health and Retirement Study, we document how care in both directions varies with family resources. Adult children whose households received grandparental childcare are 16 percentage points more likely to provide later long-term care. This association persists across income and education groups, after controlling for observed characteristics, and in comparisons between children of the same parent. Grandparental childcare is more common in higher-resource families, while its association with later care is stronger among lower-income parents. We develop a quantitative model of overlapping generations with limited commitment: families share resources and exchange care only while both generations are willing to participate. The model shows how reciprocity, saving, resource sharing, and the protection of bequests from long-term care costs connect care decisions across life stages. Earlier childcare shapes subsequent support, while anticipated care in old age influences childcare today. Subsidizing purchased childcare benefits young adults but reduces the family resources available to current grandparents, leaving them worse off on average even without policy taxes. Extending eligibility to grandparental care reverses grandparents' average losses and roughly doubles long-run lifetime welfare gains at the same public spending target.

Keywords: childcare subsidies, long-term care, intergenerational transfers, reciprocity, limited commitment, family insurance.

JEL Classification Numbers: E21, E62, D15, H31, J13, J14.

*University of Southern California, Los Angeles, CA 90089-1422, United States.

†Federal Reserve Bank of Cleveland, Cleveland, OH 44114, United States. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

‡Tilburg University, Tilburg School of Economics and Management, Tilburg, The Netherlands.

1 Introduction

Childcare subsidies can ease young families' resource constraints and help parents combine work with raising children. Their effects, however, extend to a third generation. Grandparents often provide childcare and may later rely on their adult children for support in old age. Making purchased care more affordable changes the value of this relationship to both generations: it improves young adults' alternatives to grandparental help, but can also change the resources and care their parents receive. How does this intergenerational relationship shape the welfare effects of childcare policy? Can a subsidy be designed to support young families while also benefiting the grandparents who help them?

This paper combines evidence on reciprocal care with a quantitative model of family decision-making to answer these questions. Our central finding is that a subsidy restricted to purchased childcare can raise young adults' average welfare while reducing that of incumbent parents, even when neither generation bears its financing cost. The parent's loss arises through changes in the family relationship, rather than only through taxation. Extending eligibility to grandparental care within the same fiscal spending target reverses the parent's average loss and roughly doubles long-run welfare gains for young entrants. Accounting for family care therefore changes both the incidence of childcare support and the design that delivers larger gains.

We begin with evidence from the Health and Retirement Study (HRS), linking earlier grandparental childcare to subsequent long-term care (LTC) within parent–child pairs. Adult children whose households received childcare assistance are 16 percentage points more likely to provide care to their parents later. This association is present across parent-income and child-education groups and persists after controlling for observed characteristics, including in comparisons between children of the same parent. The relationship is stronger among lower-income parents, even though grandparental childcare itself is more common in higher-resource families. These patterns distinguish the prevalence of care from the strength of the connection between its two directions. The structural analysis uses these patterns jointly to quantify the incentives for family care and study their implications for policy.

Our model follows overlapping generations through adulthood: an individual first interacts with an older parent and subsequently becomes a parent in a new care relationship. When young adults have children, cooperation allows them to receive grandparental childcare and share resources with their parents. Later, when those parents may need LTC, cooperation allows the adult child to provide informal care. Both generations choose consumption and saving, and can instead rely on their own resources and care alternatives. Agreements must remain acceptable to each generation at each stage. Limited commitment thus permits resource sharing without assuming that a parent can secure an enforceable promise of future support in return for childcare today.

The model links the two care stages in both directions. Earlier childcare raises the child's history-dependent guilt from declining subsequent care, while saving and bargaining determine the resources and terms with which the family enters old age. Anticipated support later in life also affects the parent's willingness to provide childcare today. Resource sharing, conditional guilt, and the protection of family resources and bequests from care expenses and nursing-home risk can sustain cooperation, without assuming that family members directly value one another's utility.

Our structural identification strategy draws on the joint pattern of care across histories and permanent-resource groups. Across parent–child permanent income groups, we target grandparental childcare, LTC provision without earlier childcare, and the additional LTC associated with earlier childcare. Matching these outcomes jointly disciplines the baseline incentives sustaining care and the strength of conditional guilt. Moments on wealth, receipt of consumption-floor transfers, and paid childcare further discipline the resources and alternatives available to each generation. The calibrated model then studies how a subsidy changes participation, resource sharing, saving, and later care. This links the observed care patterns to policy incidence: incumbent parents can be affected through both the allocation within the current childcare agreement and the support available in old age.

Counterfactual experiments establish that an intervention at one care stage can alter cooperation at the other. Making purchased childcare free, without changing LTC preferences or technology, reduces grandparental childcare by 9.05 percentage points and later informal LTC by 0.95 point. The better childcare alternative changes initial participation and carries forward through care histories, assets, and bargaining positions. Experiments that remove conditional guilt or improve the formal-LTC alternative show the reverse connection: changing the value of later family support also changes current childcare agreements. These responses imply that the consequences of childcare policy cannot be assessed from the immediate substitution between market and family care alone. They also motivate a separate welfare analysis, since less informal care need not mean that the family members who previously supplied it are better off.

We evaluate a paid-childcare voucher financed by a progressive earnings tax, with childcare spending targeted at one percent of benchmark taxable earnings. We compare this policy with an externally financed version that retains the voucher and removes policy taxes and refundable tax transfers. Household decisions and stationary distributions adjust in both cases. We distinguish lifetime welfare for young entrants from the remaining-lifetime welfare of their incumbent parents, who did not receive the subsidy when young. For the incumbent comparison, we hold the initial family-state distribution at its benchmark and allow subsequent decisions to follow the policy regime. Long-run entrant welfare additionally incorporates changes in the stationary distribution of initial resources.

The paid-only subsidy produces positive average welfare gains for young adults under both

financing arrangements, but the gains do not extend to their current parents. With progressive financing, incumbent parents experience an average consumption-equivalent welfare loss of 0.57 percent; the loss remains 0.42 percent when the subsidy is externally financed. By improving the child's alternative to grandparental childcare, the policy changes which families cooperate and how resources are divided within the relationship. Parents can consequently supply less childcare yet lose access to consumption and support associated with the benchmark agreement. The persistence of the loss without policy taxes identifies the importance of this family channel. The result is not that every parent loses, but that a policy benefiting the younger generation on average can impose a loss on the older generation even without a direct fiscal burden.

We then extend the voucher to grandparental childcare while retaining support for purchased care and the same spending target. A cooperating family receives the voucher in its shared budget, with its allocation still determined by the agreement between generations. Spreading the budget across more families reduces the amount available per recipient, but also directs support to the family care arrangement. Under progressive financing, this redesign raises long-run entrant consumption-equivalent welfare gains from 0.14 to 0.29 percent and turns the incumbent parent's average loss into a gain of 0.52 percent. Current young adults also gain more, and the ranking of the two designs persists without policy taxes. Broader eligibility both adds resources to cooperating families and helps sustain care relationships. Its advantage therefore reflects the allocation of public support as well as the dynamic response of informal care.

Our analysis contributes to several strands of the existing literature. First, it builds on work in which families provide insurance and exchange resources across generations. This literature studies longevity risk sharing, parental support during labor-market shocks, and transfers associated with saving and housing (Kotlikoff and Spivak, 1981; Kaplan, 2012; Barczyk and Kredler, 2014; Brandsaas, 2025). Strategic bequests and self-enforcing family arrangements provide foundations for intergenerational exchange without relying exclusively on altruism (Bernheim et al., 1985; Cigno, 1993, 2006). Related work studies reciprocal childcare and eldercare (Lefebvre et al., 2025). Fu et al. (2026) use variation in childcare-subsidy generosity to document responses in grandparental childcare and expected future LTC, and study cross-sector spillovers in a dynamic noncooperative model. We focus on how these care links shape self-enforcing resource allocations and the distribution of policy gains across current and future generations.

Second, our model connects to quantitative analyses of informal LTC, saving, and public and private insurance (Barczyk and Kredler, 2018; Ko, 2022; Barczyk et al., 2023; İmrohoroğlu and Zhao, 2018; Mikoshiba, 2025). Most closely, Mommaerts (2025) develops a limited-commitment model of an elderly parent and an adult child and shows that insurance covering family care can improve welfare. We incorporate the earlier childcare relationship, so that parents' resources, later bargaining positions, and children's willingness to provide LTC reflect care exchanged before

old-age needs arise. This extension makes childcare policy relevant for the welfare of incumbent parents even though they receive no paid-care voucher. It also allows us to evaluate support for grandparental childcare in terms of its effects on both generations over their remaining lives.

Third, the paper connects intergenerational care to the literature on intrahousehold bargaining and resource allocation. Evidence on child benefits shows that the identity of the recipient can affect spending within a household (Lundberg et al., 1997). Dynamic collective models show how commitment and changes in family members' alternatives shape consumption, saving, and labor supply (Mazzocco, 2007; Mazzocco et al., 2014; Voena, 2015); family allocation also matters for the design of taxation (Gayle and Shephard, 2019). We apply these insights to parents and adult children whose care needs occur at different stages of life. A childcare subsidy changes the child's position within that relationship and can redistribute its gains even when prices and policy taxes are held fixed. This provides a mechanism through which policy affects a generation outside its direct eligibility rule.

Finally, our policy analysis contributes to research on childcare subsidies, labor supply, and child development (Ho and Pavoni, 2020; Bastani et al., 2020; Chaparro et al., 2020). Cardia and Ng (2003) already show that grandparental time transfers matter for childcare policy and that subsidizing grandparents' care can be valuable. We add the connection to subsequent LTC and the welfare consequences of limited commitment. The resulting design comparison evaluates who benefits from childcare support, including the family provider, as well as how much care is purchased. It shows that recognizing informal care can improve the use of a given childcare budget while preserving average gains for the young adults the policy is intended to help.

The paper proceeds as follows. Section 3 presents the empirical evidence. Sections 4 and 5 describe the model and calibration. Section 6 examines the links between care stages, and Section 7 defines the welfare comparisons. Sections 8 and 9 study the paid-childcare subsidy and broader eligibility. Section 10 concludes.

2 Related Literature

Models of intergenerational transfers differ along two dimensions: the degree of commitment within the family and the motives for providing support. These dimensions are distinct. Altruism can motivate transfers without securing future cooperation, while mutually beneficial exchange and insurance can sustain resource sharing without altruism. This distinction helps locate our model within the literature.

Full-commitment and pooled-family approaches allocate resources without requiring each generation to renew its participation at every date and state. Dynastic altruism motivates grandparental childcare in Cardia and Ng (2003). With two-sided altruism, İmrohoroğlu and Zhao (2018) study

how LTC risk and demographic change affect family insurance and saving: parents and children pool resources and maximize a joint objective. Mikoshiba (2025) uses a related framework to evaluate Japanese LTC insurance. These models capture intergenerational resource flows while abstracting from whether each generation would continue to accept the family allocation. The insurance motive itself need not be altruistic: Kotlikoff and Spivak (1981) show how self-interested family members can benefit from pooling longevity risk.

Noncooperative models derive transfers from the equilibrium choices of separately optimizing family members. Altruism motivates support in models of transfers, coresidence, and housing (Barczyk and Kredler, 2014; Kaplan, 2012; Brandsaas, 2025), whereas strategic bequests and reciprocal care provide exchange motives (Bernheim et al., 1985; Lefebvre et al., 2025). Models of LTC and intergenerational care combine altruism with bequest or care-related incentives (Barczyk and Kredler, 2018; Barczyk et al., 2023; Ko, 2022; Fu et al., 2026). These approaches can include bargaining over current arrangements, as in Barczyk et al. (2023), while leaving future family transfers unenforceable. Thus, the absence of commitment does not imply either the absence of transfers or the absence of bargaining.

Limited-commitment models make the sustainability of cooperation an explicit restriction on the allocation. In dynamic collective models with limited commitment, both parties must prefer the agreement to their alternatives (Mazzocco, 2007). In the intergenerational setting, Cigno (1993, 2006) develops self-enforcing family constitutions, with altruism introduced in the later analysis. Mommaerts (2025) combines limited commitment, parental altruism, and child guilt in a model of LTC and insurance. This framework allows family resources to be shared while the terms of cooperation respond to changes in individual alternatives.

Our model combines limited commitment with care exchanged in opposite directions over the life cycle. Resource sharing, bequest protection, and history-dependent guilt support cooperation without assuming that either generation directly values the other's utility. Earlier childcare changes the history, assets, and bargaining position with which the family enters the LTC stage; expected later support also affects current childcare agreements. Both the existence of an agreement and its allocation of resources are endogenous. This combination is central to the policy results: improving the child's paid-care alternative can reduce the incumbent parent's welfare, while extending support to grandparental care can improve the gains available to both generations.

3 Empirical Motivation

Families exchange care in both directions over the life cycle: parents help their adult children raise children, and adult children subsequently help their aging parents. We use the Health and Retirement Study (HRS) to document how these transfers are linked within parent–child relationships

and how the link varies with family resources. Four patterns motivate the model. Earlier grandparental childcare predicts later informal long-term care (LTC), including in comparisons across children of the same parent. Childcare is more common among families with higher permanent-income, whereas informal LTC follows a different socioeconomic gradient. Finally, the association between the two transfers is weaker toward the top of the resource distribution. These patterns distinguish the prevalence of family care from the extent to which later care is associated with earlier assistance.

3.1 Data and Measurement

We combine the RAND HRS respondent and family files with information on helpers to construct one observation per parent–child pair. The construction uses survey waves from 1994 through 2022, excluding 2020, and links two non-overlapping phases. The childcare history covers healthy parents in the age-40–70 window, with an adult child ages 20–40 who has at least one child. The later-life history covers parents ages 70 and older who are partially or fully retired and report at least one limitation in activities of daily living (ADLs) or instrumental activities of daily living (IADLs). We record whether the parent ever provided childcare to the particular child's household in the first phase and whether that child or the child's spouse ever provided ADL or IADL assistance in the second. Help with money management is not added to the care outcome. Both measures describe transfers observed over a phase rather than annual care probabilities.

We use children's years of education and parents' average household retirement income as proxies for permanent resources. Children are grouped into 0–12, 13–15, and 16 or more years of schooling. Parent income combines pensions, annuities, and Social Security retirement benefits received by the respondent and spouse, expressed in 2022 dollars and averaged over the later-life observation period. It is a household total, without an equivalence-scale adjustment. Low-, middle-, and high-income parents are defined using the 50th and 74th percentiles of this measure in the base sample. We hold these cutoffs fixed across all tables. Education and retirement income summarize persistent differences in resources, but are not measures of exogenously assigned income.

The base sample contains 6,535 parent–child pairs. Our main LTC statistics use the 3,693 pairs with a nonmissing child-specific help indicator; the broader sample also includes histories coded as requiring no help, which are assigned zero in an alternative outcome. We follow the programs' exclusion of histories carrying the spouse-only-care flag.[1] Statistics are unweighted, so each observed pair receives equal weight. Appendix A provides the sample construction, variable definitions, treatment of missing values, and estimation details.

[1]The exclusion is applied to the flag stored for each parent–child history. It should not be interpreted as restricting the sample to unmarried parents or as imposing a uniform exclusion across all children of every parent. Appendix A describes its implementation.

3.2 Stylized Fact 1: Earlier Childcare Predicts Later Informal Care

Parents are more likely to receive informal LTC from a child whose household previously received grandparental childcare. Table 1 shows a later-care rate of 38.7 percent among pairs with earlier childcare, compared with 22.7 percent among pairs without it. The difference is 16.0 percentage points. The relevant pattern is the connection between the two directions of assistance within a relationship, rather than simply the overall frequency of either transfer.

Table 1. Earlier Childcare and Later Informal Care

Earlier childcare	Care receipt	Standard error	Observations
No	0.227	(0.009)	2,319
Yes	0.387	(0.013)	1,374
Difference	0.160	(0.016)	

Notes: Care receipt is the share of parent–child pairs in which the child or the child's spouse provided ADL or IADL help during the parent's later-life observation period. Earlier childcare refers to the same child's household. The sample contains 3,693 pairs with observed help status. The difference is Yes minus No. These descriptive standard errors treat pairs as independent; regression standard errors in Table 2 allow dependence within parents. Source: HRS; authors' calculations.

This raw comparison pools both within- and between-parent variation. Shared preferences, geographic proximity, and relationship quality could make both transfers more likely. To examine the association conditional on observed characteristics, we estimate linear probability models of the form

$$L_{ij} = a + \beta C_{ij} + X'_{ij}\gamma + u_{ij}, \tag{1}$$

where L_{ij} indicates later informal care from child j or the child's spouse, and C_{ij} indicates earlier childcare provided by parent i to that child's household. In the parent fixed-effects specifications, a is replaced by a_i. These specifications compare children of the same parent and absorb parent characteristics common to those children.

Table 2 shows that the positive association survives both observable controls and parent fixed effects. With controls, the pooled estimate is 13.5 percentage points. The corresponding within-parent estimate is 17.9 percentage points: among children of the same parent, those who received childcare are more likely to provide later care. This comparison removes differences across parents, although it does not eliminate child-specific differences in attachment, availability, or other determinants of both transfers.

We also report instrumental-variables estimates using the average number of children of the

Table 2. Childcare and Subsequent Informal Care: Regression Estimates

	Least squares		Parent fixed effects		Instrumental variables	
	(1)	(2)	(3)	(4)	(5)	(6)
Earlier childcare	0.160	0.135	0.209	0.179	0.593	0.113
	(0.017)	(0.017)	(0.023)	(0.025)	(0.300)	(0.220)
Constant	0.227	0.106			0.066	0.112
	(0.010)	(0.061)			(0.113)	(0.086)
Child controls	No	Yes	No	Yes	No	Yes
Parent controls	No	Yes	Absorbed	Absorbed	No	Yes
Parent fixed effects	No	No	Yes	Yes	No	No
Observations	3,693	3,324	3,229	2,817	3,693	3,324
Parent clusters	1,481	1,440	1,017	933	1,481	1,440
First-stage F					12.6	16.5

Notes: The dependent variable indicates later ADL or IADL help from the child or the child's spouse. Standard errors clustered by parent are in parentheses. Child controls include gender, college education, birth-order indicators, number of children, marital status, and full-time employment. Parent controls include gender, college education, race, household income, functional limitations, self-reported health, and marital status. Parent controls are absorbed by the fixed effects in columns (3)–(4). The fixed-effects estimator drops singleton parent groups; missing controls account for additional sample differences. Columns (5)–(6) report two-stage least squares, using the sibling-fertility instrument described in Appendix A.4; they do not include parent fixed effects. The childcare coefficients in columns (1)–(4) have $p < 0.01$; the corresponding p-values in columns (5)–(6) are 0.048 and 0.608. Source: HRS; authors' calculations.

focal child's eligible siblings during the childcare phase. The proposed mechanism is competition for grandparents' time: more grandchildren in other children's households may reduce the childcare available to the focal child. The first-stage relationship is negative, with excluded-instrument F-statistics of 12.6 and 16.5 on the corresponding estimation samples. Appendix A.4 gives the exact construction, including the restriction to siblings represented in the eligible childcare sample.

The instrumental-variables results are substantially less precise. The estimate falls from 59.3 percentage points without controls to 11.3 percentage points with controls, and the latter is not statistically distinguishable from zero. Moreover, a causal interpretation requires siblings' fertility to affect the focal child's later caregiving only through earlier childcare. This condition could fail if siblings' children also affect the allocation of eldercare across siblings or reflect shared family preferences. We therefore view the evidence as establishing a strong conditional association consistent with reciprocity, while leaving the mechanisms behind it to the structural analysis.

3.3 Stylized Fact 2: Childcare Is More Common in Higher-Resource Families

Grandparental childcare is not concentrated among families with the lowest measured resources. Panel A of Table 3 shows that childcare provision rises from 34.5 percent in the low-parent-income group to 45.2 percent in the high-income group. Panel B shows a similar distinction by child education: the share is 33.4 percent among children with at most 12 years of schooling, compared with approximately 45.5 percent in both higher-education groups.

Table 3. Grandparental Childcare by Parent Income and Child Education

Group	Share	Standard error	Observations	Sample share (%)
Panel A. Parent retirement income				
Low	0.345	(0.008)	3,265	50.0
Middle	0.430	(0.013)	1,569	24.0
High	0.452	(0.012)	1,701	26.0
All	0.393	(0.006)	6,535	100.0
Panel B. Child education				
0–12 years	0.334	(0.008)	3,224	49.8
13–15 years	0.455	(0.013)	1,559	24.1
16 or more years	0.456	(0.012)	1,686	26.1
All	0.395	(0.006)	6,469	100.0
Panel C. Parent income and child education jointly				
Parent income		0–12 years	13–15 years	16 or more years
Low		0.302	0.416	0.417
		(0.010)	(0.018)	(0.021)
Observations		1,940	721	566
Middle		0.395	0.480	0.445
		(0.018)	(0.026)	(0.023)
Observations		716	383	454
High		0.368	0.495	0.497
		(0.020)	(0.023)	(0.019)
Observations		568	455	666

Notes: Entries are unweighted shares of parent–child pairs with grandparental childcare during the childcare phase. Parent income is average household retirement income in 2022 dollars. Low, Middle, and High use cutoffs at the 50th and 74th percentiles of the 6,535-pair base sample; these cutoffs are held fixed throughout Section 3. Panels B and C exclude 66 pairs with missing child education. Standard errors in parentheses treat pairs as independent. Sample shares refer to the observations in the corresponding panel. Appendix A gives construction details. Source: HRS; authors' calculations.

The joint distribution in Panel C shows that these gradients are not solely a consequence of grouping one generation while ignoring the other. Within each parent-income group, children with schooling beyond high school are more likely to have received childcare. The parent-income gradient is also positive within the two higher-education groups, although it is not monotone among children with at most 12 years of schooling. Childcare reaches 49.7 percent in the high-parent-income, college-educated-child group, compared with 30.2 percent in the low-income, low-education group.

Grandparental childcare also represents a substantial regular care arrangement. Complementary evidence from the 2019 National Household Education Survey shows that, among children whose main regular relative caregiver is an unpaid grandparent, grandparents provide 21.39 hours of care per week and more than half of all nonparental care hours for 69.14 percent of children. Moreover, 69.77 percent use no reported household-paid center or nonrelative care. The majority-provider pattern holds across young-household income groups. Appendix B describes the sample and reports these comparisons. This evidence supports our focus on grandparental care as an important family resource: changes in the affordability of purchased alternatives can affect the value of the care relationship and the allocation of resources within it.

Economically, a time transfer depends on both the parent's ability to supply care and the child's value of receiving it. Greater parental resources may make providing care more affordable, while a higher opportunity cost of the child's time may increase the value of grandparental help. The observed gradients do not identify these channels separately. They establish that an explanation based only on low-income families being unable to purchase childcare is incomplete.

3.4 Stylized Fact 3: Informal LTC Follows a Different Resource Gradient

The positive resource gradient in childcare does not carry over to informal LTC. With parent-income groups held fixed, Panel A of Table 4 shows care receipt of 29.5, 28.8, and 25.8 percent across the low-, middle-, and high-income groups. The pattern by child education is non-monotone: care provision peaks among children with 13–15 years of schooling, at 33.9 percent, compared with 26.3 percent among those with at most 12 years and 29.6 percent among those with at least 16 years.

Panel C confirms that the middle-education group has the highest care rate within each parent-income group. High-income parents also receive less care than low-income parents in every child-education category. Thus, families that are more likely to exchange childcare are not uniformly more likely to exchange eldercare.

This difference is economically plausible because the two transfers solve different problems. Grandparental childcare relaxes the younger household's time constraint. Later-life care instead depends on the parent's care needs, the availability of formal substitutes, and the adult child's cost of supplying time. Higher parental resources can make formal alternatives more accessible,

Table 4. Informal LTC by Parent Income and Child Education

Group	Share	Standard error	Observations	Sample share (%)
Panel A. Parent retirement income				
Low	0.295	(0.010)	2,222	60.2
Middle	0.288	(0.016)	763	20.7
High	0.258	(0.016)	708	19.2
All	0.286	(0.007)	3,693	100.0
Panel B. Child education				
0–12 years	0.263	(0.010)	1,995	54.4
13–15 years	0.339	(0.016)	827	22.5
16 or more years	0.296	(0.016)	847	23.1
All	0.288	(0.007)	3,669	100.0
Panel C. Parent income and child education jointly				
Parent income		0–12 years	13–15 years	16 or more years
Low		0.270	0.360	0.317
		(0.012)	(0.022)	(0.024)
Observations		1,371	464	366
Middle		0.278	0.320	0.284
		(0.023)	(0.036)	(0.031)
Observations		371	172	218
High		0.206	0.304	0.278
		(0.025)	(0.033)	(0.028)
Observations		253	191	263

Notes: Entries are unweighted shares receiving ADL or IADL help from the child or the child's spouse, among pairs with observed help status. Parent-income groups use the same base-sample cutoffs as Table 3; parents are not re-ranked within the smaller LTC sample. Panels B and C exclude 24 pairs with missing child education, explaining their smaller total. Standard errors in parentheses treat pairs as independent. Sample shares may not sum to 100 because of rounding. Source: HRS; authors' calculations.

while higher child productivity can make informal care more costly. These forces provide possible explanations for the contrasting gradients; the descriptive tables alone do not establish their relative importance.

3.5 Stylized Fact 4: Reciprocity Is Weaker toward the Top of the Resource Distribution

To describe heterogeneity in reciprocity, we compare later-care rates between pairs with and without earlier childcare within each resource group. For group g, define

$$\Delta_g = \Pr(L_{ij} = 1 \mid C_{ij} = 1, g) - \Pr(L_{ij} = 1 \mid C_{ij} = 0, g). \tag{2}$$

This is an unadjusted difference in observed care rates, not a group-specific causal effect. It captures how strongly earlier assistance predicts later assistance within a given part of the resource distribution.

Table 5. Reciprocity by Parent Income and Child Education

	Earlier childcare: Yes		Earlier childcare: No		
Group	Care share	Observations	Care share	Observations	Difference
Panel A. Parent retirement income					
Low	0.416	750	0.233	1,472	0.183
	(0.018)		(0.011)		(0.021)
Middle	0.379	309	0.227	454	0.152
	(0.028)		(0.020)		(0.034)
High	0.327	315	0.204	393	0.123
	(0.026)		(0.020)		(0.033)
Panel B. Child education					
0–12 years	0.375	635	0.211	1,360	0.164
	(0.019)		(0.011)		(0.022)
13–15 years	0.431	364	0.266	463	0.166
	(0.026)		(0.021)		(0.033)
16 or more years	0.368	370	0.241	477	0.126
	(0.025)		(0.020)		(0.032)

Notes: Care shares refer to later ADL or IADL help from the child or the child's spouse. The difference is Yes minus No within each row. Parent-income groups use the fixed cutoffs defined in Table 3. Panel A includes 3,693 pairs; Panel B includes 3,669 pairs with observed child education. Parenthesized standard errors treat pairs as independent; those for differences are the square root of the sum of the two estimated mean variances. Differences are calculated before rounding. Source: HRS; authors' calculations.

Panel A of Table 5 shows that the reciprocity gap declines from 18.3 percentage points for low-income parents to 15.2 points for middle-income parents and 12.3 points for high-income

parents. Panel B shows a less pronounced education gradient: the gaps are similar for the two lower-education groups, at 16.4 and 16.6 points, and smaller for college-educated children, at 12.6 points. The latter group is more likely to have received childcare, but its subsequent caregiving is less strongly differentiated by whether childcare was received.

Table 6. Reciprocity within Parent-Income and Child-Education Groups

	Earlier childcare: Yes		Earlier childcare: No		
Child education	Care share	Observations	Care share	Observations	Difference
Panel A. Low parent income					
0–12 years	0.383	405	0.223	966	0.160
	(0.024)		(0.013)		(0.028)
13–15 years	0.492	189	0.269	275	0.223
	(0.036)		(0.027)		(0.045)
16 or more years	0.412	153	0.249	213	0.163
	(0.040)		(0.030)		(0.050)
Panel B. Middle parent income					
0–12 years	0.401	137	0.205	234	0.196
	(0.042)		(0.026)		(0.050)
13–15 years	0.383	81	0.264	91	0.119
	(0.054)		(0.046)		(0.071)
16 or more years	0.348	89	0.240	129	0.108
	(0.051)		(0.038)		(0.063)
Panel C. High parent income					
0–12 years	0.301	93	0.150	160	0.151
	(0.048)		(0.028)		(0.056)
13–15 years	0.351	94	0.258	97	0.093
	(0.049)		(0.045)		(0.067)
16 or more years	0.328	128	0.230	135	0.098
	(0.042)		(0.036)		(0.055)

Notes: The table uses the 3,669 pairs with observed help status and child education. Parent-income groups use the fixed base-sample cutoffs. Each difference compares pairs with and without earlier childcare within the indicated income–education cell. Standard errors and rounding follow Table 5. Source: HRS; authors' calculations.

Table 6 shows that the pattern is not a strict ranking of every cell. The largest gap, 22.3 percentage points, occurs among low-income parents with middle-education children; the gap is 19.6 points among middle-income parents with low-education children. It is 9.8 points in the

high-income, college-educated group. The cell estimates are imprecise, so these rankings should not be read as establishing statistically distinct reciprocity effects for every group.

Empirical evidence and structural interpretation. The association between earlier childcare and later LTC survives controls for observed characteristics and parent fixed effects, which compare children of the same parent. Additional specifications controlling for geographic proximity also preserve the positive association. The sibling-fertility IV estimates point in the same direction, although the controlled estimate is imprecise and a causal interpretation requires the exclusion restriction discussed above. Evidence from other settings reinforces the relevance of reciprocity. Using SHARE data from European countries, Brugiavini et al. (2013) find that earlier grandchild care predicts subsequent care receipt from adult children, including after controlling for reported attitudes toward family caregiving. Exploiting variation in U.S. childcare-subsidy generosity, Fu et al. (2026) find declines in grandparental childcare and in grandparents' expectations of receiving future LTC from their daughters, consistent with a link between care at the two stages.[2] Together, these findings motivate a structural analysis of reciprocity without equating the observed care-history gap with its causal magnitude. Our objective is to discipline the strength of conditional guilt jointly with the economic incentives sustaining family care. The model must reproduce how often parents provide childcare, how much LTC occurs without earlier childcare, and how much higher LTC is following childcare across parent-income and child-education groups. Structural identification draws on these joint restrictions, as detailed in Section 5.4, and provides a framework for assessing how childcare policy changes care and resource sharing across generations.

4 Structural Model

We develop an overlapping-generations model in which families exchange care over the life cycle. An older parent can provide childcare to an adult child, who may subsequently provide long-term care (LTC) to the parent. Cooperation allows the two generations to pool resources, but each must prefer the arrangement to non-cooperation. Limited commitment therefore links care at the two dates without making future care contractible. Earlier childcare affects later cooperation through saving, bargaining power, and history-dependent guilt. We study a stationary equilibrium at given prices and fiscal rules.

[2]The policy-based LTC outcome in Fu et al. (2026) is expected future assistance rather than realized care; its estimated response is significant at the 10 percent level.

4.1 Environment and Timing

Demographics. Each individual spends two periods as an adult child and two periods as an older parent. A parent–child relationship lasts two periods: a childcare stage, indexed by $t = 1$, and an LTC stage, indexed by $t = 2$. At the end of the second stage, the parent dies and the child forms a new relationship with their own adult child. Superscripts y and o denote the child and parent within a relationship; a prime denotes the next family cycle. Cooperation and non-cooperation in stage t are denoted by C_t and N_t. Cohort size is constant, and the two family stages have equal population mass.

Endowments and states. Individuals differ in permanent productivity, $\eta^i > 0$, and asset holdings. Productivity is fixed over an individual's life and transmitted across generations according to $\Pi_\eta(d\eta^y \mid \eta^o)$. An incoming adult child's assets are drawn from $F_A(da_1^y \mid \eta^y)$. Conditional on the child's productivity, this endowment is independent of parental wealth and is externally endowed. The parent's initial wealth is determined by saving and inheritance in the preceding family cycle.

The risk-free asset pays a gross return $1 + r$, the wage per efficiency unit is w, and labor earnings are taxed at rate τ. The parent receives Social Security income $SS(\eta^o)$. The stage-1 economic state is

$$s_1 = (a_1^y, a_1^o, \eta^y, \eta^o).$$

After stage-1 saving, a fraction $\delta \sim F_\delta$ of the parent's assets is destroyed:

$$\tilde{a}_2^o = (1 - \delta)a_2^o, \qquad 0 < \delta < 1. \tag{3}$$

The child's saving is unaffected. The parent's LTC need, $\lambda \in [0, 1)$, is drawn from $F_\lambda(d\lambda \mid \eta^o)$, which includes a mass at zero. With $h \in \{C_1, N_1\}$ denoting the childcare history, the stage-2 economic state is

$$s_2 = (a_2^y, \tilde{a}_2^o, \eta^y, \eta^o, \lambda, h).$$

The incoming bargaining weight and the current relationship shock enter the bargaining problem as additional arguments.

Information and timing. At the beginning of each stage, the pair observes a common transitory relationship, or love, shock $q_t \sim F_{q,t}$ and decides whether to cooperate. Under C_1, the parent provides childcare, and the pair jointly chooses consumption, the child's labor supply, and each individual's saving. Under N_1, a paid-childcare offer $K \sim F_K$ is realized and observed by both individuals. The parent chooses saving in anticipation of the child's saving, work, and childcare decisions, and also chooses whether to work. Thus, the stage-1 participation decision precedes the

paid-care offer.

Asset destruction and LTC need are realized before stage-2 bargaining. After stage-2 saving, nursing-home entry is realized, the parent's estate passes to the child, and the next family cycle begins. The love shocks, paid-care offer, and asset destruction are mutually independent and independent of productivities and LTC need. Conditional on parent productivity, LTC need is independent of these transitory shocks. Nursing-home risk depends on LTC need and the care arrangement. All shocks and choices relevant to the current family relationship are observed by both members.

4.2 Preferences, Care, and Bequests

Individuals discount future utility by $\beta \in (0,1)$ and have period utility

$$u(c,\ell) = \frac{(c^{\gamma}\ell^{1-\gamma})^{1-\sigma}}{1-\sigma}, \qquad 0<\gamma<1, \quad \sigma>0, \quad \sigma \neq 1. \tag{4}$$

Each individual has one unit of time. Consumption and leisure must be positive, and borrowing is not permitted.

Under C_1, the parent devotes $\kappa^{gp} \in (0,1)$ units of time to childcare and does not work. The child has no childcare time requirement. Under N_1, the child either provides care using $\kappa^{y} \in (0,1)$ units of time or purchases care at cost $K>0$. Paid care is available whenever the child's budget permits it. The parent's labor supply is $n_1^o \in \{0, \bar{n}^o\}$, where $0<\bar{n}^o<1$. In stage 2, the parent does not work. Cooperation requires the child to provide λ units of informal care; under non-cooperation, the parent purchases formal care at cost $p^{LTC}\lambda$.

The love shock enters each individual's utility only under cooperation. In the childcare stage, both individuals also receive a productivity-dependent payoff $U^{cc}(\eta^y,\eta^o)$, so cooperative flow utility is

$$u(c_1^i,\ell_1^i) + U^{cc}(\eta^y,\eta^o) + q_1, \qquad i \in \{y,o\}. \tag{5}$$

We allow U^{cc} to be positive or negative. Reciprocity enters through a guilt cost when the child declines to provide LTC:

$$G(\lambda,h,\eta^y,\eta^o) = \lambda\Gamma(h,\eta^y,\eta^o). \tag{6}$$

This term is subtracted from the child's utility under N_2. A positive Γ represents a guilt cost, the psychic cost of declining care. Its sign and its difference across childcare histories are unrestricted, allowing the strength and direction of reciprocity to be disciplined by the data.

Nursing-home risk and inheritance. Let $J \in \{0,1\}$ indicate nursing-home entry after stage-2 choices, with probability $p_{NH}^k(\lambda)$ under informal care ($k=I$) or formal care ($k=F$). Both family

members face the same probabilities. For parental terminal saving a, the estate remaining after nursing-home expenses is

$$e(a, J) = \max\{a - JE_{NH}, 0\}, \qquad E_{NH} \geq 0. \tag{7}$$

The parent derives warm-glow utility from this estate according to

$$b(a) = \phi \frac{(a + \kappa_B)^{\gamma(1-\sigma)}}{1-\sigma}, \qquad \phi \geq 0, \quad \kappa_B \geq 0, \tag{8}$$

and incurs utility loss $\bar{u}_{NH}$ upon nursing-home entry.

The child inherits the entire surviving estate under either care arrangement. Let $\Psi_1(A, \eta)$ denote the expected value of entering the next family cycle as a parent with assets A and productivity η. Thus, the child's continuation depends on own saving plus the realized inheritance. Below, $\mathbb{E}_J^k$ denotes expectation over nursing-home entry conditional on current need λ and care arrangement k. Care arrangements affect inherited wealth through saving and nursing-home risk.

4.3 Limited Commitment and Bargaining

For a candidate child bargaining weight $\theta \in (0, 1)$, let $V_t(s_t; \theta)$ denote the cooperative objective and $V_t^i(s_t; \theta)$ the value to individual i at the resulting allocation. These individual values exclude the current love shock. The outside option, $Z_t^i(s_t)$, is the value of non-cooperation; in stage 1 it is evaluated before the paid-care offer is known. Cooperation at weight θ is acceptable if

$$V_t^y(s_t; \theta) + q_t \geq Z_t^y(s_t), \qquad V_t^o(s_t; \theta) + q_t \geq Z_t^o(s_t). \tag{9}$$

Because q_t is additive, it affects participation but does not affect the allocation conditional on θ.

Let ϑ_t be the incoming weight. If both participation constraints hold at ϑ_t, the pair retains that weight. Otherwise, the weight adjusts by the smallest amount needed to satisfy both constraints. If no weight satisfies them, the pair does not cooperate. The same realized q_t applies to both individuals at every candidate weight.

Every new family relationship begins at $\vartheta_1 = \bar{\theta}$. Following C_1, the selected weight θ_1 becomes the incoming weight in stage 2; following N_1, bargaining begins at $\bar{\theta}$. Non-cooperation in stage 1 therefore leaves open the possibility of cooperation in stage 2. Let θ_t denote the selected weight when cooperation occurs. The individual value implied by bargaining is

$$\Omega_t^i(s_t, q_t; \vartheta_t) = \begin{cases} V_t^i(s_t; \theta_t) + q_t, & \text{under } C_t, \\ Z_t^i(s_t), & \text{under } N_t, \end{cases} \qquad i \in \{y, o\}. \tag{10}$$

We describe household decisions by backward induction, taking Ψ_1 as given, and then impose consistency across generations.

4.4 The Long-Term-Care Stage

Cooperation. For a given weight θ, the pair chooses $x_2 = (c_2^y, c_2^o, n_2^y, a_3^y, a_3^o)$ to solve

$$V_2(s_2;\theta) = \max_{x_2} \Big\{ \theta \Big[u(c_2^y, 1 - n_2^y - \lambda) + \beta \mathbb{E}_J^I \big[\Psi_1(a_3^y + e(a_3^o, J), \eta^y) \big] \Big] + (1-\theta) \Big[u(c_2^o, 1) + \beta \mathbb{E}_J^I \big[b(e(a_3^o, J)) - J\bar{u}_{NH} \big] \Big] \Big\}, \tag{11}$$

subject to

$$\begin{aligned} c_2^y + c_2^o + a_3^y + a_3^o = (1+r)(a_2^y + \tilde{a}_2^o) + SS(\eta^o) \\ + (1-\tau) w \eta^y n_2^y, \end{aligned} \tag{12}$$

and $c_2^i > 0$, $a_3^i \geq 0$, and $0 \leq n_2^y < 1 - \lambda$. The two bracketed expressions, evaluated at the optimum, give V_2^y and V_2^o. Hence $V_2 = \theta V_2^y + (1-\theta) V_2^o$. Conditional on the weight and economic resources, earlier childcare affects this stage through participation incentives.

Non-cooperation. The parent's resources net of formal care are

$$R_2^o = (1+r)\tilde{a}_2^o + SS(\eta^o) - p^{LTC}\lambda.$$

An external support program provides the transfer

$$T_2^o(R_2^o) = \max\{\underline{c}^o - R_2^o, 0\}, \qquad \underline{c}^o > 0. \tag{13}$$

A parent with $R_2^o < \underline{c}^o$ consumes $\underline{c}^o$ and leaves no assets. Otherwise, the parent chooses consumption and terminal saving freely subject to the budget. The outside option is

$$Z_2^o(s_2) = \max_{c_2^o, a_3^o} \left\{ u(c_2^o, 1) + \beta \mathbb{E}_J^F \left[b(e(a_3^o, J)) - J\bar{u}_{NH} \right] \right\}, \tag{14}$$

subject to

$$\begin{aligned} c_2^o + a_3^o = R_2^o + T_2^o(R_2^o), \qquad c_2^o > 0, \quad a_3^o \geq 0, \\ R_2^o < \underline{c}^o \quad \Longrightarrow \quad (c_2^o, a_3^o) = (\underline{c}^o, 0). \end{aligned} \tag{15}$$

The threshold determines eligibility for support; it is not a minimum-consumption constraint for ineligible parents. In particular, at $R_2^o = \underline{c}^o$ the parent receives no transfer and may save. Support is

available only under N_2. Denote the resulting saving policy by $a_3^{o,F}(s_2)$.

Taking the parent's saving as given, the child solves

$$Z_2^y(s_2) = \max_{c_2^y, n_2^y, a_3^y} \Big\{ u(c_2^y, 1 - n_2^y) - G(\lambda, h, \eta^y, \eta^o) \\ + \beta \mathbb{E}_J^F \big[\Psi_1(a_3^y + e(a_3^{o,F}(s_2), J), \eta^y) \big] \Big\}, \tag{16}$$

subject to

$$\begin{aligned} &c_2^y + a_3^y = (1+r)a_2^y + (1-\tau)w\eta^y n_2^y, \\ &c_2^y > 0, \qquad a_3^y \geq 0, \qquad 0 \leq n_2^y < 1. \end{aligned} \tag{17}$$

The parent's saving decision is independent of the child's choices, but affects the child's continuation through inheritance. The guilt cost affects the child's willingness to cooperate; conditional on N_2, it is additive and does not change the child's allocation.

4.5 The Childcare Stage

Continuation values. For stage-1 saving (a_2^y, a_2^o) and history h, the next economic state is

$$s_2^h(a_2^y, a_2^o) = (a_2^y, (1-\delta)a_2^o, \eta^y, \eta^o, \lambda, h).$$

Throughout this subsection, $\mathbb{E}_1$ integrates over asset destruction, LTC need conditional on η^o, and the next love shock, holding current productivities and saving fixed. Individual continuation values are given by Ω_2^i.

The joint cooperative objective uses the stage-2 value evaluated at the weight selected by subsequent bargaining:

$$\Omega_2(s_2, q_2; \vartheta) = \begin{cases} V_2(s_2; \theta_2) + q_2, & \text{under } C_2, \\ \vartheta Z_2^y(s_2) + (1-\vartheta) Z_2^o(s_2), & \text{under } N_2. \end{cases} \tag{18}$$

Here θ_2 is selected from incoming weight ϑ. Thus, the cooperative continuation is evaluated using the renegotiated weight, while non-cooperation is evaluated using the incoming weight. This recursive bargaining criterion is distinct from individual lifetime utility.

Cooperation. For a given weight θ, the pair chooses $x_1 = (c_1^y, c_1^o, n_1^y, a_2^y, a_2^o)$ to solve

$$V_1(s_1; \theta) = \max_{x_1} \Big\{ \theta u(c_1^y, 1 - n_1^y) + (1-\theta) u(c_1^o, 1 - \kappa^{gp}) + U^{cc}(\eta^y, \eta^o) \\ + \beta \mathbb{E}_1 \big[\Omega_2(s_2^{C_1}(a_2^y, a_2^o), q_2; \theta) \big] \Big\}, \tag{19}$$

subject to

$$\begin{aligned} c_1^y + c_1^o + a_2^y + a_2^o = (1+r)(a_1^y + a_1^o) + SS(\eta^o) \\ + (1-\tau) w \eta^y n_1^y, \end{aligned} \tag{20}$$

and $c_1^i > 0$, $a_2^i \geq 0$, and $0 \leq n_1^y < 1$. Although resources are pooled, asset ownership matters because only parental saving is exposed to destruction and each individual's assets determine their subsequent outside option.

At the chosen allocation, individual values are

$$\begin{aligned} V_1^y(s_1;\theta) &= u(c_1^y, 1-n_1^y) + U^{cc}(\eta^y, \eta^o) \\ &\quad + \beta \mathbb{E}_1\big[\Omega_2^y(s_2^{C_1}(a_2^y, a_2^o), q_2; \theta)\big], \end{aligned} \tag{21}$$

$$\begin{aligned} V_1^o(s_1;\theta) &= u(c_1^o, 1-\kappa^{gp}) + U^{cc}(\eta^y, \eta^o) \\ &\quad + \beta \mathbb{E}_1\big[\Omega_2^o(s_2^{C_1}(a_2^y, a_2^o), q_2; \theta)\big]. \end{aligned} \tag{22}$$

These values, augmented by q_1, determine participation. Because the joint continuation in (18) uses the subsequent bargaining weight, V_1 need not equal $\theta V_1^y + (1-\theta) V_1^o$.

Non-cooperation. After observing K, the parent and child play a sequential saving game in which the parent moves first. Let $d \in \{0, 1\}$ indicate paid childcare. For a proposed parental saving choice a_2^o, the child's best-response value is

$$\begin{aligned} W_1^y(s_1, K; a_2^o) = \max_{d, c_1^y, n_1^y, a_2^y} \Big\{ & u(c_1^y, 1 - n_1^y - (1-d)\kappa^y) \\ & + \beta \mathbb{E}_1\big[\Omega_2^y(s_2^{N_1}(a_2^y, a_2^o), q_2; \bar{\theta})\big] \Big\}, \end{aligned} \tag{23}$$

subject to

$$\begin{aligned} c_1^y + a_2^y + dK &= (1+r) a_1^y + (1-\tau) w \eta^y n_1^y, \\ c_1^y > 0, \qquad a_2^y \geq 0, \qquad & 0 \leq n_1^y < 1 - (1-d)\kappa^y. \end{aligned} \tag{24}$$

Denote the saving response by $a_2^{y,N}(s_1, K; a_2^o)$. Both self-provided and paid childcare lead to history N_1; neither changes permanent productivity.

Anticipating this response, the parent solves

$$\begin{aligned} \widehat{Z}_1^o(s_1, K) = \max_{c_1^o, a_2^o, n_1^o} \Big\{ & u(c_1^o, 1 - n_1^o) \\ & + \beta \mathbb{E}_1\big[\Omega_2^o(s_2^{N_1}(a_2^{y,N}(s_1, K; a_2^o), a_2^o), q_2; \bar{\theta})\big] \Big\}, \end{aligned} \tag{25}$$

subject to

$$\begin{aligned} c_1^o + a_2^o &= (1+r)a_1^o + SS(\eta^o) + (1-\tau)w\eta^o n_1^o, \\ c_1^o > 0, \qquad a_2^o &\geq 0, \qquad n_1^o \in \{0, \bar{n}^o\}. \end{aligned} \tag{26}$$

Let $a_2^{o,N}(s_1, K)$ be the optimal parental saving policy. The child's value at the resulting allocation is

$$\widehat{Z}_1^y(s_1, K) = W_1^y(s_1, K; a_2^{o,N}(s_1, K)).$$

Conditional on parental saving, the parent's work decision does not affect the child's response.[3]

Because participation is decided before the offer is observed, the relevant outside options are

$$Z_1^y(s_1) = \mathbb{E}_K[\widehat{Z}_1^y(s_1, K)], \tag{27}$$

$$Z_1^o(s_1) = \mathbb{E}_K[\widehat{Z}_1^o(s_1, K)]. \tag{28}$$

Both expectations incorporate the offer-contingent outcome of the saving game. Both individuals also anticipate possible cooperation in stage 2, starting from $\bar{\theta}$ and with reciprocity evaluated at history N_1.

4.6 Stationary Equilibrium

The intergenerational transition closes the model. After the parent dies, the current child enters the next family cycle as a parent with

$$a_1^{o'} = a_3^y + e(a_3^o, J), \qquad \eta^{o'} = \eta^y. \tag{29}$$

Terminal saving follows the policies associated with the realized stage-2 arrangement. The new adult child's productivity and initial assets are drawn from Π_η and F_A. Each new relationship begins at $\bar{\theta}$, with a fresh love shock and no inherited care history. The continuation value is therefore

$$\Psi_1(A, \eta) = \mathbb{E}_{\eta', a, q|\eta}\left[\Omega_1^o\big((a, A, \eta', \eta), q; \bar{\theta}\big)\right], \tag{30}$$

where $\eta' \sim \Pi_\eta(\cdot \mid \eta)$, $a \sim F_A(\cdot \mid \eta')$, and $q \sim F_{q,1}$. This is the current child's own continuation value in the role of parent. Uncertainty about paid childcare is already incorporated in the next relationship's outside option.

[3]Exact saving ties are resolved in favor of lower saving for each individual. At a given saving pair, ties favor self-provided childcare and no parental work. The child's best response is selected independently of the parent's payoff.

Definition (Stationary equilibrium). Take as given preferences, the prices w, r, and p^{LTC}, the fiscal rules τ, SS, and $\underline{c}^o$, and the exogenous endowment and shock processes. A stationary equilibrium consists of time-invariant consumption, labor supply, saving, and care policies; cooperation and bargaining rules; value functions V_t, V_t^i, Z_t^i, Ω_t^i, and Ψ_1; and distributions μ_1 and μ_2, such that:

1. *Household decisions.* For each state and candidate bargaining weight, cooperative allocations solve (11) and (19), subject to the corresponding budget, time, and borrowing constraints. Non-cooperative stage-2 policies solve the parent's and child's problems. Non-cooperative stage-1 policies constitute a parent-leader Stackelberg equilibrium: the child chooses a best response to parental saving, and the parent optimizes anticipating that response.

2. *Participation and bargaining.* Cooperation occurs whenever some weight satisfies both participation constraints in (9). The incoming weight is retained whenever feasible and otherwise adjusts according to the minimum-change rule. Individual values evaluate the resulting allocations and continuation outcomes. New relationships and stage-2 bargaining after N_1 begin at $\bar{\theta}$.

3. *Consistency of continuation values.* The function Ψ_1 used in stage-2 decisions satisfies (30), with inherited wealth determined by (29). All expectations use the same policies, bargaining rules, and shock probabilities that generate realized outcomes.

4. *Stationarity of the population distribution.* The distribution μ_1 describes stage-1 economic states before the love shock, and μ_2 describes stage-2 economic states and incoming bargaining weights before the stage-2 love shock. Both have unit mass. Let T_1 be the population transition generated by stage-1 bargaining, offer-contingent decisions, asset destruction, and LTC need. Let T_2 be the transition generated by stage-2 bargaining and saving, nursing-home entry, inheritance, and the arrival of the next adult child. Then

$$\mu_2 = T_1(\mu_1), \qquad \mu_1 = T_2(\mu_2). \tag{31}$$

   These transitions retain the dependence between parental wealth and productivity and, within a family cycle, between care histories, assets, and bargaining weights.

Aggregate quantities are obtained by integrating the policy functions over the two stage distributions, with equal weights on the two stages. Prices and baseline fiscal rules are specified externally; the fiscal adjustment used in the policy experiments is described in Section 8.

5 Calibration

The calibration has two steps, following the approach of Gourinchas and Parker (2002). The first specifies preferences, prices, institutional parameters, and the distributions of productivity, initial wealth, and care needs. The second jointly disciplines the remaining parameters using care arrangements across parent–child productivity groups, wealth holdings, receipt of means-tested support, and paid-childcare use and expenditure.

5.1 Units and Predetermined Inputs

Period length and normalization. One model period represents fifteen years. We normalize the wage per efficiency unit and the average hourly wage to $w = \bar{w} = 1$, and each individual's time endowment to $\bar{\ell} = 1$. Let $\bar{y}$ denote average earnings over a model period. Monetary costs and transfers that are tied to aggregate earnings are expressed as shares of $\bar{y}$; hourly care costs are expressed relative to $\bar{w}$. The discount factor is $\beta = 0.860$, and the net risk-free return is $r = 3\%$ per model period.

The Social Security schedule is

$$SS(\eta^o) = 40\% \, \bar{y} \, \eta^o.$$

Since mean productivity is normalized to one, the pension of a parent with mean productivity equals 40% of average earnings.

Preferences and care costs. We set the consumption weight in the consumption–leisure composite to $\gamma = 0.550$ and the curvature parameter to $\sigma = 2.000$. Holding leisure fixed, the implied curvature with respect to consumption is $1 - \gamma(1 - \sigma) = 1.550$. The consumption-support threshold is written as $\underline{c}^o = \xi_c \bar{y}$, where ξ_c is determined in the second step.

The price of formal LTC is $p^{LTC} = 90\% \, \bar{w}$. Nursing-home entry can reduce the parent's terminal estate by up to $E_{NH} = 20\% \, \bar{y}$. Grandparental and self-provided childcare each require 18.5% of the caregiver's time endowment. Under non-cooperation, a parent who works supplies $\bar{n}^o = 15\% \, \bar{\ell}$; participation in the labor market remains endogenous. New family relationships, and stage-2 bargaining following N_1, begin at the child weight $\bar{\theta} = 0.500$. Table 7 summarizes the predetermined inputs.

Table 7. Predetermined calibration inputs

Parameter	Description	Value
$w = \bar{w}$	Wage normalizations	1
$\bar{\ell}$	Time endowment	1
$\bar{y}$	Average earnings over one model period	normalization
β	Discount factor per model period	0.860
r	Net return per model period	3%
γ	Consumption weight in utility	0.550
σ	Curvature of the utility composite	2.000
ρ	Intergenerational productivity persistence	0.392
σ_ϵ	Log-productivity innovation SD	0.344
σ_q	Common love-shock SD	1
$\kappa^{gp} = \kappa^y$	Childcare time requirement	$18.5\%\,\bar{\ell}$
$\bar{n}^o$	Parent hours conditional on work	$15\%\,\bar{\ell}$
$\bar{\theta}$	Initial child bargaining weight	0.500
p^{LTC}	Formal LTC price	$90\%\,\bar{w}$
$SS(1)$	Pension at mean productivity	$40\%\,\bar{y}$
E_{NH}	Maximum nursing-home asset expense	$20\%\,\bar{y}$

A model period is fifteen years. Costs and transfers shown as shares of $\bar{y}$, $\bar{w}$, or $\bar{\ell}$ retain those normalizations throughout the paper. Monetary parameters without an explicit normalization are in model units. The fixed bequest parameters are reported in Table 9.

Productivity and initial wealth. Permanent productivity follows an intergenerational process in logs:

$$x^y = \rho x^o + \epsilon, \qquad \epsilon \sim N(0, \sigma_\epsilon^2), \qquad \rho = 0.392, \quad \sigma_\epsilon = 0.344. \tag{32}$$

We choose $\rho = 0.392$ so that the three-state discretization generates a parent–child permanent-productivity correlation of approximately 0.340. As its empirical counterpart, we use the national parent–child income rank–rank slope of 0.34 reported by Chetty et al. (2014). This is an empirical analogue for intergenerational persistence rather than an equality of estimands: the model statistic is a Pearson correlation in productivity levels, whereas the empirical statistic is a rank–rank slope in income. We approximate this process with three ordered types and normalize productivity to have stationary mean one. If x_j is a log-productivity type and $\pi_{\eta,j}$ its stationary probability, then

$$\eta_j = \frac{\exp(x_j)}{\sum_k \pi_{\eta,k} \exp(x_k)}.$$

The transition probability $\Pi_\eta(k, j)$ gives the probability of child type k conditional on parent type j, and the stationary probabilities satisfy $\pi_\eta = \Pi_\eta \pi_\eta$.

Mean initial assets of incoming adult children are zero, 15% of $\bar{y}$, and 30% of $\bar{y}$ for the low-

, middle-, and high-productivity types, respectively. These endowments determine F_A and are implemented on the asset grid to preserve each conditional mean. Parental wealth is determined endogenously by saving and inheritance through the stationary equilibrium.

5.2 Distributional and Preference Specifications

Relationship shocks and childcare preferences. The common love shock is normally distributed,

$$q_t \sim N(\mu_{q,t}, \sigma_q^2), \qquad \sigma_q = 1,$$

independently across stages and family relationships. The two means, $\mu_{q,1}$ and $\mu_{q,2}$, enter the second-stage calibration. The productivity-dependent childcare payoff in (5) is

$$U^{cc}(\eta^y, \eta^o) = \frac{\chi_{cc}}{1-\sigma}\left[(\eta^y)^{\omega_{cc}}(\eta^o)^{1-\omega_{cc}}\right]^{1-\sigma}. \tag{33}$$

The parameter ω_{cc} determines the relative importance of child productivity, and χ_{cc} determines the scale of the payoff. With $\sigma = 2.0$ and $\chi_{cc} > 0$, the payoff is negative and increases toward zero as the productivity composite rises. This specification allows the utility consequences of grandparental childcare to vary systematically across families.

Reciprocity. Let η_L and η_H denote the lowest and highest productivity levels. We summarize productivity by the score

$$\zeta(\eta) = 2\frac{\log \eta_H - \log \eta}{\log \eta_H - \log \eta_L} - 1, \qquad \zeta_y = \zeta(\eta^y), \quad \zeta_o = \zeta(\eta^o).$$

The low-, middle-, and high-productivity types have scores 1, 0, and -1. The guilt coefficient in (6) is parameterized as

$$\begin{aligned}\Gamma(h, \eta^y, \eta^o) = 100\Big[&\alpha_N + \beta_y\zeta_y + \beta_o\zeta_o \\ &+ \mathbf{1}\{h = C_1\}\big(\Delta_\alpha + \Delta_y\zeta_y + \Delta_o\zeta_o\big)\Big].\end{aligned} \tag{34}$$

The first three coefficients determine the level and productivity gradients after N_1. The remaining coefficients allow earlier childcare to change both the level and the gradients. All six coefficients are unrestricted in sign, allowing the data to discipline the magnitude and direction of reciprocity. The factor of 100 expresses the coefficients per percentage point of LTC need.

Paid childcare and asset loss. Paid-childcare offers are lognormally distributed:

$$\log K \sim N\left(\log \mu_{cc} - \frac{\sigma_{cc}^2}{2},\, \sigma_{cc}^2\right). \tag{35}$$

Here $\mu_{cc} = \mathbb{E}[K]$ is the arithmetic mean offer, rather than the mean of log offers. The mean μ_{cc} and the standard deviation of log offers, σ_{cc}, are jointly calibrated to the share of families outside unpaid relative care choosing paid care and their average childcare expenditure relative to average earnings. Both childcare costs and earnings cover the full fifteen-year period, from birth to age 15. Selection into paid care makes the unconditional mean offer distinct from average expenditure among purchasers.

The fraction of parental saving destroyed before stage 2 follows a truncated lognormal distribution:

$$\log \delta \sim N(\mu_A, \sigma_A^2) \quad \text{conditional on } 0 < \delta < 1. \tag{36}$$

Both paid-care offer parameters, the two asset-loss parameters, and the nursing-home utility loss enter the second-stage calibration.

5.3 Long-Term Care and Nursing-Home Risk

An auxiliary model fitted to the HRS summarizes late-life care histories in terms of care intensity, caregiver identity, institutional residence, and mortality. These histories inform the care inputs used in the structural model. The auxiliary care and residence categories are distinct from the structural model's four LTC-intensity states.

Auxiliary care model. Let S_ν^A index five states at survey wave ν: community residence without reported care, community residence with care, nursing-home residence, death recorded in a nursing home, and other recorded death. Let M_ν indicate partnership, and let R_ν identify the primary caregiver as a child, a nonfamily helper, or a spouse. The primary category is the one supplying the largest number of reported care hours; its hours are denoted by H_ν.

We estimate a multinomial logit for S_ν^A. Regressors include the two lagged states, lagged age and age squared, and interactions of the first-lag state with lagged age, care hours and hours squared, partnership, and caregiver category. Partnership is modeled with a binary logit, and caregiver identity among community care recipients with a multinomial logit. These equations use corresponding first-lag covariates, with current partnership also entering the caregiver equation.

The hours equation is estimated among community care recipients. Regressors include lagged state, lagged age and its square, state interactions with lagged age and the first two powers of lagged log hours, and current partnership. We set the transformed hours variable to zero when reported

hours are zero. The fitted coefficients and residual variance determine the conditional distribution of simulated hours.

Simulated care histories. We draw 500,000 initial records with replacement from complete initial-wave observations, preserving the joint distribution of age, care state, partnership, hours, and caregiver identity. Each path contains $T_A = 12$ observations, with age increasing by two years between observations. Death is absorbing. At subsequent observations, the simulation draws the care state, partnership, hours, and caregiver category in that order.[4]

For $\nu \geq 2$, simulated monthly hours are

$$H_\nu = \mathbf{1}\{S_\nu^A = 2\} \exp\left[\min\left\{(X_\nu^H)' b^H + \widehat{\sigma}_H \epsilon_\nu^H,\ 7.705\right\}\right], \qquad \epsilon_\nu^H \sim N(0,1), \tag{37}$$

where X_ν^H contains the hours-equation regressors and $\widehat{\sigma}_H$ is its residual root mean squared error. Innovations are independent across individuals and observations. Initial hours are taken from the sampled record. Subsequent hours are zero outside the community care state.

Care intensity and state probabilities. Nonspousal care includes hours supplied by children and nonfamily helpers. The intensity associated with a simulated path is

$$F_{\text{nsp}} = \frac{\displaystyle\sum_{\nu=1}^{T_A} H_\nu \mathbf{1}\{S_\nu^A = 2,\ R_\nu \in \{1,2\}\}}{480 B_A}, \qquad T_A = 12, \quad B_A = 8. \tag{38}$$

The monthly time endowment is 480 hours. Each observation represents a 24-month block, so the common factor of 24 cancels from the numerator and denominator. This construction sums twelve biennial care blocks and normalizes by a sixteen-year time endowment. Spousal care and nursing-home hours are excluded.

Zero nonspousal care defines the zero-intensity group. Positive values are divided at the 70th and 90th percentiles, and within-group means determine the low-, medium-, and high-intensity representatives. Let p_A be the simulated share of paths with community care or live nursing-home residence at any observation. The assigned probabilities are

$$(p_0^A, p_L^A, p_M^A, p_H^A) = (1 - p_A,\ 70\%\, p_A,\ 20\%\, p_A,\ 10\%\, p_A). \tag{39}$$

The incidence measure includes spouse-primary community care and live nursing-home stays. Consequently, the population used to assign state probabilities differs from the population with

[4]The simulation uses the first-lag terms from the fitted state equation and omits its second-lag terms without reestimating the retained coefficients. Partnership is held at its previous value after death.

positive nonspousal hours used to construct the intensity bins.

Nursing-home risk. For each positive-intensity group, the simulation distinguishes paths in which children provide at least as many total hours as nonfamily helpers from paths dominated by nonfamily care. It reports both the probability of any nursing-home observation and the probability of death recorded in a nursing home. The zero-intensity group pools paths without nonspousal care. These conditional frequencies describe associations between care histories and institutional outcomes.

Table 8 reports the structural care inputs: the time requirement in each LTC state, its probability conditional on parent productivity, and the nursing-home probability under each care arrangement. The unconditional distribution of LTC need averages the conditional columns using the stationary distribution of parent productivity. Household decisions and model moments retain the full conditional distributions.

Table 8. LTC needs and nursing-home risk

		LTC-state probability			NH probability	
LTC state	Time	Low	Middle	High	Informal	Formal
None	0.0	30.3	37.8	41.1	35.9	35.9
Low	1.8	48.8	43.5	41.2	40.6	47.5
Medium	11.0	13.9	12.4	11.8	37.3	45.8
High	43.0	7.0	6.2	5.9	36.5	43.9

All entries are percentages. Time is measured relative to the model-period time endowment. Low, middle, and high columns refer to parent productivity. LTC-state probabilities are normalized to sum to one within each productivity group before use. Nursing-home probabilities condition on LTC need and care arrangement and are not allowed to vary with either family member's productivity. The "None" row denotes a zero model LTC time requirement; it does not rule out nursing-home entry through the separately calibrated transition. Displayed values are rounded.

Nursing-home risk is identical across care arrangements when LTC need is zero and lower under informal care in each positive-need state. These differences are imposed as calibrated inputs; they are not separately identified causal effects of care provision.

5.4 Minimum-Distance Calibration

Parameters and targets. Let $\xi_c = \underline{c}^o/\bar{y}$ denote the consumption-support threshold relative to average earnings. The second-stage specification determines sixteen parameters jointly:

$$\Theta = (\mu_{q,1}, \mu_{q,2}, \omega_{cc}, \chi_{cc}, \alpha_N, \beta_y, \beta_o, \Delta_\alpha, \Delta_y, \Delta_o, \bar{u}_{NH}, \mu_A, \sigma_A, \mu_{cc}, \sigma_{cc}, \xi_c).$$

The remaining inputs, including the return, parental hours conditional on work, and bequest parameters, are held fixed.

Let D_{jk} denote a cell with parent productivity j and child productivity k, each ordered from low to high. Within each of the nine cells, the calibration targets

$$m_{1,jk} = \Pr(C_1 \mid D_{jk}), \tag{40}$$

$$m_{2,jk} = \Pr(C_2 \mid N_1, \lambda > 0, D_{jk}), \tag{41}$$

$$m_{3,jk} \equiv \Delta_{jk}^{LTC} = \Pr(C_2 \mid C_1, \lambda > 0, D_{jk}) - \Pr(C_2 \mid N_1, \lambda > 0, D_{jk}). \tag{42}$$

Childcare rates are measured over all families. The LTC moments condition on positive care need and the indicated childcare history. The history gap is a difference in conditional care frequencies; it incorporates selection into earlier cooperation as well as history dependence in preferences and bargaining. The LTC rate following C_1 equals $m_{2,jk} + m_{3,jk}$ and therefore supplies no additional independent target.

These three groups provide 27 care moments. We add two wealth moments at the beginning of stage 2:

$$\begin{aligned} m_A &= \frac{\mathbb{E}[\tilde{a}_2^o]}{\mathbb{E}[a_2^y]} = 1.36, \\ m_A^{HH} &= \frac{\mathbb{E}[\tilde{a}_2^o \mid \eta^y = \eta_H, \eta^o = \eta_H]}{\mathbb{E}[a_2^y \mid \eta^y = \eta_H, \eta^o = \eta_H]} = 1.49. \end{aligned} \tag{43}$$

Parental assets are measured after asset destruction, consistent with the timing in Section 4. These are ratios of mean assets, both unconditionally and within the high-productivity parent–child group.

Another moment matches the empirical share of older adults receiving Supplemental Security Income (SSI) to the model share receiving a positive consumption-support transfer:

$$m_{SSI} = \Pr(T_2^o > 0) = 4.5\%. \tag{44}$$

The denominator is the full stage-2 parent population. A positive transfer requires non-cooperation and own resources strictly below $\underline{c}^o$. Finally, two moments discipline paid childcare. The first is the share of families choosing paid care within the non-cooperative stage-1 regime:

$$m_{cc} = \Pr(d = 1 \mid N_1) = 30\%, \tag{45}$$

where $d = 1$ denotes the adult child's choice of paid care. The denominator includes both paid-care users and families providing their own childcare. This is the paid-care margin within the non-cooperative alternative valued by Z_1, not the paid-care share among all families. The second is

average expenditure conditional on choosing paid care, normalized by average earnings:

$$m_K = \frac{\mathbb{E}[K \mid d = 1, N_1]}{\bar{y}} = 19\%. \tag{46}$$

Here K is the effective family childcare cost over the same fifteen-year period as $\bar{y}$. The revised specification therefore uses 32 moments to discipline sixteen parameters. The empirical construction and interpretation of the targets are discussed alongside the fit below.

Identification. The 27 care moments discipline three distinct features of family behavior: how often parents provide childcare, how often children provide LTC without earlier childcare, and how much more likely LTC is following childcare. Matching these outcomes jointly is more informative than matching the history gap alone. Childcare participation restricts the gains from the initial agreement and the selection of families into the two histories. LTC after N_1 restricts the incentives sustaining care without prior childcare, including baseline guilt, resource sharing, and protection of the estate. The additional care following C_1 then disciplines conditional guilt together with the differences in assets and bargaining positions generated by earlier cooperation. The observed history gap is therefore not equated directly with a preference parameter.

The substantial provision of LTC after N_1 adds a particularly useful restriction. The model must sustain care among families that did not previously exchange childcare while generating a higher care rate among those that did. Raising a general willingness to cooperate affects care after both histories; the three Δ coefficients in equation (34) directly change the child's LTC-stage disagreement payoff only after C_1. Matching both conditional care rates therefore helps distinguish baseline guilt from its history-dependent increment within the model. These effects are nevertheless interconnected: anticipated LTC influences current childcare, and the guilt parameters can change saving and selection into either history. Identification consequently draws on the joint response of all three sets of moments.

The nine parent–child productivity cells provide further restrictions. Parent resources affect the value of formal-care alternatives and the gains from sharing consumption, while child productivity also changes the opportunity cost of care. The model must reproduce care levels and history gaps across these economic environments using the same parsimonious preference specification: three coefficients describe baseline guilt and three describe its history-dependent increment, rather than separate parameters for every cell. The parent and child gradients help discipline their respective productivity slopes. Wealth, consumption-floor, and paid-childcare moments further restrict saving, the cost of market care, and the outside options entering both generations' participation constraints.

By matching care provision across histories and productivity groups, the calibration disciplines the incentives underlying families' responses to childcare policy. A subsidy restricted to purchased

childcare improves the child's alternative to grandparental care and can change both participation and the terms of cooperation. Because childcare agreements also allocate consumption and saving, the parent can be affected at the childcare stage through resource sharing and subsequently through changes in assets, bargaining positions, and care history. The policy experiments trace these responses jointly to assess how the gains from childcare support are distributed across generations and whether extending eligibility to grandparental care can improve welfare within the same spending target.

Criterion. Holding the first-step inputs fixed, let $m_j(\Theta)$ denote stationary model moment j and m_j^d the corresponding calibration target. The minimum-distance criterion is

$$Q(\Theta) = \sum_{j=1}^{32} \left[m_j(\Theta) - m_j^d\right]^2. \tag{47}$$

This is an identity-weighted criterion. Care rates, history gaps, SSI receipt, and paid-childcare use enter as shares; wealth moments enter as asset ratios, and childcare expenditure as a share of average earnings. The units therefore determine the relative scale of the moment deviations.

Using fixed weights avoids estimating a weighting matrix from the same sample as the target moments. Correlation between sampling error in moments and estimated weights can worsen finite-sample performance, as shown by Altonji and Segal (1996) for covariance-structure models. This consideration motivates the weighting choice; it does not imply that equal weights are asymptotically efficient.

5.5 Model Fit

The current benchmark captures the main qualitative patterns of intergenerational care, although fit varies across productivity cells. We first discuss the 27 behavioral moments and then compare the additional wealth, consumption-support, and paid-childcare moments. All model values in this subsection use the benchmark with the 21-node conditional-mean quadrature for paid-childcare offers.

Behavioral moments. The minimum-distance specification includes nine childcare rates, nine LTC cooperation rates following N_1, and nine reciprocity gaps. Panel (c) of Figure 1 displays the nine childcare cells, and panel (c) of Figure 3 displays the nine reciprocity cells. Panel (c) of Figure 2 displays the nine directly targeted rates $\Pr(C_2 \mid N_1, \lambda > 0, D_{jk})$. Its one-way panels instead pool the C_1 and N_1 histories and provide complementary descriptive comparisons.

Table 9. Benchmark parameter values

Parameter	Description	Value
Parameters in the second-step estimation		
$\mu_{q,1}$	Mean love shock, childcare stage	0.320
$\mu_{q,2}$	Mean love shock, LTC stage	-0.689
ω_{cc}	Child productivity weight in childcare payoff	0.783
χ_{cc}	Scale of childcare payoff	0.463
α_N	Baseline guilt: intercept after N_1	0.034
β_y	Baseline guilt: child-productivity slope	-0.011
β_o	Baseline guilt: parent-productivity slope	0.034
Δ_α	History-dependent guilt: intercept	0.075
Δ_y	History-dependent guilt: child slope	0.013
Δ_o	History-dependent guilt: parent slope	0.014
$\bar{u}_{NH}$	Nursing-home utility loss	0.001
μ_A	Location of log asset loss before truncation	-0.350
σ_A	SD of log asset loss before truncation	0.400
μ_{cc}	Mean paid-childcare offer	0.600
σ_{cc}	SD of log paid-childcare offers	1.182
ξ_c	Consumption threshold relative to $\bar{y}$	2.500%
Predetermined bequest parameters		
κ_B	Warm-glow bequest shifter	0.012
ϕ	Warm-glow bequest strength	0.132

Values describe the parameter vector used for the benchmark model-fit comparisons and are rounded to at most three digits after the decimal point. The monetary parameters μ_{cc} and κ_B are in model units. The consumption threshold is $\underline{c}^o = \xi_c \bar{y}$.

The model reproduces the broad increase in childcare cooperation with family productivity, but the fit is uneven across individual cells. For high-productivity parent–child pairs, childcare cooperation is 49.418% in the model and 49.700% in the data. For low-productivity pairs, the corresponding values are 34.441% and 30.206%, so the model overstates cooperation in that cell. Other mixed-productivity cells account for additional discrepancies.

The pooled LTC comparisons are close in the aggregate but miss some of the cross-cell heterogeneity. Conditional on positive LTC need, the aggregate cooperation rate is 27.611% in the model and 28.800% in the data. For example, among children in the middle productivity group, the pooled rate is 26.779% in the model and 33.900% in the data. These pooled comparisons should not be interpreted as the direct fit of the conditional-after-N_1 moments. In the directly targeted cells, for example, $C_2 \mid N_1$ is 21.076% in the model and 22.300% in the data for the low–low pair, and 19.084% versus 23.000% for the high–high pair.

The model generates a positive reciprocity gap in all nine productivity cells and reproduces

its broad decline with productivity. For low-productivity parent–child pairs, the gap is 18.298 percentage points in the model and 16.015 points in the data. For high-productivity pairs, it is 12.235 points in the model and 9.850 points in the data. Thus, the model captures the direction of the history gradient, while retaining nontrivial cell-level discrepancies.

Wealth, consumption support, and paid childcare. Table 10 reports the remaining comparisons. The wealth targets are calculated from the 2022 Survey of Consumer Finances (SCF), using the ratio of mean family net worth at ages 75–90 to that at ages 45–60. This ratio is 1.360 overall and 1.490 among college graduates. We use the latter as the empirical counterpart of the model's high-productivity parent–child ratio. These are comparisons across age groups rather than observations on linked parents and adult children. Appendix C.1 describes the sample, weighting, and construction. The corresponding model ratios are 1.362 and 1.653.

The model consumption-support receipt rate is 4.478%, compared with the maintained target of 4.500%. The empirical counterpart is the SSI receipt rate among U.S. adults aged 65 or older in December 2014 (Social Security Administration, 2019, Table 3.C5). Its model counterpart is the share of the full stage-2 parent population receiving a positive consumption-floor transfer.

For paid childcare, the empirical comparisons documented in the appendix and the maintained targets in the model files are not identical. The empirical construction gives an approximate paid-care use rate of 31.500% among families without primary unpaid relative care, whereas the maintained target is 30.000%. The model rate conditional on N_1 is 32.627%. The empirical expenditure comparison is 13.200% of average earnings, whereas the maintained target used with the paid-offer distribution is 19.000%; the model value is 18.787%. Table 10 reports both columns rather than relabeling one as the other. The empirical constructions are described in Appendices C.2 and C.3.

Table 10. Model fit: additional calibration targets

Moment	Data comparison	Model
Parent-to-child asset ratio, all families	1.360	1.362
Parent-to-child asset ratio, high-productivity pairs	1.490	1.653
Consumption-floor share (%)	4.500	4.478
Paid-childcare use conditional on N_1 (%)	31.500	32.627
Mean paid-care cost among users (% of $\bar{y}$)	13.200	18.787

The wealth comparisons use SCF mean net worth for ages 75–90 relative to ages 45–60, overall and among college graduates; see Appendix C.1. Model asset ratios use stage-2 parental assets after asset destruction relative to child assets, overall and within high-productivity pairs. The SSI comparison is the December 2014 receipt rate among adults aged 65 or older (Social Security Administration, 2019, Table 3.C5). Paid-care use conditions on N_1 in the model, and mean paid-care cost additionally conditions on $d = 1$. The distinction between the empirical comparisons and the maintained paid-care targets is discussed in the text.

Aggregate implications. Table 11 reports aggregate care arrangements, which receive no separate weight in the criterion. Childcare cooperation is 41.492% in the model, compared with 39.500% in the data. Among families with positive LTC need, informal care is provided to 27.611% of parents in the model and 28.800% in the data. The share cooperating at both dates is 14.988%, compared with 14.500% in the data. The aggregate reciprocity gap is 14.823 percentage points in the model and 16.067 points in the data.

Panels (a) and (b) of Figures 1–3 show patterns by parent and child productivity separately. These comparisons are related to the underlying cell moments but are not separately targeted. In Figure 2, the one-way panels pool childcare histories, while panel (c) displays the nine conditional-after-N_1 LTC targets directly.

Table 11. Model fit: aggregate care arrangements

Outcome	Data	Model	Model − Data
All families			
Childcare cooperation, $\Pr(C_1)$	39.500	41.492	1.992
Conditional on positive LTC need			
Cooperation at both dates, C_1, C_2	14.500	14.988	0.488
Childcare cooperation only, C_1, N_2	22.800	26.282	3.482
LTC cooperation only, N_1, C_2	14.300	12.623	-1.677
Non-cooperation at both dates, N_1, N_2	48.400	46.107	-2.293
LTC cooperation, $\Pr(C_2)$	28.800	27.611	-1.189
LTC cooperation after C_1, $\Pr(C_2 \mid C_1)$	38.874	36.316	-2.558
LTC cooperation after N_1, $\Pr(C_2 \mid N_1)$	22.807	21.493	-1.314
Reciprocity gap	16.067	14.823	-1.244

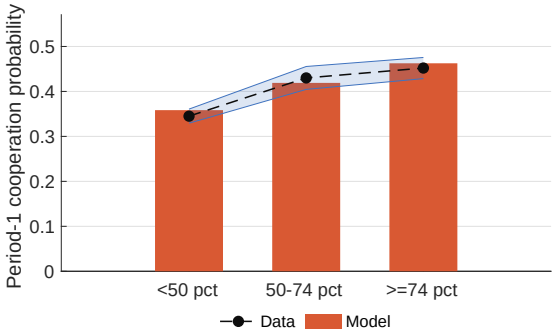
Care rates and joint shares are in percent. Reciprocity gaps and model–data differences are in percentage points. Every row in the second panel conditions on $\lambda > 0$; the conditional LTC rates additionally condition on the indicated childcare history. Conditional rates are calculated from the joint shares before rounding. The data are constructed as in Section 3. These aggregate moments receive no separate weight in the 32-moment calibration criterion.

6 Dynamic links between childcare and long-term care

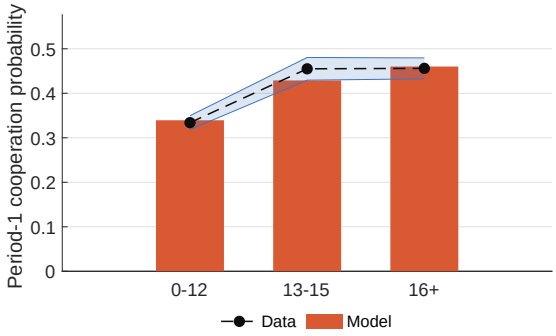
Childcare and long-term care are linked by the same family's decisions about participation, resource sharing, and saving. Grandparental childcare changes the history, assets, and bargaining position with which the family enters the LTC stage. In turn, both generations anticipate that stage

Figure 1. Childcare cooperation by parent and child productivity

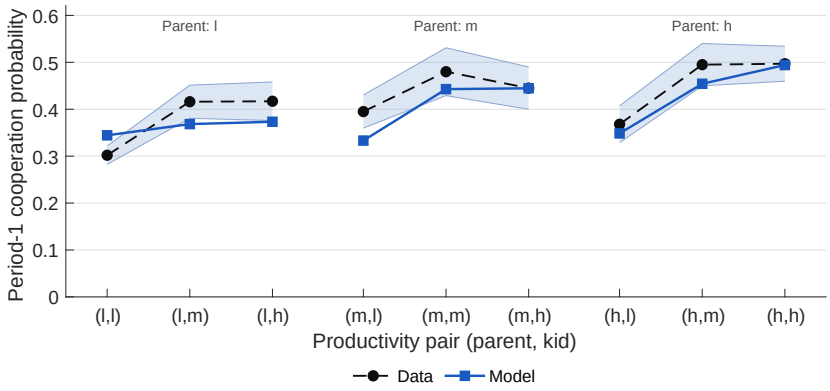
(a) Parent productivity



(b) Child productivity



(c) Joint parent–child productivity



Childcare cooperation is measured over all families. Panel (c) displays the nine directly targeted productivity cells; panels (a) and (b) are not separately targeted. Empirical construction follows Section 3; the shaded bands in panel (c) show 95% confidence intervals based on the reported unclustered standard errors.

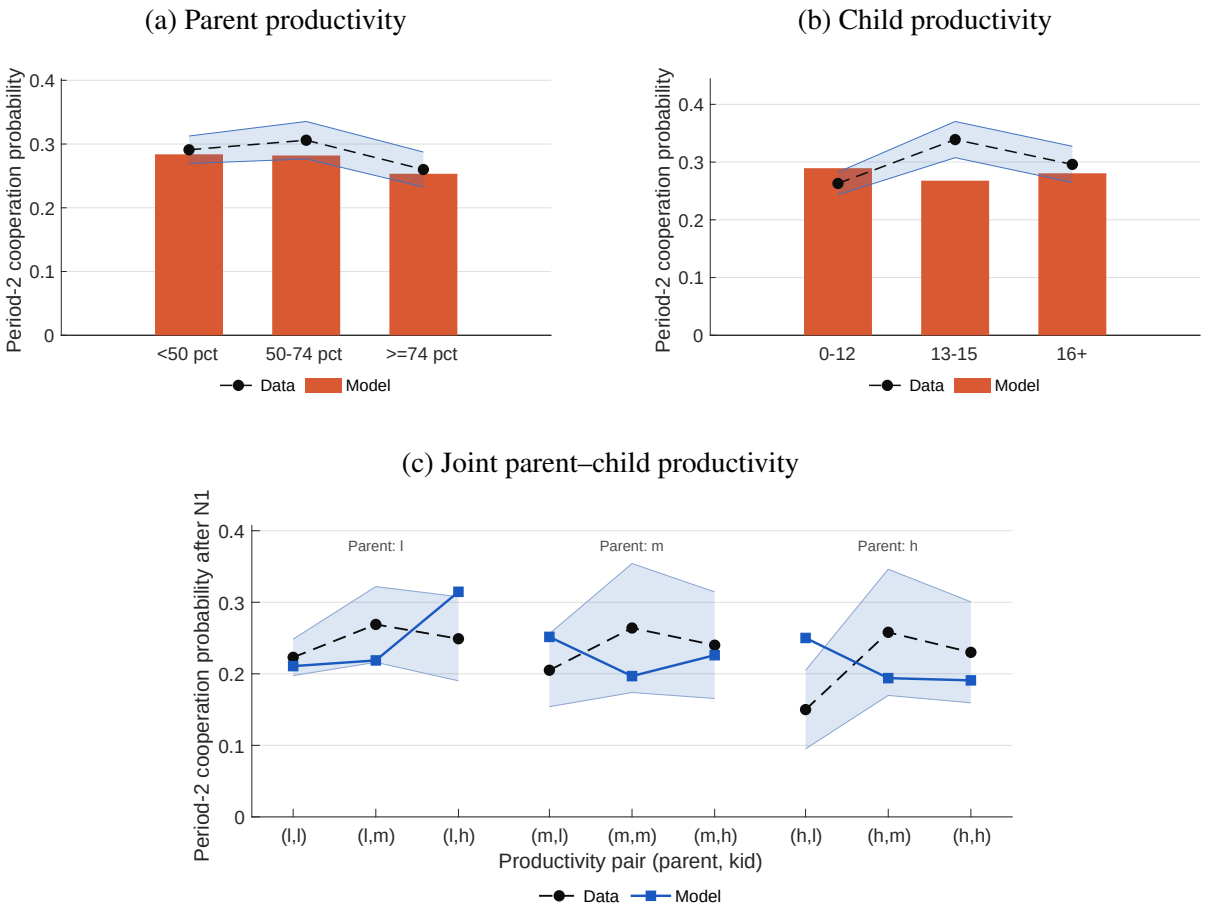
when deciding whether to cooperate today. Limited commitment makes these connections important: neither generation can require the other to preserve an agreement once its outside option changes. Earlier childcare can encourage subsequent care, but it cannot secure an enforceable promise of support.

We examine these links using four counterfactuals. Two change the incentives to cooperate when childcare is needed; two change the relative value of informal and formal LTC. Each experiment resolves household decisions and the stationary distribution at fixed prices. Throughout these experiments, an effect is

$$100 \times (\text{benchmark outcome} - \text{counterfactual outcome}),$$

measured in percentage points. A positive entry therefore means that the feature present in the benchmark raises the outcome. Participation, saving, bargaining, and the distribution of resources

Figure 2. LTC cooperation by parent and child productivity

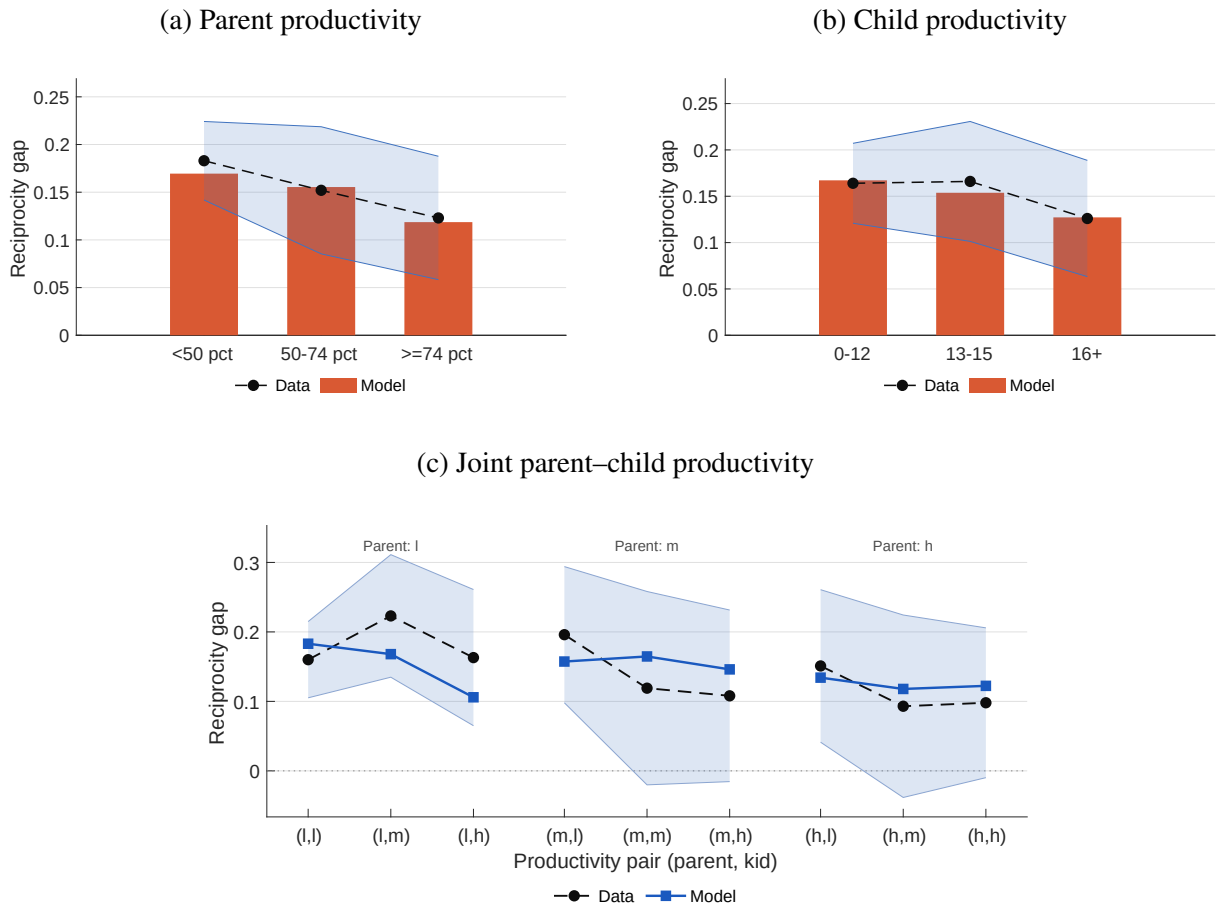


All panels condition on positive LTC need. Panels (a) and (b) pool the C_1 and N_1 histories and are complementary one-way comparisons. Panel (c) additionally conditions on N_1 and displays the nine directly targeted productivity cells. Empirical construction follows Section 3; the shaded bands in panel (c) show 95% confidence intervals based on the reported unclustered standard errors.

respond jointly to each intervention. The experiments therefore reveal how a change at one care stage is transmitted through the relationship, rather than assigning additive shares to independent mechanisms.

Across the experiments, we track cooperation at both care stages, the relationship between earlier childcare and later LTC, and parental bargaining weights. In the benchmark, 41.49 percent of period-1 families cooperate in childcare (C_1), and 27.61 percent of parents with positive care needs receive informal LTC (C_2). Among these parents, those who previously provided childcare are 14.82 percentage points more likely to receive LTC from their child; we denote this history gap by R. The gap reflects both history-dependent guilt and differences in family characteristics, assets, and bargaining positions across the two childcare histories. Table 12 reports the effects of each experiment by child productivity. The accompanying figures also report the parent's average Pareto weight among cooperating families; changes in this average reflect both changes in the terms of

Figure 3. Reciprocity by parent and child productivity



The reciprocity gap is the difference in LTC cooperation between families with childcare histories C_1 and N_1, conditional on positive LTC need. Panel (c) displays the nine directly targeted gaps; panels (a) and (b) are not separately targeted. Empirical construction follows Section 3; the shaded bands in panel (c) show 95% confidence intervals based on the reported unclustered standard errors.

cooperation and changes in which families cooperate.

6.1 Childcare incentives and later family support

The first two experiments distinguish the value of grandparental care from the attractiveness of purchasing care outside the family arrangement. This distinction matters for incidence. A change in cooperative utility affects the gains available to both generations, whereas cheaper paid care directly improves the young adult's alternative to cooperation. Both changes can influence later LTC through the initial care history, saving, and bargaining.

Productivity dependence of childcare utility. The first experiment evaluates the productivity composite in the childcare utility term at its normalized reference value of one. It retains the term's

Table 12. Determinants of family cooperation

	C_1		C_2		R	
Child productivity	Low	High	Low	High	Low	High
Benchmark level	33.936	46.009	28.943	28.057	16.717	12.720
Productivity gradient in childcare utility	-12.416	5.792	-1.859	0.759	-0.300	-0.114
Positive paid-childcare cost	9.277	5.961	1.215	0.310	2.369	1.204
History-dependent guilt	6.556	-1.854	5.736	4.577	16.758	10.315
Informal-care advantages in formal-LTC cost and nursing-home risk	7.522	2.041	5.271	2.311	2.592	1.636

Notes: Benchmark levels are percentages; counterfactual entries are benchmark-minus-counterfactual changes in percentage points. C_1 uses all period-1 families. C_2 and R condition on positive LTC need, with $R = \Pr(C_2 \mid C_1, \lambda > 0) - \Pr(C_2 \mid N_1, \lambda > 0)$. Low and high refer to child productivity. The counterfactuals, in order, make the period-1 childcare utility term productivity-neutral; set realized paid-childcare offer costs to zero; remove the child's history-dependent guilt component associated with C_1; and make formal LTC free and give it the informal-care nursing-home risk.

level at that reference and the remaining model parameters. Since this payoff enters both individuals' cooperative values, the direct intervention changes the value of C_1. Non-cooperative values can still respond indirectly through saving and continuation values.

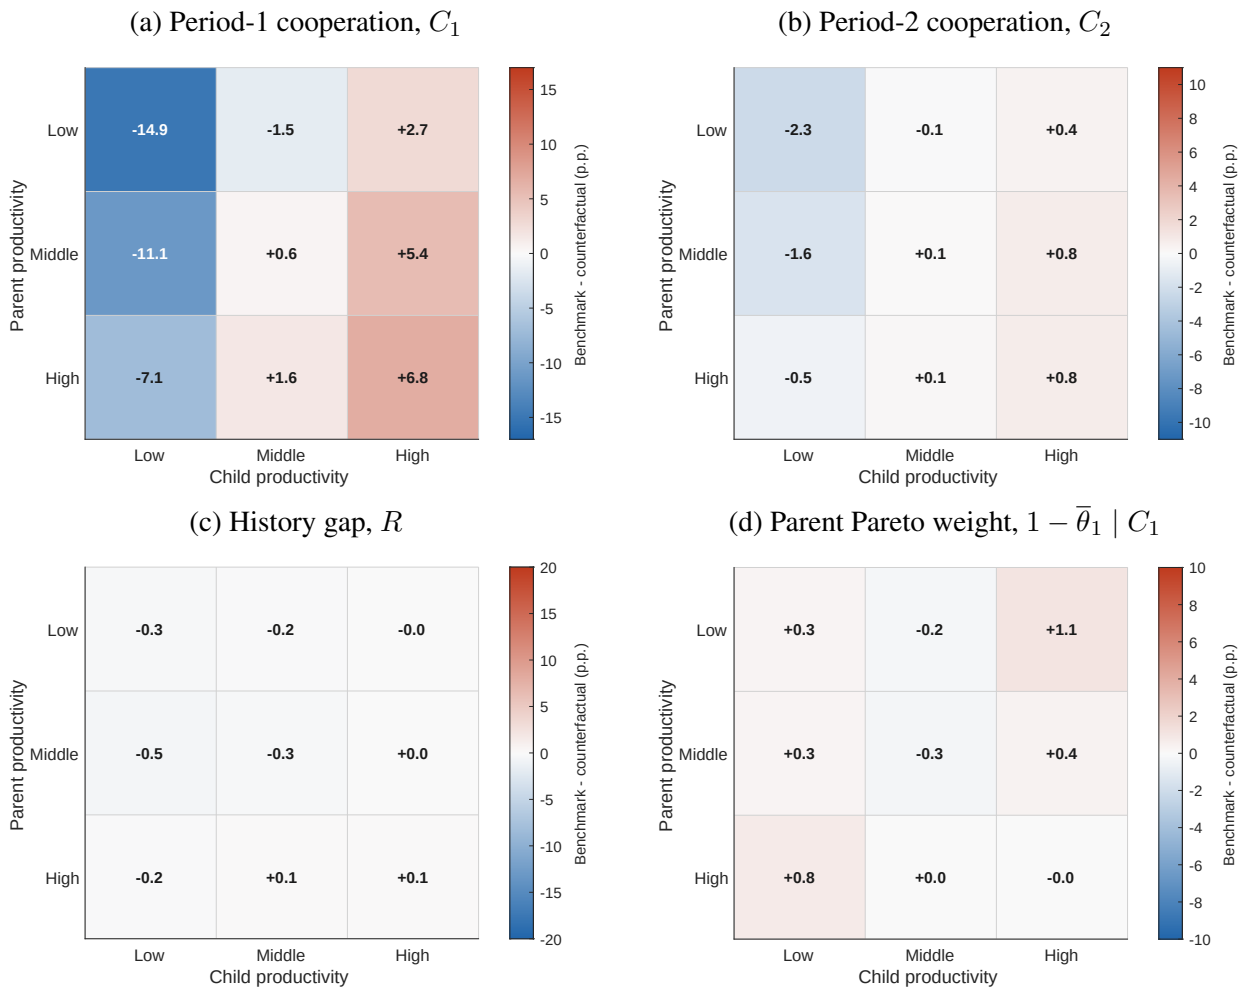
The productivity dependence is important for the allocation of grandparental care across families. Relative to the productivity-neutral specification, it lowers C_1 by 12.42 percentage points for low-productivity children and raises it by 5.79 points for high-productivity children. The aggregate effect is only -1.41 points because these opposing responses largely offset. Thus, the small aggregate change conceals a substantial reallocation of grandparental care across families.

These different childcare histories also shape later support. The benchmark productivity dependence lowers C_2 by 1.86 points for low-productivity children and raises it by 0.76 point for high-productivity children, despite no direct change in LTC preferences or technology. Earlier cooperation affects the child's subsequent disagreement payoff and the incoming bargaining weight, while saving determines both generations' resources and outside options. The corresponding history-gap effects are much smaller, at -0.3 and -0.11 point. The main dynamic response here is therefore a change in the prevalence of later care across productivity groups, rather than a large change in its association with earlier childcare.

The parent's average weight among C_1 cooperators changes by only -0.09 point in the aggregate, with cell effects ranging from -0.35 to $+1.05$ points. Figure 4 thus shows substantial changes in participation alongside modest changes in the average terms of selected agreements. Because both individuals receive the childcare utility payoff, a change in that payoff can alter their willingness to cooperate without requiring a large aggregate shift in bargaining weights. The distinction between participation and the terms of agreement becomes especially relevant when the intervention instead

improves one generation's outside option.

Figure 4. Productivity dependence of period-1 childcare utility



Each cell is 100 times the benchmark minus the productivity-neutral counterfactual, in percentage points. Rows index parent productivity and columns index child productivity, from low to high. Panel (a) uses all period-1 families. Panels (b) and (c) condition on positive LTC need. Panel (d) reports the average realized Pareto weight on the parent conditional on C_1, over all future LTC states. The other three parent- and child-weight panels appear in Appendix E.

The paid-childcare outside option. The second experiment sets the price of every realized paid-childcare offer to zero, retaining the self-care option and its time requirement. It improves the alternative available to the child after N_1. Since participation is decided before the offer is observed, expected access to cheaper care also enters the initial bargaining problem. The intervention does not directly change LTC preferences or technology and has no financing tax. It is therefore a mechanism experiment, distinct from the policy funded voucher studied below.

Making paid childcare free reduces grandparental childcare in every productivity cell. Equivalently, positive paid-care prices in the benchmark raise aggregate C_1 by 9.05 percentage points, with

effects of 9.28 points for low-productivity children and 5.96 points for high-productivity children. Cell effects range from 4.66 to 11.43 points. A more attractive market alternative reduces the child's willingness to accept some family arrangements. Cooperation can then require different allocations, and some families no longer reach an agreement acceptable to both generations.

The terms of selected agreements move in the same aggregate direction. Relative to free paid care, positive prices raise the parent's average weight among C_1 cooperators by 0.46 point. Thus, when paid care becomes free, the parent's conditional weight falls and the child's rises on average. This is consistent with the child requiring more favorable terms when the outside option improves. The response varies across families: cell effects range from -0.28 to $+2.3$ points, reflecting both renegotiation and selection into cooperation. Figure 5 reports these differences; Appendix E provides the other bargaining-weight responses.

The effects extend beyond the childcare stage. Positive paid-care costs raise aggregate C_2 by 0.95 point and the history gap by 1.42 points. The effects on later care conditional on C_1 and N_1 are $+0.56$ and -0.86 point, respectively. Making paid care free therefore does not lower every conditional care rate. It changes the prevalence and composition of childcare histories as well as subsequent choices. Fewer initial agreements reduce the prevalence of the history associated with a higher guilt cost of declining LTC, while changes in saving and inherited bargaining positions also affect later support.

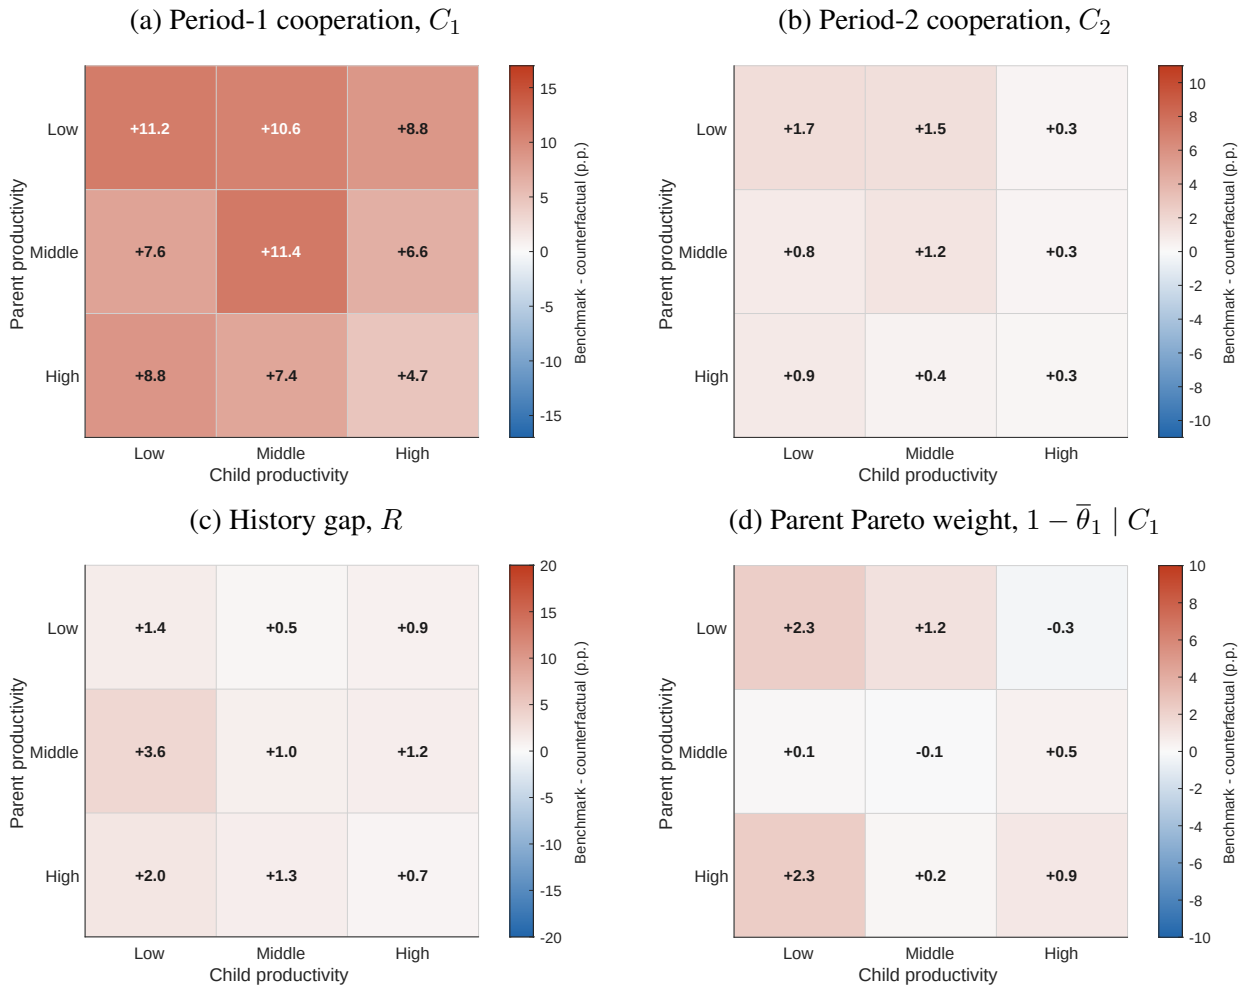
This experiment connects the mechanism analysis to the policy question. The child can benefit from a better childcare alternative while the parent loses access to a favorable family arrangement or receives less favorable terms within it. Those consequences can persist into the parent's LTC stage even when the change in the overall care rate is modest. The parent's welfare therefore depends on the resources and support available across both stages, not simply on the amount of childcare supplied. Sections 7 and 8 evaluate those consequences directly.

6.2 Future care incentives and current childcare agreements

The next two experiments establish the reverse connection. They change incentives at the LTC stage, but both generations anticipate those changes when deciding whether to cooperate and how much to save in the childcare stage. The resulting first-period responses show why the value of grandparental care cannot be assessed independently of the family's expected later arrangement.

History-dependent guilt. The history-dependent guilt experiment removes the effect of earlier childcare on the child's guilt cost of declining LTC. It sets the intercept and both productivity slopes of the C_1-minus-N_1 increment in the guilt coefficient to zero, while retaining the baseline guilt component. Removing the increment improves the child's disagreement payoff after C_1. Earlier

Figure 5. The cost of paid childcare



Each cell is 100 times the benchmark minus the zero-price paid-childcare counterfactual, in percentage points. Rows index parent productivity and columns index child productivity, from low to high. Panel (a) uses all period-1 families. Panels (b) and (c) condition on positive LTC need. Panel (d) reports the average realized Pareto weight on the parent conditional on C_1, over all future LTC states; changes in this conditional mean combine bargaining and selection into cooperation. The other three parent- and child-weight panels appear in Appendix E.

childcare still affects later resources and the incoming bargaining weight, so the two histories remain economically distinct.

The largest response occurs where the preference change acts directly. The benchmark increment raises later care conditional on C_1 by 13.33 percentage points, compared with an effect of -0.31 point conditional on N_1. It raises aggregate C_2 by 5.35 points and the history gap by 13.65 points. Without the history-dependent guilt component, the history gap remains 1.18 points. The residual gap reflects differences in assets, incoming weights, and selection across the two histories; removing one preference term does not make those histories comparable in all other respects.

Anticipating the later guilt cost changes the willingness to enter a childcare agreement, with different effects across families. The benchmark history-dependent guilt component raises C_1 by

6.56 points for low-productivity children but lowers it by 1.85 points for high-productivity children. Cell effects range from -4.57 to $+7.77$ points, and the aggregate increase is 2.42 points. A parent values the prospect of later support, while the child anticipates both the cost of supplying care and the guilt cost of declining it. These considerations enter opposite sides of the current participation problem. The net effect depends on resources, opportunity costs, and the allocation needed to keep both individuals willing to cooperate. Reciprocity therefore supports later care without necessarily encouraging initial cooperation in every family.

The parent's average weight among C_2 cooperators changes by only 0.19 point in the aggregate, with effects between -1.06 and $+0.91$ points across cells. Figure 6 shows a large care response alongside small, mixed changes in these conditional weights. This pattern is consistent with reciprocity sustaining agreements that would otherwise fail a participation constraint. The earlier-care history can thus make later support sustainable even without a large change in the average bargaining weight among cooperating families.

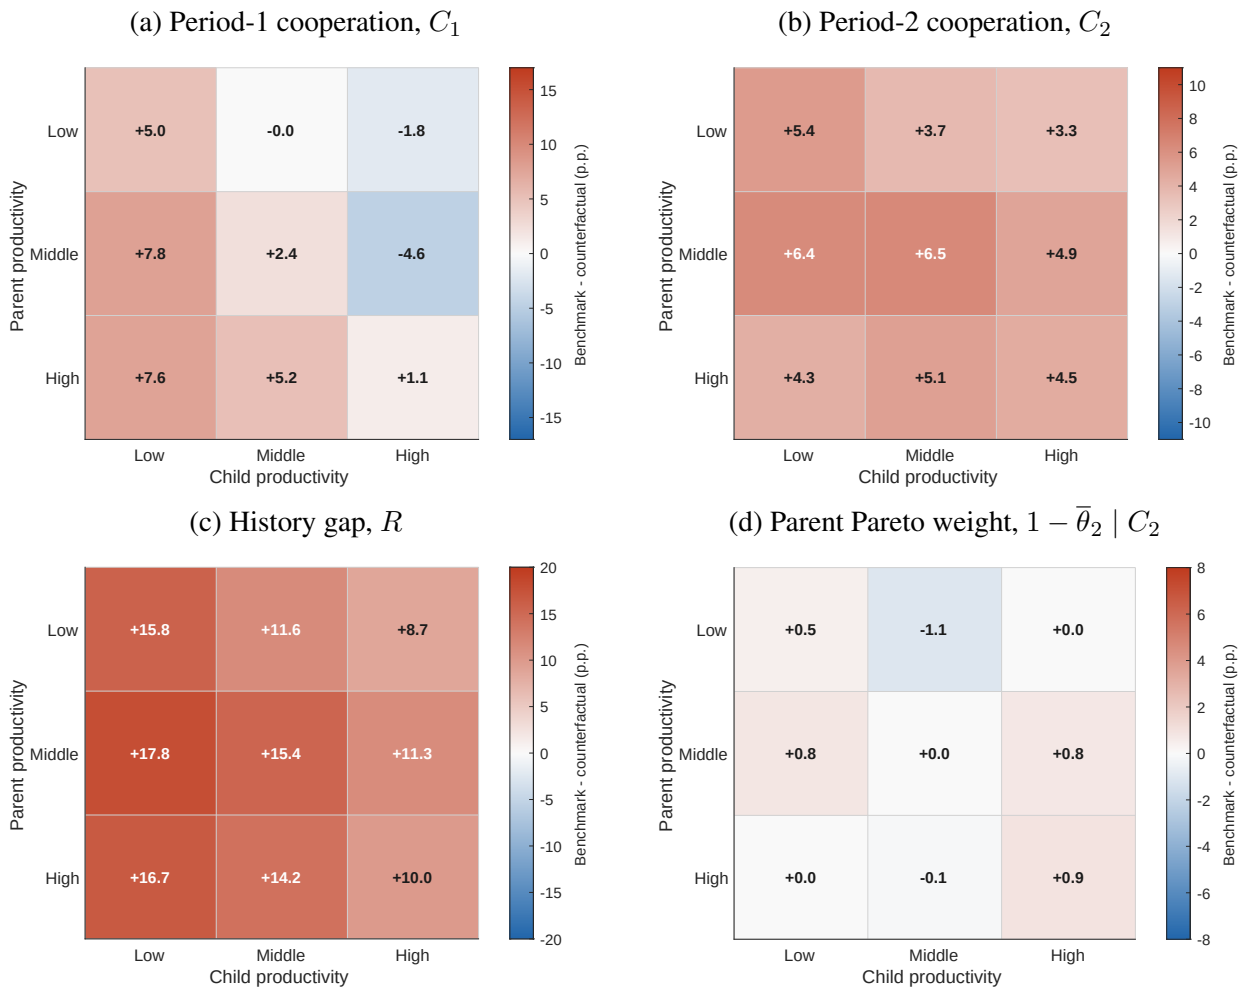
Formal-care cost and nursing-home risk. Informal LTC avoids the formal-care bill and carries a lower nursing-home probability in the benchmark. The final experiment removes both advantages. Formal LTC under N_2 becomes free to the household and receives the informal-care nursing-home probability in both individuals' values and the realized inheritance transition. Nursing-home disutility and estate depletion conditional on entry remain.

A better formal-care alternative reduces the relative value of informal care. The parent faces a smaller bill and lower institutionalization risk outside the family agreement, while the child anticipates the implications for the surviving estate and future resources. The benchmark advantages raise aggregate C_2 by 3.32 percentage points. The increase is 5.27 points for low-productivity children and 2.31 points for high-productivity children. The larger response among low-productivity children is consistent with care expenditure and protection of family resources carrying greater weight when resources are limited.

The anticipated LTC arrangement has a substantial effect on earlier childcare. Although the intervention does not directly change childcare prices or current childcare utility, the benchmark LTC advantages raise aggregate C_1 by 4.76 points. The effects are 7.52 points for low-productivity children and 2.04 points for high-productivity children. Both individuals consider the later resources and care available under cooperation and non-cooperation. Changes in those continuation values alter the childcare allocations they are willing to accept today.

Unlike the reciprocity intervention, this experiment directly changes the formal-care alternative after either childcare history. The benchmark advantages raise later care conditional on C_1 by 4.28 points and conditional on N_1 by 1.67 points, increasing the history gap by 2.61 points even though

Figure 6. History-dependent guilt

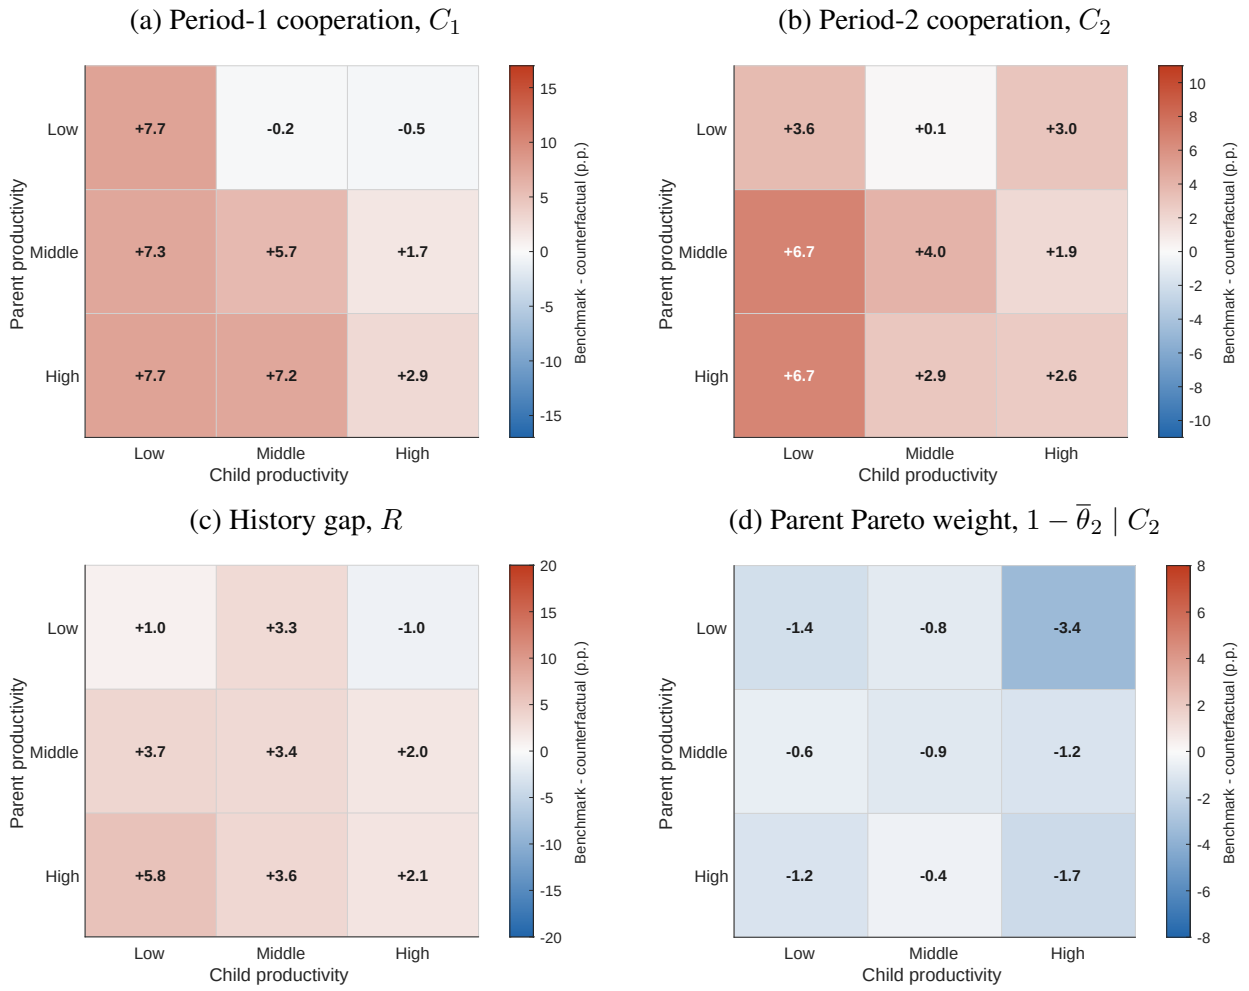


Each cell is 100 times the benchmark minus the counterfactual without the history-dependent guilt component, in percentage points. Rows index parent productivity and columns index child productivity. Panel (a) uses all period-1 families; panels (b) and (c) condition on positive LTC need; and panel (d) additionally conditions on C_2. The history gap need not be zero when the explicit history term is removed because the two histories remain endogenously selected. The remaining three parent- and child-weight panels appear in Appendix E.

the reciprocity coefficients are unchanged. The parent's average weight among C_2 cooperators falls by 1 point. Figure 7 shows that this bargaining response varies across productivity pairs. Together, the participation and bargaining responses show that a better late-life alternative changes both the willingness to cooperate and the terms on which care is exchanged.

The four experiments show that care incentives operate in both directions over the family relationship. Childcare preferences and market alternatives change initial participation and the history, resources, and bargaining position relevant for later LTC. Reciprocity and formal-care alternatives change later support and, through anticipation, the childcare agreements made earlier. Appendix F reports a further experiment that restricts resource pooling at both dates and also extends

Figure 7. Formal-LTC cost and nursing-home risk advantages of informal care



Each cell is 100 times the benchmark minus the counterfactual with free formal LTC and equalized nursing-home risk, in percentage points. Rows index parent productivity and columns index child productivity. Panel (a) uses all period-1 families. Panels (b) and (c) condition on positive LTC need, and panel (d) additionally conditions on C_2. The experiment changes formal-care cost and nursing-home risk jointly. The remaining three parent- and child-weight panels appear in Appendix E.

consumption-floor eligibility to qualifying cooperative LTC parents. Its effects combine those two changes.

The paid-childcare experiment is particularly relevant for policy. A subsidy can improve the young adult's market alternative while changing the resources and support available to the parent. A decline in family care is not itself a welfare loss, and a lower conditional parental weight does not measure the size of any loss. The next sections therefore evaluate individual welfare and ask whether subsidizing purchased childcare benefits the young adult at the expense of the incumbent parent, and whether supporting grandparental care as well changes that incidence.

7 Welfare

The preceding experiments show how changes in participation incentives at one care stage affect family arrangements at the other. These responses also redistribute the gains from cooperation. Under limited commitment, better paid-care options can benefit the adult child while reducing the parent's access to shared resources or later support. A fall in informal care therefore does not by itself establish a welfare loss, and a gain for the subsidized generation need not extend to its current parent. We evaluate individual lifetime utility and distinguish entrant welfare from the remaining-lifetime welfare of incumbent parents.

Lifetime welfare. Let $w_p(x)$ denote an individual's expected lifetime utility at entry into adulthood, conditional on initial state x, under regime p; $p = 0$ denotes the benchmark. The initial state includes both generations' assets and productivities. The value $w_p(x)$ follows the individual through two stages as an adult child and two subsequent stages as a parent. Flow utility is discounted by $1, \beta, \beta^2$, and β^3, and own terminal warm-glow and nursing-home utility by β^4. Welfare uses individual payoffs, rather than the Pareto-weighted objective that determines cooperative allocations. It includes consumption, leisure, and all care-related and terminal utility. The love shock contributes $\mathbb{E}[q\mathbf{1}\{C\}]$, which depends on which families cooperate under each regime.

Let ν_0 be the benchmark stationary distribution of entry states and ν_p the stationary entry distribution under regime p. We report

$$W_p^{\mathrm{CS}} = \int w_p(x)\, d\nu_0(x), \qquad W_p^{\mathrm{LR}} = \int w_p(x)\, d\nu_p(x), \qquad W_0 = \int w_0(x)\, d\nu_0(x). \tag{48}$$

The common-start measure evaluates the policy for entrants drawn from the benchmark distribution. Only their initial-state weights are held fixed: saving, bargaining, care, inheritance, and all subsequent outcomes follow the policy regime. The long-run measure instead evaluates entrants drawn from the policy's stationary distribution, allowing their initial family resources to adjust as well. Comparing these measures identifies the effect of changing entry-state weights; both already include dynamic responses within the individual's lifetime.

Consumption-equivalent variation. We express each welfare comparison as a proportional change in consumption at every stage of the benchmark life cycle. Write $W_0 = S_0 + D_0$, where

$$S_0 = \int \mathbb{E}_0\left[\sum_{j=0}^{3} \beta^j u(c_j^0, \ell_j^0) \,\middle|\, x\right] d\nu_0(x), \qquad D_0 = W_0 - S_0. \tag{49}$$

Scaling benchmark consumption by $1+g_p$ leaves leisure, care arrangements, event probabilities, and terminal estates fixed. With $u(c,\ell) = (c^\gamma \ell^{1-\gamma})^{1-\sigma}/(1-\sigma)$, the consumption-equivalent variation (CEV) satisfies

$$(1+g_p)^{\gamma(1-\sigma)} S_0 + D_0 = W_p, \qquad g_p = \left[1 + \frac{W_p - W_0}{S_0}\right]^{1/[\gamma(1-\sigma)]} - 1, \tag{50}$$

Here W_p is either W_p^{CS} or W_p^{LR}, and the expression in brackets must be positive. Both comparisons use the same benchmark welfare and consumption–leisure component.

This conversion holds the reference allocation fixed, including terminal bequests. Actual reform bequests remain endogenous, and their effects on inheritance and warm-glow utility enter W_p in full. Group-specific CEVs aggregate welfare and the benchmark consumption–leisure component within the group before applying the conversion. The reported share of states with $w_p(x) > w_0(x)$ uses benchmark entry-state weights and complements the average common-start welfare comparison.

Intergenerational incidence. In the childcare experiments, the directly affected young adult is the common-start entrant: their CEV includes the current childcare decision and their own subsequent stages as child and parent. It is therefore identical to the common-start CEV, whether or not they purchase paid care. The incumbent-parent comparison follows that young adult's current parent from the childcare stage through LTC and the terminal payoff. It holds the parent's initial state distribution at the benchmark and allows all remaining outcomes to respond under the policy regime. These two comparisons distinguish the subsidized generation's lifetime gain from the effect on its current parent. Their different horizons and benchmark consumption references also mean that the two CEVs cannot be added.

8 Paid childcare and intergenerational welfare

8.1 Policy design and financing

Subsidizing purchased childcare can relax young households' resource constraints and make market work more attractive. These benefits are central to the analysis of childcare policy (Ho and Pavoni, 2020; Bastani et al., 2020). They are also present here: young adults cannot borrow, and purchasing care competes with consumption and saving for limited current resources. The additional consideration is that purchased care substitutes for grandparental childcare. Improving the young adult's market alternative changes the terms on which the family is willing to share resources and exchange care. A policy that benefits the young adult can therefore harm the incumbent parent, who supplies childcare today and may need support later. This intergenerational margin connects

the time transfers studied by Cardia and Ng (2003) to evidence on the consequences of childcare arrangements for older women's economic security (Truskinovsky, 2023).

We first consider a voucher available only for purchased childcare. If the family does not cooperate in the childcare stage and the young adult purchases care at gross offer K, the voucher divides the bill into

$$K^{\text{private}} = \max\{K - v_c, 0\}, \qquad K^{\text{public}} = \min\{v_c, K\}. \tag{51}$$

An inexpensive offer can be covered in full. Neither self-provided care nor grandparental childcare is subsidized, and eligibility does not imply that a household purchases care. This design improves the alternative to a cooperative childcare agreement without adding resources directly to that agreement.

The reform is financed by a Heathcote–Storesletten–Violante tax schedule applied to each worker's gross labor earnings y. We denote its scale by Λ_T to distinguish it from the informal-care time requirement:

$$y^{\text{net}} = \Lambda_T y^{1-p_T}, \qquad T(y) = y - \Lambda_T y^{1-p_T}. \tag{52}$$

The progressivity parameter is $p_T = 0.8 \times 10^{-2}$ and the scale is $\Lambda_T = 98.5 \times 10^{-2}$. The schedule applies to the young adult's earnings at both stages and to the parent's earnings when working in the childcare stage. Taxation remains individual when resources are pooled. Pensions, assets, inheritances, and intrafamily transfers are outside the tax base. At sufficiently low positive earnings, the schedule provides a refundable transfer. These transfers are separate from the old-age consumption floor, whose eligibility rule is unchanged.

The reform targets childcare spending equal to one percent of benchmark taxable earnings:

$$B = 0.010 Y_0, \qquad R_T - G_c - \Delta F - \Delta P \simeq 0, \tag{53}$$

Here R_T is revenue net of refundable tax liabilities, G_c is voucher expenditure, and ΔF and ΔP are changes in consumption-floor and prescribed-pension expenditure.

We consider this experiment under both general equilibrium (GE) and partial equilibrium (PE). In GE, household decisions, government budget, and the stationary distribution are recomputed, while wages, the interest rate, and care-price processes remain fixed. PE retains the exact accepted voucher, removes both generations' policy taxes and refundable tax transfers, and recomputes behavior and the stationary distribution. It requires external financing. The comparison shows how the subsidy operates without its financing scheme; the difference between GE and PE also includes the behavioral and distributional responses to that scheme.

Welfare is measured as in Section 7. Common-start welfare holds the benchmark distribution of entry states fixed; long-run welfare also incorporates the reform's stationary entry distribution. Both

include all subsequent changes in care, saving, bargaining, and inheritance.

8.2 Entrant welfare and the allocation of care

The paid-care voucher raises average entrant welfare under both financing arrangements. Common-start CEV is 0.08 percent in GE and 0.6 percent in PE; long-run CEV is 0.14 and 0.71 percent, respectively. The change in entry-state weights therefore raises measured CEV by 0.06 percentage point in GE and 0.11 point in PE. Under benchmark weights, 50.4 percent of entry states gain in GE, compared with 100 percent in PE. The positive average gain with financing is consequently much smaller and less widely shared than the gain under external financing.

The policy primarily changes childcare arrangements. Paid-care use among noncooperating young adults rises from 32.63 percent to 41.06 percent in GE and 42.16 percent in PE. Grandparental childcare falls from 41.49 percent of families to 40.22 and 40.43 percent. By lowering the private cost of purchased care, the voucher makes the child less dependent on grandparental help. This improves access to market care while reducing the number of families sharing resources within a childcare agreement.

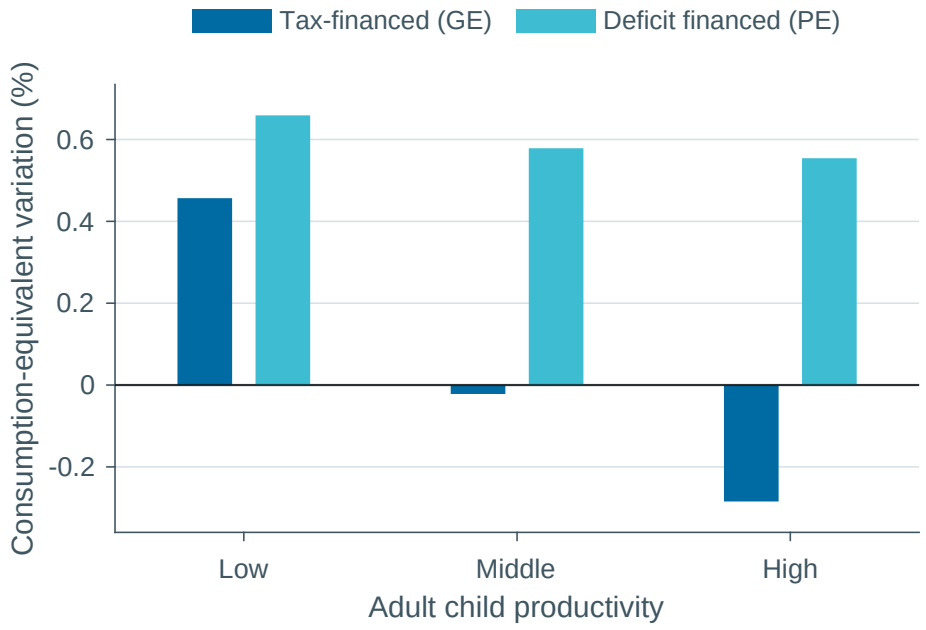
As in the above counterfactual experiments, the response extends beyond the subsidized stage. Informal LTC conditional on positive need falls somewhat and the reciprocity gap narrows. Earlier cooperation changes the child's later disagreement payoff, but also saving and the bargaining position carried into the LTC stage. The responses therefore reflect the combined dynamic consequences of changing the childcare alternative, rather than an isolated reciprocity effect. Appendix Figure D2 reports the changes across parent–child productivity pairs.

The aggregate gain also conceals substantial differences across productivity groups. Figure 8 reports common-start CEVs of 0.46, -0.02, and -0.29 percent for low-, middle-, and high-productivity young adults in GE. In PE, all three groups gain, by 0.66, 0.58, and 0.55 percent. Progressive financing and its induced adjustments thus change both the level and the distribution of welfare gains. The tax-financed reform benefits low-productivity young adults on average while imposing a loss on the high-productivity group.

8.3 Gains for young adults and losses for incumbent parents

A positive entrant CEV does not establish that both generations in the current family gain. The directly affected young adult is eligible for the voucher now and is followed through the remainder of life, whether or not paid care is purchased. The incumbent parent enters the comparison at the current grandchild-care stage, having received no voucher when young, and is followed through LTC and terminal utility. This distinction allows the same reform to improve entrant welfare while

Figure 8. Ex ante welfare from the paid-childcare policy by adult-child productivity



Notes: Bars report common-start CEVs in percent for low-, middle-, and high-productivity adult children. "Tax-financed (GE)" includes the progressive schedule in equation (52); "Zero-tax (PE)" retains the accepted voucher but requires external financing. Group utility and its benchmark consumption–leisure reference are first aggregated with benchmark group weights and then converted to a CEV. Prices are fixed in both comparisons.

reducing the welfare of a generation already in the parental role.

Let U_{Y1} and U_{Y2} denote flow utility in the childcare and LTC stages as a young adult, and U_{O1} and U_{O2} the corresponding flows as a parent. Each contains consumption–leisure utility and the applicable love and care terms; T_O is own terminal warm-glow and nursing-home utility. Under permanent regime r, the directly affected young adult's remaining-lifetime welfare is

$$W_Y^r = \int \mathbb{E}_r\left[U_{Y1}^r + \beta U_{Y2}^r + \beta^2 U_{O1}^r + \beta^3 U_{O2}^r + \beta^4 T_O^r \mid x\right] d\nu_0(x). \tag{54}$$

This is the common-start entrant measure in equation (48). The direct-young CEV therefore repeats that measure rather than identifying a treatment effect conditional on purchasing care.

For the incumbent parent, μ_{O1}^0 is the benchmark family-state distribution immediately before the current childcare-stage love shock. Remaining-lifetime welfare is

$$W_P^r = \int \mathbb{E}_r\left[U_{O1}^r + \beta U_{O2}^r + \beta^2 T_O^r \mid s\right] d\mu_{O1}^0(s). \tag{55}$$

The same initial parent population is evaluated in every regime, while subsequent choices and transitions respond to the reform. A working parent can pay tax or receive a refundable tax transfer in GE but receives no childcare voucher. In PE, the parent receives neither a voucher nor a policy

tax transfer and pays no policy tax. Any welfare effect in that comparison is transmitted through the family allocation and its continuation.

For each role $j \in \{Y, P\}$, the benchmark consumption–leisure reference is

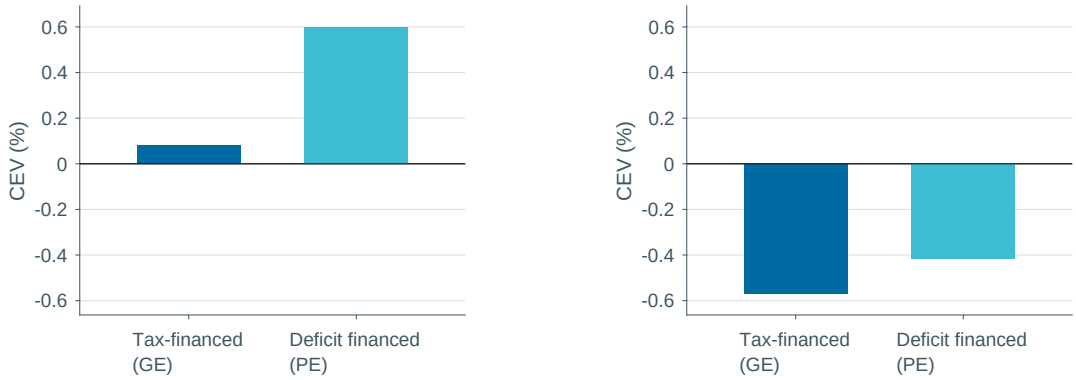
$$S_{j0} = \int \mathbb{E}_0 \left[\sum_{t \in \mathcal{H}_j} \beta^{t-t_j} u(c^0_{jt}, \ell^0_{jt}) \mid s \right] d\mu^0_j(s), \qquad D_{j0} = W^0_j - S_{j0}, \tag{56}$$

where $\mathcal{H}_j$ contains the remaining consumption stages for role j, t_j is its first remaining stage, $\mu^0_Y = \nu_0$, and $\mu^0_P = \mu^0_{O1}$. The resulting CEV is

$$g^r_j = \left[1 + \frac{W^r_j - W^0_j}{S_{j0}} \right]^{1/[\gamma(1-\sigma)]} - 1. \tag{57}$$

Each measure uses its own remaining horizon and consumption reference. The parent and young-adult CEVs describe incidence across generations; they are not additive components of aggregate welfare.

Figure 9. Welfare incidence of the paid-childcare policy across family roles



Notes: The left panel reports the directly affected young adult's remaining-lifetime CEV, which equals the common-start ex ante CEV. The right panel reports the incumbent parent's CEV using the fixed benchmark parent distribution in equation (55). CEVs are percentages and use each role's own consumption–leisure reference. Both comparisons hold prices fixed; PE requires external financing.

Figure 9 presents the central result. The young adult gains 0.08 percent in GE and 0.6 percent in PE, whereas the incumbent parent loses 0.57 and 0.42 percent. Removing financing reduces the parental loss somewhat but does not reverse it: a subsidy directed at the young adult harms the parent even when neither generation pays for the subsidy through the policy tax.

This loss is present across all reported parent productivity margins. Grouping parents by the adult child's productivity, CEVs are -0.42, -0.53, and -0.88 percent in GE and -0.31, -0.44, and -0.51 percent in PE. Appendix Figure D1 also reports the grouping by the parent's own productivity.

Negative marginal averages do not imply that every parent loses, but they show that the aggregate result is not confined to a single productivity group.

8.4 Why the parental loss persists without financing

The parental loss persists without policy taxes because the voucher strengthens the child's outside option within a relationship that also allocates consumption and saving. Under limited commitment, cooperation may then end or require terms more favorable to the child. As Section 6.1 illustrates, these changes can carry forward through saving, history-dependent guilt, and bargaining at the LTC stage. Low-productivity parents provide a useful example. In PE, their average CEV loss is 0.67 percent, the largest across parent productivity groups (Appendix Figure D1). Within this group, parents paired with low-productivity children experience declines in their mean conditional Pareto weight of 0.46 percentage point in C_1 and 0.26 point in C_2, among the largest declines across productivity pairs (Appendix Figure D3). These conditional averages reflect both changes in bargaining terms and changes in which families cooperate. Thus, less grandparental childcare can release parental time while leaving parents worse off through reduced resources and support within the family relationship.

Next, Section 9 asks whether recognizing grandparental care in the subsidy's eligibility rule can instead benefit both generations within the same target budget.

9 Designing childcare subsidies around family care

9.1 Extending eligibility to grandparental care

The preceding experiment identifies a design problem. A subsidy restricted to purchased childcare improves the young adult's alternative to a family agreement without adding resources to that agreement. The incumbent parent can consequently lose even without a financing tax. Section 6 explains why this matters beyond the current childcare decision: the terms on which a family cooperates today affect saving and later care, while the anticipated value of that later relationship helps determine whether both generations will cooperate today. We now ask whether recognizing grandparental care in the eligibility rule can improve welfare for both generations within the same childcare spending target.

Grandparental childcare uses a costly input—the parent's time—even though it does not generate a market transaction. Work on childcare subsidies emphasizes market care and parental labor supply (Ho and Pavoni, 2020; Bastani et al., 2020), while related studies document grandparental time transfers and their labor-market costs (Cardia and Ng, 2003; Truskinovsky, 2023; Backhaus and

Barslund, 2021). In our model, this time transfer is also part of a relationship through which the generations share resources and may exchange care later. The broader policy therefore supports the cooperative childcare arrangement as well as purchases of paid care.

Let s_1 denote the childcare arrangement and K the paid-care offer. The broader eligibility rule is

$$b^{\text{broad}}(s_1, K) = \begin{cases} v_b, & s_1 = C_1, \\ \min\{v_b, K\}, & s_1 = N_1 \text{ and the young adult buys care}, \\ 0, & s_1 = N_1 \text{ and the young adult supplies care}. \end{cases} \tag{58}$$

Under cooperation, the full voucher enters the family's common budget. It supports grandparental care without assigning the payment to either generation: consumption and saving remain determined by bargaining subject to both participation constraints. Following non-cooperation, the voucher continues to reimburse purchased care up to its cost, and self-provided childcare remains unsubsidized. Thus, this is an extension to cooperative grandparental care, rather than a universal payment independent of the care arrangement.

Both designs target $B = 0.01Y_0$. They use the same progressive-tax functional form and progressivity parameter in the tax-financed comparison, with the fiscal scale adjusted separately for each allocation. Childcare outlays are

$$\begin{aligned} G_c^{\text{broad}}(v_b) &= v_b \Pr(C_1) + \mathbb{E}[\min\{v_b, K\}\mathbf{1}\{N_1, \text{ paid care}\}], \\ G_c^{\text{paid}}(v_c) &= \mathbb{E}[\min\{v_c, K\}\mathbf{1}\{N_1, \text{ paid care}\}]. \end{aligned} \tag{59}$$

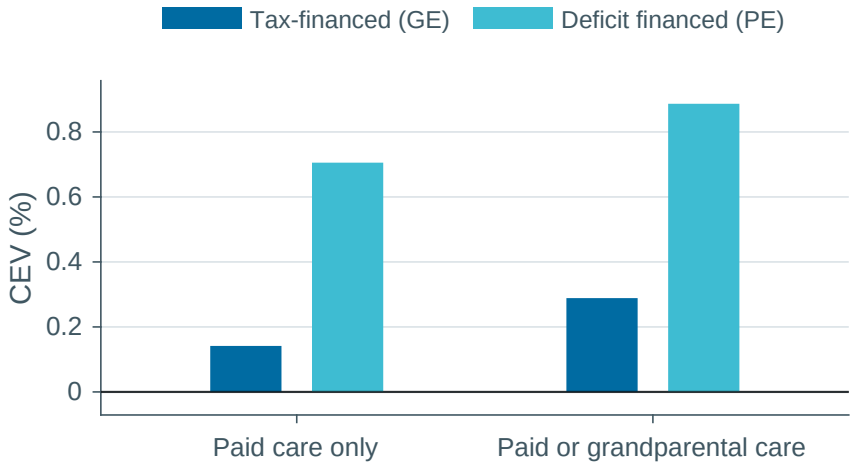
Extending eligibility requires spreading the target budget across more families. Realized tax-financed voucher spending is approximately $1.03B$ and $1.05B$, under this broader design policy and the aforementioned narrow paid childcare subsidy design, respectively. The comparison therefore holds the target budget fixed, rather than realized spending exactly fixed. Nor does a common gross voucher target impose identical net fiscal costs: changes in consumption-floor and prescribed-pension spending also enter the public budget.

Under the broad tax-financed design, 67.5 percent of childcare spending enters cooperating families' common budgets and 32.5 percent reimburses paid care. The corresponding partial equilibrium zero-tax shares are 67.66 and 32.34 percent. Eligibility therefore changes both sides of the participation comparison. The policy still makes purchased care more affordable, but it also adds resources to the cooperative arrangement. This can make an allocation acceptable to both generations without requiring the parent to bear the full adjustment to an improved outside option for the child. Cooperation and the division of resources nevertheless remain endogenous.

9.2 Welfare across generations

Figures 10–12 compare the designs under tax financing (GE) and zero policy taxes (PE). As in Section 8.1, prices remain fixed in both comparisons. PE retains each design's accepted voucher, reoptimizes household decisions and the stationary distribution, and requires external financing.

Figure 10. Long-run ex ante welfare under alternative childcare-voucher designs



Notes: Bars report long-run entrant CEVs in percent using each regime's stationary entrant distribution. Both designs target the same childcare-policy budget. GE includes progressive financing; PE is externally financed. Prices are fixed in both comparisons.

The broader design raises long-run entrant CEV from 0.14 to 0.29 percent with tax financing and from 0.71 to 0.89 percent without policy taxes. Its advantage is 0.15 and 0.18 percentage point, respectively. Broader eligibility therefore improves average lifetime welfare under both financing regimes. The welfare ranking holds both with the specified progressive financing and with policy taxes removed.

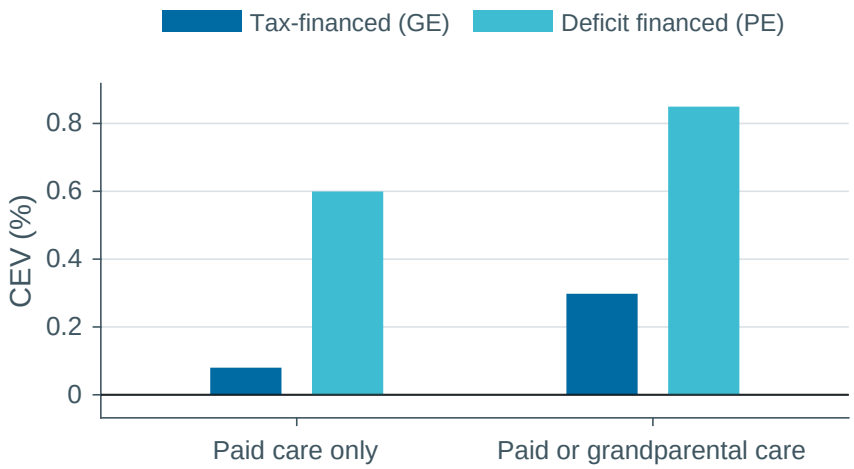
For the current young adult, broad eligibility raises common-start CEV from 0.08 to 0.3 percent in GE and from 0.6 to 0.85 percent in PE. These gains use the same benchmark distribution of initial states and include young adults who never use a voucher. They are population welfare comparisons, rather than effects conditional on receiving the subsidy. The lower payment for purchased care is more than offset, on average, by extending support to the family arrangement and by the resulting changes over the life cycle.

Tax-financed common-start CEVs are 0.87, 0.13, and -0.22 percent for low-, middle-, and high-productivity children. Without policy taxes, all three groups gain on average, with CEVs of 1.26, 0.7, and 0.56 percent.

The incumbent-parent result is the clearest consequence of the eligibility change. Paid-only eligibility lowers parent CEV by 0.57 percent in GE and 0.42 percent in PE. Broad eligibility instead

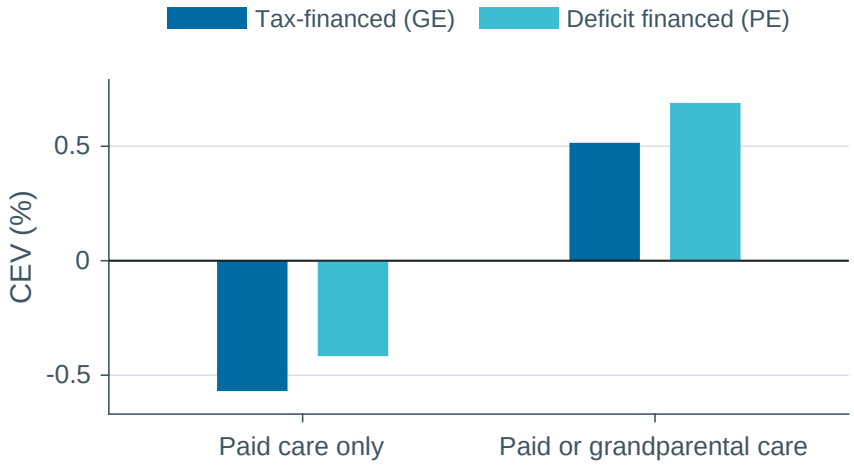
raises it by 0.52 and 0.69 percent. The improvement relative to the paid-only design is 1.08 and 1.11 percentage points, reversing the parent's average loss under either financing regime.

Figure 11. Young-adult welfare under alternative childcare-voucher designs



Notes: Bars report the directly affected young adult's remaining-lifetime CEV using the common benchmark entry distribution. This is the common-start CEV, not a CEV conditional on receiving or using a voucher. GE includes progressive financing; PE is externally financed.

Figure 12. Current-parent welfare under alternative childcare-voucher designs



Notes: Bars report the incumbent parent's remaining-lifetime CEV, holding fixed the benchmark parent distribution. Under broad eligibility, a cooperating family receives the voucher in its common budget, so the parent is not an untreated nonbeneficiary. GE includes progressive financing; PE is externally financed.

This comparison also changes the parent's policy exposure. Under paid-only eligibility, the parent receives no voucher and is affected through taxes and the family relationship. Under broad eligibility, a cooperating parent's family receives resources directly in its shared budget. This

change in exposure is central to the design result: supporting the care arrangement gives the parent access to public resources through the shared budget, whose allocation is still governed by limited commitment. The improvement extends across the marginal productivity groups. Grouping parents by their own productivity gives gains of 1.06, 0.31, and 0.17 percent in GE and 1.09, 0.56, and 0.37 percent in PE. These results contrast with the negative marginal parent CEVs under paid-only eligibility.

The broader design also reverses the decline in childcare cooperation. For the same initial incumbent-parent population, broad eligibility raises childcare cooperation by 0.58 percentage point relative to the benchmark under tax financing, whereas paid-only eligibility lowers it by 1.38 points. The same pattern holds without policy taxes. In the stationary comparison, later informal care also increases modestly relative to paid-only eligibility, consistent with the dynamic links in Section 6. Appendix Figure D4 reports the broader policy's GE changes in childcare cooperation, informal LTC, and the history gap relative to the benchmark across parent–child productivity pairs. Extending support to grandparental care therefore both adds resources to the family arrangement and helps sustain cooperation.

At the same spending target, including grandparental care roughly doubles long-run GE CEV relative to paid-only eligibility, from 0.14 to 0.29 percent, while turning incumbent parents' average welfare loss into a gain.

10 Conclusion

Family care links the effects of childcare policy to support in old age. We document a strong association between earlier grandparental childcare and later care from adult children, and show how limited commitment connects these transfers through reciprocity, saving, and resource sharing. Subsidies restricted to purchased childcare raise young adults' average welfare but leave incumbent parents worse off, even without policy taxes. Under progressive financing and the same spending target, extending eligibility to grandparental care roughly doubles long-run consumption-equivalent welfare gains for young entrants and turns parents' average losses into gains.

These findings motivate research on the optimal joint design of childcare and long-term care policy. Childcare assistance, support for family caregivers, and formal LTC coverage may complement or substitute for one another through their effects on family cooperation. An optimal policy analysis would choose benefit levels, eligibility, and financing while accounting for welfare across generations. Studying reform transitions would also clarify how to support families whose earlier care and saving decisions were made under different policy rules.

References

Altonji, J. G. and Segal, L. M. (1996). Small-sample bias in gmm estimation of covariance structures. *Journal of Business & Economic Statistics*, 14(3):353–366.

Backhaus, A. and Barslund, M. (2021). The effect of grandchildren on grandparental labor supply: Evidence from europe. *European Economic Review*, 137:103817.

Barczyk, D., Fahle, S., and Kredler, M. (2023). Save, spend, or give? a model of housing, family insurance, and savings in old age. *The Review of Economic Studies*, 90(5):2116–2187.

Barczyk, D. and Kredler, M. (2014). Altruistically motivated transfers under uncertainty. *Quantitative Economics*, 5(3):705–749.

Barczyk, D. and Kredler, M. (2018). Evaluating long-term-care policy options, taking the family seriously. *The Review of Economic Studies*, 85(2):766–809.

Bastani, S., Blomquist, S., and Micheletto, L. (2020). Child care subsidies, quality, and optimal income taxation. *American Economic Journal: Economic Policy*, 12(4):1–37.

Bernheim, B. D., Shleifer, A., and Summers, L. H. (1985). The strategic bequest motive. *Journal of Political Economy*, 93(6):1045–1076.

Brandsaas, E. E. (2025). Illiquid homeownership and the bank of mom and dad. Finance and Economics Discussion Series 2025-094, Board of Governors of the Federal Reserve System.

Brugiavini, A., Buia, R. E., Pasini, G., and Zantomio, F. (2013). Long-term care and reciprocity: does helping with grandchildren result in the receipt of more help at older ages? In Börsch-Supan, A., Brandt, M., Litwin, H., and Weber, G., editors, *Active ageing and solidarity between generations in Europe: First results from SHARE after the economic crisis*, pages 369–378. De Gruyter, Berlin and Boston.

Cardia, E. and Ng, S. (2003). Intergenerational time transfers and childcare. *Review of Economic Dynamics*, 6(2):431–454.

Chaparro, J., Sojourner, A., and Wiswall, M. J. (2020). Early childhood care and cognitive development. Technical report, National Bureau of Economic Research.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Cigno, A. (1993). Intergenerational transfers without altruism: Family, market and state. *European Journal of Political Economy*, 9(4):505–518.

Cigno, A. (2006). A constitutional theory of the family. *Journal of Population Economics*, 19(2):259–283.

Fu, C., Ko, A., and Yoon, C. (2026). Intergenerational time transfers: Altruism and strategic motives within the family. Working paper.

Gayle, G.-L. and Shephard, A. (2019). Optimal taxation, marriage, home production, and family labor supply. *Econometrica*, 87(1):291–326.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Ho, C. and Pavoni, N. (2020). Efficient child care subsidies. *American economic review*, 110(1):162–199.

İmrohoroğlu, A. and Zhao, K. (2018). The Chinese saving rate: Long-term care risks, family insurance, and demographics. *Journal of Monetary Economics*, 96:33–52.

Kaplan, G. (2012). Moving back home: Insurance against labor market risk. *Journal of Political Economy*, 120(3):446–512.

Ko, A. (2022). An equilibrium analysis of the long-term care insurance market. *The Review of Economic Studies*, 89(4):1993–2025.

Kotlikoff, L. J. and Spivak, A. (1981). The family as an incomplete annuities market. *Journal of Political Economy*, 89(2):372–391.

Lefebvre, M., Pestieau, P., and Schoenmaeckers, J. (2025). Grandchild care and eldercare. a quid pro quo arrangement. *Economic Modelling*, 146:107006.

Lundberg, S. J., Pollak, R. A., and Wales, T. J. (1997). Do husbands and wives pool their resources? evidence from the United Kingdom child benefit. *Journal of Human Resources*, 32(3):463–480.

Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies*, 74(3):857–895.

Mazzocco, M., Ruiz, C., and Yamaguchi, S. (2014). Labor supply and household dynamics. *American Economic Review*, 104(5):354–359.

Mikoshiba, M. (2025). Universal insurance with in-kind transfers: The welfare effects of long-term care insurance in Japan. RIETI Discussion Paper 25-E-030, Research Institute of Economy, Trade and Industry.

Mommaerts, C. (2025). Long-term care insurance and the family. *Journal of Political Economy*, 133(1):1–52.

Social Security Administration (2019). Annual statistical supplement, 2019. Table 3.C5. SSI receipt among the population aged 65 or older; the calibration target uses the December 2014 observation.

Truskinovsky, Y. (2023). The unintended consequences of informal childcare subsidies for older women's retirement security. *Journal of Human Resources*, 58(3):893–928.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332.

Online Appendix

A HRS Data

This appendix documents the construction underlying Section 3. The analysis combines the RAND HRS respondent and family files with respondent-level helper information, links respondents to their children across waves, and reduces the resulting histories to one observation per respondent–child pair.

A.1 Observation Windows and Care Indicators

Linkage and eligibility. Respondent and family records are merged by respondent identifier and survey wave, retaining linked child records, including stepchildren. The construction excludes the 2020 wave. The childcare phase applies the age-40–70 eligibility flag (`flag_age4070`), requires no reported ADL or IADL difficulty, and retains adult children ages 20–40 with at least one child of their own. In the retained analysis data, parents are younger than 70 during this phase. The later-life phase applies the age-70-and-older flag, requires partial or full retirement and at least one ADL or IADL difficulty, and retains adult children ages 18 and older. The two histories are merged by respondent and child identifiers. The observed childcare and later-life windows do not overlap in the retained sample.

Childcare. Let $\mathcal{T}_{ij}^{C}$ denote the eligible childcare observations for pair (i, j). We define

$$C_{ij} = \max_{t \in \mathcal{T}_{ij}^{C}} c_{ijt}, \tag{A1}$$

where c_{ijt} is the RAND child-specific indicator for grandchild care, stored as `kkdcare` in the construction program. The maximum is taken over nonmissing reports; a history without a valid report remains missing. The underlying RAND measure is subject to the HRS 100-hour childcare screen.[5] Consequently, C_{ij} records whether the particular child's household ever received reported grandchild care; it does not measure care hours, a continuous arrangement, or the share of all childcare supplied by grandparents. The resulting variable is `ever_childcare_kid`.

Later informal care and the no-help category. For each eligible later-life observation, we combine the indicators for ADL and IADL help from the child and the child's spouse. A positive

[5] See the RAND HRS Family Data 2022 documentation, Section 5C, "Kid Uses Parents for Childcare." The education variable is documented in Section 5A.

report on any of these margins counts as care. Help with money management is not added as a separate component. Taking the maximum over valid later-life reports yields `khlp_adliadl_sp`, the main outcome L_{ij}.

The programs also construct `khlp_adliadl_zero_sp`, which assigns zero to the special no-help category before aggregation. Other missing responses are not automatically converted to zero. The main outcome leaves histories without a valid child-specific help report missing, whereas the alternative includes the additional no-help histories. Both versions retain the later-life functional-limitation requirement. Thus, the distinction between them concerns the coding of help, rather than the age or health eligibility criteria.

Spouse-care exclusion. The wave-level construction identifies records in which the main-helper information points to the respondent's spouse or partner, with exceptions for reported child help. It removes flagged records and carries forward a flag into the collapsed histories. The final table programs retain only histories for which `max_flag_spouse_only` is missing. This is the implemented history-level restriction: the stored flag need not be identical across all children of a respondent, so it does not uniformly remove every child of every respondent who has a flagged history.

Analysis samples. The merged file contains 8,454 pairs. Requiring nonmissing childcare and the alternative care indicator, followed by the spouse-care exclusion, leaves the 6,535-pair base sample. The main care outcome is observed for 3,693 pairs. Missing child education accounts for the smaller samples in tables that use education. Table A1 records these steps.

Table A1. Empirical Sample Construction

Sample	Parent–child pairs	Parents
Merged histories	8,454	3,338
Nonmissing childcare and alternative care indicator	6,739	2,947
After spouse-care exclusion: base sample	6,535	2,836
With observed child education	6,469	2,819
Observed main care outcome	3,693	1,481
With observed child education	3,669	1,478

Notes: The last two rows are subsets of the base sample. Parents are distinct respondent identifiers; a respondent can contribute several child observations. Regression samples additionally depend on control availability and, for parent fixed effects, the exclusion of singleton groups. Source: HRS; authors' calculations.

These are samples of linked observed histories, rather than repeated cross-sections of all grand-

parents or all older people. Both care measures are extensive-margin indicators over unbalanced observation windows. Their levels therefore reflect the duration and availability of observed histories as well as care behavior; they should not be interpreted as annual population rates.

A.2 Permanent-Resource Proxies and Fixed Income Groups

Child education. We use `kid_yrseduc`, derived from RAND's `KAEDUC` measure of the child's last reported years of education. The categories are 0–12, 13–15, and 16 or more years. In the base sample with observed education, the respective shares are 49.8, 24.1, and 26.1 percent. These are schooling categories, not observed child-earnings quantiles. The main results do not use the alternative education construction included elsewhere in the code folder.

Parent retirement income. For respondent i in eligible later-life wave t, the household measure is

$$R_{it} = \left(P_{it} + P_{it}^{s} + S_{it} + S_{it}^{s}\right) \frac{\mathrm{CPI}_{2022}}{\mathrm{CPI}_{t}}, \tag{A2}$$

where P denotes pension and annuity income and S denotes Social Security retirement, spouse, or widow benefits; superscript s denotes the respondent's spouse. The code uses interview-year CPI values and $\mathrm{CPI}_{2022} = 292.655$. Components are combined using row sums, so missing components contribute zero to the constructed total. The income proxy $\overline{R}_i$ averages R_{it} over distinct eligible respondent–wave observations, avoiding repeated weighting of a wave merely because several children are observed.

The variable used to form the groups is `HHpension_income`. It excludes the additional SSI and Social Security disability components stored in the alternative income measure and is not divided by 1.6 for couples. The table program copies this household total into a variable whose name ends in `_eqv`; that name does not describe the measure actually used. Because the average is constructed from later-life observations, it should not be described as income measured entirely before care outcomes.

Group boundaries. The percentile boundaries are chosen to approximately match the shares of the three child-education groups. Rounding the first group's share and the cumulative share of the first two groups to whole percentages gives 50 and 74. Let q_{50} and q_{74} be the corresponding empirical quantiles of $\overline{R}_i$ in the 6,535-pair base sample. Each pair receives equal weight in calculating these quantiles. We define

$$G_i = \begin{cases} 1, & \overline{R}_i < q_{50}, \\ 2, & q_{50} \le \overline{R}_i < q_{74}, \\ 3, & \overline{R}_i \ge q_{74}. \end{cases} \tag{A3}$$

The cutoffs are approximately $25,572.47 and $38,896.05 in 2022 dollars; classification uses their full precision. These boundaries are held fixed when restricting to observed care outcomes or observed child education. Consequently, income-group shares need not remain 50, 24, and 26 percent in those subsets.

For comparison, re-estimating the quantiles within the smaller main-care sample gives cutoffs of approximately $21,709.96 and $33,111.83. Under that alternative classification, the care shares are 0.291, 0.306, and 0.260, with standard errors 0.011, 0.015, and 0.014 and sample sizes 1,845, 887, and 961. These statistics describe different group memberships. Section 3 uses the fixed base-sample boundaries so that its marginal and joint tables refer to the same parent-income groups.

A.3 Moment Construction and Sampling Uncertainty

For a binary outcome D_{ij} and an eligible group g, the reported share and descriptive standard error are

$$\widehat{p}_g = \frac{1}{n_g} \sum_{(i,j)\in g} D_{ij}, \qquad \widehat{\mathrm{se}}(\widehat{p}_g) = \sqrt{\frac{\widehat{p}_g(1-\widehat{p}_g)}{n_g - 1}}. \tag{A4}$$

The denominator n_g includes only pairs with a nonmissing outcome and the grouping variables required by that table. Childcare moments use the base sample; main LTC and reciprocity moments use pairs with observed L_{ij}. Education-based moments additionally require observed schooling. For example, the low-income cells in the joint LTC table sum to 2,201 pairs, compared with 2,222 in the parent-income marginal, because 21 low-income pairs lack child education.

For reciprocity, means are calculated separately by childcare status and differenced using unrounded values:

$$\widehat{\Delta}_g = \widehat{p}_{g,1} - \widehat{p}_{g,0}, \qquad \widehat{\mathrm{se}}(\widehat{\Delta}_g) = \sqrt{\widehat{\mathrm{se}}(\widehat{p}_{g,1})^2 + \widehat{\mathrm{se}}(\widehat{p}_{g,0})^2}. \tag{A5}$$

Independent rounding of means and differences can produce a discrepancy of 0.001 when subtracting displayed entries. The descriptive standard errors reproduce the programs' convention of treating pairs as independent. They are not survey-design-adjusted or parent-clustered tests of differences across resource groups. All regression standard errors, by contrast, are clustered by parent.

The treatment of the no-help category materially affects levels. Under the alternative outcome, care receipt is 0.133 among the 3,965 pairs without earlier childcare and 0.207 among the 2,570 pairs with earlier childcare; the corresponding descriptive standard errors are 0.005 and 0.008. The unadjusted gap is 0.074. With controls and parent fixed effects, the estimated association is 0.119 with a clustered standard error of 0.016. Thus, the positive within-parent association also appears when the additional no-help histories are included, although this specification describes a broader

sample.

A.4 Regression Controls and Instrument Construction

Controls and fixed effects. Parent controls are gender, college education, a Black-race indicator, mean household income, indicators for any ADL and IADL difficulty, mean self-reported health, and an indicator for being married or partnered during the later-life period. Child controls are gender, an indicator for at least 16 years of schooling, birth-order indicators, mean number of children, and indicators for ever being married and ever working full time during the later-life period. The birth-order variable ranks observed children by birth year within each parent's retained base-sample records. It is not necessarily the child's rank among all children ever born to that parent. Several controls summarize the later-life period and should not be interpreted as uniformly predetermined characteristics.

Parent fixed effects absorb parent characteristics common across the included children. The estimator drops parents with only one usable child observation. The fixed-effects regressions therefore contain 3,229 pairs without child controls and 2,817 with them, compared with 3,693 and 3,324 in the corresponding pooled regressions.

Sibling-fertility instrument. Let $\mathcal{J}_i^C$ be the set of distinct children represented in parent i's eligible childcare records, after the age, health, and positive-fertility restrictions. For each $k \in \mathcal{J}_i^C$, let n_{ik}^C be the maximum reported number of that child's own children over those records. With $J_i^C = |\mathcal{J}_i^C|$, the instrument is

$$Z_{ij} = \begin{cases} \dfrac{\sum_{k \in \mathcal{J}_i^C,\, k \neq j} n_{ik}^C}{J_i^C - 1}, & J_i^C > 1, \\ 0, & J_i^C = 1. \end{cases} \tag{A6}$$

It is stored as `mean_nr_kids_sibsCC`. This is a leave-one-out average of siblings' maximum observed fertility within the eligible childcare sample, rather than an average over all living siblings or a wave-by-wave mean. The zero assigned when $J_i^C = 1$ is a coding convention; it does not establish that the focal child has no siblings outside that sample.

Columns (5)–(6) of Table 2 estimate two-stage least squares, respectively without controls and with the full parent and child control sets. Table A2 reports first stages on the exact corresponding IV samples. The excluded-instrument F-statistic is the squared parent-clustered t-statistic on Z_{ij}. Instrument relevance does not establish the exclusion restriction: siblings' fertility may also affect their availability for later care or correlate with unobserved family characteristics.

Table A2. First-Stage Regressions for the Sibling-Fertility Instrument

	Without controls	With controls
Sibling-fertility instrument	-0.0256	-0.0334
	(0.0072)	(0.0082)
Excluded-instrument F	12.6	16.5
Observations	3,693	3,324

Notes: The table regresses earlier childcare on the instrument and the indicated controls using the corresponding IV outcome sample. Standard errors are clustered by parent. The excluded-instrument F-statistic is the square of the clustered t-statistic. Source: HRS; authors' calculations.

B Grandparental Care as a Regular Childcare Arrangement

The HRS histories in Section 3 establish the prevalence of care and its connection across generations. Here we examine the role of grandparents within young children's regular childcare arrangements. Evidence from the National Household Education Survey (NHES) shows that grandparents commonly supply the majority of nonparental care hours, often without purchased care alongside them. These patterns hold across young-household income groups. Additionally, evidence from the Survey of Income and Program Participation (SIPP) also identifies grandparents as the primary regular provider for a majority of preschool grandparent-care users with employed mothers.

Data and measurement. We use the 2019 NHES Early Childhood Program Participation public-use file, which covers children from birth through age six who are not yet in kindergarten.[6] The survey records usual weekly hours in regularly scheduled relative, nonrelative, and center-based care. Our sample consists of 1,226 children whose main regular relative caregiver is a grandparent and whose grandparent arrangement charges no fee to the household or another payer. Selection identifies the main provider *among relatives*; it does not require grandparents to supply the most care across all arrangements. Whether they provide a majority of nonparental hours is therefore an outcome of the analysis.

Nonparental hours include the main and additional arrangements within each of the three care categories. They exclude parents' own care. Purchased outside care denotes a reported positive household payment for the main center or nonrelative arrangement; nonrelative care includes home-based providers and need not be licensed formal care. All estimates use the child-level survey weights and official imputations. Standard errors use the 80 replicate weights and the delete-group

[6]National Center for Education Statistics, 2019 NHES public-use data, Early Childhood Program Participation; public-use codebook; and NHES:2019 Data File User's Manual. All NHES estimates in this appendix are authors' calculations from the public-use microdata.

jackknife procedure described in the survey documentation.[7] These are child-level comparisons of regular care, rather than measures of total parenting time or estimates conditional merely on having a grandparent available.

The contribution of grandparents. Table B1 shows that unpaid grandparents provide 21.39 hours of care per week on average. They supply more than half of all regular nonparental hours for 69.14 percent of sampled children and account for 64.31 percent of nonparental hours across the sample. Thus, grandparental help is the dominant nonparental input for a substantial majority of these children. Its relevance to parents' time constraints is also reflected in respondents' assessments: among those for whom the work-hours question applies, 83.95 percent report that the grandparent arrangement covers the hours needed for work well or very well.[8]

The prevalence of grandparental care without purchased outside services is equally important. Almost seven in ten children use no reported household-paid center or nonrelative care, and these children receive 77.12 percent of all grandparent care hours in the sample. The majority finding is not driven by treating missing fees as zero. Detailed fees are observed for the main provider in each category, but not for every additional arrangement. Even when we count payments by any payer and classify *all* additional relative, center, and nonrelative arrangements with unknown fees as fee charging, 57.81 percent of children remain in arrangements without any fee-charging regular care. This conservative group accounts for 64.22 percent of grandparent hours. Restricting the sample to ages zero through five leaves the headline comparisons essentially unchanged.[9]

Differences across young households. Table B2 groups children by their own household's current annual income. These categories differ from the grandparent permanent-income groups used in the HRS analysis. Mean weekly grandparent care is 23.65, 21.22, and 19.8 hours across the three groups; the corresponding shares receiving more than half of nonparental hours from the grandparent are 78.36, 73.81, and 58.11 percent. Grandparents therefore supply both substantial hours and the majority of nonparental care for most users in every group. Reliance on grandparents without purchased outside care is particularly pronounced in the lower-income group, where 84.79 percent report no household-paid center or nonrelative care. Purchased care alongside grandparental care

[7] Grandparent care is identified by RCWEEK=1 and RCTYPE=1; no fee is RCFEE=2. Grandparent hours are RCHRS. Total nonparental hours, CAREHOURX, sum the main and additional relative, nonrelative, and center arrangements. Positive household payment is identified by CPCOST¿0 with weekly center care, or NCCOST¿0 with weekly nonrelative care. The weight is FEWT, with replicates FEWT1–FEWT80.

[8] This calculation uses the 1,150 observations with an applicable answer to RCCVRWK; responses "well" and "very well" are combined.

[9] For ages zero through five, mean grandparent hours are 21.38 and the share receiving more than half of nonparental hours from the grandparent is 69.17 percent.

Table B1. The Role of Unpaid Grandparents in Regular Childcare

	Estimate	Std. error
Panel A. Grandparents' contribution		
Usual grandparent care (hours per week)	21.39	0.56
Children receiving more than half of nonparental hours from the grandparent (%)	69.14	1.78
Grandparent share of pooled nonparental hours (%)	64.31	1.83
Panel B. Care without purchased outside services		
Children with no reported household-paid center or nonrelative care (%)	69.77	1.58
Share of grandparent hours received by these children (%)	77.12	1.67
Children with no regular center or nonrelative care, paid or free (%)	60.15	1.9
Conservative lower bound: children with no fee-charging regular care of any type (%)	57.81	1.79
Children in the sample	1,226	

Notes: The sample contains children not yet in kindergarten whose main regular relative caregiver is an unpaid grandparent. Estimates use child-level survey weights; standard errors use the 80 jackknife replicate weights. Nonparental hours include all reported regular arrangements and exclude parents' own care. The pooled share divides total weighted grandparent hours by total weighted nonparental hours. Household-paid outside care refers to the main center and nonrelative providers. The hours share in Panel B weights children by their grandparent care hours; it is not the grandparent share within mixed arrangements. The conservative bound counts fees paid by any payer and treats every additional relative, center, and nonrelative arrangement with unobserved fees as fee charging. Hours measure care received by children, not unique grandparent-person hours. Source: NHES 2019; authors' calculations.

becomes more common at higher incomes, but grandparental care remains a major input even in households with income above $100,000.

Independent evidence from SIPP. The Census Bureau's 2011 childcare tables provide a separate comparison for children under five living with employed mothers. They report 3.45 million children receiving regular grandparent care and 2.36 million for whom grandparents are the primary or joint-primary arrangement. Using the published counts in thousands, the latter represent 68.39 percent of grandparent-care users. Primary care is the arrangement supplying the most weekly hours, with ties included. The accompanying report records 23 weekly hours of grandparent care among users with employed mothers.[10] Although the surveys differ in their populations and definitions, both identify grandparents as a principal regular care provider for a majority of their users.

[10]U.S. Census Bureau, *Who's Minding the Kids? Child Care Arrangements: 2011—Detailed Tables*, Tables 1A and 2A. Both counts use children under five with employed mothers: 3,445,000 receive grandparent care and 2,356,000 have a primary or joint-primary grandparent arrangement. The primary share is calculated as $2{,}356/3{,}445$. See also Lynda Laughlin (2013), *Who's Minding the Kids? Child Care Arrangements: Spring 2011*, Figure 2. SIPP records parental care during the reference parent's work or school time, rather than all parenting throughout the week.

Table B2. Reliance on Grandparental Care by Young-Household Income

	Annual income of the child's household		
	$50,000 or less	$50,001–$100,000	Above $100,000
Usual grandparent care (hours per week)	23.65	21.22	19.8
	(1.16)	(0.91)	(0.74)
Grandparent provides more than half of nonparental hours (%)	78.36	73.81	58.11
	(2.66)	(3.14)	(3.21)
No reported household-paid center or nonrelative care (%)	84.79	72.79	55.72
	(2.77)	(2.62)	(2.74)
Children	328	414	484

Notes: The sample and care definitions are the same as in Table B1. Income is current annual income of the child's household, in nominal dollars, grouped from the survey's reported income categories. Estimates use child-level survey weights. Jackknife standard errors are in parentheses. Source: NHES 2019; authors' calculations.

Implications for the family-care margin. The evidence establishes that grandparental childcare often supplies a substantial block of regular care, with most users relying on grandparents for the majority of nonparental hours and many purchasing no outside care. It gives empirical content to the model's comparison between a family agreement that provides childcare and alternatives using the young adult's own time or purchased services. A subsidy to purchased care changes the value of those alternatives. Under limited commitment, it can consequently affect the resources allocated to the grandparent within an ongoing agreement, as well as whether the family cooperates. The welfare mechanism thus operates through the terms of resource sharing in addition to changes in care provision. The prevalence and intensity of grandparental care documented here motivate accounting for that intergenerational margin when evaluating childcare policy.

C Calibration

C.1 Construction of wealth ratios

The two wealth targets in our structural estimation use the 2022 Survey of Consumer Finances (SCF) summary extract.[11] The unit of observation is the SCF family, and age and education refer to

[11] Board of Governors of the Federal Reserve System, 2022 Survey of Consumer Finances, summary extract and documentation.

its reference person. We measure wealth using NETWORTH, total financial and nonfinancial assets less liabilities, expressed in 2022 dollars. This measure includes housing and business wealth.

We compare families whose reference person is aged 75–90 with those aged 45–60. Both endpoints are included, following the underlying calculation. The first target includes all education groups. The second restricts each age group to college graduates, defined by EDCL=4. The remaining education categories are less than high school (EDCL=1) and high school or some college (EDCL=2 or 3).

The SCF provides five imputations, or implicates, of missing information. For age band A, education group E, and implicate m, let $I_{im}(A, E)$ indicate membership in the relevant cell. Using the final analysis weight $w_{im} = \texttt{WGT}_{im}$, we calculate

$$\widehat{W}_{A,E}^{(m)} = \frac{\sum_i w_{im} I_{im}(A,E)\, \texttt{NETWORTH}_{im}}{\sum_i w_{im} I_{im}(A,E)}, \qquad \widehat{W}_{A,E} = \frac{1}{5}\sum_{m=1}^{5} \widehat{W}_{A,E}^{(m)}.$$

The all-education mean is computed directly from all eligible families, rather than as an unweighted average of education-cell means. After combining the implicate-specific means, we construct

$$\widehat{R}_E = \frac{\widehat{W}_{[75,90],E}}{\widehat{W}_{[45,60],E}}.$$

Thus, the targets are ratios of combined means, rather than averages of implicate-specific ratios, family-level wealth ratios, or ratios of medians.

Sample	Ages 45–60	Ages 75–90	Older/younger
All education levels	1,177,804	1,599,411	1.35796
College graduates	2,236,291	3,334,992	1.49130

Mean net worth is reported in dollars and rounded for presentation; ratios use unrounded means. Rounding the ratios to two decimal places gives the calibration targets 1.36 and 1.49.

These cross-sectional comparisons approximate the model's stage-2 parent-to-child wealth ratios. The college-graduate ratio serves as a proxy for the high-productivity parent–child group. The SCF calculation does not identify linked parents and adult children or condition jointly on both generations' education. In the model, the corresponding numerator is parental wealth after asset destruction, $\tilde{a}_2^o = (1-\delta)a_2^o$, and the denominator is the adult child's wealth, a_2^y. The empirical and model moments therefore share a ratio-of-means construction, while their conditioning populations are approximate counterparts.

C.2 Construction of the paid-childcare share

We construct an approximate family-level target for $\Pr(d = 1 \mid N_1)$ from three arrangements: market/provider care (P), unpaid relative care (R), and self-provided or residual care (S). Families without external care remain in the denominator. Among nonparental arrangements, primary care is determined by weekly hours, excluding regular school and standalone extracurricular activities. Ties favor market/provider care, followed by unpaid relatives. Paid relatives enter P; unpaid nonrelatives and older-child self-care enter S. Formal programs enter P even when fully subsidized, making this category broader than positive household payment.[12]

For ages 0–4, the 2019 NHES Early Childhood Program Participation survey gives primary-care shares of 30.104% in centers, 17.799% with relatives, 8.399% with nonrelatives, and 1.795% in tied arrangements. Fee-paying shares are 25.4153% among relative-primary and 92.0428% among nonrelative-primary arrangements; 91.5537% of ties enter P. Hence

$$p_{0-4} = .30104 + .17799(.254153) + .08399(.920428) + .01795(.915537) \simeq .44002,$$
$$r_{0-4} = .17799(1 - .254153) + .01795(1 - .915537) \simeq .13427.$$

For school-age children, the original synthesis underlying the benchmark uses the 2005 NHES After-School Programs and Activities (ASPA) age profile. Its (p, r) inputs are $(.2675, .1172)$ for ages 5–11 and $(.1431, .0957)$ for ages 12–15. Both are multiplied by $s = .28/.3663 = .76440$, anchoring the original ages-5–12 external-care rate to the 2019 NSECE's 28% rate. The latter covers ages 6–12 and five-year-old kindergarteners using regular nonparental care outside school for at least five hours weekly.

We approximate each family's arrangement by that of its youngest child. The 2019 CPS ASEC gives 12.754, 12.011, and 6.224 million primary families with youngest own child aged 0–4, 5–11, and 12–15, respectively. The associated weights and care shares are:

Youngest child's age	Family weight	p_a	r_a	$1 - p_a - r_a$
0–4	.41157	.44002	.13427	.42571
5–11	.38759	.20448	.08959	.70594
12–15	.20084	.10939	.07315	.81746

Aggregating gives $P = \sum_a \omega_a p_a = .28232$ and $R = \sum_a \omega_a r_a = .10468$. Mapping N_1 to the complement of primary unpaid relative care yields

$$\widehat{m}_{cc} = \frac{P}{1 - R} = \frac{.28232}{1 - .10468} = .3153.$$

[12] Sources: NHES public-use data, ECPP 2019 and ASPA 2005; NSECE, OPRE Report 2024-081; and 2019 CPS ASEC family data.

C.3 Paid-childcare expenditure

The expenditure moment is mean family childcare expenditure conditional on purchasing care, divided by reference earnings:

$$m_K = \frac{\mathbb{E}[K \mid d = 1, N_1]}{\bar{y}}.$$

The numerator excludes families choosing self-provided care. The denominator is reference average earnings, rather than the earnings of paid-care users. Accordingly, this moment differs from the average of household-specific expenditure-to-income ratios. Since K covers the family, per-child expenditures require an adjustment for the number of children.

A published tabulation of the 2020 CPS Annual Social and Economic Supplement (ASEC), referring to calendar year 2019, reports annual childcare expenditure of $7,602 per household making payments. The measure covers care for children under age 15 while household members work. Census Table PINC-05 reports mean annual individual earnings of $57,765 among people with earnings for the same reference year.[13] Using average individual earnings as the empirical normalization gives

$$\widehat{m}_K^{\mathrm{CPS}} = \frac{7,602}{57,765} = .1316.$$

Thus, this comparison expresses a household childcare bill relative to average individual earnings.

To express both flows over one model period, let $\bar{k}_t$ denote the representative annual childcare bill and $\bar{e}_t$ reference annual earnings. The comparable fifteen-year ratio is

$$\frac{\sum_{t=1}^{15} \bar{k}_t}{\sum_{t=1}^{15} \bar{e}_t}.$$

Treating the annual cross-sectional means as constant equivalents throughout the period gives

$$\frac{15 \times 7,602}{15 \times 57,765} = \frac{114,030}{866,475} = .1316.$$

This conversion assumes that the observed annual bill represents an average year of the model's paid-care regime. It does not identify a longitudinal expenditure profile or account for changes in care use as children age.

Survey definitions also differ across the two childcare targets. The participation construction includes formal care with no household payment, whereas the CPS expenditure mean conditions on positive payments. Their mapping into a single paid-care regime is therefore approximate. Finally, the expenditure target applies to costs selected by purchasers, rather than the unconditional mean offer $\mu_{cc} = \mathbb{E}[K]$. The latter is determined jointly with offer dispersion to match participation and conditional expenditure.

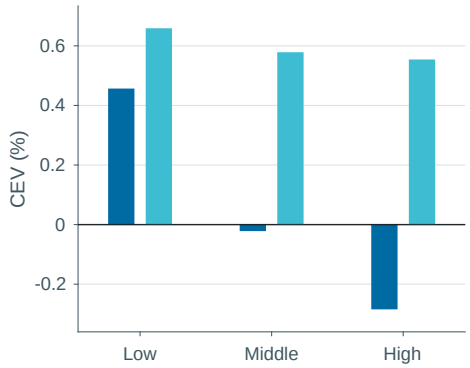
[13] Childcare expenditure: Committee for Economic Development of The Conference Board (2022), *The Economic Role of Paid Child Care in the U.S.*, Part 1. Earnings: U.S. Census Bureau, 2020 CPS Table PINC-05, work experience and earnings in 2019, total population with earnings.

D Policy welfare and behavior by productivity

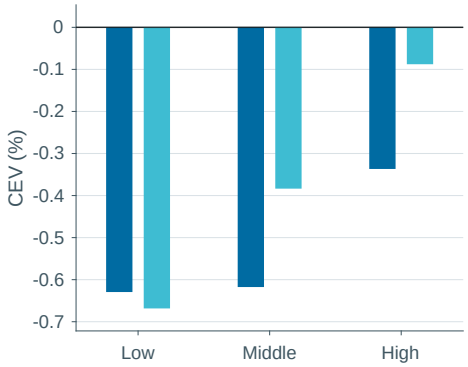
Figure D1 reports the role-specific welfare effects behind the aggregate incidence results in Section 8. The young-adult panels condition in turn on the adult child's and the incumbent parent's productivity. The parent panels measure the incumbent parent's own remaining-lifetime CEV, using the same fixed benchmark parent population as Figure 9. These are marginal productivity averages: a negative group average need not imply that every joint parent–child cell loses.

Figure D1. Role-specific welfare incidence by child and parent productivity

(a) Young-adult welfare, by child productivity



(b) Incumbent-parent welfare, by parent productivity



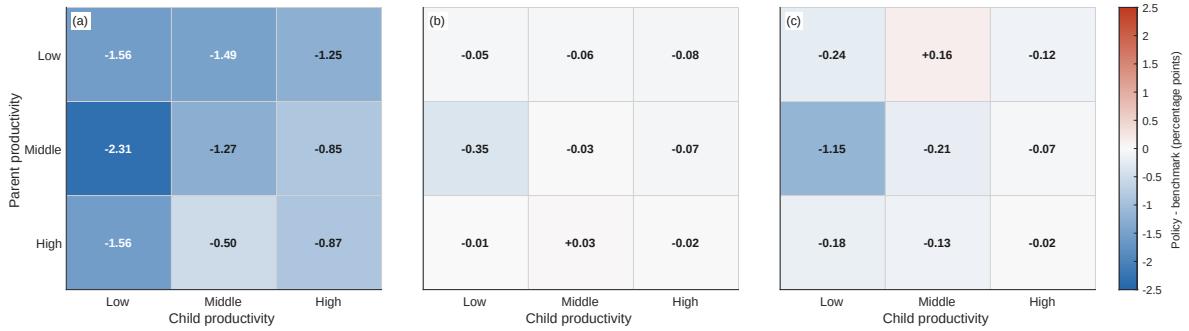
Panels report the young adult's CEV by child productivity and the incumbent parent's CEV by own productivity. Each panel compares Tax-financed (GE) and Zero-tax (PE). Utilities and benchmark consumption–leisure references are aggregated with benchmark group weights before the CEV inversion. The upper-left panel repeats Figure 8.

The following heatmaps show how the paid-childcare voucher changes family behavior across all nine productivity pairs. The cooperation outcomes are available by productivity pair for the stationary GE comparison. The bargaining comparison uses the fixed incumbent-parent cohort and reports the parent's conditional weight in both C_1 and positive-need C_2 agreements under GE and PE.

E Additional bargaining-weight responses

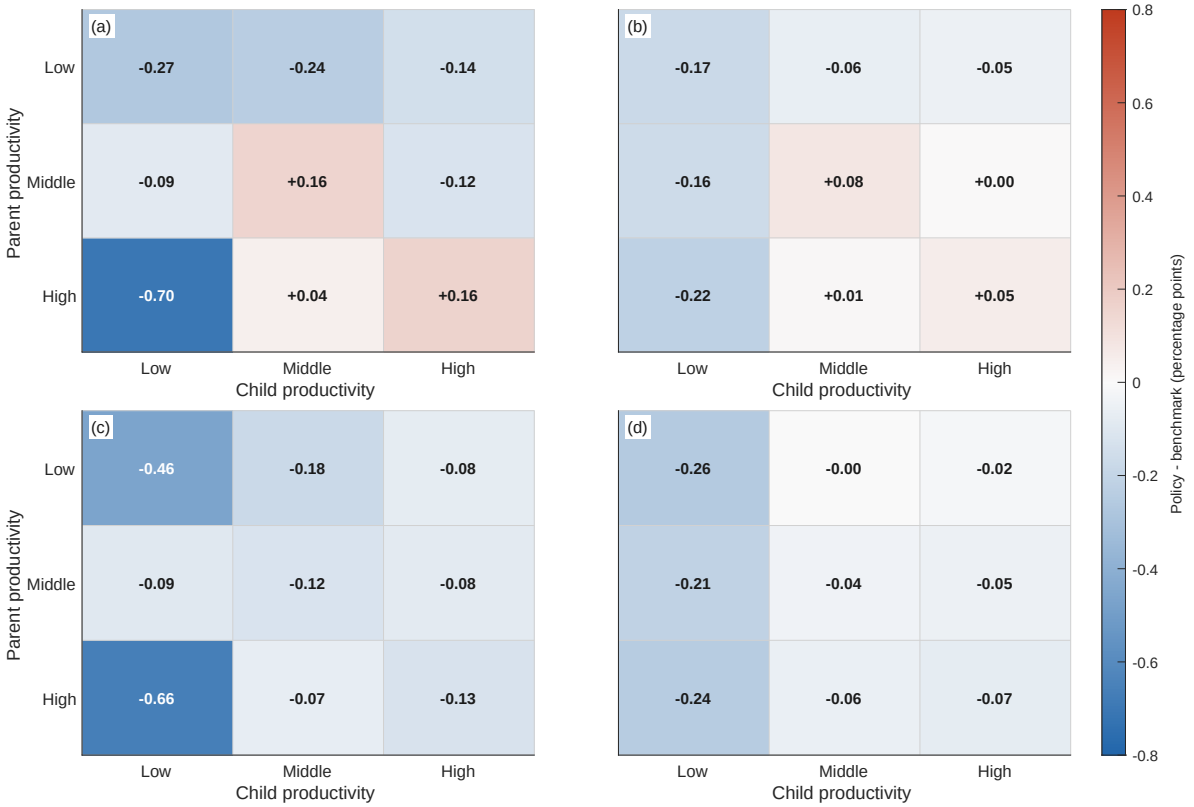
Figures E1, E2, E3, and E4 report the three bargaining-weight maps omitted from each main-text mechanism figure. For the period-1 experiments, the main text reports the parent's weight conditional on C_1; for the period-2 experiments, it reports the parent's weight conditional on C_2 and positive LTC need. The corresponding parent and child effects at a given date are exact negatives

Figure D2. Paid-childcare policy effects on family cooperation by productivity pair



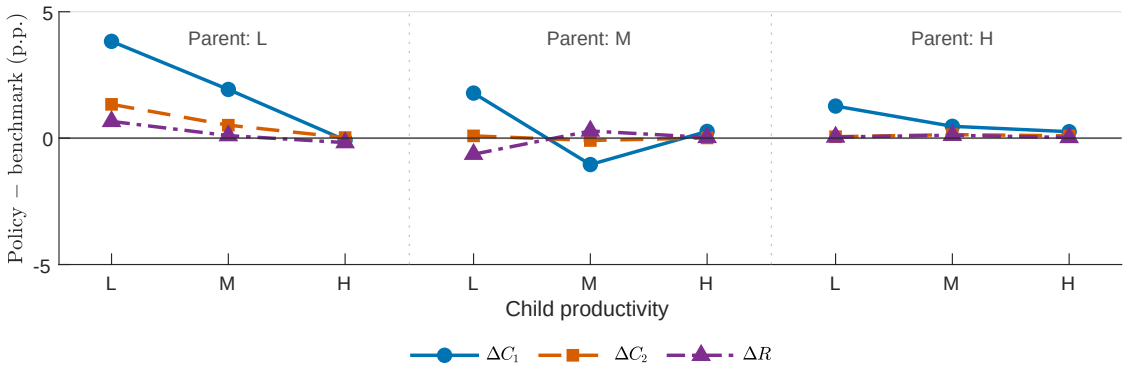
Each cell is the GE policy outcome minus the benchmark outcome, in percentage points. Rows index parent productivity and columns index child productivity, from low to high. The C_1 panel uses all period-1 families; the C_2 and reciprocity panels condition on positive LTC need. The reciprocity gap is $R = \Pr(C_2 \mid C_1, \lambda > 0) - \Pr(C_2 \mid N_1, \lambda > 0)$. These are stationary-equilibrium comparisons and combine behavioral, selection, and distributional responses.

Figure D3. Paid-childcare policy effects on the incumbent parent's bargaining weight by productivity pair



Each cell is the policy outcome minus the benchmark outcome, in percentage points. Panels (a) and (b) report GE changes in the parent's average realized Pareto weight in C_1 and C_2, respectively; panels (c) and (d) report the corresponding PE changes. The C_1 outcome is $1 - \theta_1$ conditional on a continuing C_1 agreement. The C_2 outcome is $1 - \theta_2$ conditional on positive LTC need and a continuing C_2 agreement. All four panels use the fixed benchmark incumbent-parent cohort. Rows index parent productivity and columns index child productivity. Because the set of cooperating families is endogenous, these conditional changes combine renegotiation with selection and are not fixed-family treatment effects.

Figure D4. Broader childcare subsidy: changes in family cooperation by productivity pair (GE)



The figure reports the tax-financed (GE) comparison. Lines show policy-minus-benchmark changes, in percentage points, in childcare cooperation (C_1), informal LTC (C_2), and the history gap (R). The three groups correspond to low-, middle-, and high-productivity parents. Within each group, child productivity varies from low (L) to middle (M) to high (H).

because their Pareto weights sum to one. All conditional means nevertheless reflect endogenous selection into cooperation and, at date 2, renegotiation.

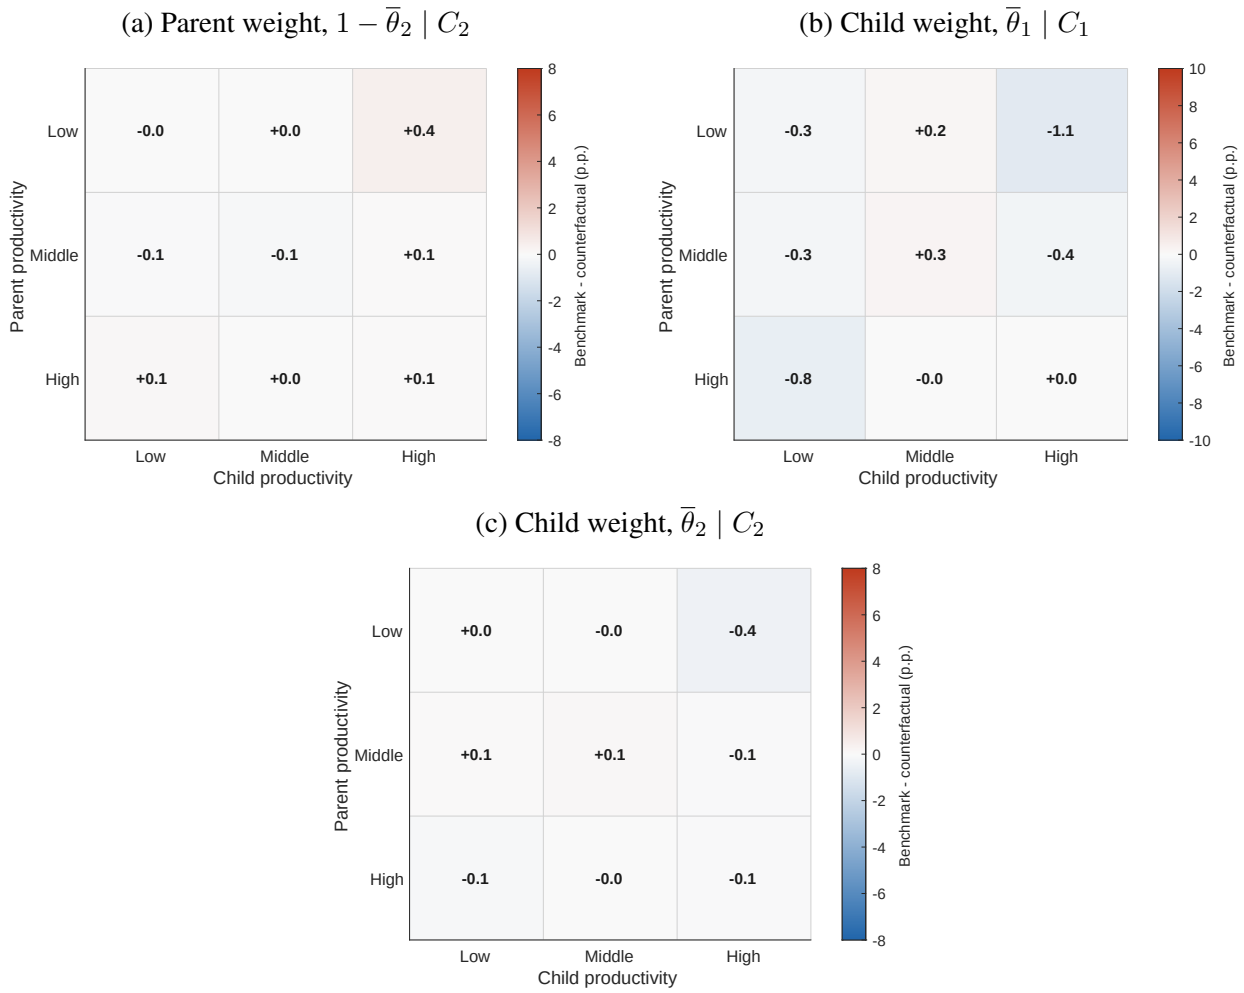
F Resource pooling and cooperative support eligibility

This cross-cutting experiment imposes separate current budgets during cooperation at both dates. One generation can no longer finance the other's current consumption or saving, although care, limited commitment, bargaining, inheritance, and each person's saving decision remain. In this implementation, an eligible parent can also receive consumption-floor support while cooperating under separate budgets; under pooled resources, the floor applies only in non-cooperation. The comparison therefore combines the removal of private resource pooling with a change in public-support eligibility.

Relative to this separate-budget specification, the pooled benchmark raises aggregate C_1 by 7.588 percentage points. The effect is positive for each child-productivity group, at 7.131, 8.493, and 6.075 points from low to high. Across the nine productivity pairs, the effects range from -0.044 to 11.316 points. Flexible resource allocation can therefore expand the set of first-period states in which both generations' participation constraints can be satisfied, although the response is essentially zero in one cell.

The second-period comparison has the opposite aggregate sign. The pooled benchmark lowers aggregate C_2 by 4.915 points; the effects are -2.720 points for low-productivity children and

Figure E1. Additional bargaining-weight responses: productivity dependence of period-1 childcare utility



Each cell is 100 times the benchmark minus the productivity-neutral counterfactual, in percentage points. Rows index parent productivity and columns index child productivity. The C_1 panel conditions on C_1 over all future LTC states. The C_2 panels condition on positive LTC need and C_2. The parent-weight C_1 response appears in Figure 4.

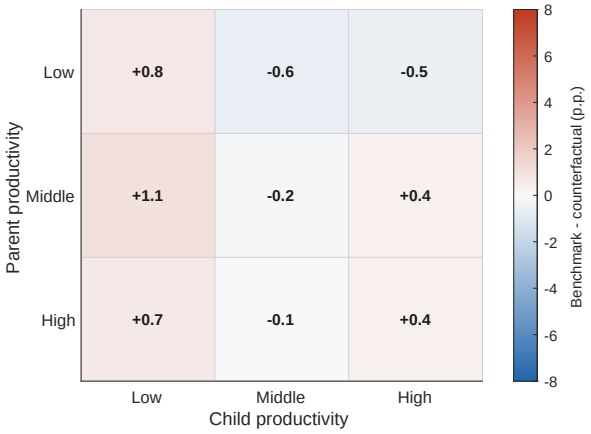
-5.039 points for high-productivity children. The nine cell effects range from -9.845 to $+2.513$ points, so this aggregate decline is not uniform. Pooling raises the aggregate history gap R by 0.844 point, but the child-productivity effects range from -0.886 to $+2.266$ points. Figure F1 reports the full productivity-pair responses.

Pooling also lowers the aggregate conditional parent Pareto weight by 3.255 points within C_1 and by 4.052 points within positive-need C_2 agreements. The C_1 response is heterogeneous across productivity pairs, whereas the parent-weight response within C_2 is negative in all nine cells. Figure F2 reports both parent weights and their child-weight complements. Because these are means among families that cooperate under each regime, they combine bargaining and selection and are not changes for a fixed set of families.

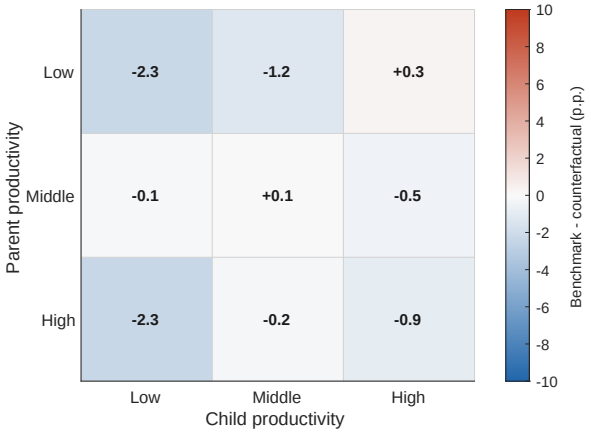
These results should not be read as a causal marginal effect of allowing a transfer at a fixed state. Separate budgets change saving and asset ownership, both generations' subsequent disagreement

Figure E2. Additional bargaining-weight responses: the cost of paid childcare

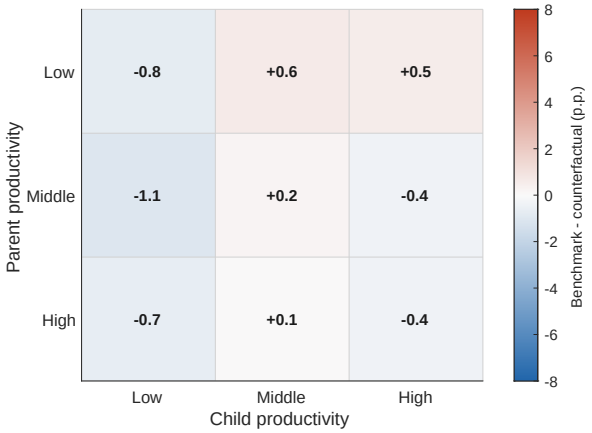
(a) Parent weight, $1 - \overline{\theta}_2 \mid C_2$



(b) Child weight, $\overline{\theta}_1 \mid C_1$



(c) Child weight, $\overline{\theta}_2 \mid C_2$

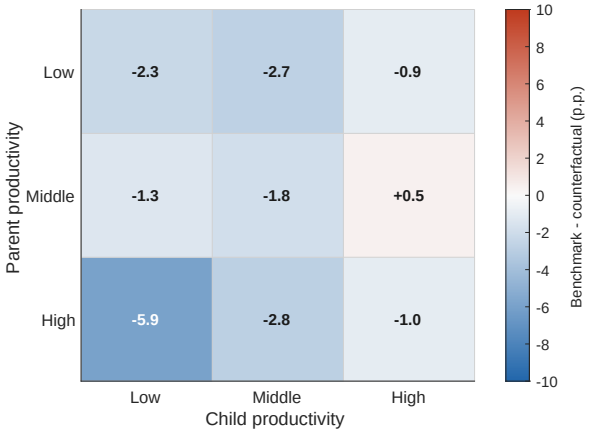


Each cell is 100 times the benchmark minus the zero-price paid-childcare counterfactual, in percentage points. Rows index parent productivity and columns index child productivity. The C_1 panel conditions on C_1 over all future LTC states. The C_2 panels condition on positive LTC need and C_2. The parent-weight C_1 response appears in Figure 5.

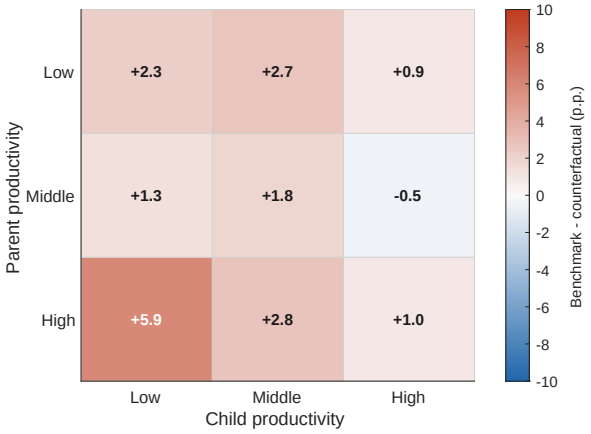
values, and the stationary distribution. Extending the consumption floor to qualifying cooperative parents can also raise C_2 directly in the counterfactual. The experiment does not separately identify that public-support channel from private resource pooling.

Figure E3. Additional bargaining-weight responses: history-dependent guilt

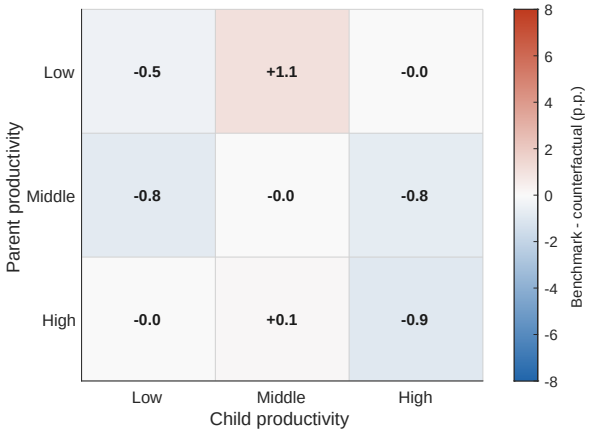
(a) Parent weight, $1 - \overline{\theta}_1 \mid C_1$



(b) Child weight, $\overline{\theta}_1 \mid C_1$



(c) Child weight, $\overline{\theta}_2 \mid C_2$

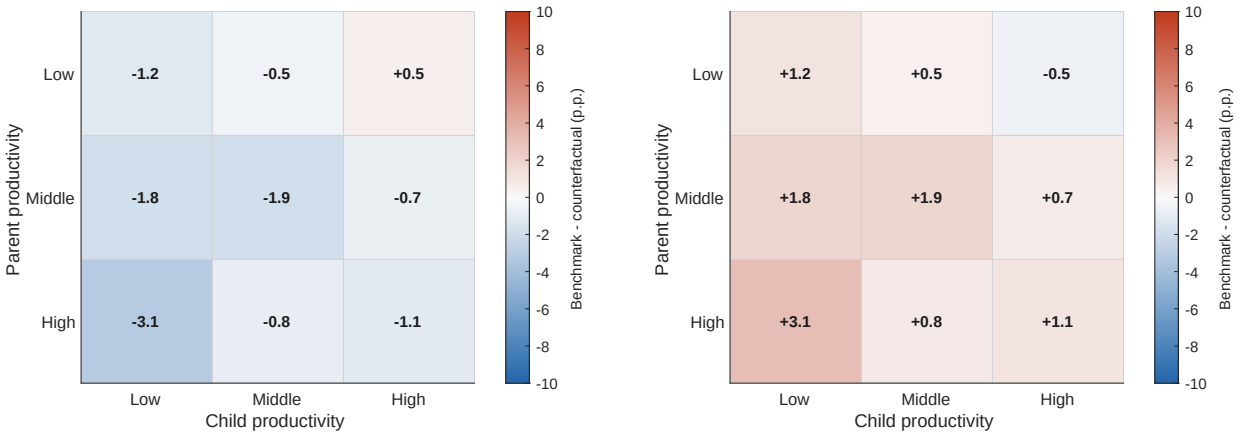


Each cell is 100 times the benchmark minus the counterfactual without the history-dependent guilt component, in percentage points. Rows index parent productivity and columns index child productivity. The C_1 panels condition on C_1 over all future LTC states. The C_2 panel conditions on positive LTC need and C_2. The parent-weight C_2 response appears in Figure 6.

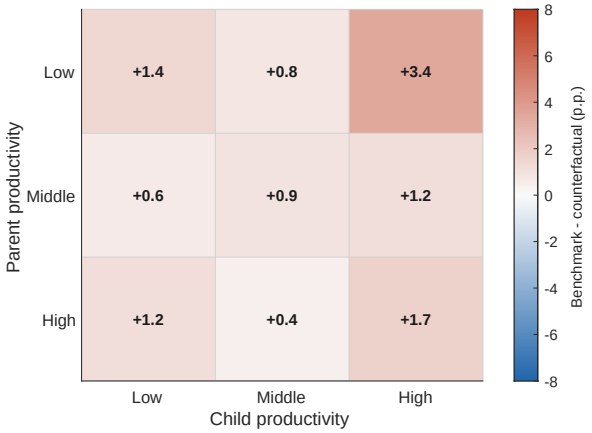
Figure E4. Additional bargaining-weight responses: formal-LTC cost and nursing-home risk

(a) Parent weight, $1 - \overline{\theta}_1 \mid C_1$

(b) Child weight, $\overline{\theta}_1 \mid C_1$



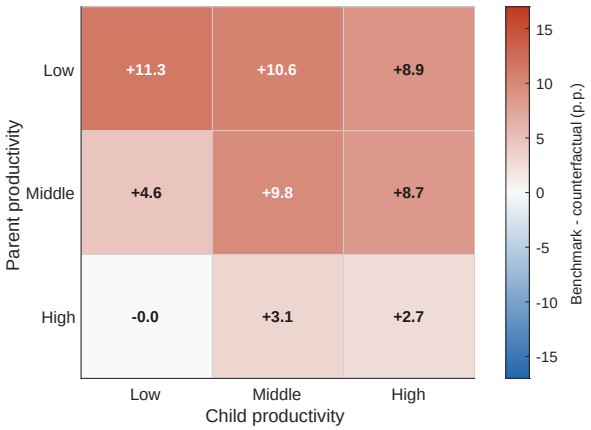
(c) Child weight, $\overline{\theta}_2 \mid C_2$



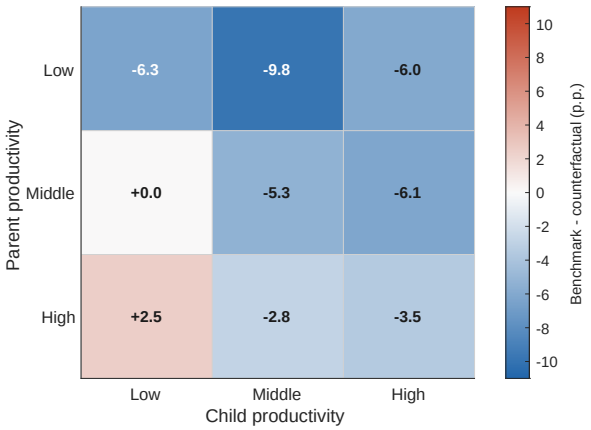
Each cell is 100 times the benchmark minus the counterfactual with free formal LTC and equalized nursing-home risk, in percentage points. Rows index parent productivity and columns index child productivity. The C_1 panels condition on C_1 over all future LTC states. The C_2 panel conditions on positive LTC need and C_2. The parent-weight C_2 response appears in Figure 7.

Figure F1. Resource pooling and cooperation by productivity pair

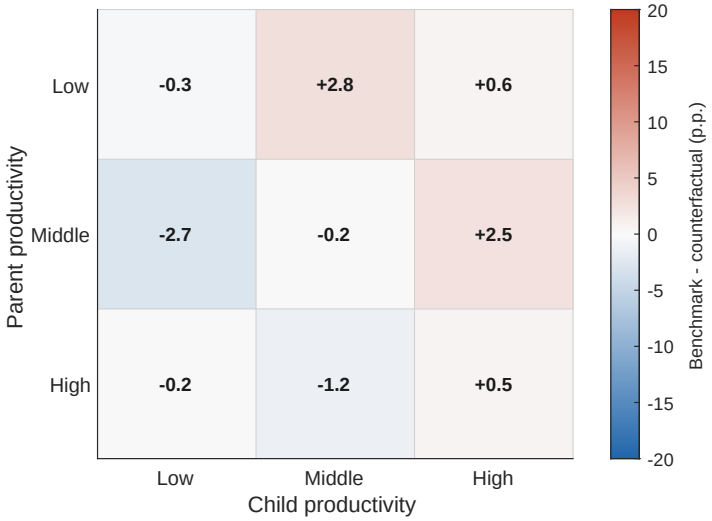
(a) Period-1 cooperation, C_1



(b) Period-2 cooperation, C_2

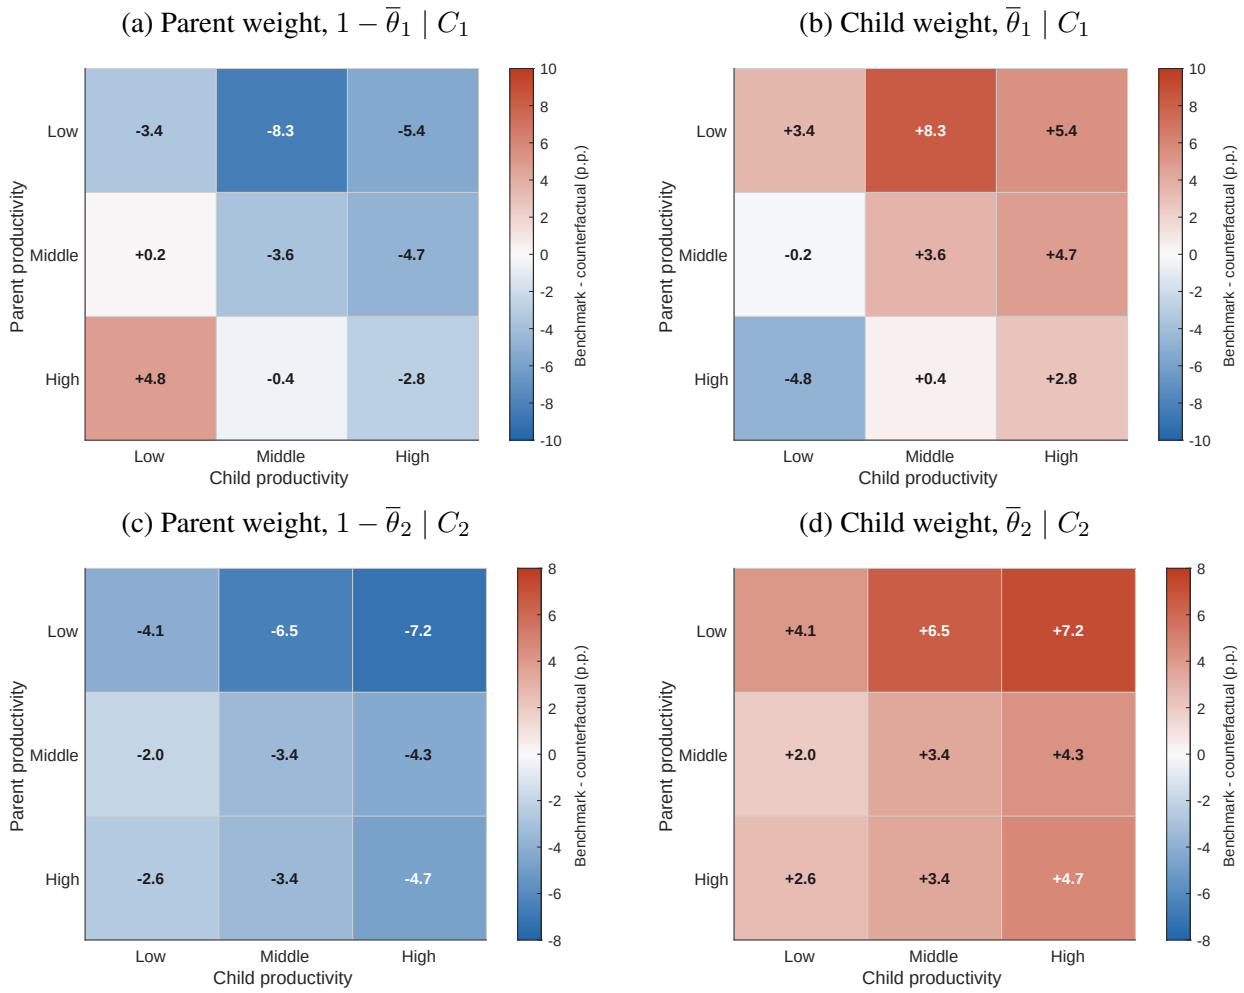


(c) History gap, R



Each cell is 100 times the pooled benchmark minus the separate-budget counterfactual, in percentage points. Rows index parent productivity and columns index child productivity, from low to high. Panel (a) uses all period-1 families. Panels (b) and (c) condition on positive LTC need. The history gap is $R = \Pr(C_2 \mid C_1, \lambda > 0) - \Pr(C_2 \mid N_1, \lambda > 0)$. All panels include the stationary-distribution response.

Figure F2. Resource pooling and bargaining weights by productivity pair



Each cell is 100 times the pooled benchmark minus the separate-budget counterfactual, in percentage points. Rows index parent productivity and columns index child productivity. Panels (a) and (b) condition on C_1 over all future LTC states. Panels (c) and (d) condition on positive LTC need and C_2. Parent and child effects at a given date are exact negatives because their Pareto weights sum to one. The conditional means reflect endogenous selection into cooperation and, at date 2, renegotiation.