Intergenerational Time Transfers: Altruism and Strategic Motives within the Family *

Chao Fu† Ami Ko‡ Chamna Yoon§

March 2026

Abstract

Extended families provide a large share of both childcare for young children and long-term care (LTC) for elderly parents. These intergenerational time transfers are often bidirectional over the life cycle, raising the question of whether they reflect altruism or strategic exchange. We study the motives underlying these transfers and their policy implications. Using data from the Health and Retirement Study and plausibly exogenous variation in childcare costs, we document evidence consistent with implicit family contracts in which early childcare assistance is later reciprocated through elderly care. Motivated by these findings, we build and estimate a dynamic non-cooperative model of an extended family that chooses care arrangements, labor supply, savings, and bequests under both altruistic and strategic motives. The estimated model shows that strategic motives play an important role in shaping the intertemporal linkage between childcare and LTC. These dynamics generate policy spillovers across care sectors. Childcare subsidies crowd out grandparental childcare and reduce the future LTC provided by adult children, whereas policies that condition old-age benefits on past childcare provision strengthen family-based care in both stages. Overall, the results highlight how dynamic family incentives shape caregiving and mediate the effects of care policies.

*We thank Hanming Fang, Rasmus Lentz, Hamish Low, Andrew Shephard, Garance Genicot, and seminar and conference participants at ASSA, the Federal Reserve Bank of Chicago, the World Congress, SEA, Penn State, and Georgetown for helpful comments.

†University of Wisconsin-Madison and NBER, cfu@ssc.wisc.edu.

‡Georgetown University, ami.ko@georgetown.edu.

§Seoul National University, chamna.yoon@snu.ac.kr.

1 Introduction

Care for dependents is a central economic activity within families. Even in economies with well-developed markets for childcare and long-term care (LTC), a large share of care continues to be provided informally by family members. For example, in the U.S. Health and Retirement Study (HRS) data, about one-third of working women with young children obtain childcare provided by their parents, and nearly two-thirds of retirees with LTC needs receive assistance from their middle-aged children. More strikingly, among retirees with LTC needs, those who previously helped their adult children with childcare are 26% more likely to receive LTC provided by these adult children.[1] These patterns suggest that families allocate time-intensive care across generations and across different stages of the life cycle, raising the question of how these care decisions are linked within the family.

This paper studies the motives underlying intergenerational time transfers of childcare and LTC. We ask whether these transfers are primarily driven by altruism or whether they reflect strategic behavior across life stages. Under altruism, family members provide care because they value each other's well-being. Under strategic motives, care provided at one stage of life influences future behavior. For example, grandparents may help with childcare to induce LTC from their adult children later in life. Understanding the relative importance of these motives is important for interpreting observed family behavior and for the optimal design of public policies: when strategic motives are strong, policies targeting one care sector can indirectly affect the other through strategic family dynamics.

We begin by documenting a set of empirical facts about intergenerational time transfers. Leveraging cross-state and over-time policy variation in the Child Care and Development Fund (CCDF), which subsidizes childcare for low-income families, we examine how CCDF generosity affects childcare and LTC among low-education women with young children. Comparing this group with otherwise-similar women with older children and controlling for state and year fixed effects, we find that higher CCDF generosity reduces grandparental childcare by 37.6 percentage points (ppts), consistent with altruism. Meanwhile, it lowers grandparents' expectation of receiving future LTC from their daughters by 31.7 ppts. These responses are unlikely to reflect higher future opportunity costs of LTC provision for daughters: we find that more generous CCDF reduces daughters' current labor supply. These results suggest a reciprocal link between the two time transfers.

[1] Authors' calculations from HRS data, 1996–2014.

Motivated by these data patterns, we develop and estimate a dynamic model of intergenerational care provision within the family, consisting of two active players, a retired grandparent (G) and a parent (P), and a passive child. The model has three stages: a first stage when childcare is needed, an intermediate stage, and a third stage when LTC may be needed. Throughout all stages, P makes labor supply decisions (full-time, part-time, nonworking), recognizing the value of accumulated experience for future wages; G, who faces health and mortality risks, makes consumption and savings decisions. In addition, in the first stage, P chooses the amount of childcare to purchase from the market, while G chooses the amount of childcare to provide, with P being responsible for the remaining childcare needs. In the third stage, a living G might have LTC needs, in which case P chooses the amount of informal LTC to provide, with the remaining need purchased by G from the formal LTC market.

Guided by our design-based findings, our model accounts for both altruistic and strategic motives for intergenerational time transfers. Altruism is embedded by allowing P and G to care about each other's well-being, governed by altruistic parameters that vary across families of different unobserved types. This unobserved heterogeneity leads to a pure correlation between family-provided childcare and LTC. Building on insights from Becker (1993), our model entertains the possibility that P may feel guilty for not providing LTC to a needy G who helped with childcare in the past; anticipating this, G may strategically provide childcare to increase the chance of receiving informal LTC when needed. This strategic motive links the two time transfers causally. In addition, P may provide LTC to reduce G's medical expenses, thereby increasing bequests.

We use three datasets to estimate the model in two steps. First, similar in spirit to Bonhomme et al. (2022), we estimate the conditional distribution of family altruistic types via K-modes clustering outside the model, using measures of parent–child closeness from Add Health. Given the distribution of types, we estimate the remaining model parameters via indirect inference, combining information from the HRS (G's assets, LTC arrangements, and the correlation between the two time transfers) and the Survey of Income and Program Participation (SIPP) (P's employment, income, and childcare arrangements).

To identify parameters critical for policy implications, i.e., parameters governing the relative importance of altruistic versus strategic motives, we exploit the dynamic nature of the model combined with exclusion restrictions. Utilizing the variation in childcare prices faced by similar households, we extract two types of identifying information. First, using childcare prices as an instrument for G-provided childcare, the associated change in P-provided LTC is informative of

P's guilt for not "paying back." Second, how G-provided childcare responds to changes in childcare prices reveals G's altruism toward P. To identify the extent to which P's informal LTC decisions are driven by a desire to protect bequests from formal care costs, we exploit variation in care provided by P across G's wealth.

The estimates reveal a pronounced asymmetry in altruism across generations, with grandparents displaying strong altruism toward their adult children but weaker altruism in the reverse direction. Consequently, adult children's care decisions are shaped mainly by guilt and bequest incentives. The estimates also indicate that grandparents strongly prefer family-provided LTC over formal LTC. As a result, grandparental childcare provision and bequests are driven not only by altruism but also by the strategic goal of securing informal LTC later in life.

Using the estimated model, we analyze the effects of counterfactual care-related policies, with particular attention to spillovers across care sectors. In the childcare sector, we consider a policy that removes the employment requirement for receiving a subsidy. In the LTC sector, we examine a reform that raises the elderly's consumption floor guaranteed by Medicaid. Finally, we study a policy that explicitly links the two sectors by expanding Medicaid benefits for grandparents in proportion to the childcare they provided earlier in life. To facilitate comparison, each reform is calibrated so that the present-discounted value of the direct fiscal cost in the targeted sector is the same.

Our findings reveal substantial cross-sector spillovers. Removing the work requirement for the childcare subsidy increases reliance on formal childcare and crowds out grandparental care. As parents receive less childcare support from grandparents early in life, they become less reciprocally committed to providing informal LTC later, shifting elderly care toward formal providers. In addition, the policy weakens work incentives for parents, increasing government transfers and reducing tax payments. Taken together, these spillovers almost double the direct fiscal cost of the reform.

In contrast, the unconditional expansion of the Medicaid consumption floor generates only modest spillovers. Although the policy weakens grandparents' incentives to strategically provide childcare in exchange for informal LTC, altruistic motives limit the decline in grandparental time, and the resulting changes in childcare use and parental labor supply are small. Consequently, the overall fiscal cost remains close to the direct cost of the reform.

Finally, tying Medicaid benefits to grandparents' earlier childcare provision strengthens the family's role in both sectors. By rewarding family-provided childcare, the policy increases grandparental involvement, reduces reliance on formal childcare, and encourages greater labor supply

among parents. As a result, it lowers childcare subsidy spending and increases tax revenue, making it the least costly of the three reforms.

Taken together, these results underscore that policies targeting one care sector can generate sizable unintended consequences in the other through dynamic family interactions over the life cycle.

Related Literature There is a long-established literature on intergenerational transfers that focuses on two foundational motives: pure altruism and strategic exchange. The altruism model, most famously articulated in the "Rotten Kid Theorem" by Becker (1974), posits that an altruistic parent's financial transfers are sufficient to align the incentives of all family members. In contrast, the strategic view, exemplified by strategic bequest motives in Bernheim et al. (1985), models transfers as an explicit exchange in which parents use the threat of disinheritance to purchase services, such as care, from their children. In his Nobel Prize lecture (Becker, 1993), Gary Becker argued that parents concerned about old-age support can invest in their children by instilling "feelings of guilt, obligation, duty, and filial love," which effectively commit children to provide support later in life. This concept, later formalized as a rational manipulation of preferences (Becker et al., 2016), provides a powerful explanation for non-financial transfers. Building on this insight, we embed both altruistic and strategic motives in our model and quantify their relative importance in explaining two-way intergenerational time transfers.

More recently, a growing literature uses structural models to understand families' role in caregiving, focusing either on childcare provision by grandparents (Cardia and Ng, 2003; Bick, 2016; García-Morán and Kuehn, 2017) or on LTC provision by adult children (Barczyk and Kredler, 2018; Ko, 2022; Mommaerts, 2025). Given the highly bi-directional nature of these two types of caregiving within a family, we contribute to this literature by providing a unified framework to study both types of transfers in a dynamic setting. Our policy simulations suggest that policy designs could be more effective by accounting for the intrinsic and endogenous linkage between these two types of transfers.

Finally, our paper also relates to studies on the relationship between female labor supply and caregiving. A vast literature documents the "child penalty," showing that the availability and cost of childcare are first-order drivers of mothers' labor force participation and career trajectories (Angrist and Evans, 1998; Kleven et al., 2019). Within this literature, some studies have highlighted the crucial role of grandparental childcare in easing this constraint and enabling mothers to work

(Compton and Pollak, 2011; Anstreicher and Venator, 2024). A separate but analogous literature documents significant negative impacts of informal LTC on the careers of prime-age women (Skira, 2015; Cil and Tuncbilek, 2018). Building on the insights from both literatures, we incorporate female labor supply in our model and study how it is affected by childcare and LTC responsibilities.

The remainder of the paper is organized as follows. Section 2 presents descriptive evidence. Section 3 introduces the model. Section 4 describes the data used to estimate the model. Section 5 presents the estimation strategy and results. Section 6 uses the estimated model to conduct counterfactual policy experiments. Section 7 concludes.

2 Descriptive Evidence

This section documents the empirical patterns that motivate our analysis and model. We first describe how childcare and LTC are allocated within families over the life cycle. We then use variation in childcare costs to provide suggestive evidence on the relationship between early childcare assistance and subsequent LTC provision.

2.1 Summary Statistics

Our primary data source is the Health and Retirement Study (HRS), a nationally representative biennial panel of older Americans starting in 1992. We use waves 3 to 12 (years 1996 to 2014) and focus on respondents who were 50 or older in wave 3, remained alive, and had no missing interviews during the sample period. We refer to these HRS respondents as G (grandparent). We consider their adult children, denoted as P, who were 21 or older in wave 3, remained alive and unenrolled in school throughout the sample period. Respondents without any adult children in wave 3 are excluded. All questions in the HRS are answered by G, including those related to intergenerational transfers and the characteristics of P.

How prevalent are intergenerational transfers? Table 1 summarizes transfers at the G-P-wave level.[2] Time transfers are prevalent. Among P–G pairs in which the youngest child of P is under age five, 29% report that G helps P with childcare.[3] Conditional on G having LTC needs,

[2] Online Appendix Table 4.1 summarizes transfers at the G-wave level. For example, 63% of Gs with LTC needs receive informal care from at least one adult child.

[3] In HRS, we do not observe the ages of P's children, but we do observe the number of children P has. We infer new births from changes in the reported number of children and use this information to impute the age of the youngest child.

Table 1: Summary Statistics on Transfers

	Mean	SD
G provides childcare to P	0.10	0.30
: P has a young kid	0.29	0.46
P provides informal LTC to G	0.04	0.20
: G has LTC needs	0.26	0.44
G's annual financial transfer to P	671.87	6514.59
P's annual financial transfer to G	52.30	718.77
P included as beneficiary of G's life insurance	0.27	0.45
Unique G	11567	
Unique G-P pairs	34394	
Unique G-P-wave combinations	222647	

Notes: Statistics are computed from G-P-wave–level data.

26% of P–G pairs report that P provides informal LTC to G.[4] In contrast, financial transfers are relatively modest, especially inter-vivos transfers: the mean annual financial transfer from G to P is \$672, and that from P to G is \$52. Financial transfers are more frequent post G's death: 27% of Ps are listed as beneficiaries of G's life insurance.

Who provides and receives time transfers? Table 2 summarizes the characteristics of Ps in the entire sample (Column 1), Ps who receive childcare assistance from G (Column 2), and Ps who provide LTC to disabled G (Column 3). Compared to an average adult child in the sample, Ps who receive childcare assistance are younger, more likely to be married, and more likely to live close to their parent. They also have more children on average, consistent with the presence of young children driving demand for grandparental childcare. In contrast, Ps who provide LTC to their disabled elderly parent tend to be older, more likely to be female, and live near their parent. The average age of P who receive childcare assistance is about 16 years lower than that of P who provide LTC, underscoring the distinct life-cycle stages associated with these time transfers.

Table 3 summarizes the characteristics of Gs in the full sample (Column 1), Gs who provide childcare to any of their adult children with young kids (Column 2), and disabled Gs who receive informal LTC from any of their adult children (Column 3). Relative to an average elderly respondent in the sample, Gs who provide childcare are younger and have higher levels of assets and income, and about 71% of these helping Gs are retired. In contrast, Gs who receive LTC are substantially older and predominantly female, with women accounting for 79% of disabled Gs receiving care from

[4]We classify G as having LTC needs if they have activities of daily living (ADLs) limitations or cognitive impairment and receive some form of LTC, including formal services or informal family care.

Table 2: P Characteristics in the HRS

	(1) All P		(2) P receiving childcare help from G		(3) P providing LTC to G	
	Mean	SD	Mean	SD	Mean	SD
Female	0.49	0.50	0.54	0.50	0.64	0.48
Age	45.41	9.70	36.73	5.62	52.93	10.61
Education (years)	13.33	2.45	13.96	2.27	13.10	2.55
Married	0.66	0.47	0.81	0.39	0.60	0.49
Coresides with G	0.06	0.24	0.06	0.23	0.28	0.45
Lives within 10 mi to G	0.37	0.48	0.55	0.50	0.73	0.44
# Children	1.96	1.55	2.38	1.19	2.13	1.70
Observations	222647		9222		7430	

Notes: Sample is at P-wave level. Column (1) includes all Ps. Column (2) includes Ps whose youngest kid is age 5 or younger and who receive childcare help from G. Column (3) includes Ps who provide care to their G with long-term care needs.

Table 3: G Characteristics in the HRS

	(1) All G		(2) G providing childcare help to any P		(3) G receiving LTC from any P	
	Mean	SD	Mean	SD	Mean	SD
Female	0.59	0.49	0.55	0.50	0.79	0.41
Age	72.11	9.70	65.37	6.48	81.39	10.14
Education (years)	11.88	3.35	12.42	3.09	10.06	3.82
Monthly LTC hrs from children	8.16	57.00	2.05	28.55	95.47	176.42
Monthly LTC hrs from daughters	6.02	49.52	1.65	27.10	70.57	159.01
Retired	0.83	0.37	0.71	0.45	0.99	0.09
Assets (in thousand)	524.78	1337.49	753.43	2257.35	200.33	727.21
Annual income (in thousand)	66.30	279.75	89.82	152.64	29.77	102.27
# Children	2.91	1.80	3.25	1.86	3.33	2.14
Have LTC needs	0.23	0.42	0.12	0.33	1.00	0.00
Observations	76437		7819		5698	

Notes: Sample is at G-wave level. Column (1) includes all Gs. Column (2) includes Gs who provide childcare to any of their children with young kids. Column (3) includes Gs with long-term care needs that receive informal long-term care from any of their children.

Table 4: Relationship between Intergenerational Time Transfers

	(1)		(2)		(3)	
	P provides LTC		G provided childcare		P provides LTC	
G provided childcare	0.102***	(0.036)				
Childcare price was high			0.030**	(0.013)	0.076	(0.063)
Observations	4344		4370		4344	
Adjusted R^2	0.375		0.615		0.372	
Mean of dependent variable	0.387		0.240		0.387	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level and are in parentheses. The sample consists of female P - female G pairs in which G has long-term care needs. Not all coefficients shown. See Table 4.2 for full estimation results.

their children.[5] Notably, among all care hours received by G, 70.6 of the 95.5 average monthly hours (74%) are provided by daughters, indicating that informal LTC provision is concentrated among female children. Taken together, these patterns indicate that intergenerational time transfers, especially LTC, in our sample occur predominantly between women. We therefore focus on female G–female P pairs in the remainder of the paper.

2.2 Relationship between Childcare and LTC Provision

To assess the relationship between the two types of time transfers, we use a subsample of disabled Gs with LTC needs within the female G–female P sample described in Section 4.1. We start by presenting the correlation between these two time transfers. Column (1) of Table 4 shows that, conditional on G's demographics, health, and resources, and on P's characteristics and proximity to G, Gs who provided childcare in the past are 10.2-ppt more likely to receive informal LTC from Ps (the sample mean is 38.7%).[6]

The positive correlation can arise from unobserved cross-family heterogeneity in altruism and/or within-family reciprocity. For further suggestive evidence of the latter explanation, we exploit plausibly exogenous variation in childcare costs, arising from two different sources, as shifters of G-provided childcare.

2.2.1 Leveraging Over-time Variation in Childcare Prices

Our first source of variation is the over-time variation in childcare prices. Using the Consumer Expenditure Survey Public Use Microdata, we find that inflation-adjusted average childcare prices

[5] This is consistent with women's higher LTC risk and their greater likelihood of outliving a spouse and thus relying on children rather than spousal care (Chang and Ko, 2026).

[6] See Table 4.2 for the full list of coefficient estimates.

rose sharply in 2005 (Online Appendix Figure 4.1). Therefore, we classify Ps with young children (aged 5 or younger) in 2004 and earlier as facing low childcare prices, and those with young children in 2005 and later as facing high prices. Using this variation relies on the assumption that, conditional on household observables, the distribution of household unobservable types is stable before and after 2005 within our sample period.

Column (2) of Table 4 shows that, among observationally-similar households, cohorts of Ps facing high childcare prices are 3-ppt more likely to receive childcare assistance from G. Given the assumption mentioned in the previous paragraph, this positive correlation is consistent with G being altruistic toward P. Column (3) shows that a positive but noisy reduced-form relationship between P's LTC provision and the childcare price level (low vs. high) they faced when their children were young.

2.2.2 Leveraging Variation in CCDF

Our second source of variation is the cross-state and over-time differences in the generosity of the Child Care and Development Fund (CCDF), the primary program providing childcare subsidies to low-income families in the U.S. For example, in 2010, CCDF supported about 5.5% of children aged 0–12, typically covering around 75% of childcare costs (Guner et al., 2020). Although relative-provided care is eligible for reimbursement, most CCDF-supported children attend formal or licensed care. Reimbursement rates for informal care are much lower, and relative care providers must meet health and safety requirements that can be burdensome.[7] While states must meet federal guidelines, such as income eligibility below 85% of the state median and qualifying activities like work or education, they retain broad discretion in setting specific parameters.

We compile state-level CCDF policy data for 2009–2014 from reports by the Administration for Children and Families.[8] As a measure of generosity of CCDF for each state-year, we construct the variable $Policy_{st} \in [0, 1]$, which is the average of five binary indicators: (1) whether job search qualifies as an eligible activity, (2) whether no minimum work hours are required, (3) whether TANF income is exempt from eligibility calculations, (4) whether TANF families are exempt from copayments, and (5) whether families below 100% of the federal poverty line are exempt from copayments.[9] Notably, while all components of the generosity index expand access to subsidized

[7] In 1999, only about 15% of CCDF-subsidized children were cared for by relatives (Ho, 2015). By 2016, 24 states required CPR certification, 22 required First Aid certification, and 21 mandated annual inspections for relative providers.

[8] https://www.acf.hhs.gov/opre/project/child-care-and-development-fund-ccdf-policies-database.

[9] Because HRS is biennial, CCDF policies enter the analysis only for the 2010, 2012, and 2014 interview

formal childcare, their effects on labor supply can go in opposite directions. For instance, components such as (1) and (2) may weaken work incentives, whereas component (3) may strengthen them. Over the sample period, 22 states experienced at least one change in CCDF generosity as measured by $Policy_{st}$ and there are 26 state–year changes in total.

Because HRS questions are retrospective and CCDF policy data begin in 2009, we use the 2010, 2012, and 2014 waves of HRS. Given its means-tested nature, CCDF benefits mostly low-income families.[10] Therefore, we restrict the sample to low-education Ps, who do not have a high school diploma.[11] We further limit the sample to female G–P pairs residing in the same state, who account for 86.4% of female G-P pairs with low-education Ps. We define the treatment group as Ps whose youngest child is not yet in school and therefore face the largest childcare expenses and potentially benefit from CCDF the most. We compare this group with otherwise-similar Ps with older children, through the following equation

$$Y_{ist} = \beta Treat_{it} \times Policy_{st} + \alpha Treat_{it} + \gamma Policy_{st} + \eta_2 Other_{st} + \eta_1 X_{ist} + \lambda_t + \theta_s + \varepsilon_{ist}, \qquad (1)$$

for G-P pair i living in state s and year t. The dependent variables include (1) whether P works full-time, (2) whether G provides childcare assistance, and (3) whether G *expects* P to provide informal LTC in the future when G becomes disabled.[12] The variable $Treat_{it}$ indicates the treatment group, $Policy_{st} \in [0, 1]$ captures the generosity of CCDF, $Other_{st}$ includes state-level controls (the unemployment rate, minimum wage, state EITC supplement) in year t, X_{ist} includes demographics of G and P in year t, and λ_t and θ_s are year and state fixed effects, respectively. The parameter of main interest is β, which examines how the generosity of CCDF affects the treatment group, as defined above.

Table 5 presents the results. Column (1) shows that more generous CCDF reduces labor supply

years. Over these years, individual CCDF policy features change infrequently when considered in isolation: we observe changes in (1) 8 states for job search eligibility, (2) 3 for minimum work hour requirements, (3) 1 for TANF income treatment, (4) 9 for TANF copayment exemptions, and (5) 3 for copayment exemptions below 100% of the federal poverty line. We therefore summarize generosity using an index that aggregates across these dimensions.

[10] In 2012, the average income of subsidy recipients was about $20,000 (Guner et al., 2020).

[11] Online Appendix Table 4.3 reports results when Ps of all education levels are included. The estimates for care-provision outcomes are qualitatively similar but considerably noisier, as expected, because the broader sample includes many Ps with incomes too high for CCDF programs to be relevant.

[12] To measure expectations about future informal care, the HRS asks: "Suppose in the future, you needed help with basic personal care activities like eating or dressing." Respondents are then asked whether they have anyone who would be willing and able to help over a long period of time. Those who answer affirmatively are subsequently asked to report the relationship(s) of these potential helpers.

Table 5: Childcare Subsidies and Intergenerational Time Transfers

	(1) P works full time		(2) G provides childcare		(3) G believes P will provide LTC	
Treat=1	0.208**	(0.093)	0.093**	(0.046)	0.063	(0.074)
Generosity of CCDF	-0.268	(0.260)	0.082	(0.177)	-0.386	(0.359)
Treat=1 × Generosity of CCDF	-0.438*	(0.247)	-0.376***	(0.117)	-0.317*	(0.190)
Observations	688		715		709	
Adjusted R^2	0.069		0.188		0.148	
Mean of dependent variable	0.390		0.080		0.353	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level and reported in parentheses. The sample consists of female P–female G pairs where P has less than a high school education. The treatment group consists of Ps whose youngest child is not yet in school.

among low-education female Ps. As mentioned above, different components of the CCDF generosity measure can have opposite effects on P's labor supply. During our sample period, there are more changes in CCDF's work-discouraging components than its work-encouraging components. Consistently, the negative estimated coefficient indicates that, on net, the policy margins embedded in our CCDF generosity index operate through work-disincentive channels.

Column (2) of Table 5 shows that more generous CCDF programs substantially reduce grandparent-provided childcare for the treatment group: moving from the least to the most generous program lowers this rate by 37.6 ppts among families with young children. Because greater CCDF generosity effectively reduces the out-of-pocket price of childcare, this result mirrors the earlier finding that G-provided childcare increases when childcare prices are higher (Table 4, Column 2). Without excluding other explanations, these findings are consistent with the hypothesis that G's caregiving decisions are partly motivated by altruism.

Consistent with the earlier results in Table 4, where higher past childcare prices are associated with greater provision of LTC, Column (3) of Table 5 shows that, among the treatment group, G's expectations of receiving informal LTC decline by 31.7 ppts as CCDF programs move from the least to the most generous. Given the finding in Column (1) that program generosity *reduces* P's labor supply, the result in Column (3) is unlikely driven by higher opportunity costs of P's LTC provision. Instead, it suggests a causal link between intergenerational time transfers.

2.3 Interpretation and Implications for the Model

Taken together, results from our descriptive analysis suggest that intergenerational care decisions are dynamically linked. Families that exchange childcare and LTC appear to form persistent relationships over time, and exogenous reductions in the need for grandparental childcare weaken subsequent family-provided LTC. These patterns are difficult to reconcile with a purely altruistic model. Instead, they are consistent with a setting in which current care provision affects future behavior through strategic or reciprocal motives. These findings motivate the structural framework developed in the next section, which allows us to quantify the relative importance of altruistic and strategic motives and to evaluate how policies affecting one care margin propagate through the family over the life cycle.

3 Model

In a dynamic non-cooperative game framework, we study two-way time transfers of childcare and LTC in an extended family, consisting of two active players, a retired grandparent (G) and a parent (P), and a passive child. The model has three stages: a first stage when childcare is needed, a third stage when LTC may be needed, and a stage in between. Throughout all stages, P makes labor supply decisions; G, who faces health and mortality risks, makes consumption and savings decisions. In addition, in the first stage, P chooses the amount of childcare to purchase from the market, G chooses the amount of childcare to provide, with P being responsible for the remaining childcare needs. In the third stage, a surviving G might face LTC needs, in which case, P chooses the amount of informal LTC to provide, with the remaining need purchased by G from the formal LTC market. The two-way transfers can arise from both altruistic motives and two types of strategic motives. First, P may provide LTC to reduce G's formal LTC expenses, thereby increasing bequests. Second, P may feel guilty for not providing LTC to a needy G who helped with childcare in the past; foreseeing this, G may provide childcare strategically to increase the chance of receiving informal LTC when needed. Appendix A specifies functional forms we use.

3.1 Primitives

A household is characterized by $\left(X_1, \chi, H_0^P, a_0^G\right)$. The vector X_1 consists of 3 subvectors: X_1^P (P's demographics), X_1^G (G's demographics, initial health, and retirement income), and X_1^{PG} (household-level observables). The variable χ is unobserved by the researcher, which captures

permanent heterogeneity across households in terms of their altruism; we model χ as an unobserved discrete type that is correlated with X_1. Finally, H_0^P is P's initial work experience and a_0^G is G's initial asset level. Over time, G's health (part of X_t) evolves exogenously, while H_t^P and a_t^G evolve endogenously.

3.1.1 Timing and Decisions

The model has three stages, starting from the child's birth. Stage 1 corresponds to the first five years of the child's life. Stage 3 corresponds to years when the grandparent, if still alive, may need LTC (G is aged above 77 up to 95). Stage 2 corresponds to years in between. To ease exposition, we present the model as if all stages have the same length; our empirical analysis accounts for stage-specific lengths (Online Appendix 1). At each stage, P and G move sequentially, with P moving first. The specific timing and decisions to be made in these stages are as follows.

Stage 1. The child requires Q hours of care, which can be provided by P, G, and/or the market in the form of paid childcare. P observes her Stage-1 choice-specific shocks and chooses her work hours $h_1^P \in \{0, h^{PT}, h^{FT}\}$ and the hours of market-provided childcare $q^M \in \{0, q^{PT}, q^{FT}\}$, corresponding to none, part-time, and full-time hours respectively. Next, G observes her choice-specific shocks related to childcare provision $\left(q^G\right)$ and chooses $q^G \in \{0, \ldots, Q - q^M\}$ and her asset level a_1^G.

Stage 2. G experiences a health shock. With probability $\lambda\left(X_1^G\right)$, G survives with an updated health status, governed by transition function $\Gamma(X_t^G | X_{t-1}^G)$; with probability $1 - \lambda\left(X_1^G\right)$, G dies and leaves bequest a_1^G to P. After the realization of G's survival and health status and P's Stage-2 labor-supply shocks, P chooses her work hours h_2^P. Next, surviving G chooses her asset level a_2^G.

Stage 3. G experiences another health shock. With probability $\lambda\left(X_2^G\right)$, G survives with an updated health status, which determines G's weekly hours of LTC needs $N\left(X_3^G\right) \in \{0, ..., N_{\max}\}$; with probability $1 - \lambda\left(X_2^G\right)$, G dies and leaves bequest a_2^G to P. After the realization of G's LTC needs and P's Stage-3 choice-specific shocks, P chooses work hours h_3^P and her LTC provision $n^P \in \left\{0, ..., N\left(X_3^G\right)\right\}$ if G needs it. Next, surviving G pays the out-of-pocket expense for $N\left(X_3^G\right) - n^P$ hours of formal LTC and chooses her asset level (a_3^G). At the end of period 3, G dies and leaves bequest a_3^G to P.

3.1.2 Constraints

Constraints for P. We assume that the parent does not save and uses all her after-tax income to fund her consumption and, if $t = 1$, her expenses on market-provided childcare. P's after-tax labor *and* non-labor income $W\left(h_t^P, H_{t-1}^P, X_t^P\right)$ depends on her work hours h_t^P, past work experience H_{t-1}^P, and demographics X_t^P.[13] At the beginning of next period, P's experience evolves such that

$$H_t^P = \widetilde{H}(H_{t-1}^P, h_t^P). \tag{2}$$

In the first stage, P's out-of-pocket expense on paid childcare depends on the sticker price of childcare p^M, P's usage of this service q^M, and childcare subsidies P may be eligible for given her income $W(\cdot)$, labor supply h_t^P, and demographics X_t^P. P's consumption in Stage t is given by

$$c_t^P = \begin{cases} Y\left(h_t^P, W\left(h_t^P, H_{t-1}^P, X_t^P\right), X_t^P, q^M, p^M\right) \text{ for } t = 1, \\ W\left(h_t^P, H_{t-1}^P, X_t^P\right) \text{ for } t = 2, 3, \end{cases} \tag{3}$$

where $Y(\cdot)$ is P's after-tax income net of out-of-pocket childcare expenditure. We adopt a broad notion of taxes that encompasses both tax liabilities and transfer payments such as Supplemental Nutrition Assistance Program (SNAP); net taxes can be negative for low-income households. Note that $W(\cdot)$ enters $Y(\cdot)$ because it determines both P's total available resources and her eligibility for SNAP and childcare subsidies. Similarly, h_t^P and X_t^P affect not only income $W(\cdot)$ but also eligibility for these programs, and hence enter $Y(\cdot)$ directly.

The parent allocates her endowed time $\mathcal{T}$ to work (h_t^P), childcare $\left(q^P\right)$, elderly care $n^P \in \left\{0, ..., N\left(X_3^G\right)\right\}$, and leisure (l_t^P), such that

$$l_t^P = \mathcal{T} - h_t^P - I\left(t = 1\right) q^P - I\left(t = 3\right) n^P, \tag{4}$$

$$\text{with } q^P = Q - q^M - q^G, \tag{5}$$

i.e., at $t = 1$, P provides all childcare hours not covered by paid care or G.

[13] Given that P's income process is not a main focus of our paper, we model it as a simple function of age, education, experience, and labor supply.

Constraints for G. The grandparent's consumption $\left(c_t^G\right)$ and savings $\left(a_t^G\right)$ decisions are subject to the following constraints:

$$c_t^G = \max\left\{(1+r)a_{t-1}^G + y^G - a_t^G - I\left(t=3\right)m\left(N-n^P\right),\ \underline{c}^G\right\} \tag{6}$$

$$a_t^G \geq 0 \tag{7}$$

$$a_t^G = 0 \text{ if } (1+r)a_{t-1}^G + y^G - I\left(t=3\right)m\left(N-n^P\right) \leq \underline{c}^G, \tag{8}$$

In 6, a_{t-1}^G is G's assets at the beginning of Stage t, on which G earns interest at rate r, $y^G \in X_t^G$ is G's retirement income. These resources are used to fund G's consumption c_t^G and savings next period a_t^G, and if $t=3$, the cost of formal LTC services $m\left(\cdot\right)$. In the case of large LTC expenditure needs, G is provided a consumption floor $\underline{c}^G$ by implicit governmental or societal insurance. We put two constraints on G's asset choice. First, G cannot borrow (7); second, G has to run down her asset to zero before she can resort to the consumption floor (8).

The grandparent splits her time $\mathcal{T}$ between leisure l_t^G and childcare $q^G \in \{0,..,Q-q^M\}$ in $t=1$, and spends all her time on leisure in other periods, such that

$$l_t^G = \mathcal{T} - q^G I\left(t=1\right). \tag{9}$$

3.1.3 Preferences

P's Flow Utility. We specify P's flow utility as

$$u^P(c_t^P, l_t^P, h_t^P) + I\left(t=1\right)\begin{pmatrix}\theta_q \frac{q^P}{Q} + \\ \theta_K K\left(q^M, q^G, q^P\right)\end{pmatrix} - I(t=3)\begin{pmatrix}\psi\left(N(X_3^G), n^P\right)\times \\ \left(\omega_\chi^P + \rho\frac{q^G}{Q}\right)\end{pmatrix}. \tag{10}$$

Common across all stages, P derives utility $u^P(\cdot)$ from her own consumption c_t^P, leisure l_t^P, and a fixed cost of labor force participation $(h_t^P > 0)$. In Stage 1, P derives additional utility θ_q from time spent caring for her child and utility θ_K from her child's skill development $K(\cdot)$, the latter being a function of childcare arrangement.[14] In Stage 3, the grandparent derives disutility $\psi(\cdot)$ from using formal LTC services, where ψ is a function of the severity of her condition $N(X_3^G)$ and formal LTC hours $N(X_3^G) - n^P$.[15] P internalizes this disutility to a degree that depends on 1)

[14] Although only appearing in Stage-1 utility, θ_K captures the sum of P's discounted utility from child development across future periods.

[15] While the existing literature consistently finds this disutility to be positive, i.e., retirees dislike formal LTC relative to family-provided care, we do not impose this sign ex ante and instead allow ψ to take either

her type-specific altruism parameter, ω_χ^P, and 2) G's past help with childcare from Stage 1. The latter, governed by the parameter ρ, captures the guilt or indebtedness P feels toward G for G's time transfer.[16]

G's Flow Utility. We specify G's flow utility as

$$u^G(c_t^G, l_t^G) + \omega_\chi^G \left(u^P(.) + I\left(t=1\right) \left(\theta_q \frac{q^P}{Q} + \theta_K K\left(q^M, q^G, q^P\right) \right) \right) - I\left(t=3\right) \psi\left(N(X_3^G), n^P\right). \quad (11)$$

Common across all stages, the grandparent, if alive, derives utility $u^G(.)$ from her own consumption and leisure and from P's flow utility weighted by her type-specific altruism parameter ω_χ^G. As explained earlier, in Stage 3, G derives disutility $\psi(.)$ from using formal LTC services. If this disutility is sufficiently large, G will have strategic incentives to induce informal LTC provision from P to avoid long hours of formal LTC. She can do so through two channels. First, by accumulating savings, which strengthens P's incentive to provide informal LTC in order to protect the inheritance from being depleted by formal care expenses. Second, by providing childcare in Stage 1, which activates the guilt mechanism and increases P's willingness to reciprocate with informal LTC.

3.2 Value Functions

In this subsection, we specify the value functions and solve the family's problem by backward induction. Readers primarily interested in the model's implications rather than its technical details may skip to Section 3.3.

Recall that G's survival probability at the beginning of $t = 2, 3$ is $\lambda\left(X_{t-1}^G\right)$. Conditional on G's survival, we denote the transition probability from X_t^G to X_{t+1}^G by $\Gamma\left(X_{t+1}^G | X_t^G\right)$. We denote by Ω_t the common state variables faced by P and G at Stage t; they face their own choice-specific shocks when applicable. Notably, the first mover P's Stage-t choices are G's state variables at t.

3.2.1 Stage 3

Common state variables Ω_3 include X_3, χ, H_2^P (P's work experience), a_2^G (G's asset), and q^G (G's childcare provision in Stage 1).

sign in the structural estimation.

[16] Alternatively, ρ can be interpreted more broadly as a reduced-form representation of the future consequences of reneging on an implicit intergenerational arrangement, including the possibility that failing to support one's parent weakens similar support from one's own children. This interpretation is related to the "demonstration effect" emphasized by Cox and Stark (2005), whereby adults may care for their parents partly to shape the future behavior of the next generation.

Surviving G's Problem. Given Ω_3 and P's Stage-3 choices, G chooses her optimal asset level a_3^G, yielding the following value function:

$$V_3^G\left(\Omega_3, c_3^P, h_3^P, n^P\right) = \max_{a_3^G} \left\{ \begin{array}{c} u^G(c_3^G, \mathcal{T}) + \omega_\chi^G u^P\left(c_3^P, \mathcal{T} - h_3^P - n^P, h_3^P\right) \\ -\psi(N(X_3^G), n^P) + \beta\omega_\chi^G B\left(a_3^G\right) \end{array} \right\} \quad (12)$$
$$s.t.\ (6) - (8).$$

The first three components of (12) are G's flow utility as specified in (11). The last component reflects the fact that G dies at the end of Stage 3 and leaves a bequest, providing P with inheritance utility $B\left(a_3^G\right)$ discounted at rate β. Being altruistic, G internalizes this utility to the degree of ω_χ^G.

P's Problem. If G survives, P chooses labor supply and LTC provision, $z_3^P \equiv \left(h_3^P, n^P\right)$, yielding the value function $V_3^P(\cdot)$ given by

$$V_3^P\left(\Omega_3, \epsilon_3^P\right) = \max_{z_3^P} \left\{ \begin{array}{c} u^P(c_3^P, \mathcal{T} - h_3^P - n^P, h_3^P) - \left(\omega_\chi^P + \rho\frac{q^G}{Q}\right)\psi(N(X_3^G), n^P) \\ +\beta B\left(a_3^G\right) + \epsilon_3^P\left(z_3^P\right) \end{array} \right\} \quad (13)$$
$$s.t. \qquad (3),$$

where $\epsilon_t^P(.)$ is an i.i.d. choice-specific shock.

If G dies at the beginning of Stage 3, P receives G's bequest a_2^G and chooses labor supply h_3^P, yielding the following value function $\widetilde{V}_3^P(\cdot)$ given by

$$\widetilde{V}_3^P\left(\Omega_3, \epsilon_3^P\right) = \max_{h_3^P} \left\{ u^P(c_3^P, \mathcal{T} - h_3^P, h_3^P) + B\left(a_2^G\right) + \epsilon_3^P\left(h_3^P\right) \right\} \quad s.t. \quad (3). \quad (14)$$

3.2.2 Stage 2

Common state variables Ω_2 include X_t, χ, H_1^P, a_1^G, and q^G.

Surviving G's Problem. G chooses optimal asset a_2^G , yielding the value function

$$V_2^G\left(\Omega_2, c_2^P, h_2^P\right) = \max_{a_2^G} \left\{ \begin{array}{c} u^G(c_2^G, \mathcal{T}) + \omega_\chi^G u^P\left(c_2^P, \mathcal{T} - h_2^P, h_2^P\right) \\ +\beta(1 - \lambda\left(X_2^G\right))\omega_\chi^G B\left(a_2^G\right) \\ +\beta\lambda\left(X_2^G\right) \sum_{X_3^G} \Gamma\left(X_3^G | X_2^G\right) E_{\epsilon_3^P}\left[V_3^G\left(\Omega_3, c_3^P, h_3^P, n^P\right)\right] \end{array} \right\} \quad (15)$$
$$s.t.\ (2), (6) - (8).$$

With probability $1-\lambda\left(X_2^G\right)$, G dies at the beginning of Stage 3 and leaves bequest a_2^G to P (the second row of (15)). With probability $\lambda\left(X_2^G\right)$, G survives in Stage 3 (the third row of (15)). In this case, G's expected value accounts for two types of uncertainties. First, with transition probability $\Gamma\left(X_3^G|X_2^G\right)$, G's health evolves to a new status. Second, G takes expectations over shocks to P's Stage-3 choices, because $V_3^G(\cdot)$ depends on these choices.

P's Problem. If G survives, P's value function is given by

$$V_2^P\left(\Omega_2,\epsilon_2^P\right) = \max_{h_2^P}\left\{\begin{array}{c} u^P(c_2^P,\mathcal{T}-h_2^P,h_2^P)+\epsilon_2^P\left(h_2^P\right) \\ +\beta(1-\lambda\left(X_2^G\right))E_{\epsilon_3^P}\left[\widetilde{V}_3^P\left(\Omega_3,\epsilon_3^P\right)\right] \\ +\beta\lambda\left(X_2^G\right)\sum\limits_{X_3^G}\Gamma\left(X_3^G|X_2^G\right)E_{\epsilon_3^P}\left[V_3^P\left(\Omega_3,\epsilon_3^P\right)\right] \end{array}\right\} \tag{16}$$
$$s.t. \quad (2),(3).$$

Parallel to (15), the second row in (16) describes the case where G dies at the beginning of $t=3$, and the third row describes the case where G survives in $t=3$.

If G dies at the beginning of $t=2$, P receives G's bequest a_1^G and chooses labor supply h_t^P for the remaining periods $t=2,3$, yielding the value function

$$\widetilde{V}_2^P\left(\Omega_2,\epsilon_2^P\right) = \max_{h_2^P}\left\{\begin{array}{c} u^P(c_2^P,\mathcal{T}-h_2^P,h_2^P)+B\left(a_1^G\right)+\epsilon_2^P\left(h_2^P\right)+ \\ \beta E_{\epsilon_3^P}\max_{h_3^P}\left\{u^P(c_3^P,\mathcal{T}-h_3^P,h_3^P)+\epsilon_3^P\left(h_3^P\right)\right\} \end{array}\right\} \tag{17}$$
$$s.t. \quad (2),(3).$$

3.2.3 Stage 1

Common state variables Ω_1 include X_1,χ, H_0^P and a_0^G.

G's Problem. Given Ω_1, P's Stage-1 choices $\left(c_1^P,h_1^P,q^M\right)$, and her choice-specific shocks to childcare provision $\epsilon_1^G\left(q^G\right)$, G chooses her asset level and childcare provision q^G:

$$V_1^G(\Omega_1,c_1^P,h_1^P,q^M,\epsilon_1^G) = \max_{a_1^G,q^G}\left\{\begin{array}{c} u^G(c_1^G,\mathcal{T}-q^G)+\omega_\chi^G\left[u^P\left(c_1^P,l_1^P,h_1^P\right)+\theta_q\frac{q^P}{Q}+\theta_K K(.)\right]+\epsilon_1^G(q^G) \\ +\beta(1-\lambda(X_1^G))\,\omega_\chi^G B(a_1^G) \\ +\beta\lambda(X_1^G)\sum_{X_2^G}\Gamma(X_2^G|X_1^G)\,E_{\epsilon_2^P}\left[V_2^G(\Omega_2,c_2^P,h_2^P)\right] \end{array}\right\}$$
$$s.t.\ (2),(4),(5),(6)-(8). \tag{18}$$

The first row of (18) describes G's flow utility, including her choice-specific shock. With probability $1-\lambda(X_1^G)$, G dies in Stage 2 and leaves bequests to P. With probability $\lambda(X_1^G)$, G survives in Stage 2; in this case, G forms expectations over her Stage-2 health shocks and shocks to P's Stage-2 choices.

P's Problem. P chooses labor supply and paid childcare hours $z_1^P = \left(h_1^P, q^M\right)$, where each choice z_1^P is associated with a preference shock $\epsilon_1^P(z_1^P)$, yielding the value function:

$$V_1^P(\Omega_1, \epsilon_1^P) = \max_{z_1^P} \left\{ E_{\epsilon_1^G} \left[\begin{array}{c} u^P(c_1^P, l_1^P, h_1^P) + \theta_q \frac{q^P}{Q} + \theta_K K(q^M, q^G, q^P) + \epsilon_1^P(z_1^P) \\ +\beta(1-\lambda(X_1^G))\, E_{\epsilon_2^P}\left(\widetilde{V}_2^P(\Omega_2, \epsilon_2^P)\right) \\ +\beta\lambda(X_1^G) \sum_{X_2^G} \Gamma(X_2^G | X_1^G)\, E_{\epsilon_2^P}\left[V_2^P(\Omega_2, \epsilon_2^P)\right] \end{array} \right] \right\} \quad (19)$$
$$s.t.\ (3) - (5).$$

Because G moves after P and G's choices affect P's flow and future utilities, P takes expectations over G's choice shocks ϵ_1^G, hence the operator $E_{\epsilon_1^G}(\cdot)$. Inside $E_{\epsilon_1^G}(\cdot)$, (19) describes P's flow utility (row 1), and her expected continuation value in Stage 2 if G dies (row 2) or survives (row 3).

3.3 Discussion

Non-Cooperative Framework. In modeling the family dynamics, both cooperative and non-cooperative frameworks serve as approximations; we choose the latter for two reasons. First, the distribution of assets between family members is indeterminate in cooperative models with full commitment. Given that Medicaid eligibility depends only on the beneficiary's wealth, cooperative models suggest that children should own all of their parents' assets, which is at odds with our data. Second, the majority of parent-grandparent pairs in our sample belong to separate households, where information asymmetry and transaction costs can be high, making cooperation more difficult. As shown in Lundberg and Pollak (1993) and Castilla and Walker (2013), in such cases, intrahousehold allocations can default to a non-cooperative equilibrium.

Sequential Move. Within each stage, we assume that players move sequentially. Modeling a simultaneous-move game would cause substantial complications (e.g., multiple equilibria) without gains in realism. Furthermore, we assume that P is the first mover within each stage for the following reasons. First, P's choices of labor supply and the use of childcare services are likely subject to constraints such as workplace and childcare service schedules. In our model, these choices are

made in anticipation of G's childcare provision; this is arguably more plausible than assuming that these choices adjust to the realized G-provided childcare (which is subject to preference shocks). Second, given the potentially strategic motive for G-provided childcare, a model assuming that G moves first may overstate the extent to which G can influence P's decision to provide LTC later. By modeling P as the first mover, we are more conservative in drawing our counterfactual policy implications.

Abstraction of P's Savings Decision. We include G's savings decision in our model because it is critical for our research question. First, G's savings determine bequests, one of the strategic levers linking childcare and LTC provision. Second, G's asset level governs her eligibility for means-tested programs (e.g., Medicaid) that provide consumption floors for the elderly, a key policy parameter we examine. In contrast, P's savings do not play an analogous role in the mechanisms we study. Financial transfers from P to G are rare and negligible in the data, and childcare subsidies depend on P's income, which we model, rather than on assets. Therefore, we abstract away from P's savings decision, which would introduce another continuous endogenous state variable, to maintain tractability.[17] We acknowledge that this abstraction prevents us from studying how P's assets interact with caregiving decisions, an interesting question that we leave for future work.[18]

Abstraction of Location Decisions. The model abstracts away from location decisions that determine residential proximity between P and G. This choice is highly correlated with time transfers, and families may move closer to each other during care-needed stages of their lives (e.g. Anstreicher and Venator, 2024). To the extent that families with stronger preferences for intergenerational support are more likely to locate near each other and also more likely to provide time transfers, our model would interpret these families as being more altruistic. This simplification affects how one might interpret the altruism parameters in our model, but it does not affect the policy implications we derive: changing policy incentives in one care sector will spill over to the other sector through within-family strategic dynamics.

[17] Furthermore, while the HRS provides detailed information on respondents' wealth, who are treated as G in our model, it does not provide information on the assets of respondents' children.

[18] If children were able to accumulate assets, they could buffer earnings losses associated with reducing labor supply to care for a disabled parent. However, existing evidence suggests that caregiving is not associated with systematically different asset dynamics. For example, Skira (2015) finds no descriptive evidence that caregivers experience differential changes in assets relative to non-caregivers.

4 Estimation Data

4.1 HRS

From the HRS 1996–2014 data introduced in Section 2, we build an estimation sample by imposing the following restrictions. Our model considers interactions between one G and one P. Table 3 shows that informal long-term care is provided mostly to female Gs. For each household in the HRS, we therefore choose a female G. Regarding P selection, we consider daughters. This is based on the fact that daughters account for most of the informal long-term care hours provided by children. Applying these gender restrictions results in 5,082 Gs and 9,483 Ps.

For each G, we pick one P, following the procedure described in Online Appendix 2. During this procedure, we pick the P with the highest probability of becoming the major caregiver for G among P's for whom we have credible information about G's childcare provision, resulting in 2,426 G's and 3,341 P's.

Finally, we exclude G who is not retired or has assets in the top 1.5% of the distribution. The final sample includes 2,368 families and 16,064 family-wave observations. Tables 4.4 and 4.5 in the Online Appendix provide summary statistics of the families in the HRS estimation sample.

4.2 SIPP

While the HRS is informative about how childcare and long-term care transfers within a family are linked, it does not contain detailed information on childcare arrangements. For example, it lacks measures of the intensive margin of grandparent-provided childcare and information on the use of formal childcare services. Furthermore, G's reports on P's employment and earnings can be very noisy. To complement HRS, we use the Survey of Income and Program Participation (SIPP), which provides information on income, employment and a wide variety of subjects. Nationally representative panels were drawn annually from 1984 to 1993, and there were usually several panels in progress simultaneously. Starting in 1996, they had larger non-overlapping panels drawn every few years. Prior to the 2014 panel, data were collected at 4-month intervals ("waves"), and most panels consisted of between 8 and 16 waves.

SIPP consists of "core" questions that are asked every wave and Topical Module (TM) questions that are not repeated in each wave. TM questions gather specific information on a wide variety of subjects, including child care arrangements. We construct the SIPP TM sample using waves

from 2001, 2004 and 2008 panels where TMs included questions on child care arrangements.[19] A designated parent is selected from each household to answer questions about care arrangements for each and every child. In married couples, the mother is the designated parent. Among TM respondents, 97% are women, and we exclude the remaining 3% from the sample. We further restrict to respondents with children of ages 0-5 years who are not yet in school. As our model considers only one grandchild, in cases with multiple children, we first select the child for whom the grandparent provides the most hours of care, and then we choose the youngest child to break the ties. The SIPP TM sample consists of 15,136 parents and the number of parent-wave observations is 21,250.

Information about P's employment and earnings is asked in "core" questions, which were asked in all waves for panels 2001, 2004 and 2008. We restrict to female respondents aged between 21-64. The SIPP Core sample consists of 111,765 respondents and the number of respondent-wave observations is 865,484. We use this sample to estimate P's income function and to construct P's employment moments. Table 4.6 in the Online Appendix reports the summary statistics of the respondents in the SIPP Core sample. Compared with P's in the HRS sample, the mean age is lower by one year at 42.5. The share working full-time is 0.52, which is lower than the value of 0.63 observed in the HRS estimation sample. This difference may partly reflect that the cutoff for full-time work is higher in the SIPP (35 hours per week) than in the HRS (30 hours per week). The labor force participation rate is $0.52 + 0.16 = 0.68$, which is also lower than that in the HRS sample ($0.63 + 0.12 = 0.75$) as reported in Table 4.4 of the Online Appendix.

4.3 Add Health

To estimate the distribution of unobserved altruism type, we use the publicly available survey data of the Add Health study, administered by the National Longitudinal Study of Adolescent to Adult Health. It is a longitudinal study of a nationally representative sample of adolescents in the U.S. and asks questions about how children and parents perceive their relationship. Online Appendix 3 provides details on the data.

[19] While TM questions for child care arrangements were asked to earlier panels, we do not use them due to inconsistencies in questions. For example, care arrangements were asked only to parents working or in school.

5 Estimation Strategy and Identification

5.1 Empirical Definitions

Time Horizon. We consider three values of P's age when the first stage of the model starts: 25, 29 and 34 years old. These are the first three quartiles of the maternal age distribution in the SIPP sample when their child is born. The grandparent G is assumed to be 29 years older than P, which is the mean age difference in the HRS estimation sample.[20] The first stage lasts 5 years which captures the time until free public education starts. The third stage starts when G's mean age is 77 which is when the risk of developing LTC needs starts to become substantial. This implies that the second stage lasts 14 years. The third stage lasts until G dies, and we assume G dies with probability one once they reach the age of 95.

G's Income. We closely follow De Nardi et al. (2025) in specifying G's income process. We model log income as the sum of a flexible age profile and a permanent income (PI) component. We first estimate household fixed effects from a regression of log income on age polynomials. To identify PI, we rank households by their mean residual across waves and fit a fifth-order polynomial in the PI percentile. We then assign households to terciles of the estimated fixed effect distribution for discretization. Appendix A provides more details.

P's Income. P's income process is estimated using SIPP Core data. Labor income is modeled as a function of hours worked (part-time or full-time), work experience, and education. Non-labor income is modeled as a function of education and age. We impute federal income tax liabilities using the Internet TAXSIM program provided by NBER, using the 2005 tax schedule. Tax liabilities are calculated as a function of age, labor income, non-labor income, and child age.

Government Transfers. We incorporate government transfer programs for both P and G. For P, in stage 1, working mothers with household income below the threshold $\bar{y}$ are eligible for CCDF, which subsidizes 75% of formal childcare expenditures, as in Guner et al. (2020). We estimate $\bar{y}$ via indirect inference to match the 5.5% recipiency rate reported in Guner et al. (2020). In all stages, P may also receive SNAP (food stamp) benefits if her pre-transfer household income falls below the federal eligibility limit. Eligible households receive benefits up to a maximum annual amount $\bar{b}_{SNAP}$, which we estimate via indirect inference; benefits phase out at the statutory rate

[20] Grandparents are therefore either 54, 58 or 63 years old when the model starts.

of 30 cents per dollar of pre-transfer income. For G, we estimate the consumption floor $\underline{c}^G$ via indirect inference, using the Medicaid recipiency rate in the HRS data (approximately 13%) as an informative moment.

Care Prices. As described in Section 2, we use the Consumer Expenditure Survey Public Use Microdata to measure cross-time variation in inflation-adjusted childcare prices. Low formal childcare prices are set at \$4.3 per hour, which represent the mean costs between years 1996-2004. High childcare prices are set at \$5.5, which are the mean costs between years 2005-2014. For information about formal LTC prices, we use Genworth's Annual Cost of Care Survey and set the hourly rate of formal care at \$46. We incorporate variations in the cost of LTC based on whether the respondent (G) in the HRS sample has daughters other than the matched child (P). Conditional on having long-term care needs, the average nights spent in a nursing home is higher by 37% when the selected child is the only daughter that the respondent has. We therefore reflect this variation by increasing the cost of LTC for HRS respondents with just one daughter relative to those with multiple daughters.

Care Hours. The total required childcare is 84 hours per week. In the SIPP sample, among individuals that use formal child care for more than 35 hours per week, the majority report using the service 40 hours. For those that use less than 35 hours, the average hour is 17. We therefore allow formal childcare hours to be 0, 17 or 40 hours per week. Likewise, we allow grandparental childcare hours to be 0, 12, and 40 hours per week, where 12 is the mean among those that report receiving less than 35 hours of help. G in $t = 3$ can be healthy, lightly disabled or severely disabled. For the lightly disabled state, G needs 15 hours of care, while in the severe state, 40 hours of care are needed. When G is disabled, P can choose to provide 0%, 50% or 100% of required LTC hours.

Other Values. The total available time $\mathcal{T}$ for both P and G is set at 124 hours per week. P's part-time and full-time work requires 15 and 40 hours per week, respectively. The annual real interest rate is set at 3% as in Brown and Finkelstein (2008).

5.2 Parameters Estimated or Set Outside of the Model

G's Health Risks. We use the HRS data to estimate G's health and mortality risks outside the model. G is assumed to be healthy in the first stage. In the second stage, conditional on being alive, G's health can be either good or bad. We use self-reported health in the HRS to measure

this. In the third stage, conditional on being alive, G can be healthy, lightly disabled or severely disabled. We treat an individual (1) as being healthy if the individual has zero or one limitation carrying out activities of daily living (ADL); (2) as lightly disabled if the individual has two or three ADL limitations without cognitive impairment; and (3) as severely disabled if the individual has four or more ADL limitations or cognitive impairment.[21] G's health transition probabilities follow a Markov process where the next period's health depends on the current health and permanent income. We adopt a logit function and recover the transition by maximum likelihood estimation.

Child Skill Production Function. Rather than allowing the parent to have direct preferences over a specific combination of childcare arrangements, $q = \left(q^M, q^G, q^P\right)$, we discipline the model using a child skill production function $K\left(q^M, q^G, q^P\right)$, such that P's preferences over different care arrangements arise through its effect on child skill development. The function $K(\cdot)$, however, is neither the focus of this paper nor credibly identifiable using our data. Therefore, we borrow estimates from Bernal and Keane (2011), who use variation in welfare policy rules to estimate the effects of cumulative childcare inputs relative to maternal time on cognitive outcomes of young children (ages 4-6). Consistent with Flood et al. (2022), who document that relative-provided childcare has the lowest average quality among non-parental care arrangements, Bernal and Keane (2011) show that relative-provided care (predominantly provided by grandparents) reduces test scores relative to maternal care. Specifically, one additional quarter of relative-provided care reduces test scores by about 0.75% relative to maternal care, while one additional quarter of formal care increases test scores by about 0.47%. We specify our $K(\cdot)$ function accordingly (details in Appendix A).

Type Distribution. We estimate the type distribution $\Xi(\chi|X_1)$ using measures of parent-child closeness from Add Health. We largely follow the approach in Bonhomme et al. (2022), with two differences. First, because the Add Health survey measures are categorical, we use K-modes clustering instead of K-means clustering. Second, these measures are obtained from Add Health rather than from our main estimation sample, which prevents us from identifying types at the household level as in Bonhomme et al. (2022). Instead, using variables available in both datasets

[21] The HRS asks respondents whether they have a difficulty carrying out each of five ADLs, which are bathing, dressing, eating, getting in/out of bed, and walking across a room. The HRS also provides a cognitive score based on several tests including word recall, subtraction, backward number counting, object naming, date naming, and president naming. We treat an individual as cognitively impaired if the individual's test score is in the bottom 10%.

(G's and P's education) as conditioning variables, we nonparametrically estimate the distribution of types $\Xi(\chi|X_1)$. In Online Appendix 3, we provide details on this procedure.

5.3 Structural Estimation via Indirect Inference

Given the estimated distribution of types conditional on household observables, we estimate the rest of model parameters (Θ) using the indirect inference estimation method. The estimated vector of structural parameters is the solution

$$\widehat{\Theta} = \text{argmin}_{\Theta} \left\{ [\widehat{\beta}(\Theta) - \overline{\beta}]' W [\widehat{\beta}(\Theta) - \overline{\beta}] \right\},$$

where $\overline{\beta}$ is our chosen set of auxiliary model parameters computed from data, $\widehat{\beta}(\Theta)$ is the corresponding auxiliary model parameters obtained from a large dataset simulated from the model, and W is a weighting matrix. In particular, the vector Θ to be estimated in Step 2 includes parameters governing 1) P's and G's consumption-leisure preference function $u^P(\cdot)$, $u^G(\cdot)$, 2) P's preferences over time spent with the child and child development, 3) G's disutility from using formal LTC, 4) household-type specific altruism parameters ω_χ^P and ω_χ^G, and 5) parameters governing strategic motives: P's guilt parameter ρ and utility from bequests.

5.4 Identification

We provide an overview of how we identify the rest of the model, given $\Xi(\chi|X_1)$; these identification arguments guide our choice of the auxiliary models. As shown in Lewbel (2000), static discrete choice models with unobserved types can be identified semiparametrically if there exists a variable, with a large support, that affects choice-specific payoffs but is excluded from the type distribution. Given our assumption that households' χ types are permanent, our identification relies on the dynamic nature of our model and the panel data structure, combined with the following exclusion restrictions: Conditional on X, households of different types face the same (1) prices and (2) survival probabilities and health status transition probabilities for G. Given (1), the sensitivity of households' choices with respect to market care prices provides information on how households trade off consumption against the (dis)utility from using different care arrangements. Given (2), households of the same unobserved type, due to different realizations of G's health shocks, will face different choice-specific payoffs. However, the supports of these excluded variables are bounded; we therefore impose functional form assumptions.

With these assumptions and the dynamic nature of the model, we discuss how we identify households' preference parameters. Although all parameters are jointly identified, some parts of the data are more informative of certain parameters than others; we argue our identification accordingly, with particular focus on how we differentiate altruistic motives (governed by the type-specific ω_χ) and strategic motives (governed by the guilt parameter ρ and P's utility from bequests).

Guilt Parameter (ρ). This parameter governs the causal effect of G-provided childcare on P's subsequent LTC provision. An OLS correlation between these transfers conflates the unobserved heterogeneity in altruism with the causal guilt mechanism, as more altruistic families naturally provide more of both types of care. We separate these channels by using market childcare prices as an instrument for G-provided childcare (q^G). Given the assumption that prices are uncorrelated with unobserved types, conditional on X, the IV estimand isolates the causal reciprocity effect: an exogenously-shifted q^G should not induce variation in n^P if $\rho = 0$. Similarly, given that G is forward-looking, the correlation between LTC price and q^G, the latter being the time "lent out," provides information on the strength of strategic motives.

G's Altruism (ω_χ^G). Price variation is also informative about G's altruistic motives. When P faces a higher childcare price, a more altruistic G would respond by providing more childcare to reduce P's childcare costs. Conditional on the same expected health evolution (given X_t), a more altruistic G would also save more to leave a larger bequest while providing more childcare assistance. To strengthen the identification of altruistic heterogeneity χ, we also exploit the information embedded in G's choices over time, in particular the intertemporal correlations between G's childcare provision in Stage 1 and G's wealth accumulation in Stages 2 and 3.

P's Altruism (ω_χ^P) and Bequest Motive. P's LTC provision in families where G did not provide childcare in the past and where G has limited assets is informative about P's altruism. In such cases, the scope for reciprocal and bequest motive is weak, so P's decision to provide LTC primarily reflects her altruistic concern for G.

P's consumption–leisure preferences are identified from her labor supply in the absence of caregiving responsibilities, i.e., when G is healthy in Stage 3 (a conditionally exogenous event). Observing labor supply in these periods allows us to recover the value of P's leisure and the opportunity cost of her time implied by the wage function.

Given these parameters, P's LTC decision reflects a comparison between the opportunity cost

of providing care and its benefits. The opportunity cost consists of the value of P's leisure time and forgone earnings. The benefits include the savings in G's formal LTC spending, potential increases in bequests, and the reduction in guilt governed by ρ. Therefore, P's decision to allocate time to LTC is informative about her altruism (ω_χ^P). In particular, if $\omega_\chi^P = 0$, P would not provide LTC when the opportunity cost exceeds these benefits.

Other Parameters. Given the identification of P's preference for leisure and the pre-estimated child skill function $K(\cdot)$, P's Stage-1 childcare choices identify her direct utility from caregiving (θ_q) and weight on child development (θ_K). Given parameters governing altruistic, bequest, and guilt motives, all of which are identified out of the co-variation of prices and choices, the extent to which G uses childcare and bequests to induce P's LTC assistance provides useful information on G's distaste for formal LTC: the larger the disutility from formal LTC, the more G would want to induce P's LTC provision. The Medicaid recipiency rate is particularly informative about the consumption floor for G, while the labor force participation of low-education P is informative about welfare program benefits.

5.5 Auxiliary Models

Reflecting the identification argument in Section 5.4, we target the following auxiliary models.

Aux 1: G's Childcare Supply Response (q_i^G) to Childcare Prices (p_i^M). The coefficient α_1 from the following OLS regression is informative about G's altruism:

$$q_i^G = \alpha_1 p_i^M + \alpha_2 X_i + \varepsilon_i. \tag{20}$$

Aux 2: P's LTC Provision Response (n_i^P) to G's Past Childcare Provision (q_i^G). The OLS coefficient γ_1^{OLS} from (21) and the reduced-form IV estimate γ_1^R from (22) are informative for identifying guilt and unobserved heterogeneity in family types:

$$n_i^P = \gamma_1^{OLS} q_i^G + \gamma_2^{OLS} X_i + \varepsilon_i, \tag{21}$$

$$n_i^P = \gamma_1^R p_i^M + \gamma_2^R X_i + u_i. \tag{22}$$

Aux 3: Persistence in G's Choices. The coefficients β_1 and θ_1 from the following regressions

are informative for identifying G's permanent altruism types:

$$\ln(a_{i3}^G) = \beta_1 \ln(a_{i2}^G) + \beta_2 X_{i3} + \varepsilon_i, \tag{23}$$

$$\ln(a_{i3}^G) = \theta_1 q_{i1}^G + \theta_2 X_{i3} + u_i. \tag{24}$$

Here, a_{it}^G denotes G's asset holdings in stage t.

Aux 4: G's Wealth Accumulation Moments. We use the first moments of G's wealth in each stage, conditional on G's retirement income.

Aux 5: P's LTC Provision. To help identify P's altruism (ω_χ^P), we use P's LTC provision when G did not provide childcare and has limited assets. To help discipline P's inheritance motive, we use informal LTC provision rates by G's wealth and LTC need severity.

Aux 6: Childcare Arrangement Moments. We use G's childcare provision rates by P's education level and G's retirement income from the HRS. Information on formal childcare use comes from the SIPP. Because grandparental childcare is measured differently across the two datasets and therefore not directly comparable in levels, we focus on the composition of non-grandparental care. In particular, we use the share of formal childcare hours out of total non-grandparental care hours, $q^M/(q^M + q^P)$, by P's education level.

Aux 7: P's Labor Supply Moments. We use full-time and part-time employment rates in each stage by P's education level.

5.6 Estimation Results

Parameter Estimates. Using K-modes clustering, we classify families into two types (A and B) based on the Add Health data, which we detail in Online Appendix 3. Type A families exhibit systematically stronger parent–child relationships than Type B families. Table 6 shows the share of Type A households by parent-child education levels, which ranges from 69% to 74%.

Table 7 reports the estimates of the structural parameters obtained via indirect inference. G's altruism toward P is 0.92 for Type A families (the type exhibiting stronger parent-child relationships in Add Health); it is substantially lower, at 0.20, for Type B families. The type-weighted mean is 0.71, within the range of estimates from previous studies in which parents are homogeneous in

Table 6: Share of Type A Families by Education Group (%)

Mom\Child	HS or Less	Some College	College +
HS or Less	72.6	69.6	74.2
Some College +	69.3	71.4	69.4

Notes: The table reports the percentage of Type A families conditional on the education of the mother and the child, based on K-modes clustering using Add Health data. Type A families are the more altruistic type, exhibiting systematically stronger parent–child relationships than Type B families.

their altruism toward their children.[22] P's altruism toward G is estimated to be only 2.3% of G's altruism, underscoring pronounced asymmetry in intergenerational altruism.[23] This asymmetry implies that care transfers from P to G are not primarily motivated by direct concern for G's utility. Instead, P's transfers are driven by a strong desire to avoid guilt, which is estimated to be large, as well as by financial incentives to protect bequests from being depleted by formal LTC expenses, captured by P's weight on bequest utility.

We find that P derives a nontrivial amount of utility (θ_q) from spending time with her children. Given that P's contemporaneous utility is measured in log of consumption, the estimate of 2.6 implies that, for a mother with median first-stage annual consumption ($15,594), the willingness to pay for one additional hour per week (52 additional hours per year) with her child is 3.05% of her annual consumption.[24] The weight on the child's skill (θ_K) is estimated to be 0.76. Given that child skill is measured in log of test scores, this estimate implies that the willingness to pay for a one-standard-deviation improvement in the child's log test scores is approximately 1.17% of her consumption for a mother with median first-stage consumption, which is relatively low. However, we note that θ_K reflects a combination of preferences and beliefs. Rather than interpreting this finding as mothers placing a low value on their children's skill development, a more plausible explanation is that mothers hold biased beliefs about the skill production function, as suggested by previous studies (Cunha et al., 2013, 2022).[25]

[22] For example, parents' altruism parameter is estimated to be 0.32 in Lee and Seshadri (2019), 0.60 in Ko (2022), and 0.86 in Caucutt and Lochner (2020).

[23] In an alternative specification, we allow for type-specific scaling factors ($\frac{\omega_X^P}{\omega_X^G}$). The difference between the two scaling factors is economically and statistically zero.

[24] Notice that because we normalize the productivity of mothers' time to zero in the children's skill production function, θ_q captures both mothers' direct utility from childcare time and the indirect utility from her child's skill.

[25] For example, an overly optimistic (pessimistic) view of the productivity of grandparental (formal) care would lead to its overuse (underuse), which would be reflected in a low estimated value of θ_K. Given our data, it is impossible to distinguish between true preferences and beliefs; however, this does not affect our policy implications, because our counterfactual simulations in Section 6 alter neither preferences for nor beliefs about child skill development.

Table 7: Parameter Estimates from Indirect Inference

	Estimate
Panel A: P's Preference Parameter	
P altruism scaling factor ($= \frac{\omega_X^P}{\omega_X^G}$)	0.023
P leisure utility	1.590
P direct utility from child care	2.599
P weight on child's skill	0.755
P guilt when child care received	2.517
P weight on bequest utility	1.071
P FT employment cost (HS or less)	0.209
P FT employment cost (some college)	0.005
P FT employment cost (college +)	-0.006
P additive PT employment cost	0.624
P logit shock scale (t=1)	0.506
P logit shock scale (t=2)	0.843
P logit shock scale (t=3)	0.816
Panel B: G's Preference Parameter	
G altruism (Type A)	0.916
G altruism (Type B)	0.196
G leisure utility	1.533
G formal LTC disutility (light disability)	0.047
G formal LTC disutility (severe disability)	1.625
G logit shock scale (t=1)	0.135
Panel C: Others	
Medicaid annual consumption floor	5013
Childcare subsidy annual income threshold	14313
SNAP annual maximal benefit	3849

Notes: The table reports estimates from indirect inference. Logit shock scale parameters are reported in annual utility units. Flow utility parameters are annualized, except for the weight on the child's skill, which reflects a stock.

G's distaste for formal LTC is modest when disability is mild but becomes large in the severe state. For example, for a G with the median annual consumption, the willingness to pay to avoid one additional hour per week in formal care over a year when she has severe disability is 3.98% of her consumption. This pattern is intuitive, as severe disability is more likely to result in institutionalization in a nursing home, a form of formal LTC that is well documented to generate substantial disutility for retirees (Barczyk and Kredler, 2018; Mommaerts, 2025; Ko, 2022). To reduce the risk of nursing home entry, which occurs with positive probability for every G when viewed from Stage 1, G has strategic incentives to provide childcare earlier in life to activate the guilt mechanism on P's side and/or to accumulate savings that can be used as bequests to elicit informal LTC from P.

Model Fit. Appendix B reports tables and figures showing the model fit. Table B1 indicates that the model also performs well in matching childcare and LTC arrangements. It reproduces the increasing gradient in formal childcare use by P's education, though the levels are modestly overstated. The model closely matches P's LTC provision rates by P's education, while overpredicting part-time provision when G has light disability. The model matches the Medicaid recipiency rate well, while slightly underpredicting the childcare subsidy recipiency rate.

Table B2 shows that the model captures the main patterns in P's labor supply. It reproduces the rise in full-time employment over time and the higher full-time employment rates for more educated individuals. Some discrepancies remain, including a tendency to overpredict full-time work in later periods for lower-education groups.

Figure B1 shows that the model does a decent job of matching G's wealth over time across PI groups, although for high-PI households, the model slightly overpredicts wealth in period 2 and underpredicts it in period 3.

Table B3 reports the fit for selected targeted regression coefficients. The model reproduces the positive association between G's childcare and childcare prices, although the magnitude is somewhat underpredicted. It closely replicates the positive association between G's childcare and P's LTC provision. It also does a decent job of capturing persistence in G's choices, which is informative about her altruism types. Finally, the model replicates the negative association between P's formal childcare use and high childcare costs.

6 Counterfactual Policy Analysis

We conduct a series of counterfactual policy experiments to evaluate the impact of alternative childcare and LTC policies. To facilitate comparison across reforms, we calibrate the policy parameters so that each reform requires the same *direct* fiscal cost in the targeted sector.

6.1 Policy Design

Unconditional Childcare Subsidy (C1). The first policy, C1, replaces the existing means-tested childcare subsidy with an unconditional version that no longer ties eligibility to employment. The subsidy continues to be means-tested but removes the work requirement. The generosity of the subsidy is kept at 75%, implying that recipients continue to pay the remaining 25% of childcare expenses, as in the baseline. This design draws on evidence from Guner et al. (2020) that unconditional transfers can enhance welfare despite their negative impact on labor supply, and reflects a policy orientation that emphasizes better access and child development rather than employment incentives. We add to these findings by examining such reforms' spillover effects on the LTC sector. Under C1, government spending on childcare subsidies increases by \$4,657 per household. For the remaining reforms, we calibrate the policy parameters so that the discounted present value (PV) of the direct fiscal cost in the targeted sector is also \$4,657 per household.[26]

Medicaid Floor Expansion (C2). The second policy, C2, raises the Medicaid consumption floor for grandparents. We calibrate the consumption floor so that the resulting increase in the PV of Medicaid spending is \$4,657 per household, corresponding to a level 73% above the baseline.

Childcare-Contingent Medicaid Floor Expansion (C3). The third policy, C3, directly links the two care sectors by conditioning Medicaid generosity on grandparents' prior childcare provision. Specifically, we increase the Medicaid consumption floor for grandparents proportionately to the hours of childcare they provided in $t = 1$, calibrated so that the resulting increase in the PV of Medicaid spending equals \$4,657 per household. This implies a maximum increase of 331% over the baseline consumption floor, which applies to grandparents providing 40 hours of childcare per week (total childcare need is 84 hours per week).

[26] We use an annual discount factor of $\frac{1}{1+r}$, where r is the real interest rate, set at 3%. All present values are computed from the beginning of Stage 1.

Two aspects of C3 are worth discussing. First, it requires verification of grandparental childcare. Such monitoring appears feasible in practice: several European governments already track unpaid caregiving for the purpose of determining future social insurance benefits. For example, the Specified Adult Childcare Credits program in the U.K. grants pension credits to grandparents who provide regular childcare. Second, under C3, G has a "third" motive to provide childcare, in addition to the altruistic and strategic motives present in the baseline model: securing access to a more generous old-age social insurance program. Our baseline model assumes a motive-invariant guilt technology: P's guilt parameter is tied to the act of receiving childcare from G (and the associated internalized obligations), rather than to the perceived intention behind G's decision to provide care. Given this assumption, our counterfactual policy implications are better interpreted as short-run effects; in the long run, P's guilt "parameter" may evolve under sustained exposure to policies such as C3.

6.2 Policy Impacts

Table 8 shows the impacts of each of our counterfactual policies (Columns (1) to (3)); Column (4) reports the baseline mean of each outcome. In Panels A to D, we report policy impacts as percentage changes from the baseline, defined as the change in the average outcome divided by the average baseline level; in Panel E, we report level changes in the government budget in PV terms.

Care Arrangements. Panel A of Table 8 shows the effects on childcare arrangements. By removing the employment requirement for childcare subsidies, C1 shifts households toward market-based childcare, substituting primarily away from grandparental care rather than parental care. These effects are especially pronounced among lower-educated parents, reflecting the means-tested nature of the subsidies. The results are qualitatively consistent with the descriptive evidence from the CCDF variations in Section 2.[27] Increasing the Medicaid consumption floor (C2) indirectly reduces grandparental childcare. As the public LTC insurance option becomes more attractive, G has a lower incentive to elicit informal LTC from P, which weakens the strategic provision of childcare in $t = 1$. However, because the expansion of the Medicaid program is modest and because G is altruistic, the reduction in grandparental childcare is limited. C3 has the opposite effect: by explicitly rewarding childcare provision with more generous Medicaid benefits, C3 substantially

[27] A direct quantitative comparison is infeasible because the program generosity index used in the descriptive analysis draws on multiple policy dimensions to ensure sufficient variation over the sample period.

increases grandparental childcare. Anticipating more help from G, P reduces her usage of formal childcare; moreover, G's help also lowers parental childcare hours.

Panel B presents the effects on P's LTC provision in $t = 3$ when G is severely disabled.[28] Under C1, the decline in grandparental childcare weakens P's reciprocal obligation to provide LTC. This is consistent with the descriptive evidence in Section 2, which shows that more generous childcare subsidies reduce family caregiving across both the childcare and LTC sectors. Similarly, under C2, grandparents reduce strategic childcare provision, and Ps reciprocate with lower LTC, although the effect is smaller. C3 moves in the opposite direction and strengthens intergenerational reciprocity: P's provide more LTC hours after receiving greater childcare support from their parents.

P's Labor Supply and G's Savings. Panel C summarizes changes in parental labor supply in $t = 1$.[29] The unconditional subsidy C1 substantially reduces work hours across all education groups, consistent with Guner et al. (2020). C2 has little effect on parental work hours, reflecting the modest decline in grandparental childcare under that policy. C3, by contrast, increases work hours, particularly among low-educated parents, as greater childcare support from grandparents reduces their childcare responsibilities.

Panel D shows policy impacts on G's wealth accumulation. Under C2 and C3, with a more generous Medicaid consumption floor, G has smaller precautionary savings motives and, in later stages, larger dis-saving incentives to qualify for Medicaid, resulting in a nontrivial decrease in her wealth. In comparison, C1 has a very small but positive impact on G's savings on average.

Government Budget. Panel E reports the fiscal effects of each reform, with all items reported in PV terms. Although the PV of the direct fiscal cost of each policy is calibrated to be the same at \$4,657 per household in the targeted sector, spillovers across care sectors generate substantial differences in net government spending. C1 directly increases childcare subsidy expenditures by removing the work requirement and, indirectly, raises Medicaid spending by reducing informal LTC and increasing reliance on formal care. In addition, by weakening work incentives, C1 increases welfare transfers and reduces tax payments by P. Taken together, these spillovers almost double the direct fiscal cost: the net government spending PV increases by \$9,150 per household.

C2 directly increases Medicaid spending, and by discouraging childcare provision by G, it slightly

[28] The impact when G is lightly disabled is minimal (Table 4.7).

[29] Changes in $t = 2$ and $t = 3$ are minimal (Table 4.7), suggesting that dynamic adjustments in labor supply are limited within our model.

Table 8: Counterfactual Policy Impacts

	(1) C1 Unconditional Childcare Subsidy	(2) C2 Medicaid Floor Expansion	(3) C3 Childcare-Contingent Medicaid Floor Expansion	(4) Baseline
Panel A: Childcare arrangements (% change from baseline)				hrs/week
Grandparental childcare				
P's edu = HS or less	-23.02	-4.81	73.86	10.1
P's edu = Some college	-14.66	-3.42	35.06	12.6
P's edu = College +	-11.66	-2.71	31.13	11.7
Formal childcare				
P's edu = HS or less	47.81	0.75	-3.47	7.8
P's edu = Some college	20.29	0.38	-1.84	10.4
P's edu = College +	11.44	0.15	-1.32	13.6
Parental childcare				
P's edu = HS or less	-2.13	0.64	-10.81	66.2
P's edu = Some college	-0.43	0.64	-6.96	60.9
P's edu = College +	-0.33	0.51	-5.89	58.7
Panel B: P's LTC hours when G is severely disabled (% change from baseline)				hrs/week
P's edu = HS or less	-7.51	-2.21	26.14	17.8
P's edu = Some college	-5.02	-2.28	15.40	19.4
P's edu = College +	-4.74	-1.88	14.38	18.5
Panel C: P's work hours in $t = 1$ (% change from baseline)				hrs/week
P's edu = HS or less	-36.89	-1.29	9.62	13.6
P's edu = Some college	-21.97	-0.59	3.14	21.2
P's edu = College +	-16.61	-0.44	2.32	22.7
Panel D: G's wealth (% change from baseline)				$
$t = 1$	0.01	-0.20	-0.33	405659
$t = 2$	0.07	-1.39	-1.93	490286
$t = 3$	0.01	-2.16	-2.64	346226
Panel E: Fiscal effects per household (PV $ change from baseline)				PV $
Childcare subsidy	4657	12	-81	883
Medicaid	1371	4657	4656	27778
P's income tax	-3123	-62	367	23067
Change in net spending	9150	4731	4209	-

Notes: Columns (1)–(3) report the changes in average values relative to baseline: Panels A–D are in % terms, and Panel E is in PV $ terms. Column (4) reports the mean values in the baseline economy: weekly hours in childcare, LTC, and work, $ in G's wealth, and PV in government spending and tax revenues. Policy C1 removes the work requirement for the means-tested childcare subsidy. Policy C2 raises the Medicaid consumption floor by 73%. Policy C3 increases the baseline Medicaid consumption floor for grandparents in proportion to childcare hours provided in $t = 1$, up to a maximum increase of 331% for full-time childcare. Each reform is calibrated so that the present-discounted value of the direct fiscal cost in the targeted sector equals $4657 per household. In the baseline economy, the annual Medicaid consumption floor is $5013; childcare subsidy recipiency rate is 4%; and Medicaid recipiency rate is 13%.

increases childcare usage, leading to a modest rise in childcare subsidy expenditures. Consistent with the small decline in grandparental childcare time, labor supply responses are also limited, resulting in a modest reduction in taxes paid by P. As a result, the net fiscal cost PV of C2 is \$4,731 per household, close to its direct cost.

In contrast, C3 generates spillovers that operate in the opposite direction. By rewarding prior childcare provision, C3 encourages G to provide more childcare, which reduces P's usage of formal childcare and increases P's labor supply. The lower usage of formal childcare and the increase in P's income jointly reduce the government's childcare subsidy spending by \$81 per household (a 9.2% reduction from the baseline childcare spending). Meanwhile, the increase in P's income also leads to a higher tax revenue by \$367 per household. As a result, the net fiscal cost PV of C3 is \$4,209 per household (lower than its direct cost), making it the least costly among all three reforms.

Summary. Taken together, these counterfactual experiments underscore how childcare and LTC policies are interconnected through intergenerational family dynamics. Policies that expand access to market-based care in one sector can weaken informal care networks and increase families' reliance on social insurance programs in the other. In contrast, policies that create dynamic incentives for family-provided care can be more cost-effective by strengthening intergenerational time transfers.

7 Conclusion

This paper studies how families allocate time-intensive care across generations and over the life cycle. Using evidence from the Health and Retirement Study and cross-state variation in childcare subsidies, we document that lower childcare costs reduce reliance on grandparent-provided childcare and subsequently reduce informal LTC received by grandparents. These patterns are difficult to reconcile with purely altruistic behavior and point to an important strategic component in intergenerational care decisions.

To interpret these facts, we develop and estimate a dynamic non-cooperative model of extended families that choose care arrangements, labor supply, savings, and bequests under both altruistic and strategic motives. The estimated model indicates that strategic reciprocal considerations play a central role in linking early childcare assistance to later LTC provision. By embedding these motives in a unified life-cycle framework, the model provides a quantitative account of how care decisions today shape incentives and outcomes many years later.

These dynamics have important implications for policy evaluation. Counterfactual policy simulations reveal substantial spillover effects across the two care sectors. Policies that focus on expanding access to market care in one sector could weaken informal care networks and increase families' reliance on social insurance programs in the other care sector. In contrast, policies that provide dynamic incentives for family-provided care could strengthen intergenerational transfers and reduce total government expenditure. More broadly, these results highlight the importance of modeling families as dynamic systems in which transfers are shaped by both preferences and incentives. Accounting for strategic interactions within families is essential for understanding caregiving behavior and for designing policies that affect care provision across generations.

Our framework and analysis open several important avenues for future research. We have studied only a part of within-family dynamics; a richer framework would also include other family members (e.g., siblings) to gain a more complete picture of how extended families interact across generations and within a generation. Another extension of our framework would be to incorporate migration decisions of the family, which would endogenize residential proximity between P and G and thus the feasibility of time transfers across generations.

References

Angrist, J. D. and Evans, W. N. (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *American Economic Review*, 88(3):450–477.

Anstreicher, G. and Venator, J. (2024). To grandmother's house we go: Informal childcare and female labor mobility. Unpublished.

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

Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy*, 82(6):1063–1093.

Becker, G. S. (1993). The economic way of looking at life. *Journal of Political Economy*, 101(3):385–409.

Becker, G. S., Murphy, K. M., and Spenkuch, J. L. (2016). The manipulation of children's preferences, old-age support, and investment in children's human capital. *Journal of Labor Economics*, 34(S2):S3–S30.

Bernal, R. and Keane, M. P. (2011). Child care choices and children's cognitive achievement: The case of single mothers. *Journal of Labor Economics*, 29(3):459–512.

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

Bick, A. (2016). The quantitative role of child care for female labor force participation and fertility. *Journal of the European Economic Association*, 14:639–668.

Bonhomme, S., Lamadon, T., and Manresa, E. (2022). Discretizing unobserved heterogeneity. *Econometrica*, 90(2):625–643.

Brown, J. R. and Finkelstein, A. (2008). The interaction of public and private insurance: Medicaid and the long-term care insurance market. *American Economic Review*, 98(3):1083–1102.

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

Castilla, C. and Walker, T. (2013). Is ignorance bliss? the effect of asymmetric information between spouses on intra-household allocations. *American Economic Review: Papers & Proceedings*, 103(3):263–268.

Caucutt, E. M. and Lochner, L. (2020). Early and late human capital investments, borrowing constraints, and the family. *Journal of Political Economy*, 128(3):1065–1147.

Chang, M. and Ko, A. (2026). Intrahousehold inequality and savings in retirement. Working paper.

Cil, G. and Tuncbilek, Z. E. (2018). The effect of parental health shocks on female labor supply. *Journal of Human Resources*, 53(4):1118–1153.

Compton, J. and Pollak, R. A. (2011). Proximity and coresidence of adult children and their parents: Description and correlates. Working Paper 17120, National Bureau of Economic Research.

Cox, D. and Stark, O. (2005). On the demand for grandchildren: Tied transfers and the demonstration effect. *Journal of Public Economics*, 89(9-10):1665–1697.

Cunha, F., Elo, I., and Culhane, J. (2013). Eliciting maternal expectations about the technology of cognitive skill formation. Working Paper 19144, National Bureau of Economic Research.

Cunha, F., Elo, I., and Culhane, J. (2022). Maternal subjective expectations about the technology of skill formation predict investments in children one year later. *Journal of Econometrics*, 231(1):3–32.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2025). Why do couples and singles save during retirement? household heterogeneity and its aggregate implications. *Journal of Political Economy*, 133(3):750–792.

Flood, S., McMurry, J., Sojourner, A., and Wiswall, M. (2022). Inequality in early care experienced by us children. *Journal of Economic Perspectives*, 36(2):199–222.

García-Morán, E. and Kuehn, Z. (2017). With strings attached: Grandparent-provided child care and female labor market outcomes. *Review of Economic Dynamics*, 23:80–98.

Guner, N., Kaygusuz, R., and Ventura, G. (2020). Child-related transfers, household labour supply, and welfare. *Review of Economics Studies*, 87(5):2290–2321.

Ho, C. (2015). Welfare-to-work reform and intergenerational support: Grandmothers' response to the 1996 prwora. *Journal of Marriage and Family*, pages 407–423.

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Ko, A. (2022). An Equilibrium Analysis of the Long-Term Care Insurance Market. *Review of Economic Studies*, 89(4):1993–2025.

Lee, S. Y. and Seshadri, A. (2019). On the intergenerational transmission of economic status. *Journal of Political Economy*, 127(2):855–921.

Lewbel, A. (2000). Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables. *Journal of Econometrics*, 97(1):145–177.

Lundberg, S. and Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of Political Economy*, 101(6):988–1010.

Mommaerts, C. (2025). Long-Term Care Insurance and the Family. *Journal of Political Economy*, 133(1):1–52.

Skira, M. (2015). Dynamic wage and employment effects of elder parent care. *International Economic Review*, 56(1):63–93.

A Appendix: Functional Forms

Flow Utility. Both agents have a log–additive separable utility over consumption and leisure:

$$u^P(c_t^P, l_t^P, h_t^P) \;=\; \ln(c_t^P) \;+\; \alpha^P \ln(l_t^P) \;-\; \eta(h_t^P), \tag{25}$$

$$u^G(c_t^G, l_t^G) \;=\; \ln(c_t^G) \;+\; \alpha^G \ln(l_t^G), \tag{26}$$

where α^P and α^G are estimated leisure weights. $\eta(h_t^P)$ is P's fixed utility cost of employment:

$$\eta(h_t^P) \;=\; \begin{cases} \eta^{edu} + \eta^{PT} & \text{if } h_t^P = h^{PT}, \\ \eta^{edu} & \text{if } h_t^P = h^{FT}, \\ 0 & \text{if } h_t^P = 0, \end{cases} \tag{27}$$

with $\eta^{edu} \in \{\text{HS or less, some college, college +}\}$ being the full-time participation cost specific to each education group, and $\eta^{PT} \geq 0$ an additional cost for part-time work.

Child's Skill Production Function. Following Bernal and Keane (2011), K is expressed in log test score units. Based on their estimates, we use

$$K = 4.5 - 0.00751\, Q^G + 0.00472\, Q^M, \tag{28}$$

where Q^G and Q^M denote quarters of grandparental and formal care, respectively. Maternal time is the omitted baseline.[30]

[30] We treat part-time use as using two quarters and full-time use as using four quarters, resulting in $2T_1$ and $4T_1$ quarters, respectively, where T_1 is the years in Stage 1. The constant term is a normalization chosen

LTC Disutility. G's disutility from formal LTC, $\psi\Big(N(X_3^G), n^P\Big)$, is proportional to the share of LTC needs met by the formal market:

$$\psi\Big(N(X_3^G), n^P\Big) \;=\; \begin{cases} \psi_1 \dfrac{N(X_3^G) - n^P}{N(X_3^G)} & \text{if } N(X_3^G) = N_1 \text{ (light disability)}, \\ \psi_2 \dfrac{N(X_3^G) - n^P}{N(X_3^G)} & \text{if } N(X_3^G) = N_2 \text{ (severe disability)}, \\ 0 & \text{if } N(X_3^G) = 0 \text{ (healthy)}, \end{cases} \tag{29}$$

where $N_1 < N_2$ are the two positive LTC need levels. Allowing $\psi_1 \neq \psi_2$ captures health-dependent preferences for formal care. Although our model does not distinguish between types of formal care, this flexibility indirectly accommodates the empirical pattern that care type and disability severity are correlated: paid home care predominates in lighter disability states, whereas nursing home care becomes more prevalent in severe states. Since retirees may have substantially stronger aversion to institutional care, we allow ψ_2 to be different from ψ_1.

Bequest Utility. We model P's bequest utility as a parsimonious approximation that captures that P values inheriting wealth, and that the marginal value of the bequest diminishes with P's own resources:

$$B(a_t^G) \;=\; \frac{\theta_b}{a_d} \ln\Big(y_d^P + a_d\, a_t^G\Big), \tag{30}$$

where $a_d = (1-\beta)/(1-\beta^{T_d})$ is an annuity factor over a calibrated retirement window T_d that P will have, y_d^P is P's non-labor income, and $\theta_b > 0$ is estimated via indirect inference.

P's Experience Accumulation. The experience function is:

$$H_t^P \;=\; H_{t-1}^P \;+\; T_t\, \mathbf{1}_{h_t^P = h^{FT}} \;+\; \xi\, T_t\, \mathbf{1}_{h_t^P = h^{PT}}, \tag{31}$$

where T_t is the duration of stage t in calendar years and $\xi = 0.6$ is the experience credit per year of part-time work.

G's Retirement Income. G's retirement income y_t^G follows the fixed-effects specification of De Nardi et al. (2025):

$$\ln y_t^G \;=\; \hat{\kappa}(\text{age}_t^G) \;+\; \hat{h}(PI), \tag{32}$$

where $\hat{\kappa}(\cdot)$ is a cubic polynomial in G's age, and $\hat{h}(PI)$ is a PI group fixed effect.

B Appendix: Model Fit Tables and Figures

to match the mean of the outcome variable in Bernal and Keane (2011) and does not affect model predictions (the mean log test score is 4.5 with a standard deviation of 0.186).

Table B1: Model Fit for Care Arrangements and Program Recipiency

	Empirical	Simulated
Panel A: G-provided childcare rate		
G's income = low	0.463	0.492
G's income = mid	0.502	0.579
G's income = high	0.559	0.509
Panel B: Formal childcare hrs out of non-grandparental time		
P's edu = HS or less	0.069	0.101
P's edu = Some college	0.115	0.138
P's edu = College +	0.153	0.177
Panel C: P's LTC supply		
By P's education:		
P's edu = HS or less	0.577	0.594
P's edu = Some college	0.582	0.615
P's edu = College +	0.561	0.596
G has light disability:		
P provides 50% of required LTC hours	0.252	0.327
P provides 100% of required LTC hours	0.245	0.268
G has severe disability:		
P provides 50% of required LTC hours	0.294	0.296
P provides 100% of required LTC hours	0.305	0.315
Panel D: Program Recipiency Rates		
Medicaid recipiency rate	0.131	0.129
Childcare subsidy recipiency rate	0.055	0.043

Notes: Empirical G-provided childcare (Panel A), P-provided LTC supply (Panel C), and Medicaid receipt (Panel D) are computed using HRS data. Empirical formal childcare moments (Panel B) are computed using SIPP data. Childcare subsidy receipt is from Guner et al. (2020).

Table B2: Model Fit for P's Employment

	Full-time		Part-time	
	Empirical	Simulated	Empirical	Simulated
Panel A: Employment by P's edu and t				
P's edu = HS or less				
$t = 1$	0.330	0.263	0.155	0.205
$t = 2$	0.458	0.480	0.171	0.190
$t = 3$	0.443	0.508	0.132	0.123
P's edu = Some college				
$t = 1$	0.444	0.459	0.210	0.188
$t = 2$	0.551	0.522	0.197	0.210
$t = 3$	0.555	0.625	0.147	0.130
P's edu = College +				
$t = 1$	0.487	0.491	0.188	0.205
$t = 2$	0.567	0.558	0.197	0.203
$t = 3$	0.659	0.656	0.142	0.118
Panel B: Employment when G healthy				
$t = 3$	0.660	0.664	0.102	0.087

Notes: Employment moments by P's education (Panel A) are computed using SIPP data. Employment when G is healthy (Panel B) is computed using HRS data, where G's health status is observed.

Table B3: Model Fit for Targeted Regression Coefficients

Dependent variable	Key regressor	Empirical	Simulated
G provides childcare	Childcare price is high	0.064	0.033
P provides LTC	G provided childcare (OLS)	0.154	0.141
log(G wealth at $t = 3$)	G provided childcare	0.323	0.394
log(G wealth at $t = 3$)	log(G wealth at $t = 2$)	0.877	0.963
P uses formal childcare	Childcare price is high	-0.081	-0.074

Notes: All empirical regression coefficients except the last row are computed using HRS data. The last row is computed using SIPP data.

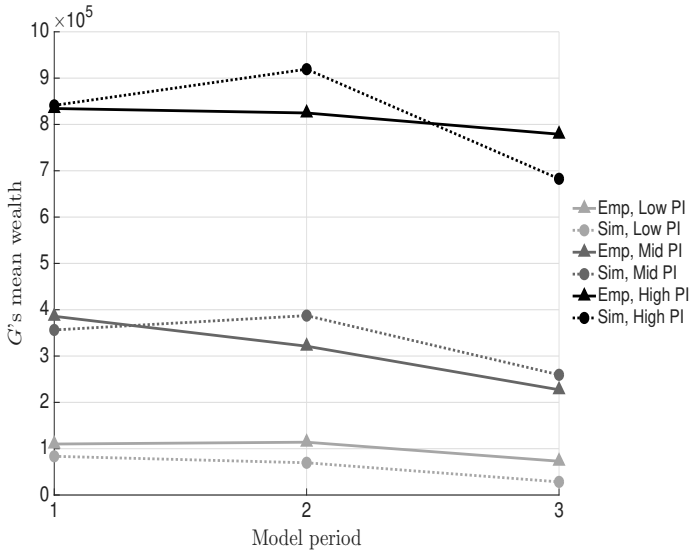


Figure B1: Model Fit for G's Wealth Evolution by Permanent Income (PI)

Notes: Darker colors represent higher PI. Solid lines with triangle markers represent empirical moments, and dashed lines with circle markers represent simulated moments.

Intergenerational Time Transfers: Altruism and Strategic Motives within the Family (Online Appendix)

Chao Fu Ami Ko Chamna Yoon

March 2026

1 Model Details: Accounting for Stage-Specific Length

As noted in the main text, stage durations differ, calibrated to match life-cycle timing in the data: Stage 1 lasts $T_1 = 5$ years (the child's first five years), Stage 2 lasts $T_2 = 14$ years (until G's mean age of 77 at the LTC onset), and Stage 3 lasts $T_3 = 18$ years (until G turns 95). In a given stage, since an action is held fixed for all T_t years, annual flow utility is scaled by the geometric sum $(1 - \beta^{T_t})/(1 - \beta)$. The exception is the utility from the child's skill in Stage 1, which is treated as a stock and is therefore not annualized. Continuation values are discounted by β^{T_t} rather than β. In addition, we use the agents' mean age within a given stage as the input to income functions and health transition matrices.

2 Data Details: HRS Sample Selection

For each G, we pick one P. If we were to pick the P who is observed to provide the most long-term care hours to G, then she is likely to be too old for us to observe whether G provided childcare help when P had young kids. In contrast, if we picked the P based solely on childcare hours from G, then we risk picking the wrong P who might not be the primary long-term caregiver for G. Our solution is a compromise between the two, involving three steps. First, using families that are observed to rely on informal long-term care, we regress the dummy for whether the child is the primary caregiving child on a vector of child characteristics that are exogenous to the model.[1] We use the estimated coefficients to predict each P's probability of being the major caregiver.

Second, we restrict the sample to P's for whom we have credible information about G's childcare provision. This is not a trivial task, as we do not observe the ages of P's children and can only

[1] To construct the dependent variable, we use the following criteria until ties are broken: (i) the number of waves where the child is reported to provide informal LTC, (ii) the number of help days, and (iii) the number of help hours. The regressors include the child's age, proximity to the parent, education, marital status, number of kids, number of siblings, and whether the child is the oldest child.

infer new births from changes in the reported number of children. We first keep P's for whom we observe their first birth.[2] Assuming they have all of their subsequent children during the sample period (mean number of children P has is around two), we can infer G's entire childcare history. For P's that enter the sample with kids already, we keep them if they ever receive childcare help from G at some point during the sample period. Applying this second-step restriction results in 2,426 G's and 3,341 P's.[3]

Third, for each G, we select the P with the highest predicted probability of being the primary caregiver computed in the first step.

3 Estimation of Unobserved Types from Add Health

To estimate the distribution of unobserved altruism types, we use data from the National Longitudinal Study of Adolescent to Adult Health (Add Health), a longitudinal survey of a nationally representative sample of U.S. adolescents. The study began in 1994–1995 with 20,745 respondents in grades 7–12 and follows them into adulthood across five waves, with the most recent wave conducted in 2015–2018 when respondents were ages 33–42.[4]

To align with our model, which focuses on interactions between adult children and their parents, we use data from Waves 3–5. We restrict attention to 2,557 individuals who are observed in all three waves. Our goal is to use information on the perceived quality of the parent–child relationship to recover heterogeneity in altruistic preferences.

We focus on survey questions that capture closeness, warmth, enjoyment of time together, and the frequency and quality of interactions with the mother. The specific questions used are listed below:

1. How close do you feel to your mom (Waves 3–5)? [close_mom]
2. How warm was your mom growing up (Wave 3)? [warm_mom]
3. How much do you enjoy your time with your mom (Wave 3)? [enjoy_mom]
4. How often do you see your mom (Wave 4)? [seeoften_mom]
5. How good is the quality of communication with your mom (Wave 4)? [gdcomm_mom]
6. How frequently do you communicate with your mom (Waves 4–5)? [frcomm_mom]

[2] We treat a case as a first birth when P's reported number of children increases from zero in the previous wave to a strictly positive value.

[3] The second-step restriction excludes G's for whom we cannot credibly impute their childcare provision to any of their daughters. That is, their daughters enter the sample already with kids, and G never provide childcare to them over the sample period. We do not use these families because if we did, we would treat G as never providing childcare, when in fact, they could have provided care before the sample period started, when P's kids were young.

[4] The five waves were conducted in 1994–1995 (Wave 1), 1996 (Wave 2), 2001–2002 (Wave 3), 2008 (Wave 4), and 2015–2018 (Wave 5). Parents were surveyed in Waves 1 and 5.

We classify individuals into two types using the K-modes clustering algorithm, which is well suited for categorical data. Analogous to K-means, K-modes partitions observations into a pre-specified number of clusters by minimizing within-cluster dissimilarity, defined here as the number of mismatches between an observation and the cluster mode. The algorithm iterates between assigning observations to clusters and updating cluster modes until convergence.

After clustering, we label the two groups based on their survey responses. One group (Type A) exhibits systematically stronger parent–child relationships across all measures, while the other group (Type B) exhibits weaker relationships. Figure 3.1 illustrates these differences. Relative to Type B, Type A respondents report feeling closer to their mothers, experiencing greater warmth during childhood, enjoying time together more, and maintaining more frequent and higher-quality interactions in adulthood. We interpret Type A as the more altruistic type.

Finally, we estimate the distribution of altruism types conditional on educational attainment, reported in Table 6 of the main text.

Figure 3.1: Survey Answers by Altruism Type



4 Additional Tables and Figures

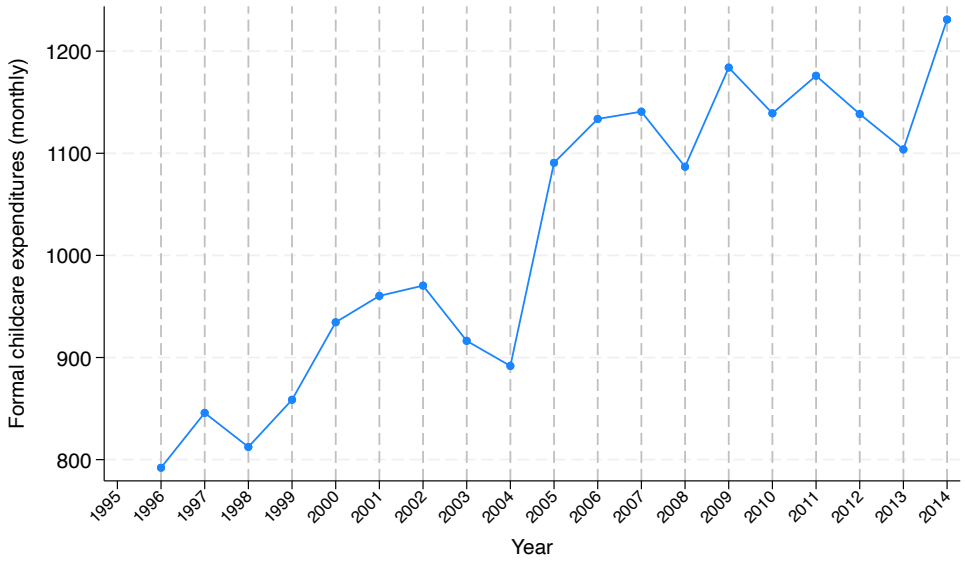


Figure 4.1: Formal Childcare Expenditures Over Time

Notes: Data = Consumer Expenditure Survey Public Use Microdata. Inflation adjusted and shown in 2019 dollars.

Table 4.1: Summary Statistics on Transfers at G-wave level

	Mean	SD
G provides childcare to any P	0.21	0.41
: G has a young grandchild	0.43	0.50
G receives informal LTC from any P	0.09	0.29
: G has LTC needs	0.63	0.48
G's mean financial transfer to P	887.18	6460.65
G's mean financial transfer from P	64.82	652.94
G has any P as a beneficiary of life insurance	0.38	0.48
Unique G	11567	
Unique G-wave combinations	76437	

Notes: Statistics are computed from G-wave–level data.

Table 4.2: Relationship between Intergenerational Time Transfers

	(1)		(2)		(3)	
	P provides LTC		G provided childcare		P provides LTC	
G provided childcare	0.102***	(0.036)				
G: Age	-0.000	(0.004)	0.021***	(0.003)	0.002	(0.004)
G: Married	-0.066	(0.051)	0.001	(0.014)	-0.066	(0.051)
G: log assets	0.002	(0.003)	0.001	(0.001)	0.002	(0.003)
G: log income	-0.005	(0.011)	0.001	(0.001)	-0.005	(0.011)
G: 0-1 ADL limitations	0.000	(.)	0.000	(.)	0.000	(.)
G: 2-3 ADL limitations	0.060***	(0.021)	-0.004	(0.008)	0.059***	(0.021)
G: 4-5 ADL limitations	0.099***	(0.028)	-0.004	(0.008)	0.098***	(0.028)
G: says health is excellent/very good	0.000	(.)	0.000	(.)	0.000	(.)
G: says health is good	0.001	(0.032)	-0.006	(0.011)	0.002	(0.032)
G: says health is fair/poor	-0.008	(0.028)	0.006	(0.011)	-0.006	(0.028)
P: Age	-0.000	(0.002)	-0.017***	(0.003)	-0.002	(0.002)
P: Years of education	-0.005	(0.007)	0.008	(0.007)	-0.004	(0.008)
P: Lives within 10 mi to G	0.376***	(0.023)	0.066***	(0.016)	0.383***	(0.023)
P: Number of children	-0.013*	(0.007)	0.015*	(0.008)	-0.012	(0.007)
P: Married	-0.010	(0.026)	0.029	(0.025)	-0.007	(0.026)
G is a Medicaid recipient	-0.044	(0.029)	-0.008	(0.009)	-0.044	(0.028)
G has private LTC insurance	-0.020	(0.042)	-0.011	(0.015)	-0.021	(0.043)
Childcare price was high			0.030**	(0.013)	0.076	(0.063)
Observations	4344		4370		4344	
Adjusted R^2	0.375		0.615		0.372	
Mean of dependent variable	0.387		0.240		0.387	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level and are in parentheses. The table reports estimates not reported in Table 4 of the main text.

Table 4.3: Childcare Subsidies and Intergenerational Time Transfers – All Education Levels

	(1)		(2)		(3)	
	P works full time		G provides child care		G believes P will help	
Treat=1	-0.013	(0.031)	0.047**	(0.024)	0.020	(0.028)
Generosity of CCDF	0.091	(0.070)	-0.039	(0.054)	-0.050	(0.097)
Treat=1 × Generosity of CCDF	0.128*	(0.070)	-0.069	(0.055)	-0.057	(0.070)
Observations	8713		8931		8793	
Adjusted R^2	0.072		0.115		0.098	
Mean of dependent variable	0.665		0.130		0.425	
Family fixed effects	Yes		Yes		Yes	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level and reported in parentheses. Compared to Table 5 of the main text, the sample includes *P*s of all education levels.

Table 4.4: Summary Statistics of P in the HRS Estimation Sample

	Mean	SD
Female	1.00	0.00
Age	45.11	8.67
Education (years)	13.68	2.19
Married	0.70	0.46
Lives within 10 mi to G	0.52	0.50
# Children	2.10	1.37
Works full-time	0.63	0.48
Works part-time	0.12	0.32
Provides informal LTC to G	0.08	0.27
Observations	16064	

Notes: See 4.1 for the sample construction.

Table 4.5: Summary Statistics of G in the HRS Estimation Sample

	Mean	SD
Female	1.00	0.00
Age	71.49	8.49
Education (years)	11.81	3.01
Provides childcare	0.26	0.44
Retired	1.00	0.00
Assets (in thousand)	406.46	571.17
Annual income (in thousand)	51.64	88.27
# Children	3.37	1.92
Have LTC needs	0.23	0.42
Observations	16064	

Notes: See 4.1 for the sample construction.

Table 4.6: Summary Statistics of P in the SIPP Core Sample

	Mean	SD
Female	1.00	0.00
Age	42.47	12.17
HS or less	0.36	0.48
Some college	0.36	0.48
College and plus	0.28	0.45
Married	0.57	0.50
Work experience (years)	11.98	9.67
Works full-time	0.52	0.50
Works part-time	0.16	0.37
Full-time earnings (per month)	2880.59	2636.01
Part-time earnings (per month)	1067.07	1488.86
Total family income (per month)	5420.59	5495.10
Observations	865484	

Notes: The sample is at the SIPP respondent-wave level. The sample is constructed using 2001, 2004 and 2008 panels. Monetary values are in 2019 dollars.

Table 4.7: Counterfactual Policy Impacts: Additional Outcomes

	(1) C1 Unconditional Childcare Subsidy	(2) C2 Medicaid Floor Expansion	(3) C3 Childcare-Contingent Medicaid Floor Expansion	(4) Baseline
Panel A: P's work hours in $t = 2$ (% change from baseline)				hrs/week
P's edu = HS or less	-0.75	-0.01	0.12	22.1
P's edu = Some college	-0.37	-0.02	0.04	24.1
P's edu = College +	-0.42	-0.00	0.08	25.4
Panel B: P's work hours in $t = 3$ (% change from baseline)				hrs/week
P's edu = HS or less	-0.78	0.04	-0.22	22.2
P's edu = Some college	-0.41	0.02	-0.14	26.9
P's edu = College +	-0.28	0.02	-0.04	28.0
Panel C: P's LTC hours when G is lightly disabled (% change from baseline)				hrs/week
P's edu = HS or less	-0.30	-0.20	0.49	6.5
P's edu = Some college	-0.35	-0.05	0.21	6.6
P's edu = College +	-0.26	0.01	0.05	6.3
Panel D: G's wealth by PI (% change from baseline)				$
Low-PI G:				
$t = 1$	-0.10	-1.27	-2.34	83544
$t = 2$	-0.59	-9.98	-19.01	69620
$t = 3$	-1.47	-21.93	-34.50	28359
Middle-PI G:				
$t = 1$	0.00	-0.26	-0.44	356328
$t = 2$	0.01	-2.36	-3.36	387266
$t = 3$	-0.22	-4.06	-5.39	259526
High-PI G:				
$t = 1$	0.03	-0.04	-0.04	841217
$t = 2$	0.15	-0.38	-0.21	919316
$t = 3$	0.17	-0.60	-0.29	682681

Notes: Columns (1)–(3) report percent changes relative to baseline. Column (4) reports the mean values in the baseline economy: LTC and work hours are in weekly terms. Policy C1 removes the work requirement for the means-tested childcare subsidy. Policy C2 raises the Medicaid consumption floor by 73%. Policy C3 increases the baseline Medicaid consumption floor for grandparents in proportion to childcare hours provided in $t = 1$, up to a maximum increase of 331% for full-time childcare. Each reform is calibrated so that the present-discounted value of the direct fiscal cost in the targeted sector equals $4657 per household. Panel C is restricted to Gs who are lightly disabled.