Insuring Labor Income Shocks: The Role of the Dynasty*

Andreas Fagereng, Luigi Guiso, Luigi Pistaferri, Marius A. K. Ring

September 8, 2025

Abstract

We provide empirical evidence on the importance of a relatively understudied channel of insurance against labor income shocks: transfers from (cash-rich) parents to (cash-short) children when the latter experience negative wage shocks. Matching population data for Norway across two generations, we establish several results. First, parents make a transfer—i.e., run down liquid assets—when adult children experience negative labor income shocks. Consistent with dynastic insurance, we observe no transfers when income shocks are positive. Second, parents' responses depend on the nature of the shock. If losses are temporary, parents dissave; if they are persistent, parents save in order to make future transfers. Parental transfers offset 43% of temporary and 27% of persistent losses. Third, insurance is lower when children have other smoothing options, like spousal labor supply, and is greater for shocks to their own child versus a child's spouse. Support also increases if the spouse's parents can contribute, suggesting "competition for attention." Lastly, insurance flows are one-way: children do not insure their parents against income losses.

Keywords: Income shocks, Insurance, Financial wealth, Intergenerational transfers.

JEL codes: D31, E21, E24, G11

*Andreas Fagereng is at BI Norwegian Business School and affiliated with CEPR and Statistics Norway; Luigi Guiso is at EIEF and affiliated with CEPR; Luigi Pistaferri is at Stanford University and affiliated with SIEPR, NBER and CEPR; Marius Ring is at the University of Texas at Austin. This article has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement 851891). Thanks to Sylvain Catherine, Alex Michaelides, and seminar participants at EIEF, Bristol University, Cambridge University, New York FED, San Francisco FED, the CEPR European Conference on Household Finance, the NBER Summer Institute conference on Household Finance, the Stone-UCL-LSE Workshop on Wealth and Capital Taxation, and the 3D-In-Macro Workshop in Oslo for comments. All errors are ours.

1 Introduction

It is notoriously difficult to rely on formal credit and insurance markets to cope with labor income shocks. This has led people and societies to resort to other mechanisms to smooth labor market adversities. A whole strand of research has studied the role of such arrangements, focusing on progressive taxation, government transfers, reliance on own savings, as well as risk sharing among spouses (see Blundell, Pistaferri, and Preston 2008; Guler, Guvenen, and Violante 2012; Ortigueira and Siassi 2013; Heathcote, Storesletten, and Violante 2014; Blundell, Pistaferri, and Saporta-Eksten 2016). All these channels have been shown to help workers smooth consumption, though to different degrees. With a few notable exceptions, there is limited evidence on what may be the most obvious source of insurance for most young workers: help from parents.[1]

In this paper, we study whether transfers from parents are an effective source of insurance against their children's labor income shocks. Parents' help is the most natural source of insurance against income shocks for one key reason: when children are in the early years of their working careers they typically have limited assets to buffer shocks, while parents on the other hand are in a phase of their life cycle where they have accumulated considerably more wealth. Their accumulated assets can be used as an effective mechanism for smoothing children's income shocks. Our paper is one of the few in the literature to shift the focus from the actions and decisions of the children (i.e., how they respond to intergenerational transfers) to those of the parents.

To study dynastic insurance against labor income shocks, we use individual-level administrative data covering the entire Norwegian population. These data allow us to reconstruct the full network of ties linking parents and children and, crucially, to observe parents' and children's assets as well as their incomes. For married children, we observe the parents of *both* spouses and can thus study which set of parents reacts to which spouse's income shock. We present a simple model of parental insurance to help guide our empirical analysis. A central, realistic feature of the model is that children have limited access to financial markets. Our model implies that altruistically motivated parents transfer resources to their child when the child faces an income realization below a certain threshold. Under plausible assumptions,

[1] There is a vast literature on the importance of intergenerational transfers and intergenerational family linkages (see e.g., Altonji et al. 1997 and Waldkirch et al. 2004). However, this literature is primarily concerned with identifying the nature of the linkages (e.g., whether motivated by altruism as in Altonji et al. 1997 or by features such as cognitive abilities and preferences that are common to parents and their offspring, as in Waldkirch et al. 2004). This literature does not study the impact of children's income shocks on the asset accumulation/decumulation decision of parents, which is the focus of this paper.

this threshold is negative. Accordingly, our theory predicts that parental transfers only occur when the child faces an income *drop*. Besides this testable asymmetry, the model also predicts that, upon observing the shock to their child, the parent's response will depend on the nature of the shock. In general, a negative income shock has two effects. First, it reduces the current resources available to the child. Second, it leads parents to a downward revision in the child's expected future income. When the shock is i.i.d., the second effect is absent and parents dissave (make a current transfer only) to help children smooth consumption across periods. But when the shock is persistent, parents expect to make transfers in the future as well, and hence need to save for "(a child's) rainy day", to paraphrase Campbell (1987), in order to be able to meet their transfer obligations in future periods. When parents are not motivated by altruism towards their children, their saving choices are independent of the child's income realizations, positive or negative. This provides a clean, powerful test of the altruism hypothesis. Hence, we study how the saving behavior of parents changes in response to negative vs. positive shocks, and in response to transitory vs. persistent income declines suffered by their children.[2]

We find strong evidence of parental insurance. First, consistent with our model, we find that parents' wealth only changes when shocks to children's income are negative. The relation between positive income shocks and parents' change in financial wealth is flat at zero. Second, parents decumulate financial wealth in response to transitory negative shocks to their children's income, and accumulate assets in response to persistent shocks. This is consistent with consumption-smoothing agents wishing to front-load the adverse effects of anticipated future transfers on their own consumption.[3]

Identification is complicated by the fact that the model predicts different parental saving responses to shocks of different sign (negative vs. positive) and to shocks of different duration (transitory vs. permanent). Under the null of the altruistic model, an OLS regression of parental saving on negative income changes suffered by their child identifies a weighted average of the responses to transitory and persistent income drops. Since the model predicts that these responses are of opposite sign, OLS estimates are thus biased toward finding no evidence of altruistic behavior. To distinguish between the response to transitory vs persistent shocks, we resort to an Instrumental Variables (IV) strategy. In particular, we

[2] Of course, alternative ways for parents to "finance" a transfer to their children would be to cut their own consumption or increase labor supply. However, short of non-standard preference considerations (such as mental accounting, etc.), these are clearly sub-optimal responses when parents have savings to draw upon. In our regression we control for parents' initial wealth. We also examine labor supply responses separately.

[3] This is also consistent with immediate, positive effects on saving behavior from higher long-term wealth tax exposure as in Ring (2019).

assume (and validate empirically) that firm (negative) shocks are among the "primitives" underlying persistent drops in income experienced by workers. On the one hand, firm-related shocks are hard to manipulate for each individual worker and hard to avoid (at least in the short run), as established in a still growing literature on the importance of pass-through of firm productivity shocks onto wages in the presence of labor market frictions. On the other hand, we show that firm-related shocks load primarily on the persistent component of workers' wages, consistent with the empirical evidence firstly established by Guiso, Pistaferri, and Schivardi (2005).[4] Hence, negative shocks to value added translate into negative long-lasting shocks to workers' earnings. Combining OLS with IV estimates allows separate identification of parent responses to the child's transitory and persistent income shocks.

We estimate an elasticity of parents' financial wealth to a persistent negative shock to their child's labor earnings of about -0.25 (i.e., parents save for a child's rainy day). For transitory shocks, we estimate a positive elasticity of around 0.39 (i.e., parents dissave). Evaluated at the median value of parental wealth and child's earnings, the corresponding marginal effects are \$0.12 extra parental saving for \$1 permanent loss in the child's earnings, and \$0.19 decrease in accumulated liquid wealth for each \$1 transitory drop in the child's earnings. These effects suggest that parents play quite a relevant role in offering protection against their children's adverse labor market outcomes. In principle, the insurance role of parents may even be understated if, in addition to their saving choices, they also modify their labor supply choices (i.e., delay retirement or re-enter employment) when the child faces income drops. However, we find no evidence for this, most likely due to the transaction/regulatory costs associated to changes in the timing of retirement or employment choices at a relatively old age.

Having established a significant role for parental insurance, we next explore potential sources of heterogeneity. An important one is whether insurance (altruism) is towards the child or towards the household unit they becomes part of after marriage. In other words: Does "blood matter"? Since we can match spouses to their parents, our data allow us to study whether parents tend to offer insurance primarily to their own offspring or whether they also activate transfers in response to drops in the income of their offspring's spouse.[5] We find that there are larger insurance responses when the shock (transitory or persistent) affects the own child: "blood matters" (although the estimates tend to be noisy). We present evidence

[4]See Guiso and Pistaferri (2020) for a survey covering evidence for other countries.

[5]If couples follow a collective utility model and drops in income lead to a reallocation of consumption within the household, parents can decide to offer insurance to their child in order to redress the increase in intra-household inequality that follows.

consistent with the idea that divorce risk reduces parents' incentive to insure against shocks to the child's spouse.

Other sources of heterogeneity may derive from different access to alternative mechanisms for smoothing shocks, which may displace parental insurance. For example, marriage may provide risk sharing opportunities, and especially so if the spouse works. We thus test whether parents' savings tend to respond less when children are married rather than single and, among married children, if the spouse is employed. Second, we show that marriage, besides opening up risk sharing possibilities among spouses, may also in principle increase the overall supply of dynastic insurance as it expands the number of parents that can activate transfers if shocks occur. Interestingly, in this context, there is little evidence of one set of parents "free riding" on the other; in fact, we find the opposite: parents provide more insurance when another set is present, perhaps due to "competition for attention". Finally, we show that while parents insure their children *vis-à-vis* labor income shocks, the reverse is not true. This finding lends support to the logic of dynastic insurance which requires that a transfer occurs only when the agent making the transfer has enough assets compared to the agent experiencing a drop in income. The different asset position over the life cycle of parents and children implies that this condition tends to be met by parents but not by children, at least on average. It is of course possible that the insurance flowing from children to parents takes a different form (time rather than money). We test whether labor supply of children respond to negative income shocks faced by parents, but find no evidence for this channel either.

Relation to the literature. Our paper is related to two strands of literature. The first studies the mechanisms that allow households to buffer labor income shocks and attempts to provide a quantitative assessment of their importance (Blundell, Pistaferri, and Preston 2008; Guler, Guvenen, and Violante 2012; Ortigueira and Siassi 2013; Heathcote, Storesletten, and Violante 2014; Blundell, Pistaferri, and Saporta-Eksten 2016). We contribute to this literature in a relatively understudied direction—namely, insurance-motivated transfers from parents—and establish the quantitative relevance of this channel. The only other papers we are aware of that emphasize parents' role in insuring children's labor income shocks are Kaplan (2012) and more recently Boar (2021) and Andersen et al. (2020).[6]

[6] Several papers have examined how altruistic behavior is affected by the child's income and probability of being liquidity constrained. Guiso and Jappelli (1991) and McGarry (1999) show that *intervivos* transfers are related to the extent of liquidity constraints faced by adult children, while bequests depend on the level of permanent income; Altonji et al. (1997) show that, when the child's income is uncertain, parents may prefer to defer *intervivos* transfers unless the child is currently liquidity constrained; Benetton et al. (2022)

Kaplan (2012) uses high-frequency US panel data to document the “boomerang” effect, i.e., children returning to the parental home following job loss. Using a calibrated model of parent-child living arrangements, he shows that the option to return to the parental home increases children’s welfare considerably, reduces their precautionary savings and consumption response to shocks, and induces riskier occupational choices. While Kaplan (2012) focuses on a form of in-kind insurance, our focus is on monetary transfers; moreover, in our case, the shock that may trigger insurance is not only a job loss but also, possibly, long-lasting income losses that children cannot buffer otherwise. Our finding that children’s access to other sources of insurance of labor income losses affects parents’ willingness to make a transfer parallels Kaplan’s result that the option of relying on parents affects children’s risk-taking behavior. Boar (2021) relies on matched parent-child pairs in the PSID to show that parents’ consumption correlates negatively with variation in permanent income risk across children’s age and occupation groups. This is consistent with parents exhibiting a precautionary saving motive in response to their children’s income risk. Unlike Boar (2021), we focus on *ex-post* transfers (after shocks occur) and relate them to the *size* and *nature* of the income loss. Importantly, while parental accumulation of precautionary savings in response to their children’s income uncertainty requires parents to have a precautionary motive—i.e., to have preferences exhibiting convex marginal utility—the transfer motive that we study *does not* require a precautionary motive: it only requires that utility is concave and that children suffer a sufficiently large income loss to activate the transfer. This is important because it implies that, provided parents have enough accumulated assets, an insurance channel may be operational even if those assets were accumulated for reasons other than to set up a buffer *vis-à-vis* children’s labor income risk, as in Boar (2021). Because *ex-post* transfers occur when income losses are sufficiently large, our paper, like Kaplan (2012), highlights the fact that parents appear to act as insurance providers of last resort. Finally, and similarly to us, Andersen et al. (2020) study post-shock transfers using administrative data from Denmark. They observe money transfers sent between accounts of a large national bank. If the sender of the transfer has an account in the same bank, they can use administrative data to discern family ties. This means that they can obtain a direct measure of monetary transfers from *own* parents, and look at how parental transfers respond to changes in children’s income and other adverse occurrences (such as unemployment, divorces, and expenditure shocks).

show that parents may relieve their children from downpayment constraints by extracting equity from their own house, which generates “dynastic home equity” effects; Cox (1987) studies whether *intervivos* transfers from parents to children are motivated by altruism or exchange motives and finds evidence supporting the latter.

Andersen et al. (2020) document that monetary transfers from parents are related to drops in the income position of the child as well as other adverse shocks, consistent with the insurance role of the dynasty.[7] However, they find that the replacement rate—the share of the child's income loss that is covered by the parental transfer—is quite small, around 7%. While our results and those in Andersen et al. (2020) are qualitatively similar, they differ in various respects. First, we estimate a significantly higher replacement rate. One potential reason for this difference is that parental transfers are unobserved in Andersen et al. (2020) if they are given as cash in order to avoid *inter vivos* gift taxation. Another reason is that Andersen et al. (2020) do not observe if parents *pay directly* for some of the children expenses. The understatement of the insurance coverage could be substantial if direct payments involve bulky expenses, such as the child's home rent, mortgage payments, or utility bills. Our methodology is free from these problems.[8] We observe the parental linkages of both spouses and the parents' wealth changes, hence capturing *all* parents' transfer sources. Furthermore, because we observe all parents we can study whether parents care only about shocks to their own children or also about shocks to their child's spouse and how the supply of insurance varies when the parents of one spouse are missing. Third, we devise a specific strategy to identify parents' response to transitory and permanent income shocks. On the other hand, we have to assume that a reduction in parental savings is causally related to the child's income shock; we do so by identifying truly idiosyncratic shocks to the child's income (i.e., dismissing concerns related to facing common shocks), and using exogenous variation stemming from firm shocks passing through the child's earnings.

[7]Andersen et al. (2020) also establish that the most relevant money transfers originate from parents; transfers from other members of the individual's social network (siblings, co-workers, school friends) are negligible. This highlights the unique role of parents as supplier of insurance, most likely because, unlike school friends or siblings (who tend to be of similar age as the child), parents are on average at the peak of their wealth accumulation trajectory.

[8]Even though there exists a registry of *inter vivos* gifts and bequests for Norway (see Ring and Thoresen, 2022), we still prefer to study parental insurance by looking at saving rather than transfers. There are several reasons for this. First, as mentioned above direct monetary transfers from parents to children miss indirect transfers that diminish parental resources (and hence change their saving) while benefiting their children (such as parents paying for the children's utility bills, rent or mortgage payments, in-kind transfers, etc.). Second, the registry does not contain monetary transfers that are exempt from reporting requirements. This is true of all transfers made after 2013 (when the inheritance tax was abolished), and of all transfers below half the so-called basic amount from the National Insurance Scheme (approximately 39,500 NOK=$7,000 in 2011) before 2013. Before 2008 there was no set amount that specified when a gift had to be reported to tax authorities, so many gifts may be missing (see Halvorsen and Thoresen, 2010). Finally, before 2013 gifts above 30,000 NOK but below 470,000 NOK (approximately $83,000 in 2011) were subject to reporting requirements but exempt from tax. This may have reduced incentives to report them (especially because, unlike most sources of income or wealth in the tax records, *inter vivos* gifts are self-reported rather than reported by a third-party).

Our paper also relates to a recent strand of macroeconomics research that studies the aggregate implications of idiosyncratic labor income shocks. For example, Bayer et al. (2019) study the aggregate implications of microeconomic uncertainty shocks through accumulation of liquid precautionary assets (see also Basu and Bundick 2017 and Leduc and Liu 2016); Schaab (2020) allows also for time variation in fundamental aggregate risk and for correlation between the latter and the micro uncertainty that individuals experience, emphasizing the tail-risk nature of microeconomic labor income shocks during recessions (as in Guvenen et al., 2014). However, in these papers the only insurance mechanism is precautionary savings. One exception is Bardoczy, 2020, who builds a macroeconomic model with incomplete markets and heterogeneous agents facing idiosyncratic and cyclical labor income risk but allows insurance among spouses, both because marriage offers unemployment risk diversification and because a spouse can increase his/her labor supply to absorb the shock to the family income. However, none of these papers allow for parental insurance. As ours and Kaplan's (2012) findings suggest, parental insurance activates precisely in response to the type of shocks to labor income - large and possibly persistent - that children find hardest to self insure and because of this may cause the strongest consumption response. Accordingly, ignoring parents' insurance role can significantly overstate the importance of microeconomic labor income shocks for macroeconomic fluctuations. The most instructive difference comes from comparing insurance provided through, say, an added worker effect (i.e., a secondary worker increasing labor supply when the primary worker faces a layoff) with insurance provided by parents (who are often retired and hence insulated against labor market risk). Aggregate shocks make the former type of insurance ineffective, while the latter is preserved. In general, our evidence suggests that a full understanding of the macroeconomic implications of microeconomic labor income risk would require modeling households not as isolated entities but as dynastically connected through transfers from one generation to the next.

The rest of the paper proceeds as follows. In Section 2, we outline a simple model and describe its basic predictions. In Section 3 we lay down our empirical strategy and discuss identification of parents' savings response to transitory and permanent shocks to their children's labor income. In Section 4 we introduce the data and discuss how we measure shocks to labor income and to firm performance. We present our main results in Section 5, while in Section 6 we show evidence of heterogeneity in parents response to children income shocks and discuss extensions. In Section 7 we put the results in perspective and conclude.

2 A basic framework

In this section, we set up a simple model of dynastic relations between parents and children (similar to Altonji et al., 1997) to isolate the main forces that induce parents to offer insurance to children experiencing labor income losses. The model provides a set of predictions that we use to guide our empirical analysis.

Parents live for three periods and interact with children in periods 1 and 2. In the initial period (period 0) parents obtain income y_0^P and save w_0^P. In period 1 they receive income $y_1^P = y_0^P$, can count on previous period savings w_0^P and save w_1^P to finance spending in the last period; they make an *inter vivos* transfer τ_1 to their child after observing the *level* and *nature* - transitory or persistent - of the shock faced by the child. In period 2 they spend w_1^P after making another transfer τ_2 to the child.[9] In both periods children are endowed with the same flow of income/cash-on-hand, a^K, which is subject in period 1 to a shock ϵ_1, and in period 2 to a shock ϵ_2, observed by both parent and child. We assume no lifetime uncertainty. Parents draw utility from helping their children when alive, but no utility from leaving bequests, which are therefore absent/omitted.

To consider a stark case for insurance, assume that children have no access to formal insurance or credit markets. This assumption captures the idea that there are financial market frictions to which parents are not subject because, being older, they have accumulated enough assets when children start facing labor market shocks. Accordingly, they can use these assets to time transfers to their children to help them smoothing consumption.[10]

We focus on the parents' optimization problem once children are born, i.e., focus on choices made in periods 1 and 2. Assuming for simplicity no discounting and zero return on savings, parents choose their current savings and transfers in the two periods to maximize expected utility, subject to the constraint that they cannot make negative transfers:

$$Max_{w_1^P,\tau_1,\tau_2} : u(y_0^P + w_0^P - w_1^P - \tau_1) + Eu(w_1^P - \tau_2) + \alpha[u(a^K + \epsilon_1 + \tau_1) + Eu(a^K + \epsilon_2 + \tau_2)] \quad (1)$$

[9] Note that transfers can come in the form of reducing the child's consumption needs instead of increasing the child's income (i.e., the parents pay for the child's rent). If parental transfers are infra-marginal, there is no practical difference between the two. But even if parental transfers were not infra-marginal, parents might still prefer "transfers as consumption" to maintain some control over the child's behavior.

[10] A less extreme assumption is that parents can invest their savings at a higher rate than children, for instance because the scale of their wealth allows to reap a higher risk-free return. If children could borrow but the borrowing rate exceeded the risk free rate, parents could profitably lend money to them.

$$s.t. \tau_1 \geq 0; \tau_2 \geq 0$$

where period utility $u(.)$ is increasing and concave. The parameter $0 \leq \alpha \leq 1$ measures the degree of parental altruism. As for the shock to the child's income, we assume that it follows an AR(1) process:

$$\epsilon_2 = \rho\epsilon_1 + \eta$$

The period 1 shock ϵ_1 is a symmetrically distributed, zero mean error with variance σ_ϵ^2; η is a zero mean, i.i.d. innovation with standard deviation σ_η; this parameter marks the importance of the innovation η for the second period shock ϵ_2. The parameter ρ measures the degree of persistence of the income shock; we assume $0 \leq \rho \leq 1$. From the parents' period 1 viewpoint, ϵ_1 is known and so is the persistent component of ϵ_2, i.e., $\rho\epsilon_1$, whereas the innovation η is not. A negative income realization ($\epsilon_1 < 0$) produces two effects. First, it reduces current resources available to the child. Second, it leads to a revision in the child's expected income. If the income process is i.i.d. ($\rho = 0$) the latter effect is absent.

To focus on the role of persistence and obtain a closed form solution for parents' saving, we assume for simplicity $\sigma_\eta = 0$. We discuss below the case where $\sigma_\eta > 0$.

2.1 Optimal transfers

The following proposition establishes the conditions under which the transfers are operative so that shocks to the child's income may affect parents' savings flows.

Proposition 1

When parents are not altruistic ($\alpha = 0$), they make no transfers ($\tau_1^* = \tau_2^* = 0$) and their saving decisions are independent of the child's income. When parents are altruistic ($\alpha > 0$), they may make transfers, depending on the realization of the child's income. In the spirit of the role of insurance, assume that in the absence of shocks, i.e., when $\epsilon_1 = 0$, parents make no transfers, that is: $u'\left(\frac{y_0^P + w_0^P}{2}\right) \geq \alpha u'(a^K)$. Under the assumption $0 \leq \alpha \leq 1$, this condition holds if the child's endowment is at least a given share of the parents' endowment. *A fortiori*, parents make no transfers when the shock is positive. Let $\bar{\epsilon} \leq 0$ denote a threshold value for the first period shock above which no transfer is made, and τ_1^* and τ_2^* the optimal transfers. Hence:

- The threshold $\bar{\epsilon}$, defined implicitly by: $u'(\frac{y_0^P+w_0^P}{2}) = \alpha u'(a^K + \bar{\epsilon})$, is decreasing in children's cash on hand, a^K, and increasing in parents' period 1 endowment, $\left(y_0^P + w_0^P\right)$, and degree of altruism α.

- Transfers occur as a function of the realization of the shocks affecting the child's income: $\tau_1^* = \tau_2^* = 0$ if $\epsilon_1 \geq \bar{\epsilon}$; $\tau_1^* > 0$ and $\tau_2^* = 0$ if $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$; and $\tau_1^* > 0$ and $\tau_2^* > 0$ if $\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$. It follows that if the shock is purely transitory ($\rho \to 0$), there is no transfer in the second period, and in the first period a transfer is only observed if $\epsilon_1 < \bar{\epsilon}$; if the shock is purely permanent ($\rho \to 1$), transfers are positive in both periods (as long as $\epsilon_1 < \bar{\epsilon}$); for intermediate values of ρ, the probability of making transfers in both periods increases with the size of the shock.

- When both transfers are positive, their level equalizes children's marginal utility of consumption in the two periods. This requires $\tau_2^* = \tau_1^* + \epsilon_1 - \epsilon_2$. If only the first period transfer is active, its optimal level equalizes parents' period 1 marginal utility to the parents' perception of children's period 1 marginal utility.

Proof: In the Appendix.

The proposition implies that, provided parents have large enough accumulated savings and care about their children (so that they internalize the child's budget constraint), they are ready to help financially constrained children smooth current consumption when children suffer a sufficiently large income loss in the current period.

In addition, if the negative shock is large and/or persistent enough, parents plan to transfer cash also in the next period; this requires compressing their current consumption to make sure that enough resources are available in period 2 to sustain their own consumption and finance the transfer to the child. Assuming both transfers are operative, their size will be increasing in the value of parents endowment and degree of altruism, decreasing in the child's endowment, and will be larger the more negative the shock. The size of the second period's transfer increases with the degree of persistence ρ.

2.2 Parents' wealth dynamics

Parental transfers triggered by shocks to the child's income affect the dynamics of parents' wealth, establishing a link between observable shocks to children's labor income and changes in their parents' wealth.[11] From the first order condition for savings w_1^P in the parents' maximization problem above, the parents' flow of savings is:

$$\Delta w_1^P = \frac{1}{2}(y_0^P - w_0^P + \tau_2^* - \tau_1^*)$$

If a current transfer τ_1^* takes place, parents decumulate wealth, whereas they save (for a "child's rainy day") if, *ceteris paribus*, a transfer τ_2^* is planned for the future period. Because optimal transfers depend on the income loss experienced by the child and by its persistence, the existence of an insurance motive can be inferred from the response of parental wealth changes to the children's observable transitory and persistent labor income shocks, which is the key implication for our empirical analysis. To obtain closed form solutions for parents' optimal transfers and wealth changes, assume period utility is CRRA $u(x) = \frac{x^{1-\delta}}{1-\delta}$ with relative risk aversion parameter $\delta > 0$. We summarize the link between parents' wealth changes and income shocks in the following proposition.

Proposition 2

If the period utility is CRRA, $u(x) = \frac{x^{1-\delta}}{1-\delta}$, the condition for observing no transfers when $\epsilon_1 = 0$ is $a^K \geq \mu\left(\frac{y_0^P + w_0^P}{2}\right)$, where $\mu = \alpha^{1/\delta}$ satisfies $0 \leq \mu \leq 1$ and is increasing in risk aversion. The shock threshold for the first period transfer to be operative is $\bar{\epsilon} = \left(\mu\left(\frac{y_0^P + w_0^P}{2}\right) - a^K\right) < 0$, and the condition for transfers to be operative in both periods is $\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$. The solution for the optimal transfers and for the dynamics of parents wealth depends then on the realized value of ϵ_1 and is given by the following:

Case	τ_1^*	τ_2^*	Parents' current saving Δw_1^P
$\epsilon_1 \geq \bar{\epsilon}$	0	0	$\Delta w_1^P = \frac{y_0^P - w_0^P}{2}$
$\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$	> 0	0	$\Delta w_1^P = \frac{y_0^P - w_0^P}{2} - \frac{\tau_1^*}{2}$
$\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$	> 0	> 0	$\Delta w_1^P = \frac{y_0^P - w_0^P}{2} - \frac{\tau_1^*}{2} + \frac{\tau_2^*}{2}$

[11] *Inter vivos* transfers are typically unobservable, particularly at high frequency. When information on *inter vivos* transfers is available, it is often either about transfers made by the donor or those received by the recipient, but rarely do researchers observe both sides of the exchange. Andersen et al. (2020) is an exception as they observe transfers made by parents to children who have checking accounts in the same bank. However, even in this case, it is unlikely that all transfers are captured. Apart from cash transfers that do not go through the bank account, parents can support their offspring by paying directly for some of their expenses, such as rent, mortgage, or utility bills. But these transfers are not observed but they affect parents' wealth dynamics. If wealth dynamics is well measured, as is in our data, it captures in principle all type of monetary transfers (direct or indirect) that affect parental savings. It may also capture in-kind transfers (e.g., provision of child care) if parents reduce their labor supply in order to make them.

Proof: In the Appendix.

Parents' wealth dynamics is invariant to the shocks to child's income when the latter are positive ($\epsilon_1 \geq \bar{\epsilon}$), in which case it is purely determined by the need to smooth their own consumption between periods 1 and 2. When the kid's income shock is purely transitory ($\rho = 0$) the third case is redundant. Parents make a positive transfer $\tau_1^* = \frac{2}{2+\mu}(\bar{\epsilon} - \epsilon_1)$ and have to dissave (relative to the first case or the case without altruism), $\Delta w_1^P = \frac{y_0^P - w_0^P}{2} - \frac{\tau_1^*}{2}$. Since $0 < \mu \leq 1$, the parents offer only partial insurance against the child's shock. The transfer is increasing in the size of the shock, the degree of altruism, the child's endowment, and decreasing in the parent's endowment. A persistent shock ($\rho > 0$) makes a period-2 transfer more likely because it reduces the size of the intermediate region $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$ in which only a period-1 transfer is made. Indeed, when the shock is purely permanent ($\rho = 1$), this intermediate case becomes redundant. Persistent shocks lead parents to save for a child's rainy day (i.e., because they expect a negative shock in the current period to repeat itself next period).[12]

Figure 1 shows graphically the relation between change in parental wealth and the period 1 shock to children's labor income under different assumptions about the degree of persistence as well as the no altruism case. Two things emerge from this picture. First, for given degree of persistence, the amount of resources that parents carry to the next period (savings) decreases with the (absolute value) size of the child's shock, since insurance requires a larger transfer from the parent to the child. Second, savings increase with the degree of persistence of the shock, since parents balance the insurance against current shock vs. the expectation of having to make a transfer in the second period as well. In fact, one can decompose the saving response to the shock into two components: the current "insurance" component (the effect of ϵ_1 in the purely transitory case, or $\rho = 0$) and the "child's rainy day" component (which comes from the fact that, with $\rho > 0$, parents expect to make a transfer in the second period if the shock is large enough in absolute value). The latter is plotted in Panel B.

In the next section we present the empirical strategy that we use to test these basic implications of the parental insurance model and discuss how we can identify parents responses to transitory and permanent shocks. Before moving to the empirical analysis, we return to the $\sigma_\eta = 0$ assumption. This amounts to ignore a precautionary savings channel (and thus

[12]When making transfers in both periods is optimal, one important question is why parents do not choose to make a single transfer in period 1 and let children manage it. In our simple model, this is because children have no access to financial markets (including access to a saving technology). In more realistic settings, this may still be optimal if the child's income in the second period is uncertain, if parents have higher returns on financial wealth due to scale effects, or if there are moral hazard considerations.

dropping the expectations operator from the problem (1)) against children's labor income shocks when parents plan their saving and transfer policies for period 2. This channel is the focus of Boar (2021).[13] *Ceteris paribus*, if we consider the case $\sigma_\eta > 0$, then parents with preferences characterized by prudence will save more (for precautionary reasons) than in the case $\sigma_\eta = 0$ – because of the fear of not having enough resources to fund a more generous transfer to their child in case of a bad realization of η. However, allowing for uncertainty about second period income of the child does not change the qualitative nature of the solution, which is the focus of our empirical analysis. First, it will still be the case that parental savings will behave asymmetrically, responding only to bad realizations of the income process of the kid; second, it will still be the case that transitory shocks require parents to dissave in order to finance a current transfer, while more persistent shocks raise the need for future transfers, and hence alter the saving decision in the opposite direction. A precautionary motive for saving alters the values of the elasticities and the thresholds at which transfers become active, but not the fundamental nature of the problem.

3 The empirical model

An empirical specification that captures the key implications of the model of Section 2 is the following:

$$\Delta w_t^P = \alpha_{Trans} \Delta y_{Trans,t}^{-K} + \alpha_{Pers} \Delta y_{Pers,t}^{-K} + \gamma \Delta y_t^{+K} + \beta_1 w_{t-1}^P + \beta_2 y_{t-1}^P + \beta_3 a_{t-1}^K + \eta_t^P, \quad (2)$$

where Δw_t^P is the log-change in parents' liquid wealth in year t, $\Delta y_{Trans,t}^{-K}$ and $\Delta y_{Pers,t}^{-K}$ denote transitory and persistent drops in the child's labor income, Δy_t^{+K} is a positive income shock, and w_{t-1}^P is beginning-of-period parent wealth.[14] We also control for parental income, y_{t-1}^P, and the child's liquid resources (i.e., cash on hand, measured as after-tax income plus financial wealth), a_{t-1}^K, both measured a full calendar year earlier.[15] This specification captures both the asymmetric response to positive vs. negative income changes as well as the distinction between temporary vs. more persistent earnings declines experienced by the child.

[13] An alternative equivalent assumption is to assume quadratic utility, and thus the absence of a precautionary saving motive.

[14] Note that $\Delta y_t^{+K} = \Delta y_t^K \times \mathbf{1}\{\Delta y_t^K \geq 0\}$, and similar notation applies to $\Delta y_{Trans,t}^{-K}$ and $\Delta y_{Pers,t}^{-K}$.

[15] Because the left hand side of equation (2) is the first difference in parents' wealth between t and t-1, to avoid attenuation bias due to measurement error in parents' liquid assets, the control for parents' lagged liquid assets on the right-hand side of equation (2) is dated t-2. Our findings are robust to changing the length of these lags.

Figure 1: The relation between parental savings and children's labor income shocks

Panel A: Parents savings and current shock

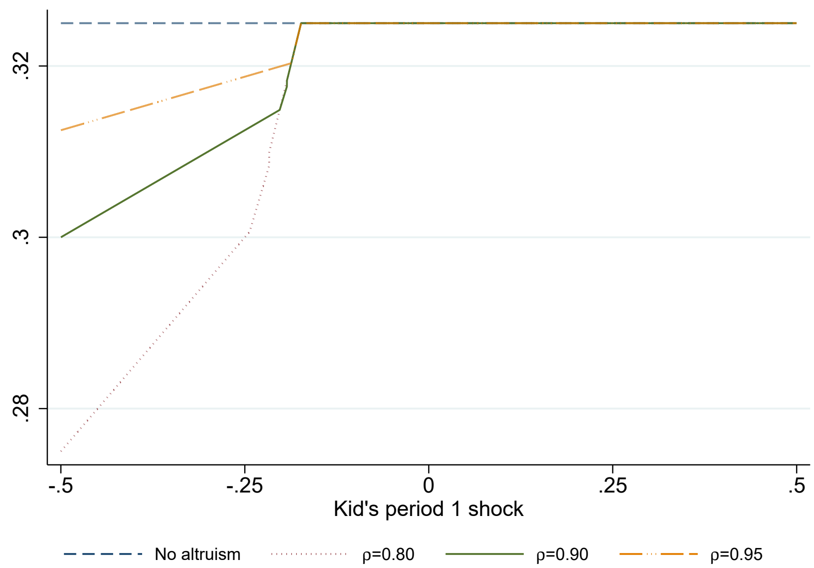

Panel B: Parents savings and persistent shock

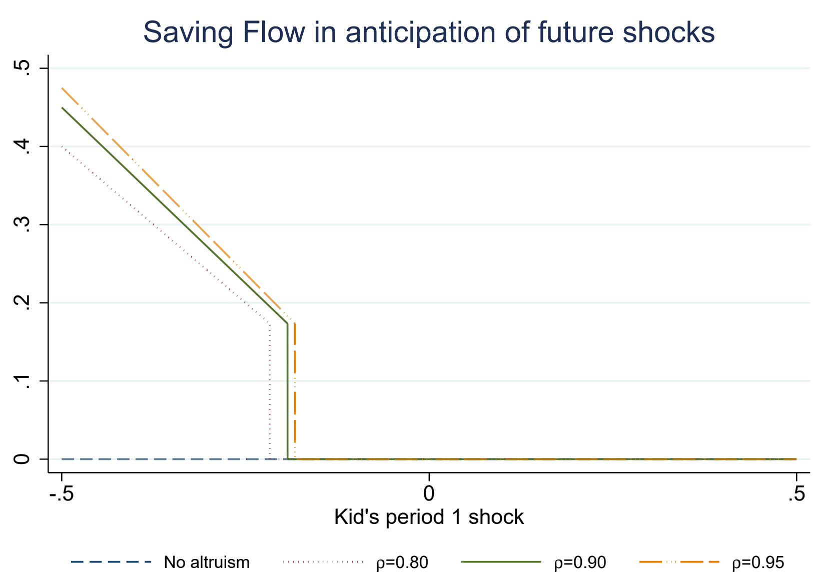

Notes: Panel A shows the relation between the change in parents wealth and the current shock to children labor income; Panel B shows the relation between parents savings and the persistent component of the current shock to children labor income. The graphs are generated under the (illustrative numerical) assumptions: $\alpha = 0.66, \delta = 3, \left(y_0^P + w_0^P\right) = 0.75, a^K = 0.5$.

The theoretical model implies that parents make a transfer, and thus dissave, when their child faces a transitory drop in income (parent saving and earnings shocks move in the same direction, or $\alpha_{Trans} > 0$), and that they save in anticipation of having to make future transfers in the future when the shock is persistent (parental saving and earnings shock move in opposite directions, or $\alpha_{Pers} < 0$). No response should be observed when children experience positive income shocks – that is, $\gamma = 0$. Finally, a stark test of no altruism is that parental saving decisions are independent of the child's income process: $\alpha_{Trans} = \alpha_{Pers} = \gamma = 0$. The theoretical model with altruism also delivers implications for the sign of parents' initial wealth ($\beta_1 < 0$) and children's cash on hand ($\beta_3 > 0$).

There are two issues in estimating (2). The first is that we do not observe transitory and persistent income changes separately, but only the overall (unexplained) change in earnings:

$$\Delta y_t^K = \Delta y_{Trans,t}^K + \Delta y_{Pers,t}^K \tag{3}$$

This is a conventional problem. The literature has addressed it by considering the restrictions that the model imposes on the covariance structure of earnings dynamics and parental savings (Blundell et al., 2008); or by combining data on expectations and realizations of earnings (Pistaferri, 2001). The additional complication we face is that while we observe whether the individual experiences an increase or a drop in her earnings ($\Delta y_t^K > 0$ or $\Delta y_t^K < 0$), we cannot determine whether a negative change is caused by a persistent or a transitory shock. We first discuss the key features of the income process, then turn to a discussion of the identification of the insurance parameters.

3.1 The income process

A recent and growing literature on rent sharing (Guiso et al., 2005; Kline et al., 2019; Card et al., 2018; Lamadon et al., 2022) shows that firm-level productivity shocks are transmitted onto wages. Suppose that firm-level productivity shocks are adequately captured by unexplained changes in value added (ΔVA_t^F) and that these are the sum of an i.i.d. innovation to a random walk process, ν_t^F, and changes in a transitory i.i.d component, κ_t^F:

$$\Delta VA_t^F = \nu_t^F + \Delta\kappa_t^F \tag{4}$$

On the worker's earnings side, unobserved changes in earnings are given by (3). Assume that the transitory component $y_{Trans,t}^K$ is i.i.d. while the permanent component, $y_{Pers,t}^K$, follows a random walk process:

$$\Delta y^K_{Pers,t} = \underbrace{\theta \nu^F_t + \widetilde{\zeta}^K_t}_{\zeta^K_t} \tag{5}$$

where ζ^K_t is the composite innovation to the random walk. This innovation originates from two distinct sources. One is the pass-through of permanent firm value added shocks onto wages (measured by the pass-through coefficient θ); the other is any other persistent earnings shocks unrelated to the firm's fortunes (captured by the term $\widetilde{\zeta}^K_t$) due to e.g., deterioration of health capital, obsolescence of general skills, etc. We thus assume that permanent shocks to the firm's value added load onto the persistent component of workers' earnings (consistent with evidence firstly established by Guiso et al., 2005), while transitory shocks to value added are "insured" (i.e., there is no pass-through onto wages). We show that these assumptions hold true in the Norwegian setting as well (see Appendix OA.3). In particular, if we allow worker wages to depend on both transitory and persistent value added shocks, we find that the coefficient that measures the pass-through of transitory value added shock onto wages is small and insignificant, while the one that measures the pass-through of persistent value added shock is an order of magnitude larger and statistically significant. Moreover, we calculate that, taking our pass-through estimates at face value, 98% of the wage variation explained by pass-through of firm value added shocks originates from the pass-through of persistent firm shocks.

3.2 Identification of insurance parameters: Symmetric case

To build intuition about the forces that, in the data, identify the parameters of interest, we start by discussing identification of the insurance parameters in the symmetric case, then discuss the (empirically relevant) asymmetric case in the next section. Omitting individual subscripts for simplicity, suppose that the relationship between parental saving and income shocks faced by the child were symmetric, a special case of specification (2):

$$\Delta w^P_t = \alpha_{Trans} \Delta y^K_{Trans,t} + \alpha_{Pers} \Delta y^K_{Pers,t} + \eta^P_t, \tag{6}$$

where we have omitted other controls for simplicity. This regression is infeasible because we do not observe $\Delta y^K_{Trans,t}$ and $\Delta y^K_{Pers,t}$ separately; we only observe their sum, Δy^K_t. Consider then running an OLS regression of Δw^P_t on Δy^K_t. It can be shown (see Appendix OA.2) that the OLS estimate is a linear combination of the two responses α_{Trans} and α_{Pers}:

$$\text{plim } \hat{\alpha}^{OLS} = \text{plim}\frac{\text{cov}(\Delta w_t^P, \Delta y_t^K)}{\text{var}(\Delta y_t^K)} = \omega_{Trans}\alpha_{Trans} + (1 - \omega_{Trans})\alpha_{Pers}, \tag{7}$$

where ω_{Trans} is the share of total earnings growth variance due to transitory shocks.

To separately identify the two parameters of interest α_{Trans} and α_{Pers}, consider an Instrumental Variables (IV) strategy that uses productivity shocks to the firm where the child is employed to isolate persistent variation in wages (as posited by equation (5)). One can show (see again Appendix OA.2) that this pins down α_{Pers}:

$$\text{plim } \hat{\alpha}^{IV} = \text{plim}\frac{\text{cov}(\Delta w_t^P, \Delta VA_t^F)}{\text{cov}(\Delta y_t^K, \Delta VA_t^F)} = \alpha_{Pers}, \tag{8}$$

Shocks to firm's value added, ΔVA_t^F, are a valid instrument for identifying α_{Pers} under three assumptions: (a) unobserved heterogeneity in parental savings is orthogonal to firm value added shocks ($E(\eta_t^P|\Delta VA_t^F) = 0$); (b) permanent value added shocks load onto the persistent component of earnings ($\theta \neq 0$); and (c) temporary value added shocks are "insured within the firm". Assumption (a) would be violated if, say, parents hold shares in the child's company. We will show that eliminating cases where this may occur produces identical results. Assumption (b) can be easily tested by looking at the power of the instrument, and is motivated economically by the fact that workers cannot avoid pass-through of firm persistent shocks onto wages, especially when labor markets are characterized by frictions, a hard-to-dispute feature of most labor markets. Assumption (c) is a typical finding in the rent sharing literature (see Guiso et al., 2005), and is empirically confirmed in our Norwegian setting (see Appendix OA.3).

To close the circle on identification in the symmetric case, one can combine (7) and (8) to show that:[16]

$$\hat{\alpha}_{Trans} = \frac{1}{\hat{\omega}_{Trans}}\hat{\alpha}^{OLS} - \frac{(1 - \hat{\omega}_{Trans})}{\hat{\omega}_{Trans}}\hat{\alpha}^{IV} \tag{9}$$

[16] A consistent estimate of ω_{Trans} can be obtained using the dynamics of the child's earnings growth (the variance and the first order autocovariance). In particular, it is easy to show that: $\hat{\omega}_{Trans} = \frac{-2\text{cov}(\Delta y_{t-1}^K, \Delta y_t^K)}{\text{var}(\Delta y_t^K)}$. It is worth stressing that in estimating ω_{Trans} we do not need to distinguish between positive and negative shocks to labor income and can thus use all the variation in the data. For the discussion that follows, a possible concern would be if the variance of shocks differs conditioning on positive vs. negative shocks. However, as we shall see, the distribution of residualized log income changes is fairly symmetric around the mean.

It is then easy to show that plim $\hat{\alpha}_{Trans} = \alpha_{Trans}$. Combining OLS and IV estimates (along with the dynamics of the worker's income process to pin down $\hat{\omega}_{Trans}$) is a transparent way of solving the problem of identifying two parameters ($\hat{\alpha}_{Pers}$ and $\hat{\alpha}_{Trans}$) using two data "moments".

3.3 Identification of insurance parameters: Asymmetric case

The relevant relationship between parental saving and child's shocks is asymmetric, as described by equation (2). Under the null of the altruistic/insurance model, and assuming that the child's labor income shock at time t is orthogonal to lagged values of his cash-on-hand and the parents' income and wealth, one can think of using the analogs of the OLS and IV regressions of the symmetric case:

$$\check{\alpha}_{OLS} = \frac{\text{cov}(\Delta w_t^P, \Delta y_t^{-,K})}{\text{var}(\Delta y_t^{-,K})} \tag{10}$$

$$\check{\alpha}_{IV} = \frac{\text{cov}(\Delta w_t^P, \Delta VA_t^{-,F})}{\text{cov}(\Delta y_t^{-,K}, \Delta VA_t^{-,F})} \tag{11}$$

where $\Delta y_t^{-,K} = \Delta y_t^K \times \mathbf{1}\{\Delta y_t^K < 0\}$ and $\Delta VA_t^{-,F} = \Delta VA_t^F \times \mathbf{1}\{\Delta VA_t^F < 0\}$. The first parameter is the OLS coefficient of a regression of parental saving on the child's income *drops*. In our model, this is what activates parental transfers. The second is the IV coefficient of a regression of parental saving on the child's income *drops* using value added *declines* as an instrument for the latter. In the spirit of the pass-through narrative, workers employed in shrinking firms should experience downward adjustment in wages (and *vice versa* for workers employed in growing firms), and these wage changes should be persistent. In our model, this is what activates parental transfers in response to persistent shocks suffered by the child. Hence, the identification logic described in the previous section still applies: the IV parameter (11) pins down the response of parents' wealth to persistent drops in the child's income (the parameter α_{Perm}), while the OLS parameter (10) is a combination of the parents' wealth response to transitory and persistent drops in the child's income.

Because of the non-linearities induced by the sign-switching terms, it is no longer possible to prove identification of α_{Perm} and α_{Trans} analytically from (10) and (11) as done in the symmetric case of the previous section.[17] However, we show that, using (10) and (11) as

[17] Several papers in the consumption insurance literature (Blundell et al. 2008; Kaplan 2012; to cite a few) use the covariance restriction that the model imposes on the joint behavior of consumption (or saving) and

auxiliary parameters in an Indirect Inference procedure, delivers virtually identical estimates of α_{Trans} and α_{Perm} than our baseline methodology. See Appendix OA.4 for the technical details.

4 Data and variables construction

In this section, we discuss our data sources, the criteria for selecting the sample we focus on, and how we measure income shocks.

4.1 The Norwegian population data

We use population data for Norway and match through family identifiers every parent to all sons and daughters (we refer to them as children), either single or married, who live in an independent domicile. We link parents and children identifiers with several administrative registries: (a) tax records containing detailed information about the individual's sources of income as well as asset holdings and liabilities; (b) balance sheet data for the private businesses owned by the individual; (c) a housing transaction registry; (d) balance sheet information for all firms individuals work for. The value of asset holdings and liabilities is measured as of December 31. While tax records typically include information about income, they rarely (if ever) contain exhaustive information about wealth. In Norway, this happens because of a wealth tax that requires taxpayers to report their asset holdings in their tax filings. From the tax records we observe labor income as well as any other income component, including income from capital and from transfers. From the same source we have information on the assets and liability holdings of each taxpayer. Assets values are available separately by asset type (deposits, bonds, mutual funds, listed stocks, non-listed stocks, real estate and private business wealth); data on liabilities are available separately for mortgages and consumer loans, and for student loans.[18]

income at different leads and lags to identify the effect of transitory and more persistent shocks on behavior. We cannot use this popular strategy here because the altruism model of Section 2 draws a net distinction between the impact of negative-transitory, negative-persistent, and positive shocks, which are not separately identified by covariance restrictions.

[18]Data on private pension wealth and other (minor) wealth components are absent. However individual pension (i.e., the equivalent of IRA accounts in the US) are quantitatively negligible (less than 1% of aggregate household gross wealth). Furthermore, liquidation of these assets is typically costly as they entail early withdrawal penalties as well as the loss of tax subsidies. Hence, they are unlikely to be relied upon to fund a transfer to own children. The other component of wealth that is missed is assets held abroad and not reported to the tax authority. However, this is unlikely to matter for our results. Alstadsæter et al.

The data we assemble have several, distinguishing useful features for the purpose of this study. First, our income and wealth data cover all individuals in the population, including people at the bottom as well as at the very top of the wealth distribution. This is important since whether the parents insurance channel is active or not depends on their wealth holdings relative to those of their children. Furthermore, population coverage allows to span both sources of parents transfers for married couples, transfers from the parents of the husband and from those of the wife (if alive).

Second, in our data set most components of income and wealth are reported by a third party. Financial assets and data on debt are reported at market value directly to the tax authority by the intermediary where the asset is held (e.g., a bank or mutual fund); similarly, labor income is reported directly by the employer to the tax agency. Housing is either reported at market value in the year in which a transaction is observed, or imputed using a hedonic regression in the years in which no transaction occurs. Finally, wealth in private business is obtained as the product of the equity share held in the firm (available from the shareholder registry) and the fiscally-relevant "assessed value" of the firm. The latter is the value reported by the private business to the tax authority to comply with the wealth tax requirements. See Fagereng et al. (2020) and Ring (2019) for more details about the wealth data. Given the prevalence of third-party reporting, the data do not suffer from the standard measurement errors that plague household surveys, where individuals self-report income and asset components and confidentiality considerations lead to censoring of asset holdings and top incomes.

Third, the Norwegian data have a long panel dimension, which allow us to observe interactions between parents and children over several periods where negative income shocks have a chance to materialize and insurance-motivated transfers to be observed. Because the data cover the whole relevant population, they are free from attrition, except the unavoidable ones arising from mortality and emigration, implying that parents-children interactions are rarely interrupted and can be followed also in case of divorce of either the young or old member of the dynasty.

Next, we discuss how we measure shocks to children's labor income and to their firms' value added, parents' savings and endowments, as well as our sample selection.

(2018) find that hidden foreign assets are mostly held by people in the top 1% of the distribution of wealth. These individuals have so much accumulated assets that missing their foreign assets is unlikely to affect the estimated ability to transfer money in case their children are confronted with a shock.

4.2 Measuring income shocks and financial holdings

We restrict the analysis to children aged 25 to 55 to focus on people who have completed their education and are not yet retired.[19] To obtain an exogenous earnings shock measure, we focus on children employed in the private sector where we have a meaningful definition of firm value added. We drop observations where parents and children work in the same industry (and hence, also those working in the same firm); we also drop children whose lagged earnings are below a threshold of approximately \$10,000 ("basic income" level in Norwegian Social Security). These selections leave us with a sample of 3 million child-parent pairs observed during the 1997–2014 time period, providing a total of more than 13 million observations, of which about 9.7 million from matches with married children.

Shocks to the firm's performance. Following Guiso et al. (2005), we measure firm performance by its value added, defined as total revenues net of operating costs, excluding depreciation and labor cost. We use value added as it is the relevant flow to remunerate labor and capital. We set value added to zero whenever it is missing during or following a bankruptcy or large collective dismissal event.[20] In order to accommodate (near-) zeros in value added when taking logs, we shift the observed measure by NOK 10,000 (in 2011 NOKs, around \$1,750). To obtain a measure of shocks to value added we purge the observed value added data (Ψ_{jt}) of the non-idiosyncratic component in firm j in year t by estimating the following process:

$$\ln \Psi_{jt} = X'_{jt}\phi + VA^F_{jt},$$

where X_{jt} is a vector of observables that captures the predictable component of firm's performance and ϕ the corresponding vector of coefficients. The vector X_{jt} includes firm fixed effects, 2-digit NACE-code industry fixed effects at the county-year level, more granular 3-digit NACE-code fixed effects at the year level, as well as the latter fixed effects interacted with lagged log revenues. Our measure of shocks to the firm value added are the first-differenced residuals from this regression after we winsorize them at ± 3 (approximately 1.5% of all observations). The resulting variable is denoted ΔVA^F. Figure 1, Panel A shows the histogram of its distribution, benchmarked against the normal distribution (green line).

Shocks to labor income. We model the log of labor earnings Y_{it} of worker i in year t,

[19] AppendixOA.5, Table OA.2 shows that our results are invariant to reducing the age interval to 25–45.

[20] We define a large collective dismissal event as one in which at least 10 employees or 10% of the firm's work force receive unemployment insurance. This addresses potential missing tax filings in the presence of large employment-relevant shocks to these firms.

in a similar vein, as a linear function of a predictable component that depends on a vector of workers observed characteristics, Z_{it} and an idiosyncratic, unpredictable component Δy_{it}:

$$\Delta \ln Y_{it} = Z_{it}'\gamma_t + \Delta y_{it}.$$

We take out predictable time-invariant components of earnings by taking first differences. Z_{it} includes a third-order polynomial in age, year-specific fixed-effects interacted with (i) family size $\times$ marital status, (ii) year-specific county$\times$field of education$\times$education level bins, and (iii) municipality of residence. Controls (ii) and (iii) remove region- and occupation-related shocks or trends that could be correlated with firm performance shocks. Similarly to the value added shocks, outliers are accommodated by censoring the estimated residual Δy_{it} at ± 3 (this winsorizes less than 1% of the earnings data). Figure 2, Panel B shows the histogram of the distribution of labor income shocks; the green line is the normal distribution. Compared to the distribution of the value added shocks, labor income shocks are much less spread out, as one would expect; both distributions appear fairly symmetric and reveal some excess kurtosis, which is more marked for labor income shocks.

Figure 2: Distribution of Value-Added Shocks and Changes in Child Labor Earnings

Panel A: Value-Added Shocks

Panel B: Log-Differenced Child Labor Earnings

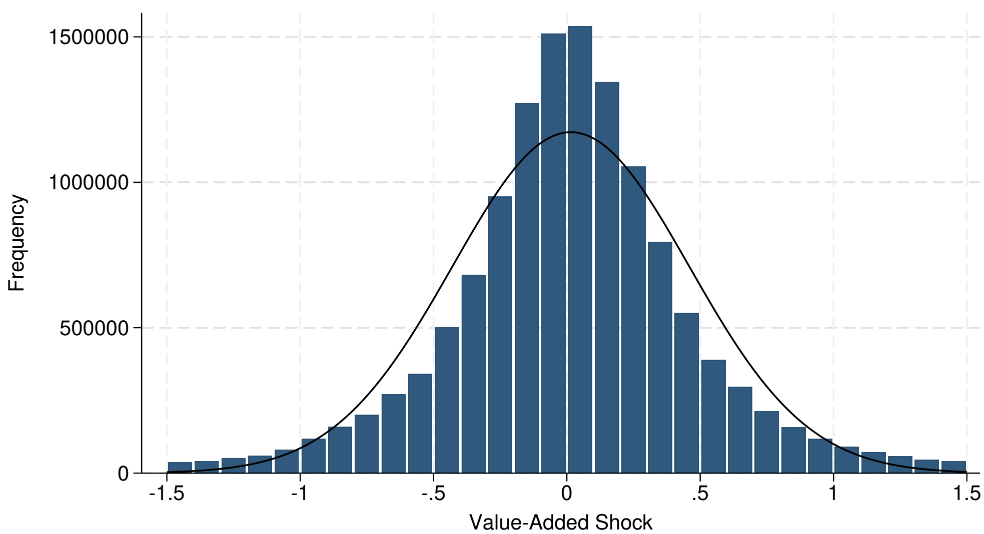

Notes: Panel A shows the distribution of the estimated first differenced value-added shocks to the child's firm employer; Panel B that of the estimated shocks to child labor earnings. The bin size is 0.1 (measured in log-differenced values). Values outside [-1.5,1.5] are omitted for readability. The continuous line is a fitted normal distribution.

Changes in Financial Wealth. We follow a similar strategy to obtain a measure of residualized parental net savings that we use to infer their insurance role. For this we use log-differenced parental financial wealth, residualized with the same variables used to model children's income growth. We also include year-specific coefficients on the lagged stock market share in parents' financial wealth to avoid potential confounding from differences in stock market exposure. We use financial wealth to focus on a measure of wealth that can be readily made available for insurance purposes. We use residuals to insulate changes that cannot be attributed to predictable shifts in portfolio composition, taxation, and the like.

Table 1 shows summary statistics for parents and children. We also present separate statistics for the subsample of children that are married. The last three rows present statistics for the main variables of interest: residualized changes to parents' financial wealth as well as children's labor earnings and firm value added shocks. Since the distribution is symmetric (Figure 2), approximately half of the sample experience some unexplained drop in earnings. The average age gap between parents and children is 27 years. As expected, parents have much more liquid wealth than their children (roughly twice as much at the average) but children have relatively more labor earnings, which is consistent with parents and children being in different phases of their life cycles.

Table 1: Summary Statistics

	N	Mean	SD	p25	p50	p75
Parents						
Financial Wealth	14,049,742	88,320	436,757	8,854	35,280	97,344
Total Labor Income	14,049,742	83,669	64,908	43,217	69,086	107,768
Labor Earnings	14,049,742	45,139	71,968	0	5,805	75,192
Age	14,049,742	66	10	58	65	73
Children						
Financial Wealth	14,049,742	47,229	306,179	5,346	16,915	45,413
Total Hh. Labor Income	14,049,742	134,994	77,244	88,079	125,314	164,767
Cash-on-hand	14,049,742	151,798	371,861	82,202	118,422	169,472
Labor Earnings	14,049,742	82,083	53,129	55,937	73,179	97,510
Age	14,049,742	39	8	32	38	45
Married	14,049,742	0.72	0.45	0.00	1.00	1.00
Married Children						
Spouse Works	10,146,890	0.95	0.23	1.00	1.00	1.00
Spouse's Parents Present	9,553,028	0.92	0.28	1.00	1.00	1.00
Parent of Spouse	10,146,890	0.50	0.50	0.00	0.00	1.00
Not divorced $(t, ..., t+3)$	8,856,782	0.87	0.34	1.00	1.00	1.00
Share of Hh. Labor Earnings	10,131,468	0.59	0.21	0.45	0.58	0.71
Residuals						
Parents' FW, Δw^P	13,821,570	0.00	0.79	-0.24	-0.03	0.19
Child's Labor Earnings, Δy^K	13,882,984	0.02	0.43	-0.09	-0.00	0.09
Value-Added Shocks, ΔVA^F	13,474,290	0.02	0.66	-0.24	0.02	0.27

Notes: Monetary values are in USD, 2011 prices - using the 2011-average exchange rate between USD and NOK in 2011 (USD/NOK= 5.607). Total labor income includes labor earnings as well as pensions and labor-related public transfers. Cash-on-hand is the sum of after-tax total income and financial wealth. "Married children" is the subset of children for whom we observe a legally married spouse. The Spouse Works dummy indicates that the spouse has non-zero labor earnings. Spouse's Parents Present indicates whether the spouse's parents are alive. Parent of Spouse indicates that the child-parent relationship goes through the spouse (i.e., children-in laws). Share of Hh. Labor Earnings is the ratio of the child's own labor earnings to their household's total labor earnings, which includes any spousal earnings.

5 Results

5.1 Graphical evidence on parental insurance

Before discussing the regression results, we provide some visual evidence regarding the key property of the insurance model: insurance from parents should only be active when children experience a drop in income and parents should not respond to positive shocks to children's income, as illustrated in Figure 1. Figure 3 plots the change in parents' residualized log liquid assets against the change in the child's residualized log labor earnings using a kernel-weighted local polynomial regression. The continuous line runs the local polynomial regression over the whole range of the distribution of the shock to the child's labor earnings. There is a clear positive relation between changes in parents' liquid wealth and children's labor income shock over the negative domain of the shock: as the latter becomes increasingly more negative, parents decumulate (or slow down accumulation of) financial assets. Over the positive domain of the shock the relation is essentially flat —that is, movements in parents' liquid savings appear unrelated to the child's labor income innovations. The asymmetry in parents' behavior is confirmed when running separate polynomial regressions for negative and positive values of the shock (the dashed line). The non-parametric empirical relation between parents' net saving behavior and children's labor income shocks tracks closely the theoretical prediction of Figure 1 and lends some *prima facie* support to the role of parents as insurance providers against labor income shocks affecting the younger members of the dynasty.[21] In fact, since we are using residuals that eliminate the influence of "common shocks" that may create spurious co-movements between changes in labor earnings of children and wealth changes of parents (such as economy-wide or geographically-concentrated effects), it would be hard to rationalize the relationship that we find - and its asymmetry - absent some form of insurance or altruistic behavior from parents to children.

To separately identify the effect of children's persistent and transitory earnings shocks on parental savings, we follow the empirical strategy discussed in Section 3. A first step is to establish the power of the instrument (first stage of the regression), i.e., to show that shocks to the value added of the firm predict shocks to the child's labor income. In particular, we want to establish that negative shocks to the performance of the firm cause negative changes

[21] Notice that in Figure 3 the relation flattens when the drop in child income is large, a feature predicted by the model (see Figure 1). This happens because when the drop in income is large parents plan to also make a transfer in the future and this requires some extra savings. Some of these extra savings come from reducing current transfers, which flattens the curve.

Figure 3: Parental Saving and Changes to Children's Labor Earnings

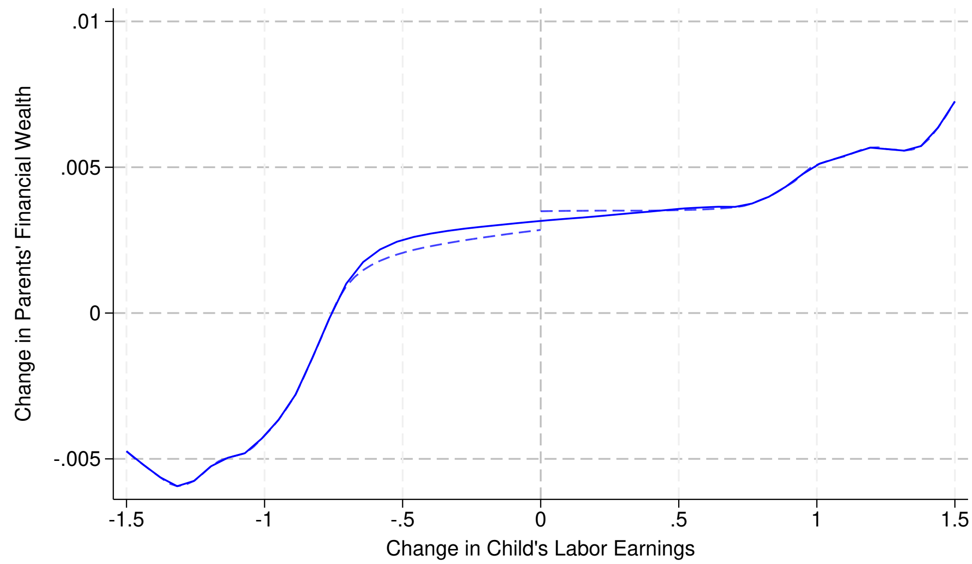

Notes: The figure shows the relationship between (residualized) changes in parents' financial wealth and changes in the child's labor earnings. The solid line is a local polynomial fit (bandwidth=0.33), excluding values exceeding 1.5 in absolute value for readability. The dashed lines perform the local polynomial fit separately for negative and positive changes in earnings.

in the child's labor income, which is what triggers parental transfers in the model. Second, we want to establish that persistent shocks to the firm performance are the drivers of the correlation between parental savings and the child's labor income drops.

Again, we start by illustrating these forces with visual evidence. Figure 4, Panel A is a depiction of the first stage. It plots a kernel-weighted local polynomial regression of negative shocks to children labor income (vertical axis) against the entire support of shocks to the firm's value added. Clearly, only *negative* shocks to firm performance appear to predict *negative* shocks to labor income.[22] Figure 4, Panel B provides a non-parametric version of the reduced form. It shows the local polynomial regression of the change in parents liquid assets (vertical axes) against the shocks to the value added of the child's employer. As in Figure 1, the solid line is the local polynomial fit on the entire support of value added shocks while the dotted line shows the relation separately for positive and negative shocks. Positive shocks to the firm's value added (say, increase in productivity due to the adoption

[22] A similar figure for positive shocks shows that positive shocks to firm value added are strong predictors of growth in earnings, see panel A of Appendix Figure OA.1; panel B of the same Figure OA.1 shows that there are no clear asymmetries: on average, workers in growing firms experience wage gains and workers in shrinking firms experience wage losses.

of new technologies) result in positive shocks to children's labor income and, as predicted by the insurance model, have no effect on parents' transfers and savings. Negative value added shocks are passed over to children's earnings only to the extent that they are permanent (say, fundamental scaling down of the firm's operations in foreign markets). The model predicts that, confronted with larger drops in children's earnings, parents save more (or dissave less), anticipating the need for future transfers. Empirically (Panel B), this prediction is borne out by the negative correlation when value added shocks are negative. The most striking aspect of this picture is that - in a world without parent-to-child insurance - there would be no reason for parental savings to respond to productivity declines experienced by the firm employing *their child*.[23] The fact that parents' savings changes when the firm employing the child is doing poorly is again *prima facie* evidence in favor of dynastic-motivated saving behavior. In the next section we present regression results that allow us to separate the effect of transitory vs. more persistent income shocks on parental savings (which the simple theoretical model of Section 2 suggests should move in opposite directions), quantify the extent of insurance, while controlling for other confounders.

5.2 Parents' savings response to children's earnings shocks

5.2.1 OLS and IV Regression Estimates

In Table 2 we show the OLS and IV regressions, along with information on the power of the instruments. These regressions are the multivariate version of equations (10) and (11), since we add a full set of controls. They are also the parametric counterparts of Figures 3 and 4 (again, allowing for controls), and they reproduce their qualitative patterns. In particular, the OLS regression (column (1) of Table 2) shows that the impact on parents' savings of *drops* in the child's income is positive (0.021) and highly significant (s.e. 0.0011). In contrast, the estimate of parents' savings response to positive innovations in children's income is virtually zero, implying that, economically, over the positive range of variation of children income shocks parents' savings are insensitive to the shocks, as already shown in

[23]Recall that aggregate shocks that may create some spurious association - say a local, large firm going bankrupt and depressing the local economy - have been netted out. Moreover, most parents are already retired, and hence mostly insulated from labor market shocks, local or otherwise. A final possibility is that parents hold stocks in their child's company, so may experience a direct decline in the value of their wealth if the company is doing poorly. This is unlikely to be a concern for two reasons. First, most parents are already retired and do not hold many risky assets in their portfolio; second, the Norwegian stock market has few listed companies. We will show that even if we eliminate parents with stocks (in listed or unlisted companies), the results remain identical.

Figure 4: The Relationship Between Value-Added Shocks, Child Earnings, and Parental Saving

Panel A: First Stage

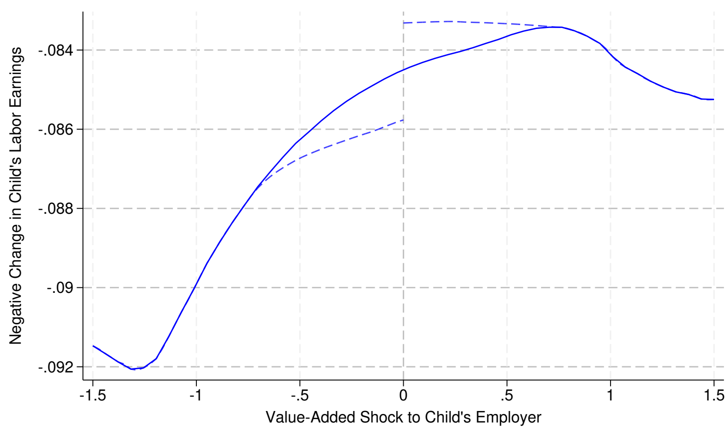

Panel B: Reduced-form Relationship

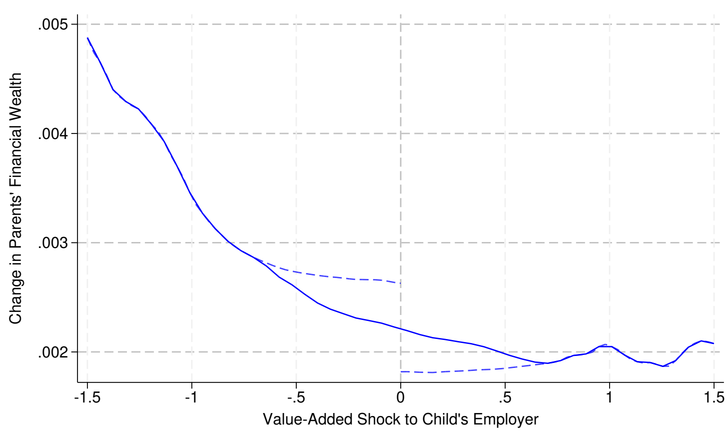

Notes: Panel A shows the first-stage relationship between negative changes in the child's labor earnings (Δy^{-K}) and value-added shocks experienced by the child's employer. Panel B shows the reduced-form relationship between changes in parents financial wealth and the child's firm value-added shocks. The solid line is a local polynomial fit (bandwidth=0.33), excluding values exceeding 1.5 in absolute value for readability. The dashed lines perform the local polynomial fit separately for negative and positive value-added shocks.

Figure 2. In Section 3.2 we have argued that - even in the symmetric case - the OLS estimates combine the responses to transitory and persistent shocks to children's income. To isolate persistent (negative) changes in earnings we use as an instrument (negative) shocks to the firm's value added, which the rent sharing literature has shown to pass-through primarily onto the permanent component of earnings. The IV estimates (column (2) of Table 2) show that persistent declines in the child's earnings correlate negatively and significantly with parents' savings (coefficient -0.24 with a s.e. of 0.08). In this IV specification, parents' savings increase slightly in response to positive income shocks, but the effect is economically negligible. In column (3) we show the first stage, where negative earnings changes are regressed onto the observables and the instrument, negative changes in the firm's value added. Drops in earnings are associated to drops in firm productivity. The first stage is powerful, with an F-statistic of 81. Hence, the lack of response to parental saving when the child experiences growth in earnings in the IV case is not due to low power of instruments.[24]

[24] Note that the first stage has fewer observations than the second stage because the parent-level regressions include all parent-child (and parent-child-in-law) pairs. In principle, a "child" could be matched with four parent households: the household of her father, the household of her mother (if she no longer lives with her father), the household of her father-in-law, and the household of her mother-in-law (if she no longer lives with her father-in-law). In the first stage we only use unique individual-firm pairs (although the practical implementation of the IV regression uses the full data matrix, which includes replicated individual-firm pairs).

Table 2: OLS, IV, and First Stage Regressions

Dep. variable	Δw^p		$\Delta y^{-,K}$
	(1)	(2)	(3)
$\Delta y^{-,K}$	0.0209***	-0.2362***	
	(0.0011)	(0.0821)	
$\Delta y^{+,K}$	-0.0002	0.0225***	
	(0.0007)	(0.0074)	
$\Delta VA^{-,F}$			0.0060***
			(0.0007)
Other controls	Yes	Yes	Yes
First-stage F-statistic			80.90
N	13,550,902	12,993,333	7,948,626

Notes: Columns (1) and (2) report the estimates of the OLS and IV regressions; in column (3) we show the first-stage of the IV specification. Regressions also include the log of parental financial wealth (dated $t-2$), the log of parental income and the log of the child's cash-on-hand (both dated $t-1$). Standard errors in columns 1-2 are clustered at the child-pair level; in column 3 are clustered at the firm level.

5.2.2 Estimates of Structural Parameters

Table 3 shows the estimates of the structural parameters of interest using the methodology described in Section 3. We estimate the variances of the earnings shocks using data on the dynamics of the child's earnings (equations (3) and (5)) and a GMM procedure. The estimated variance of transitory shocks to the child's labor income is 0.0410 (s.e. 0.0002), smaller than that of permanent shocks, estimated at 0.1129 (s.e. 0.0004). The implied share of total variance in earnings growth explained by transitory shocks is thus $\hat{\omega}_{Trans} = 0.42$. The rest of Table 3 shows estimates of α_{Trans}, α_{Pers}, and γ (equation (2)), which are key to testing the different theoretical implications of the model. In this table and in the ones that follow standard errors are obtained by the block bootstrap (where the clustering is by the parent-child pair), based on 100 replications. In column (1), parents' savings elasticity to negative transitory shocks to children's labor income (α_{Trans}) is 0.37 and highly statistically significant (standard error 0.13). In contrast, the estimate of parents' savings elasticity to negative persistent shocks is negative ($\alpha_{Perm} = -0.24$) and also statistically significant at conventional levels. Finally, our estimate of γ (the response to positive earnings shocks) is small and statistically insignificant. At the bottom of the table we present the result of a "No

Altruism" test (i.e., a test that parents' saving behavior is independent of earnings shocks faced by the child, or α_{Trans}=$\alpha_{Pers} = \gamma = 0$). The null is overwhelmingly rejected.

In Table 3 we also report the effect on parental savings of parents' lagged financial wealth and income (both in logs) as well as children's lagged (log) cash on hand (the sum of financial wealth and after-tax income) to capture the volume of resources that children can count on to smooth earnings shocks. Consistent with the model predictions, the initial stock of parents' financial wealth is negatively correlated with parents' current savings while the latter are positively correlated with the children's cash-on-hand; both are highly statistically significant.[25]

Given that we find no evidence that parents alter their saving behavior when children do well, in column (2) (and in all the other specifications we consider below) we impose $\gamma = 0$. The estimates of α_{Trans} and α_{Pers} are slightly larger in absolute value (0.39 and -0.25, respectively), and remain highly statistically significant.[26] The estimates of the other coefficients are unaffected. In column (3) we show that our estimates of α_{Trans} and α_{Pers} are virtually identical if we use (10) and (11) as auxiliary parameters in an Indirect Inference procedure (see Appendix OA.4).

Our data set includes all parent-child pairs, where a "child" can be either the parents' own child (biological or adopted) or the parents' child-in-law (if their own child is married). In column (4) we restrict the analysis to own children, and find that the results are unchanged. However, this stability may hide some heterogeneity in the child's marital status (insurance may be different for single vs. married children) as well as heterogeneity in consanguinity (insurance may be different for own children vs children-in-law). We analyze these issues more in detail below.

Overall, the evidence in Table 3 lends strong support to the dynastic channel for insuring labor income shocks. Consistent with the insurance model predictions, parents react to permanent negative shocks to their child's wages by saving more in anticipation of the transfers they plan to make in future years, when children will continue to need their support (what we call "saving for the (child's) rainy day"). If shocks are temporary, they dissave to smooth their child's consumption in the current period. Finally, if shocks are positive,

[25] It is worth noticing that the positive correlation between parents savings and children cash on hand arises if parents are altruistically linked to their children and when the transfer motive is active: transfers are of smaller size (i.e. parents dissave less) when children cash on hand are larger, giving rise to the correlation.

[26] Dropping parents with stocks in listed or unlisted companies (including companies where their child is employed) produces similar estimates: the estimates of α_{Trans} and α_{Pers} are 0.44 (s.e. 0.14) and -0.28 (s.e. 0.10), respectively.

parents do not change their saving behavior.[27]

Table 3: Parameter Estimates

	(1)	(2)	(3)	(4)
ω_{Trans}	0.4209^{***}			
	(0.0019)			
Elasticities to Shocks				
α_{Pers}	-0.2362^{***}	-0.2508^{***}	-0.2448^{***}	-0.2516^{**}
	(0.0814)	(0.0862)	(0.0920)	(0.1107)
α_{Trans}	0.3746^{***}	0.3947^{***}	0.4052^{***}	0.4039^{***}
	(0.1125)	(0.1191)	(0.1310)	(0.1522)
γ	-0.0002			
	(0.0007)			
Additional effects				
$\log(\text{par. FW})_{t-2}$	-0.0440^{***}	-0.0440^{***}	-0.0437^{***}	-0.0439^{***}
	(0.0001)	(0.0001)	(0.0002)	(0.0001)
$\log(\text{par. income})_{t-1}$	0.0253^{***}	0.0253^{***}	0.0247^{***}	0.0269^{***}
	(0.0003)	(0.0003)	(0.0003)	(0.0004)
$\log(\text{child cash-on-hand})_{t-1}$	0.0291^{***}	0.0291^{***}	0.0291^{***}	0.0283^{***}
	(0.0002)	(0.0002)	(0.0003)	(0.0003)
F-test of no altruism	158.97	112.53	195.34	129.96
(p-value)	<0.0001	<0.0001	<0.0001	<0.0001
N	13,550,902	13,550,902	13,550,902	8,690,684

Notes: In columns (1) the null hypothesis of the test of no altruism is $\alpha_{Trans} = \alpha_{Pers} = \gamma = 0$. In the other columns the null hypothesis is $\alpha_{Trans} = \alpha_{Pers} = 0$. In columns 1-3 the sample is both children and children-in-law (if present). In column 4 the sample is only own children. The block bootstrap (clustered by parent-child pair) is used to calculate the standard errors and the Wald statistics for the test of no altruism.

5.2.3 Quantification of Effects

How large are the effects we estimate? Using the estimates from column (2) of Table 3, we calculate that – faced with a 10% transitory drop in children's income – parents

[27] In principle, parents could finance transfer to their children by changing their labor supply rather than by changing their wealth accumulation decisions. However, we find no evidence that they do so. We regress the change in the parent's employment indicator and follow the same identification logic described in Section 3. The labor supply elasticity to transitory (permanent) shocks faced by the child is 0.04, with a standard error of 0.04 (-0.03, with a standard error of 0.03). This lack of response on the employment side is partly because some parents are well in their retirement stage of the life cycle, and it may be hard for them to find employment opportunities. But even for younger parents, there are important adjustment costs related to changing employment status or hours of work.

decrease their liquid assets by about 4%. Similarly, parental insurance activates in response to persistent shocks in children's labor earnings - a 10% persistent shocks faced by the child induces parents to increase their saving by 2.5% to cover a child's rainy day. To get a better sense of these figures, consider that a 10% drop in the child's labor income corresponds to roughly $8,200. If this drop is temporary, our estimates predict that parents would reduce their liquid wealth (potentially funding a direct or indirect transfer to their child) of about $3,500 (an insurance coverage ratio of 43%); if the shock is persistent, they would save an extra $2,200 to cover future transfers (an insurance coverage ratio of 27%).

A different way to quantify the effects is to compute marginal effects and, to avoid the influence of outliers, evaluate all the effects at the median. We find that the marginal effect of temporary earnings losses is 0.19 (s.e. 0.06), while the marginal effect of persistent earnings losses is -0.12 (s.e. 0.04). Hence, a $1 temporary decline in the child's earnings induces a decline in parents' savings of about 19 cents; in contrast, a $1 (perceived) persistent decline in the child's earnings induces an increase in parent's savings of about 12 cents. These are non-negligible effects, leaving us to conclude that parents provide a meaningful source of insurance even in the presence of formal insurance mechanisms (e.g., unemployment insurance).

5.3 Your parents or my parents? Checking insurance providers

In married couples that pool resources, it is immaterial which spouse suffers the adverse income shock: a dollar lost by one spouse is a dollar lost by all. In the previous sections, we have shown that an income loss induces parents to step up and insure such loss (by dissaving when the shock is temporary and by saving when it is persistent). However, there could be a violation of income pooling: a dollar lost by a daughter could be treated differently (for insurance purposes) than a dollar lost by a son-in-law.[28] Moreover, even if it exists, insurance of the child-in-law's income shocks could be contingent on the expectation that the marriage continues in the future. In this section, we provide evidence on how the insurance mechanisms studied in the previous section change when the child is married; in the next section, we study other sources of heterogeneity stemming from the demographics of the child's household (such as the presence of another set of parents who may, in principle, offer insurance against the same shock). We are able to study these issues thanks to the

[28]This may interact with violation of income pooling at the household level. If a dollar lost by a daughter weakens her bargaining position inside the marriage, her parent may want to provide insurance so as to reduce intra-household consumption reallocations. But if the same dollar is lost by the daughter's spouse and hence her bargaining position inside the marriage improves, her parent may not need to intervene since she's already better off (net of income effects).

richness of our data, which allow us to link each spouse to his or her parents and provide us with information about the wealth of each set of parents.

We regress the change in financial wealth of the parent set i, Δw_{it}^{P}, on two key covariates: Δy_{it}^{-K} (the negative shock suffered by their child), and Δy_{jt}^{K-} (the negative shock suffered by their child-in-law). We thus estimate the following variant of model (6):

$$\begin{aligned}\Delta w_{it}^{P} = {} & \alpha_{Trans,i}\Delta y_{Trans,it}^{-K} + \alpha_{Pers,i}\Delta y_{Pers,it}^{-K} + \alpha_{Trans,j}\Delta y_{Trans,jt}^{-K} + \alpha_{Pers,j}\Delta y_{Pers,jt}^{-K} \\ & + \beta_1 w_{it-1}^{P} + \beta_2 y_{it-1}^{P} + \beta_3 a_{it-1}^{K} + \eta_{it}^{P}\end{aligned}$$

The subscripts i and j index the child and the child's spouse, while the superscripts P and K continue to index the parent's and the child's household, respectively. If parents respond equally to a negative transitory (persistent) shock hitting their own child or his/her spouse, then $\alpha_{Trans,i} = \alpha_{Trans,j}$ ($\alpha_{Pers,i} = \alpha_{Pers,j}$).

Table 4 shows the results, focusing on the elasticity estimates (all regressions control for the same variables included in the main regression, namely the lagged values of the parents' financial wealth, parents' income, and child's income, all in logs). The first column reproduces the estimates of Table 3 in the restricted sample of households where both spouses' parents are present. It shows that in the sample of married children the parents' savings response is similar in magnitude to that in the sample that includes both married and single children (see Table 3).

The second column distinguishes between negative shocks hitting the parents' own child and negative shocks hitting the child-in-law. At face value, the response of parents to shocks is larger and statistically significant only if it is their own child that is hit by negative persistent or transitory shocks. If the shock hits the child's spouse, the point estimate is lower while the standard errors are unchanged, resulting in statistically insignificant responses.

The third column drops from the sample children who divorce at some point in the near future (5 years or less). Interestingly, while the size of parents response to a permanent drop in their own child's income is only slightly affected, the point estimate response to a shock to their child's spouse becomes as large (-0.33, significant at the 10% level). While the estimate remains imprecise, its value is consistent with the idea that parents may be willing to provide insurance to children-in-law's persistent shocks only insofar as they believe their children's marriage is relatively stable. This interpretation is further supported by the estimates of the elasticities to transitory negative shocks, showing little difference between own child and child-in-law when future divorcees are dropped from the sample.

Table 4: Heterogeneity in response by type of parent

	Sample		
	Married child+child-in-law (joint) (1)	Married child+child-in-law (separate) (2)	No Divorce within 5 years (3)
Elasticities to Shocks			
$\alpha_{Pers,i}$	-0.2531^{***}	-0.2858^{**}	-0.2962^{*}
	(0.1057)	(0.1416)	(0.1563)
$\alpha_{Pers,j}$		-0.2180	-0.3331
		(0.1552)	(0.2146)
$\alpha_{Trans,i}$	0.3903^{***}	0.4400^{**}	0.4562
	(0.1458)	(0.1945)	(0.2996)
$\alpha_{Trans,j}$		0.3375	0.4966^{**}
		(0.2137)	(0.2374)
Additional effects			
Income, wealth controls	Y	Y	Y
N	9,782,241	4,922,023	4,287,345

Notes: "No divorce within 5 years" is an indicator variable that takes the value 1 if the child is not observed divorcing his/her spouse within the next 5 years. All regressions consider the sub-sample of married children. The block bootstrap (clustered by parent-child pair) is used to calculate the standard errors.

5.4 Robustness

We perform a number of robustness tests. First, we repeat our main estimates focusing on a sample of children aged 25–45 instead of our baseline range of 25–55. Second, we retain only parents younger than 75, to take care of two concerns: old parents may have already decumulated their assets to finance old age consumption, or they may be cognitively unable to make discretionary transfer decisions. The estimated elasticities to both transitory and persistent shocks are smaller but qualitatively similar to those in the baseline sample, and are precisely estimated. Second, we estimate the model using a sub-sample that includes only child-parent pairs residing in different counties. This results in larger insurance estimates, most likely because parents that live in the same county as their children can provide more in-kind assistance (e.g., allowing children to move back home as in Kaplan 2012, or providing services such as child care), while those living in different counties can only offer monetary support, which is what our savings variable captures. Finally, to ensure that our findings are not driven by selection into firms after value-added shocks are realized, we use a sub-sample of children employed at their firm at the beginning of year t (i.e. January). Results remain qualitatively and quantitatively similar.

6 Extensions

In this section we consider some additional, intuitive implications of dynastic insurance. We first consider how parental insurance supply changes when children have access to alternative forms of insurance. Next, we study *reverse* insurance, namely monetary transfers from children to parents.

6.1 Heterogeneity in responses

It stands to reason that the insurance role of parents should be less relevant when the child has access to alternative forms of insurance. In Norway, social insurance is very generous but there is no variation that would allow us to test whether parents provide less insurance when the child has access to the state's safety net. In contrast, we can use demographic variation (whether a child is married or not), as well as variation within married couples (whether the child's spouse has an independent source of income), to test whether parents' need to intervene with cash transfers depends on the ability of the child to absorb earnings shocks through alternative insurance channels. Marriage can mitigate the effect of income shocks suffered by one spouse either because the other spouse can enter the labor market (an added worker effect) or can work more (if already employed). Moreover, in married couples with both sets of parents alive, the potential for insurance could be wider and hence shocks smoothed to a larger extent (see Laferrère and Wolff 2006 for a discussion). In practice, this need not be so. First, there could be free riding: each parent set may insure less waiting for the other to activate transfers. On the other hand, there may be "competition for attention" between the two sets of parents: a set may transfer money not to be outdone by the other set, losing part of the attention of the children.

We summarize the effect of being married and also having a working spouse in Table 5, which shows the estimated effects for different subsamples: single children (column 1), married children (column 2), married with working and non working spouse (columns 3 and 4, respectively) and married with only one set or two set of parents (columns 5 and 6, respectively). The first row provides the responses to persistent shocks. The second row reports the response to transitory shocks.

Comparing column (1) and column (2) - i.e. single and married children - reveals that (at least at face value) parents tend to offer more insurance to married kids (although the standard errors in the sample of singles are high because of the smaller sample size). One explanation is that, in response to very bad income shocks, single children can always return to live with their parents (an in-kind form of insurance), while for married children this is less of an option and insurance needs to be of a monetary nature. Comparing married with one set of parents and married with two sets - i.e. columns 3 and 4 - shows that, at face value, expanding the set of parents raises dynastic insurance coverage *vis-á-vis* transitory income losses (elasticity 0.42, compared to 0.36) as well as persistent ones (elasticity -0.28 compared to -0.24). Note that this is inconsistent with "free riding". We find that when parents in-law are present, parents provide *more* insurance, not less. One way to interpret this result is that it captures "competition for attention": parents offer more insurance in an attempt to sway the child's (or the couple's) attention towards them. Moving from married children with non-working spouse to married with a working spouse (columns 5 and 6) shows a rather dramatic drop in parental insurance provision. When the spouse does not work the elasticity of parents' savings to a transitory income drop suffered by the child is 1.08 (significantly different from zero at 10%);[29] when the spouse works it falls to 0.35 (significant at 10%). Similarly, the elasticity to persistent income losses fall (in absolute value) from -0.75 (significant at 10%) to -0.22 (significant at 10%). This strongly suggests that marriage is a particularly effective way of buffering shocks when both spouses work. It is only then that parents can significantly limit their role of insurance providers *vis-á-vis* their children labor market adversities.

[29]At sample means of parents financial assets and children earnings the elasticity point estimate of 1.2 implies insurance coverage greater than 1 (computed as (($\hat{\alpha}_{Trans}$×Parents financial wealth)/(Children earnings)); but we cannot reject that the elasticity equals the value (around 0.92, which at sample means just guarantees full coverage).

Table 5: Heterogeneity: marriage and spouse parents

	(1)	(2)	(3)	(4)	(5)	(6)
Child's marital status	Single	Married	Married	Married	Married	Married
Sets of parents			One	Two		
Working spouse					No	Yes
$\hat{\alpha}_{Pers}$	-0.1880	-0.2531**	-0.2402	-0.2776**	-0.7543*	-0.2225**
	(0.2120)	(0.1057)	(0.2450)	(0.1237)	(0.4320)	(0.1075)
$\hat{\alpha}_{Trans}$	0.3307	0.3903***	0.3602	0.4238***	1.0792*	0.3482***
	(0.2917)	(0.1458)	(0.3380)	(0.1710)	(0.5918)	(0.1485)
N	3,768,661	9,782,241	782,386	8,542,523	516,399	9,265,842

Note: The table shows the estimated heterogeneity in parental saving responses to income shocks. The block bootstrap (clustered by parent-child pair) is used to calculate the standard errors.

6.2 Reverse Insurance: Do Children Insure Parents?

Dynastic insurance arises from cash-rich parents transferring resources to cash-constrained children when markets for smoothing income shocks are incomplete. One implication of dynastic insurance is that it is likely to be unilateral: even if children are as altruistic towards their parents as their parents are towards them, they typically do not have enough assets to activate transfers to their parents in case the latter are hit by adverse shocks to income.[30] We can test this distinct implication of the dynastic model by reversing our empirical model. We estimate a specification equivalent to (2), but with the change in the child's liquid assets on the left hand side and the shocks to parents' income on the right hand side.

If children insure parents we would expect a positive correlation between unexpected drops in parents' earnings and the child's change in liquid assets, and no relation if parents experience positive shocks. Following the same logic illustrated in Section 2, we would expect children to save more (or dissave less) when their parents face permanent income shocks (e.g., due to disability). That is, we should observe a negative relation between the change in children's liquid assets and persistent shocks to parents' income. One caveat is that by requiring parents to be employed (so that instruments are well defined), we are implicitly focusing on relatively younger children-parent sets. This implies that children are

[30] In our sample, average financial wealth held by children is roughly half of the average amount held by parents. Even at the median, parents have twice as much financial wealth as their child. Nonetheless, children could still provide assistance to their parents in kind. Unfortunately, we do not observe these.

less likely to have accumulated assets and hence less able to provide reverse insurance in the first place.[31]

Table 6 provides the estimated (reverse) insurance parameters. While the signs of the effects are consistent with reverse insurance, the effects are economically and statistically indistinguishable from zero. This lack of evidence for children offering insurance to their parents is consistent with the story that dynastic insurance of income shocks runs mostly from parents to children, as only the former are in a position to absorb children's labor income shocks. Our findings complement Boar (2021), who shows that while parents consumption responds to the variance of children's labor income, the reverse is not true: children reduce consumption in response to their own income risk but not that of their parents. Our results are also qualitatively consistent with Gale and Scholz (1994), according to whom *inter vivos* transfers from parents to children are about ten times larger than those from children to parents.[32]

Table 6: Reverse Insurance

		(1)	(2)
Elasticities to Shocks	$\hat{\alpha}_{Perm}$	-0.0567	-0.0485
		(0.3553)	(0.4560)
	$\hat{\alpha}_{Trans}$	0.0733	0.0698
		(0.4885)	(0.6274)
N		4,573,024	3,054,665

Notes: This table considers the relationship between negative changes in parental income (Δy^{-P}) and the change in children household-level financial wealth. Parental income includes transfers and pensions, but the sample restrictions require that the main parent-household earner is not receiving any pension income. To identify $\hat{\alpha}_{Perm}$, we use value-added shocks to the parent's employer as an instrument for income changes. In a separate regression, we obtain OLS estimates by regressing changes in the child's household financial wealth on Δy^{-P}. To identify $\hat{\alpha}_{Trans}$, we combine the OLS and $\hat{\alpha}_{Perm}$ estimates as described in Section 3. Column (1) uses both children and children-in-law. Column (2) is restricted to own children only. The block bootstrap (clustered by parent-child pair) is used to calculate the standard errors.

[31] In the reverse-insurance sub-sample, the mean (median) child age is 32.4 (32) years, which contrasts with our main analysis sample where the mean (median) is 38.5 (38). Child average household financial wealth in this sub-sample is approximately NOK 168,000, which is about a third less than in the main sample. On the other hand, parents are also slightly less wealthy in this sample, which may partially offset the fact that children have less assets. Parental financial wealth is about 428,000 on average, which is 7% lower than in the main sample.

[32] We also find no evidence that children change their labor supply decision in response to parents receiving negative income shocks. We regress the change in the child's employment indicator and follow the same identification logic described in Section 3. The labor supply elasticity to transitory (permanent) shocks faced by the parent is -0.004, with a standard error of 0.012 (0.003, with a standard error of 0.008).

7 Conclusions

Our paper presents new evidence that parents serve a key role in insuring their adult children's income shocks. We find that parents deplete their savings when children experience large, negative income shocks. This relationship does not hold for positive shocks, consistent with the insurance motive for parental transfers. Importantly, we find that current transfers compete with expected future transfers. Shocks that are primarily permanent in nature trigger savings-accumulation responses. This implies that a small relationship between overall shocks (including transitory and permanent elements) may mask more substantial but offsetting responses as permanent shocks cause additional saving but transitory shocks cause parents to deplete their assets. Our methodology – combining OLS and IV estimates using plausibly exogenous variation in the child's income induced by pass-through of firm productivity shock onto wages – allows to separate the two responses and identify parents' savings elasticity to transitory and permanents drops in children income. They are both substantial, covering 27% of permanent and about 43% of transitory negative shocks to children income. Furthermore, parents activate insurance-motivated transfers only when income losses are large enough and increase transfers when the size of the loss is larger. Because of this, dynastic insurance is an effective way to attenuate the adverse effects on workers consumption of fat negative tails in the distribution of labor income growth during recessions, first documented by Guvenen et al. (2014). Hence, dynastic insurance can greatly mitigate the macroeconomic effects of idiosyncratic shocks.

We also find evidence that the estimated effects exhibits economically relevant heterogeneity. Parental insurance is not unconditional: it becomes active when children lack alternative ways of smoothing labor income shocks, particularly through marriage and access to transfers from in-laws. Furthermore, parents' willingness to supply insurance is affected by the stability of their offspring's marriage: our evidence implies that if parents anticipate a divorce, they become less cautious in insuring shocks to their child-in-law. Thus, divorce may not only destroy risk pooling opportunities among spouses—it also weakens dynastic insurance.

References

Alstadsæter, A., Johannesen, N., & Zucman, G. (2018). Who owns the wealth in tax havens? macro evidence and implications for global inequality. *Journal of Public Economics*, *162*, 89–100.

Altonji, J. G., Hayashi, F., & Kotlikoff, L. J. (1997). Parental altruism and inter vivos transfers: Theory and evidence. *Journal of Political Economy*, *105*(6), 1121–1166.

Andersen, A. L., Johannesen, N., & Sheridan, A. (2020). Bailing out the kids: New evidence on informal insurance from one billion bank transfers. *Available at SSRN 3624022*.

Bardoczy, B. (2020). Spousal insurance and the amplification of business cycles. *Unpublished Manuscript, Northwestern University*.

Basu, S., & Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, *85*(3), 937–958.

Bayer, C., Lütticke, R., Pham-Dao, L., & Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, *87*(1), 255–290.

Benetton, M., Kudlyak, M., & Mondragon, J. (2022). Dynastic home equity.

Blundell, R., Pistaferri, L., & Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, *98*(5), 1887–1921.

Blundell, R., Pistaferri, L., & Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, *106*(2), 387–435.

Boar, C. (2021). Dynastic precautionary savings. *The Review of Economic Studies*, *88*(6), 2735–2765.

Campbell, J. Y. (1987). Does saving anticipate declining labor income? an alternative test of the permanent income hypothesis. *Econometrica*, *55*(6), 1249–1273. Retrieved July 29, 2023, from http://www.jstor.org/stable/1913556

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, *36*(S1), S13–S70. https://doi.org/10.1086/694153

Cox, D. (1987). Motives for private income transfers. *Journal of Political Economy*, *95*(3), 508–546.

Fagereng, A., Guiso, L., Malacrino, D., & Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, *88*(1), 115–170.

Gourieroux, C., Monfort, A., & Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, *8*, S85–S118. Retrieved May 30, 2024, from http://www.jstor.org/stable/2285076

Guiso, L., & Jappelli, T. (1991). Intergenerational transfers and capital market imperfections: Evidence from a cross-section of italian households. *European Economic Review*, *35*(1), 103–120. https://doi.org/https://doi.org/10.1016/0014-2921(91)90106-S

Guiso, L., & Pistaferri, L. (2020). The insurance role of the firm. *The Geneva Risk and Insurance Review*, *45*(1), 1–23.

Guiso, L., Pistaferri, L., & Schivardi, F. (2005). Insurance within the firm. *Journal of Political Economy*, *113*(5), 1054–1087.

Guler, B., Guvenen, F., & Violante, G. L. (2012). Joint-search theory: New opportunities and new frictions. *Journal of Monetary Economics*, *59*(4), 352–369.

Guvenen, F., Ozkan, S., & Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, *122*(3), 621–660.

Halvorsen, E., & Thoresen, T. O. (2010). Parentsâ Desire to Make Equal Inter Vivos Transfers*. *CESifo Economic Studies*, *57*(1), 121–155. https://doi.org/10.1093/cesifo/ifq007

Heathcote, J., Storesletten, K., & Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, *104*(7), 2075–2126.

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

Kline, P., Petkova, N., Williams, H., & Zidar, O. (2019). Who Profits from Patents? Rent-Sharing at Innovative Firms*. *The Quarterly Journal of Economics*, *134*(3), 1343–1404. https://doi.org/10.1093/qje/qjz011

Laferrère, A., & Wolff, F.-C. (2006). Microeconomic models of family transfers. *Handbook of the economics of giving, altruism and reciprocity*, *2*, 889–969.

Lamadon, T., Mogstad, M., & Setzler, B. (2022). Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, *112*(1), 169–212. https://doi.org/10.1257/aer.20190790

Leduc, S., & Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, *82*, 20–35.

McGarry, K. (1999). Inter vivos transfers and intended bequests. *Journal of Public Economics*, *73*(3), 321–351. https://doi.org/https://doi.org/10.1016/S0047-2727(99)00017-1

Ortigueira, S., & Siassi, N. (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics*, *60*(6), 650–666.

Pistaferri, L. (2001). Superior Information, Income Shocks, and the Permanent Income Hypothesis. *The Review of Economics and Statistics*, *83*(3), 465–476. https://doi.org/10.1162/00346530152480117

Ring, M. A. K. (2019). Wealth taxation and household saving: Evidence from assessment discontinuities in norway. *Available at SSRN 3716257*.

Ring, M. A. K., & Thoresen, T. O. (2022). Wealth taxation and charitable giving. *CESifo Working Paper No. 9700*.

Schaab, A. (2020). Micro and macro uncertainty. *Available at SSRN 4099000*.

Waldkirch, A., Ng, S., & Cox, D. (2004). Intergenerational linkages in consumption behavior. *Journal of Human Resources*, *39*(2), 355–381.

Online Appendix for "Insuring Labor Income Shocks: The Role of the Dynasty"

by Andreas Fagereng, Luigi Guiso, Luigi Pistaferri and Marius A.K. Ring

In this Online Appendix we provide supplementary material to the article.

OA.1 Proof of Propositions

Proof of proposition 1

Under the assumption $\sigma_\eta = 0$,the parents' problem consists of:

$Max_{w_1^P,\tau_1,\tau_2} : u(y_0^P + w_0^P - w_1^P - \tau_1) + u(w_1^P - \tau_2) + \alpha[u(a^K + \tau_1 + \epsilon_1) + u(a^K + \tau_2 + \epsilon_2)]$
s.t. $\tau_1 \geq 0; \tau_2 \geq 0$

We have three first order conditions (FOC):

FOC1 (optimal parent savings): $u'(y_0^P + w_0^P - w_1^P - \tau_1) = u'(w_1^P - \tau_2)$

FOC2 (optimal first period transfer): $u'(y_0^P + w_0^P - w_1^P - \tau_1) = \alpha u'(a^K + \tau_1 + \epsilon_1) + \lambda_1$

FOC3 (optimal second period transfer): $u'(w_1^P - \tau_2) = \alpha u'(a^K + \tau_2 + \epsilon_2) + \lambda_2$

and the complementary slackness conditions:

(CSC): $\lambda_1 \geq 0$ $(\lambda_1 \tau_1 = 0)$; $\lambda_2 \geq 0$ $(\lambda_2 \tau_2 = 0)$

where λ_1 and λ_2 are the Karush–Kuhn–Tucker multipliers associated to the non-negativity constraints on the two transfers.

Consider first the non-altruistic case in which $\alpha = 0$. Making positive transfers is clearly not optimal because transfers are pure "bads" and "consuming" positive quantities of them reduces total utility, so that parents will choose the corner solutions $\tau_1^* = \tau_2^* = 0$ and their saving choices will be independent of the child's income realizations.

Consider now the altruistic case in which $\alpha > 0$. From FOC1 we get $w_1^P = \frac{1}{2}(y_0^P + w_0^P + \tau_2 - \tau_1)$, and thus an expression for saving in period 1:

$$\Delta w_1^P = w_1^P - w_0^P = \frac{1}{2}(y_0^P - w_0^P + \tau_2 - \tau_1)$$

First, we want to determine under what conditions it is optimal for parents to transfer resources to their children. There are three cases of interest: (a) $\tau_1^* = \tau_2^* = 0$; (b) $\tau_1^* > 0, \tau_2^* = 0$; and (c) $\tau_1^* > 0, \tau_2^* > 0$.[33]

From the complementary slackness conditions, the constraint on parental transfers are binding ($\tau_1^* = \tau_2^* = 0$) whenever $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$, which require that $u'(\frac{y_0^P + w_0^P}{2}) \geq \alpha u^{'}(a^K + \epsilon_t)$ (for $t = \{1, 2\}$). Given concavity of the utility function, there exists a threshold value of ϵ_t (say, $\bar{\epsilon}$), such that $u'\left(\frac{y^P + w_0^P}{2}\right) = \alpha u'\left(a^K + \epsilon_t\right)$, so that $\lambda_t = 0$. Hence, it is optimal to set transfers to zero whenever $\epsilon_t \geq \bar{\epsilon}$, i.e., if the shock is sufficiently large. In the spirit of the role of insurance, assume that in the absence of shocks, i.e., when $\epsilon_t = 0$, parents make no transfers, implying $u'\left(\frac{y_0^P + w_0^P}{2}\right) \geq \alpha u'\left(a^K\right)$. Assuming $\alpha < 1$, this condition holds if the the child's endowment is at least a certain share of the parents' endowment. If parents make no transfers when the shock is zero, *a fortiori*, they will not do so when the shock is positive. Hence, $\bar{\epsilon} \leq 0$. Totally differentiating the condition $u'\left(\frac{y_0^P + w_0^P}{2}\right) = \alpha u'\left(a^K + \bar{\epsilon}\right)$, and using concavity of $u(.)$ shows that the threshold $\bar{\epsilon}$ is decreasing in the child's endowment a^K and increasing in parents' income endowment y_0^P , initial wealth w_0 and degree of altruism α. Note also that, since $\epsilon_2 = \rho\epsilon_1$, the constraint on second-period transfers is binding ($\tau_2^* = 0$) when $\epsilon_1 \geq \frac{\bar{\epsilon}}{\rho}$.

The second case of interest ($\tau_1^* > 0, \tau_2^* = 0$) arises when the shock is negative but not large enough in absolute value, i.e., $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$. In this case the first period transfer is positive and determined by :

$$u'\left(\frac{y_0^P + w_0^P - \tau_1^*}{2}\right) = \alpha u^{'}(a^K + \tau_1^* + \epsilon_1),$$

which is increasing in parents' endowment, initial wealth, altruism and income loss and decreasing in children's cash on hand.

The final case of interest ($\tau_1^* > 0, \tau_2^* > 0$) occurs when the shock is large and negative ($\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$). In this case both transfers are operative, and their optimal value is determined by:

$$u'\left(\frac{y_0^P + w_0^P - (\tau_1^* + \tau_2^*)}{2}\right) = \alpha u^{'}(a^K + \tau_1^* + \epsilon_1)$$

$$u'\left(\frac{y_0^P + w_0^P - (\tau_1^* + \tau_2^*)}{2}\right) = \alpha u^{'}(a^K + \tau_2^* + \epsilon_2)$$

implying that parents make transfers to equalize the marginal utility of children's con-

[33] A fourth case where $\tau_1^* = 0, \tau_2^* > 0$ is a mathematical impossibility.

sumption across the two periods. This implies that $\tau_2^* = \tau_1^* + \epsilon_1 - \epsilon_2$. Using this, total differentiation of the first order conditions shows that first period transfer increases with parents' endowment, degree of altruism α, and the size of first period income loss; it decreases with children's cash on hand and the size of second period income loss, i.e., the degree of persistence. Following the same logic, the second period transfer increases with parents endowment and altruism and the degree of persistence of the shock; it decreases with the child's endowment. $\therefore$

Proof of proposition 2

Recall from FOC1 that the change in wealth in period 1 (the saving flow) is

$$\Delta w_1^P = \frac{1}{2}(y_0^P + w_0^P + \tau_2 - \tau_1)$$

Let $u(x) = \frac{x^{1-\delta}}{1-\delta}$. Then the condition for no transfer in the absence of a shock $u'(\frac{1}{2}(y_0^P + w_0^P)) \geq \alpha u'(a^K)$ is

$$a^K \geq \mu \frac{y_0^P + w_0^P}{2}$$

where $\mu = \alpha^{1/\delta} < 1$ if the risk aversion parameter is positive and finite.

The threshold for the income shocks that defines whether the transfers are active is (under the maintained assumptions):

$$\bar{\epsilon} = \mu \frac{y_0^P + w_0^P}{2} - a^K \leq 0$$

Parents response to the shocks varies depending on which transfer is active. Consider again the three relevant cases: (a) $\tau_1^* = \tau_2^* = 0$ (which occurs when $\epsilon_1 \geq \bar{\epsilon}$); (b) $\tau_1^* > 0, \tau_2^* = 0$ ($\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$); and (c) $\tau_1^* > 0, \tau_2^* > 0$ ($\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$).

Transfers not active: $\epsilon_1 \geq \bar{\epsilon}$

Setting transfers to zero in the general expression for parents wealth dynamics we have :

$$\Delta w_1^P = \frac{y_0^P - w_0^P}{2}$$

Transfers active in period 1 only: $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$

The optimal first period transfer is determined from the FOC2 in Proposition 1 after substituting for w_1^P from FOC1:

$$u'\left(\frac{y_0^P + w_0^P - \tau_1^*}{2}\right) = \alpha u'(a^K + \tau_1^* + \epsilon_1)$$

Using the CRRA utility assumption yields

$$\tau_1^* = \frac{2}{2+\mu}\left(\mu\frac{y_0^P + w_0^P}{2} - a^K - \epsilon_1\right) = \frac{2}{2+\mu}(\bar{\epsilon} - \epsilon_1)$$

Using this expression and and setting $\tau_2^* = 0$, parents wealth dynamics among those with an active transfer in period 1 is

$$\Delta w_1^P = \frac{1}{2+\mu}y_0^P - \frac{1+\mu}{2+\mu}w_0^P + \frac{1}{2+\mu}\left(a^K + \epsilon_1\right).$$

Transfers active in both periods: $\epsilon_1 < \frac{\bar{\epsilon}}{\rho}$.

From the proof of proposition 1, the optimal transfers in this case are determined by the first order conditions

$$u'\left(\frac{y_0^P + w_0^P - (\tau_1^* + \tau_2^*)}{2}\right) = \alpha u^{'}(a^K + \tau_1^* + \epsilon_1)$$

$$u'\left(\frac{y_0^P + w_0^P - (\tau_1^* + \tau_2^*)}{2}\right) = \alpha u^{'}(a^K + \tau_2^* + \epsilon_2)$$

Using the CRRA utility and solving for τ_1 and τ_2 yields:

$$\tau_1^* = \frac{1}{1+\mu}\left(\mu\frac{y_0^P + w_0^P}{2} - a^K\right) - \frac{2+\mu}{2(1+\mu)}\epsilon_1 + \frac{\mu}{2(1+\mu)}\epsilon_2$$

$$\tau_2^* = \frac{1}{1+\mu}\left(\mu\frac{y_0^P + w_0^P}{2} - a^K\right) + \frac{\mu}{2(1+\mu)}\epsilon_1 - \frac{2+\mu}{2(1+\mu)}\epsilon_2$$

Using these expressions, parents wealth dynamics is

$$\Delta w_1^P = \frac{y_0^P - w_0^P}{2} + \frac{(\epsilon_1 - \epsilon_2)}{2} = \frac{y_0^P - w_0^P}{2} + \frac{(1-\rho)}{2}\epsilon_1$$

Suppose we ignore the intermediate case $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$ (for example because ρ is large enough, so there is very little density mass in the $\frac{\bar{\epsilon}}{\rho} \leq \epsilon_1 < \bar{\epsilon}$ region). Then we can put together the two cases of interest and express the change in parents wealth (saving) as:

$$\begin{aligned}\Delta w_1^P &= \left(\frac{y_0^P - w_0^P}{2}\right) I^+ + \left(\frac{y_0^P - w_0^P}{2} + \frac{(1-\rho)}{2}\epsilon_1\right) I^- \\ &= \left(\frac{y_0^P - w_0^P}{2}\right) I^+ + \underbrace{\left(\frac{y_0^P - w_0^P}{2} + \frac{1}{2}\epsilon_1\right)}_{\alpha_T} I^- + \underbrace{\left(\frac{y_0^P - w_0^P}{2} - \frac{1}{2}E(\epsilon_2|\epsilon_1)\right)}_{\alpha_P} I^-\end{aligned}$$

where $I^+ = I(\epsilon_1 \geq \bar{\epsilon})$, and $I^- = I(\epsilon_1 < \frac{\bar{\epsilon}}{\rho})$ are indicator functions for the two intervals of interest (positive/slightly negative shocks vs. moderately/strongly negative income shocks). When income shocks are positive there is no change in parental savings (i.e., they are independent of the child's income shock). When shocks are negative, parental saving will adjust in response to the child's income shock. Moreover, the adjustment will depend on how persistent shocks are. If shocks are purely transitory ($\rho = 0$), parents will decumulate to fund a current transfer. As shocks become more persistent, the saving response will be dampened by the prospect of having to make future transfers in the future. The second row in the equation above tries to make this more explicit by writing the transitory and persistent component of the income process as ϵ_1 and $E(\epsilon_2|\epsilon_1)$, respectively. A negative shock today has two effects: an immediate effect (measured by the α_T coefficient), inducing parents to dissave in order to finance a transfer to smooth the child's consumption, and a delayed effect (the coefficient α_P) that works through altering the expectations about future income (i.e., if the shock is expected to be persistent), and inducing parents to save for a (child's) rainy day.

OA.2 Proof of equations (7) and (8)

In the symmetric case, parental wealth is (in its simplified form, omitting other controls):

$$\Delta w_t^P = \alpha_{Trans} \Delta y_{Trans,t}^K + \alpha_{Pers} \Delta y_{Pers,t}^K + \eta_t^P, \tag{OA.1}$$

This regression is infeasible because we do not observe $\Delta y_{Trans,t}^K$ and $\Delta y_{Pers,t}^K$ separately; we only observe their sum, Δy_t^K. Consider then running an OLS regression of Δw_t^P on Δy_t^K. The probability limit of the OLS estimate is:

$$
\begin{aligned}
\text{plim } \hat{\alpha}^{OLS} &= \text{plim}\frac{\text{cov}(\Delta w_t^P, \Delta y_t^K)}{\text{var}(\Delta y_t^K)} \\
&= \text{plim}\frac{\text{cov}(\alpha_{Trans}\Delta y^K_{Trans,t} + \alpha_{Pers}\Delta y^K_{Pers,t} + \eta_t^P, \Delta y^K_{Trans,t} + \Delta y^K_{Pers,t})}{\text{var}(\Delta y^K_{Trans,t} + \Delta y^K_{Pers,t})} \\
&= \text{plim}\frac{\alpha_{Trans}\text{var}(\Delta y^K_{Trans,t}) + \alpha_{Perm}\text{var}(\Delta y^K_{Perm,t})}{\text{var}(\Delta y^K_{Trans,t} + \Delta y^K_{Pers,t})} \\
&= \omega_{Trans}\alpha_{Trans} + (1 - \omega_{Trans})\alpha_{Pers}
\end{aligned}
$$

where $\omega_{Trans} = \text{plim}\frac{\text{var}(\Delta y^K_{Trans,t})}{\text{var}(\Delta y^K_{Trans,t} + \Delta y^K_{Pers,t})}$ is the share of total earnings growth variance due to transitory shocks. We have assumed that the the two shocks are orthogonal to each other and are orthogonal to idiosyncratic variation in parental wealth, η_t^P.

Consider next the Instrumental Variables (IV) regression in in which we use productivity shocks to the firm where the child is employed as an instrument for the child's earnings growth in order to isolate persistent variation in wages. In particular, we consider:

$$
\begin{aligned}
\text{plim } \hat{\alpha}^{IV} &= \text{plim}\frac{\text{cov}(\Delta w_t^P, \Delta VA_t^F)}{\text{cov}(\Delta y_t^K, \Delta VA_t^F)} \\
&= \text{plim}\frac{\text{cov}(\alpha_{Trans}\Delta y^K_{Trans,t} + \alpha_{Pers}\Delta y^K_{Pers,t} + \eta_t^P, \Delta VA_t^F)}{\text{cov}(\Delta y^K_{Trans,t} + \Delta y^K_{Pers,t}, \Delta VA_t^F)} \\
&= \text{plim}\frac{\alpha_{Trans}\text{cov}(\Delta y^K_{Trans,t}, \Delta VA_t^F) + \alpha_{Pers}\text{cov}(\Delta y^K_{Pers,t}, \Delta VA_t^F) + \text{cov}(\eta_t^P, \Delta VA_t^F)}{\text{cov}(\Delta y^K_{Trans,t}, \Delta VA_t^F) + \text{cov}(\Delta y^K_{Pers,t}, \Delta VA_t^F)} = \\
&= \alpha_{Pers}
\end{aligned}
$$

This is under the following assumptions, spelled out in the text: (a) unobserved heterogeneity in parental saving is orthogonal to firm value added shocks (plim cov $(\eta_t^P, \Delta VA_t^F) = 0$); (b) permanent value added shocks load onto the persistent component of earnings (plim cov$(\Delta y^K_{Pers,t}, \Delta VA_t^F) \neq 0$, which requires $\theta \neq 0$); and (c) temporary value added shocks are "insured within the firm" (plim cov$(\Delta y^K_{Trans,t}, \Delta VA_t^F) = 0$). In the main text we discuss the plausibility of these assumptions. In the next section we provide direct tests of assumptions (b) and (c).

OA.3 Estimating pass-through of transitory and persistent firm shocks

Suppose that the permanent component of the firm's value added follows a random walk process, so that:

$$\Delta VA_t^F = \nu_t^F \tag{OA.2}$$

and the transitory component is i.i.d., $VA_{jt,Trans}^F = \kappa_{jt}^F$, so that $\Delta VA_{jt}^F = \nu_t^F + \Delta\kappa_{jt}^F$.

Consider a more general case in which both transitory and persistent shocks to value added may pass through onto wages, i.e..

$$\begin{aligned}
\Delta y_{it}^K &= \Delta y_{it,Pers}^K + \Delta y_{it,Trans}^K \\
&= \underbrace{(\theta \Delta VA_{jt,Pers}^F + \tilde{\zeta}_{it})}_{\Delta y_{it,Pers}^K} \quad + \underbrace{(\lambda \Delta VA_{jt,Trans}^F + \Delta\varepsilon_{it})}_{\Delta y_{it,Trans}^K} \\
&= \theta \nu_t^F + \lambda \Delta\kappa_{jt}^F + \tilde{\zeta}_{it} + \Delta\varepsilon_{it}
\end{aligned}$$

We need to estimate the following parameters: $\sigma_\eta^2, \sigma_\varepsilon^2, \sigma_\zeta^2, \sigma_\nu^2, \theta$ and λ. To do so we use GMM on the following moments (clustering standard errors at the year-firm level since this is the shock that is common to all workers):

$$\begin{aligned}
E((\Delta VA_{jt}^F)^2) &= \sigma_\nu^2 + 2\sigma_\kappa^2 \\
E(\Delta VA_{jt}^F \Delta VA_{jt-1}^F) &= -\sigma_\kappa^2 \\
E((\Delta y_{it}^K)^2) &= \theta^2\sigma_\nu^2 + 2\lambda^2\sigma_\kappa^2 + \sigma_{\tilde{\zeta}}^2 + 2\sigma_\varepsilon^2 \\
E(\Delta y_{it}^K \Delta y_{it-1}^K) &= -\lambda\sigma_\kappa^2 - \sigma_\varepsilon^2 \\
E(\Delta y_{it}^K \Delta VA_{jt}^F) &= \theta\sigma_\nu^2 + 2\lambda\sigma_\kappa^2 \\
E(\Delta y_{it}^K \Delta VA_{jt-1}^F) &= -\lambda\sigma_\kappa^2 \\
E(\Delta y_{it-1}^K \Delta VA_{jt}^F) &= -\lambda\sigma_\kappa^2
\end{aligned}$$

We obtain the estimates reported in column (1) of Table OA.1. The pass-through coefficient of transitory shocks (λ) is statistically insignificant and quite small (an order of magnitude smaller than the one on the persistent component). In fact, even at face value, the amount of wage variation due to the transmission of firm shocks is almost entirely coming from the pass-through of firm shocks of persistent nature: $\frac{\theta^2\sigma_\eta^2}{\theta^2\sigma_\eta^2 + 2\lambda^2\sigma_\varepsilon^2} = 98\%$. This confirms our identifying assumption. In column (2) we estimate the model imposing $\lambda = 0$. These are the estimates we feed into the Indirect Inference code, explained below.

Table OA.1: Estimates of the joint earnings growth-value added growth process

	(1)	(2)
σ^2_ν	0.1209***	0.1204***
	(0.0001)	(0.0001)
σ^2_κ	0.1379***	0.1390***
	(0.0001)	(0.0001)
$\sigma^2_{\tilde{\zeta}}$	0.1134***	0.1130***
	(0.0008)	(0.0008)
σ^2_ε	0.0407***	0.0411***
	(0.0007)	(0.0006)
θ	0.0247***	0.0298***
	(0.0034)	(0.0023)
λ	0.0022	
	(0.0013)	

OA.4 Indirect Inference

We first simulate the income process using the estimates of $\sigma^2_\eta, \sigma^2_\varepsilon, \sigma^2_\zeta, \sigma^2_\nu, \theta$ described in the previous section:

$$\Delta VA^F_{jt} = \nu^F_t + \Delta\kappa^F_{jt}$$
$$\Delta y^K_{it} = \theta\nu^F_t + \tilde{\zeta}_{it} + \Delta\varepsilon_{it} = \zeta_{it} + \Delta\varepsilon_{it}$$

This is done for a number of observations equal to our sample size. Our structural model is one for parental wealth from specification (2):

$$\Delta w^P_t = \alpha_{Trans}\Delta\varepsilon^-_{it} + \alpha_{Pers}\zeta^-_{it} + \beta_1 w^P_{t-1} + \beta_2 y^P_{t-1} + \beta_3 a^K_{t-1} + \eta^P_t$$

where we assume that $\eta^P_t \sim N(0, \sigma^2_\eta)$. This error term is simulated as well using the residual of our OLS wealth regression to obtain an estimate of σ^2_η. The structural parameters we are interested in recovering are in the vector: $\Theta = (\alpha_{Trans}, \alpha_{Pers}, \beta_1, \beta_2, \beta_3)$. Consider the OLS and IV regressions:

$$\check{\psi}_{OLS} = (\sum x'_t x_t)^{-1} \sum x'_t \Delta w^P_t \qquad \text{(OA.3)}$$

$$\check{\psi}_{IV} = (\sum z_t' x_t)^{-1} \sum z_t' \Delta w_t^P \tag{OA.4}$$

where $x_t = (\Delta y_t^{-,K}, w_{t-1}^P, y_{t-1}^P, a_{t-1}^K)$ is the vector of controls in the OLS regression and $z_t = (\Delta VA_t^{-,F}, w_{t-1}^P, y_{t-1}^P, a_{t-1}^K)$ the vector of (excluded and included) instruments in the IV regression. Call $\check{\psi} = (\check{\psi}_{OLS}, \check{\psi}_{IV})$ the coefficients from the auxiliary regressions. Indirect Inference solves the problem:

$$\hat{\Theta}_{II} = \arg\min_{\Theta} (\check{\psi} - S^{-1} \sum_{s=1}^{S} \check{\psi}^s(\Theta))' \Omega (\check{\psi} - S^{-1} \sum_{s=1}^{S} \check{\psi}^s(\Theta))$$

where $\check{\psi}^s(\Theta)$ is an estimate of $\check{\psi}$ obtained by simulating the model for a given value of Θ. We choose $S = 20$. Standard errors are computed as:

$$var(\hat{\Theta}_{II}) = (1 + \frac{1}{S})(J'\Omega J)^{-1} J'\Omega V \Omega J (J'\Omega J)^{-1}$$

where J is the Jacobian matrix (evaluated at the estimated parameter values $\hat{\Theta}_{II}$), and V the variance matrix of the auxiliary parameters. Under regularity conditions spelled out in Gourieroux et al. (1993), $\text{plim}\hat{\Theta}_{II} = \Theta$. These estimates are reported in column (3) of Table 3.

OA.5 Additional Tables and Figures

Table OA.2: Robustness Checks

Note: This table provides the elasticities of parent financial wealth to child income shocks. In column (1) we reproduce the baseline estimates. In column (2) we focus on a sample of children aged 25-45. In column (3) we restrict the analysis to parents younger than 75. In column (4) we focus on a sample of parents and children living in different counties. Column (5) focuses on observations in which the children were employed by their employers on January 1st. The block bootstrap (clustered by parent-child pair) is used to calculate the standard errors.

$\hat{\alpha}_{Pers}$	-0.2508**	-0.1913*	-0.2143**	-0.7354***	-0.2066***
	(0.0862)	(0.1118)	(0.1083)	(0.2132)	(0.0817)
$\hat{\alpha}_{Trans}$	0.3947**	0.3088**	0.3398**	1.0522***	0.3322***
	(0.1191)	(0.1544)	(0.1498)	(0.2949)	(0.1125)
N	13,550,902	10,446,753	10,903,730	3,730,705	11,973,149

Figure OA.1: The Relationship Between Value-Added Shocks and Positive as well as Over-all Changes in Child Labor Earnings

Notes: Panel A shows the first-stage relationship between positive changes in the child's labor earnings (Δy^{+K}) and value-added shocks experienced by the child's employer. Panel A shows the first-stage relationship between over-all changes in the child's labor earnings (Δy^{K}) and value-added shocks experienced by the child's employer. The solid line is a local polynomial fit (lpoly, bandwidth=0.33), excluding values exceeding 1.5 in absolute value. The dashed lines perform the local polynomial fit separately for negative and positive value-added shocks.

Panel A: Positive Changes in Earnings

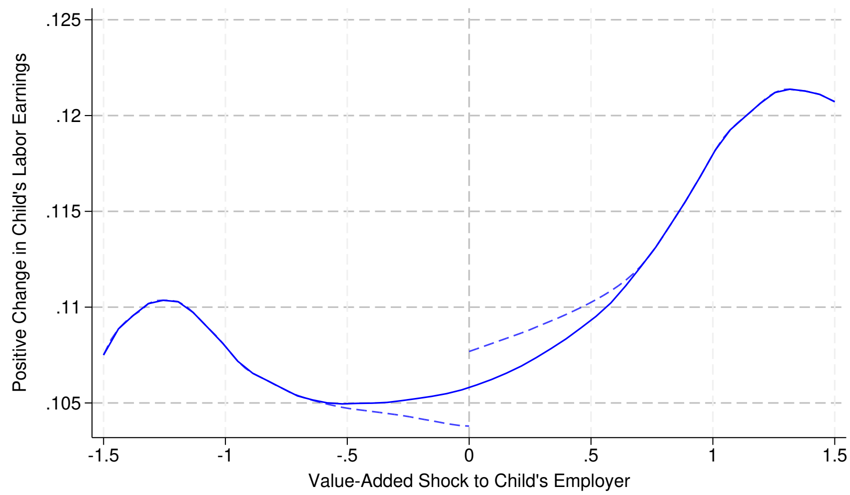

Panel B: Over-all Change in Earnings