Inter-generational Rebalancing

The 2nd SHARE Conference

Institute for Fiscal Studies, London — 14–15 May 2026

Intergenerational Altruism and Transfers of Time and Money: A Life-Cycle Perspective

Jamie Hentall-MacCuish, Uta Bolt, Eric French, and Cormac O'Dea
HEC Paris, University of Bristol, University of Cambridge, and Yale University

---

Education, earnings, and wealth persist strongly across generations, yet relatively little is known about the mechanisms through which parental background shapes children's adult outcomes. Hentall-MacCuish and his coauthors address this gap by developing and estimating a dynastic life-cycle model that jointly captures three channels of parental investment: time spent with children, educational expenditures, and direct monetary transfers in the form of inter-vivos gifts and bequests. The paper draws on the National Child Development Study, a British cohort of approximately 17,000 individuals born in a single week in 1958 and followed through their mid-fifties, allowing parental inputs to be linked directly to children's cognitive skills, educational attainment, earnings and wealth. The model is estimated by two-stage simulated method of moments, first recovering child skill production functions and then embedding them in a dynastic problem in which parents choose investments over the life cycle.

Two findings on the technology of human capital production stand out. The authors document modest complementarity between early and later parental time investments, and substantially stronger complementarity in wage production between realised skills and education, implying that the returns to skill are considerably higher for those with more education, so that investments in skills and schooling reinforce each other in the labour market. A further novel contribution concerns parental preferences. The estimates indicate that parents treat time invested in children as intermediate between work and leisure: more burdensome than pure leisure but less costly than a commensurate amount of market labour. This implies that the private return to parental time investment falls below the market wage.

The estimated model is used to assess a student loan policy that relaxes intergenerational borrowing constraints. Introduction of student loans reduces intergenerational persistence, with the largest gains to high-ability children from low-educated families who gain access to university they would not otherwise have attended. The policy has heterogeneous effects: infra-marginal children who would have attended regardless may be made worse off, since they end up borrowing to fund education their parents would have financed directly. A variant targeted at the highest-ability children improves average child welfare but does less to erode persistence, since advantaged families are better positioned to prepare their children to qualify. Participants questioned how the estimated returns to parental time investment compare to those documented in James Heckman's preschool literature, and raised the potential importance of geography and single-parent household structure, both of which Hentall-MacCuish acknowledged as natural extensions.

How Do House Prices Affect Social Mobility?

David Sturrock and Peter Levell
HM Treasury and Institute for Fiscal Studies

The ratio of average house prices to annual household earnings in the UK roughly doubled between the 1980s and the mid-2010s, while homeownership among young adults aged 30 to 36 fell from around three-quarters in 1990 to about half by the mid-2010s, with the sharpest decline among children of renters. Sturrock and Levell ask whether these trends are causally connected and, if so, through what mechanisms — a question with direct bearing on intergenerational mobility now that housing wealth has become a larger component of lifetime resources than at any time in the post-war period.

The study draws on the Longitudinal Study of Linked Censuses, the universe of housing transactions between 1995 and 2019, a newly assembled set of local-authority housing supply elasticities, and the Wealth and Assets Survey. Three cohorts, born in the 1950s, 1960s, and 1970s, are observed in childhood and again in their late twenties and thirties. To separate the causal effect of facing higher local prices from that of having wealthier parents, the authors use an instrumental variables strategy that exploits geographic variation in housing supply elasticity interacted with parental homeownership status.

The main findings are striking. A £100,000 increase in local house prices reduces the probability of homeownership by around 12 percentage points — a large affordability effect. A £100,000 increase in parental housing wealth, by contrast, has essentially no effect on whether a child becomes a homeowner. What it does affect is where children end up living and what their homes are worth: children of wealthier parents are significantly more likely to live in London and the South East, and their homes are worth roughly £10,000 to £15,000 more — an effect explained not by purchasing larger properties but by purchasing in higher-opportunity areas. The Wealth and Assets Survey confirms that wealthier parents make larger and more frequent gifts, concentrated in areas of greater house price growth.

The earnings consequences follow the same geographic logic. Among men, higher parental housing wealth raises the probability of reaching the top of the earnings distribution, with effects concentrated in business, management, media, law and cultural occupations — sectors disproportionately located in London. Effects for women run in the opposite direction, consistent with an income effect enabling earlier withdrawal from the labour force. The authors interpret these patterns through an overlapping-generations model in which moving to a high-productivity location requires a large upfront housing cost and young people face borrowing constraints; parental transfers ease the liquidity constraint and allow higher-skilled children of wealthier parents to access more productive locations. The model also implies that falling real interest rates steepen the house price gradient and amplify the role of parental wealth. Participants raised local labour market shocks as a potential confound, the interaction with public housing policy, and the role of rental markets in shaping the decision to buy. Sample selection arising from the focus on those continuously resident in England was noted as a further concern.

Parental Responses to Child Endowments

Pedro Carneiro, Neil Davies, Sebastian Espinoza Rojas, and Emma Scandolo
University College London

Whether families compensate for or reinforce differences in children's endowments is among the oldest questions in the economics of the family. Carneiro and his coauthors revisit it using an empirical strategy that exploits the random transmission of genetic material from parents to children to identify causal parental responses to child endowments. The identifying assumption follows from Mendel's laws of inheritance: conditional on parental genotype, variation in children's genotype is effectively random. The authors operationalise this using polygenic scores for educational attainment, cognition, externalising behaviour and BMI, regressing parental investment on the child's score while controlling for both parents' scores.

The data come from the Millennium Cohort Study, with genetic material collected at age 14. The sample is restricted to white families with genetic data for both biological parents, which skews the sample toward higher socioeconomic status — a limitation participants raised as a concern. Parental investment is measured through a composite index assembled from the survey's investment-related items at each wave, spanning basic care in early childhood through academic support and enrichment activities in adolescence.

The main finding is one of reinforcement: children with higher genetic propensity for educational attainment receive greater parental investment throughout childhood. The pattern that gives the study its sharper interpretation, however, is that this reinforcement is entirely absent in single-child families. Among only children, there is no detectable gradient between polygenic score and parental investment. The gradient emerges and steepens as the number of siblings increases, and — within the same child, controlled for child fixed effects — grows each time a new sibling is born. The effect is concentrated at the bottom of the endowment distribution: children with the lowest polygenic scores in multi-child families face a double disadvantage, receiving both lower genetic propensity and systematically fewer parental resources, while children in the middle are treated comparably to those at the top.

The authors consider two candidate mechanisms: parental learning about relative endowments through comparison across children, and resource scarcity that forces sharper prioritisation under tighter household budgets. An income fixed-effects analysis finds no significant change in the investment-endowment gradient with changes in household income, casting doubt on a pure resource-scarcity channel, though the authors call for a dynamic model to disentangle these mechanisms fully. Discussion raised several methodological points. Participants questioned how polygenic scores for education should be interpreted within a skill production function, noting that such scores are themselves environmentally loaded and derived from population studies conducted in different policy contexts than the MCS. The match between the Biobank reference population and the cohort sample was raised as a source of measurement error, and one participant noted recent evidence that genes transmitted through the maternal line may carry greater influence on certain traits.

The Family Multiplier: Understanding Delinquency and Parent-Adolescent Interactions

Kai Liu, Marc Chan, Sriya Iyer, and Anwen Zhang
University of Cambridge, University of Melbourne, and University of Glasgow

---

Standard models in the economics of the family treat children as passive recipients of parental investment, but for adolescents this assumption is particularly strained: teenagers make active choices about their own behaviour, and those choices feed back into how parents respond. Liu and his coauthors develop and estimate a strategic model of parent-adolescent interaction in which delinquency and parenting style are jointly determined in equilibrium. The central quantity of interest, the family multiplier, measures how much of the total effect of any environmental input on adolescent delinquency operates through parental responses, over and above the direct technological effect.

The empirical setting is the National Longitudinal Study of Adolescent to Adult Health, using the first two waves covering approximately 9,700 secondary-school students. The dataset combines school surveys, in-home interviews and parent surveys, with information on peer environments, parental behaviour and children's polygenic scores for risk tolerance, educational attainment and cognitive function. Child delinquency is a binary indicator from 13 items spanning property crime, violent crime, drug offences and status offences. Parental inputs are collapsed into binary measures of support and control.

The model treats the parent and adolescent as strategic agents in a simultaneous incomplete-information game. To separately identify the parental reaction to perceived child delinquency and the child's reaction to parental control, Liu employs two exclusion restrictions: peer delinquency among peers sharing the child's school, cohort, race and gender shifts the child's best response without entering the parental payoff, and parental community involvement and religiosity shift the parent's best response without operating through the child's payoff. The model also allows for four discrete unobserved parent-child types.

The main structural estimates show that parental control has a large and significant direct effect on adolescent delinquency, with an estimated elasticity of −0.54, while parental support has a small and insignificant effect. Crucially, parents reduce control when their child is perceived as delinquent — a "giving up" response that creates a feedback loop amplifying the initial deterioration. The family multiplier for the average parent is estimated at 1.07, but this masks substantial heterogeneity: heavily reinforcing parents generate a multiplier of 1.12, while mitigating parents yield 0.85. A structural decomposition of delinquency gaps by race, mother's education and the child's risk-tolerance polygenic score finds that roughly three-quarters of the genetic gap is attributable to the direct effect of endowment, with the remaining quarter operating through equilibrium adjustments in parent-child interactions — a finding connected to recent work on gene-environment interplay. Participants questioned whether parental control measures capture actual rather than intended influence, and whether peer effects can be separated from correlated neighbourhood shocks. Several encouraged Liu to extend the model to allow for learning and gradual adaptation over multiple periods.

Intergenerational Mobility in Complex Family Structures

Aline Bütikofer, Katrine V. Løken, and Nikodem Szewczyk
Norwegian School of Economics

Standard models of intergenerational mobility measure family background through biological or legal parentage, implicitly assuming a stable nuclear family throughout childhood. Yet in Norway 35% of children live with a single parent and 17% with a stepparent before the age of 17. Bütikofer and her coauthors ask whether the omission of stepparents biases conventional estimates of persistence and, if so, through what mechanisms stepparents come to influence children's socioeconomic trajectories. The study draws on Norwegian administrative registers covering children born between 1987 and 1993, linked to biological parents and to any stepparents who co-resided during childhood. Annual family trees are constructed for each child; permanent income is averaged over ages 30 to 36 for children and 40 to 55 for parents.

Entry into stepfamilies is not random. Biological parents in stepfamilies tend to come from lower socioeconomic backgrounds, while stepfathers who repartner earn roughly 8% to 10% more than the biological fathers they replace. To separate transmission from selection, Bütikofer exploits variation in the timing and nature of family transitions, examines outcomes determined before repartnering — including height and fifth-grade test scores — and extends the model to control for grandparents and aunts and uncles.

Standard estimates of intergenerational persistence are materially incomplete for children in stepfamilies. Accounting for stepfathers alone raises measured father-linked persistence by 28% to 31%. Once both biological and stepparents are included, total measured persistence rises by 13% to 17%. The stepfather income rank-rank slope — around 0.09 for sons, with corresponding magnitudes for daughters in education — is approximately 60% of the residual biological father slope, a substantial figure given the absence of any genetic connection. Stepmothers also matter, though their coefficients are smaller.

A key contribution is the demonstration that transmission varies systematically with relationship quality. When a stepfather has had a common child with the biological mother — a signal of commitment that predicts longer co-residence — his persistence coefficient rises by 70% to 80%. When the biological father lives in a different municipality, his coefficient falls by 13% to 19% while the stepfather's rises, consistent with the stepfather occupying the role vacated by a geographically absent father. Longer co-residence with lower-earning stepfathers produces stronger effects, pointing toward direct environmental investment rather than latent family background. Stepfather coefficients are substantially smaller for height and pre-cohabitation test scores than for adult income and education. Adding extended family members leaves stepparent coefficients largely intact, and municipality fixed effects do not materially alter the estimates.

Intergenerational Transmission of Family Influence: A Value Function Approach

James J. Heckman, César Chávez, and Rasmus Landersø
University of Chicago and the ROCKWOOL Foundation

Heckman argued that standard frameworks for measuring intergenerational mobility rest on simplifying assumptions that are not merely abstractions but active sources of distortion. The overlapping-generations model that has dominated this literature since Becker and Tomes treats the family as a single parent investing in a single child over a single period, abstracts from uncertainty and credit constraints, and invokes stationarity across cohorts. Because parents and children are born into different economic and social environments, comparing their outcomes at matched ages introduces alignment errors that bias IGE estimates in ways rarely acknowledged. A father's educational credential in 1955 carries a fundamentally different signal than the same credential in 1975.

Heckman and his coauthors propose two alternative measures of lifetime resources. The first is the expected present discounted value of future income, conditioned on the agent's information set at each period of the life cycle. The second is an approximate lifetime wealth measure derived from the household Euler equation, accounting for both expected income and the credit constraints that prevent its smooth consumption. Income realised over a lifetime is not income expected and acted upon at each stage, and it is the latter that governs investment in children. A central methodological contribution is the formal treatment of agent information sets: a sufficiently rich information set fully absorbs the predictive content of current income for future outcomes.

When lifetime measures replace snapshot proxies, two patterns emerge in opposite directions. Relative mobility is substantially lower than conventional estimates imply: the father-child IGE for expected lifetime wealth is 0.371, compared to 0.125 for wage income and 0.287 for disposable income measured at ages 30 to 35. Absolute mobility, by contrast, is substantially higher — snapshot proxies miss much of the upward movement in resources that accumulates over a working lifetime. Lifetime wealth and expected PDV are consistently stronger predictors of cognitive test scores, college attainment and years of schooling than any snapshot measures. Resources measured in ages 0 to 4 are among the most predictive.

The paper also addresses family structure and the equivalisation of resources, showing that IGE estimates vary meaningfully across nuclear, single-parent and blended families. Credit market reforms in Denmark over the sample period play a material role in movements of the IGE across cohorts. Participants questioned how the information set is identified in practice and whether the procedure is robust to additional variables. Heckman acknowledged that the test is only as good as the candidate set specified, but noted that results are stable across a wide range of specifications. The finding that strong intergenerational dependence in lifetime resources persists in Denmark — a generous welfare state with extensive redistribution — was seen as particularly significant, suggesting that the mechanisms of transmission are not easily neutralised by redistributive policy alone.

Intergenerational Divides — And What Can Be Done About Them

David Willetts
Resolution Foundation

Willetts began with demography. Cohort size, he argued, is not merely a social curiosity but an economic and political variable of the first importance: large cohorts compete more fiercely for jobs, housing and university places, and accumulate the voting weight to shape policy in their favour as they age. The baby boom generation has done precisely that. Tracking median real hourly pay by birth cohort from 1975 to 2025, each successive cohort entering the labour market since the mid-1980s has failed to replicate the wage trajectory of its predecessors. The one partial exception is at the bottom of the distribution, where the rapid rise of the National Living Wage has compressed pay differentials and closed the gap between the minimum wage and the earnings of recent graduates.

The fiscal treatment of different age groups has compounded these labour market trends. Resolution Foundation analysis shows that social security changes since 2010 have left non-pensioner households around £1,400 per year worse off by 2024–25, while pensioners are on average over £900 better off. The mechanism is largely the triple lock, which has increased the state pension far faster than either prices or earnings would have warranted, at roughly three times its initially projected cost. On wealth, the generational divide is sharpest: household wealth has risen consistently relative to national income for four decades, driven by rising house prices and the capitalisation of defined benefit pension entitlements. Baby Boomer wealth at pension age has surged compared to earlier cohorts, while younger cohorts face a housing market that has compressed their ownership rates well below those of their parents and a pension landscape dominated by defined contribution schemes that transfer longevity and investment risk onto the individual.

Willetts arranged his proposals across three domains. On earnings and skills: a stronger maintenance loan and grants offer to widen access to higher education, a reorientation of the Growth and Skills Levy toward younger workers, and — given elevated youth unemployment — a pause in the convergence of the youth minimum wage with the adult rate. On wealth: a citizen's capital grant of £10,000 to every 25-year-old, funded by a reformed inheritance tax. On pensions and housing: an end to the triple lock in favour of a smoothed earnings link, a gradual increase in the state pension age, and a combination of building more homes, reforming loan-to-income lending rules, and a targeted deposit support scheme.

Discussion ranged across the political economy of pension reform, the treatment of student debt in mortgage affordability assessments, and the potential of the Lifetime ISA. Willetts noted that UK student debt functions more like a graduate tax than a conventional loan, and that the practical question for mortgage lenders is the monthly repayment rather than the face value. On demographics more broadly, he observed that persistently low fertility in countries such as Japan, South Korea and Italy has accelerated investment in care robotics and social technology — a development that may eventually reshape the economics of elderly care in Britain, though regulatory frameworks have yet to catch up.

Spreading Wealth Across the Generations: Inheritance and Universal Basic Capital

Julian Le Grand

London School of Economics and Political Science

Le Grand opened with a distinction drawn from Michael Sherraden's Assets and the Poor: while income maintains consumption, assets change the way people think and interact with the world, inducing longer time horizons and the pursuit of longer-term goals. It is this transformation of psychology and behaviour, Le Grand argued, that makes asset ownership so consequential and so underappreciated in policy design. The empirical case rests on longitudinal evidence from the US and the UK. Individuals who own assets exhibit better health, lower mortality, greater marital stability, less domestic violence, better educational outcomes for their children, and higher savings rates when those children reach adulthood. Asset ownership at age 23 has strong documented links with full-time employment between 22 and 33, earnings at 33 for men, and health outcomes for both sexes at 33, surviving controls for income, social class and personality.

The policy vehicle Le Grand proposed is Universal Basic Capital: a grant of £10,000 to every individual at 18, funded by inheritance tax. The idea has a long pedigree — Tom Paine argued for a payment to every 21-year-old from a ground rent on land; Le Grand himself first proposed a variant in 1989; Ackerman and Alstott proposed $80,000 at 21. The most concrete British instantiation was the Child Trust Fund, introduced in 2003, which provided every child a £250 grant at birth, an additional £250 for disadvantaged children, and the option for parents to add up to £1,200 per year tax-free.

The Child Trust Fund was abolished by the coalition government after the financial crisis — a significant policy mistake, made easier because the benefits were invisible to recipients who had not yet turned 18. The American comparison is instructive: the Trump accounts announced in 2026 provide a one-off federal deposit of $1,000 for children born between January 2025 and December 2028, with families able to contribute up to $5,000 per year. The Dell Foundation has committed $6.25 billion to supplement these accounts in lower-income zip codes — a targeted overlay Le Grand found instructive.

A significant portion of the presentation addressed universality versus targeting. Le Grand argued for universal provision, citing the familiar traps of means-testing — administrative complexity, stigmatisation, and perverse incentives — while acknowledging the value of targeted top-ups for young people leaving care. On the fear that 18-year-olds would spend their grants irresponsibly, he pointed to Child Trust Fund evidence: a substantial share of recipients continued to contribute rather than withdraw immediately. He drew a clear distinction between UBC and Universal Basic Income: UBI supports ongoing consumption, while UBC operates as a springboard at a critical biographical transition. Participants asked whether shared infrastructure investments constitute collective basic capital. Le Grand accepted their value but maintained that direct asset transfers have a qualitatively different effect on individual agency that collective provision cannot replicate.

Starter Capital Accounts and Incentivised Learning

Gavin Oldham
The Share Foundation

Oldham opened by invoking Sir Keith Joseph's mission from over 50 years ago to break the cycle of deprivation — a mission that, by any honest assessment, remains unaccomplished. He cited the significant burden of student debt; 16% of Britain's 16-to-24-year-olds being unemployed; the freezing of tax thresholds imposing a stealth tax falling disproportionately on working-age earners; and the political retreat of both major parties from promoting individual ownership which has removed a policy tool which once enjoyed cross-party support.

The policy vehicle Oldham has championed for decades is a targeted starter capital account for disadvantaged young people. The original Child Trust Fund, introduced in 2002 and closed by the coalition government in 2011, was vulnerable due to its universality. Just over six million accounts were issued but, despite two-thirds having matured, many young adults remain unaware they exist. The Share Foundation has therefore proposed automatic release at twenty-one for HMRC-allocated accounts. Were this introduced now, approximately 276,000 accounts holding over £570 million would be released, of which £381 million would flow directly to low-income young adults.

Paired with the accounts for young people in care is Stepladder Plus, an incentivised learning programme in which participants earn incremental financial benefits by completing modules in financial literacy, budgeting, planning, and the Open University's Managing My Money course, with a final stage involving a mentor to secure employment, education, or training. It has demonstrated attitudinal transformation through measurable reductions in welfare dependency among adult care leavers — the group facing the sharpest transition cliff.

His proposed Child Trust Fund Mark Two differs significantly: targeted at low-income young people only, with all accounts HMRC-allocated, incentivised learning built in, and automatic release of unclaimed accounts at twenty-one. The proposal asks government only for legislation and logistical support; the accounts themselves could be philanthropically funded.

Turning to a global perspective, Oldham situated the proposal within a broader philosophy of individual ownership and disintermediation — reducing the extent to which governments make decisions affecting individuals' ability to control their own destiny. Global polarisation of wealth is significantly more acute than within the UK: the UBS Global Wealth Pyramid records that the poorest 40.7% of the world's adult population own just 0.6% of total wealth, while 1.6% own 48.1%, with young generations being particularly impacted. With OECD projecting ODA falling to 2020 levels by 2027, an international starter capital initiative could link philanthropists with structured banking support in disadvantaged countries.

Discussion focused on HMRC account mechanics, the role of local authorities in reaching care leavers, and the potential for philanthropy to substitute for government funding if the legislative framework can be secured.

Input and Output: The Opportunity Offered by Tech Wealth Participation

Heloïse Greeff
Share Alliance

Greeff began by quantifying the asymmetry that motivates her proposal. The top five US technology platforms now exceed $12 trillion in combined market capitalisation, with between 35% and 50% of that value attributable to data on intangible-asset accounting. Over the decade to 2024, those firms returned roughly 25% annually against OECD median real-wage growth of approximately 1%. The intergenerational dimension compounds this: Greeff noted that US workers aged 22–25 in AI-exposed roles saw a 16% relative fall in employment between 2022 and 2025 while older cohorts saw employment unchanged or rising. The young, in her formulation, are paying twice − supplying the data inputs to the AI economy and absorbing the displacement of its first outputs.

The four existing families of response — micropayments, data dividends, data-as-labour, and universal basic income — share a common limitation: each treats data as a commodity to be priced, taxed or shared, and none confers ownership of the productive asset it generates. Greeff's proposal, developed with Capital Economics, is to issue equity in technology companies to the users whose data and creativity build them, with the allocation tilted toward the generations most exposed to AI displacement. Contributors would continue to receive free services but would in addition accumulate equity entitling them to dividends, capital appreciation and, depending on share class, governance rights.

The Capital Economics modelling examined 10 firms across 21 countries, projecting to 2050 under three valuation methods. A fixed-income approach assigning a flat dollar amount per user proved unworkable, with cumulative allocations exceeding total market capitalisation for companies such as Alibaba and Duolingo. A revenue-share approach generated wide dispersion across company types. A market-capitalisation approach at 2% per annum emerged as the most robust, reducing inter-company disparity while keeping shareholder dilution below 30% cumulative dilution by 2035 in the base case.

The distributional finding Greeff regarded as decisive concerned age weighting. Under a base case allocating 80% of value equally and 20% as an age-weighted top-up, the 18-to-25 cohort in the US received a 20% income boost by 2050, twice that of users aged 46 and over, narrowing the intergenerational income gap by 11 percentage points. Discussion focused on the share-class question — whether non-voting shares would leave the governance ambition unrealised — and on the choice of first-mover jurisdiction. Greeff acknowledged that Meta had expressed willingness to engage, but maintained that voluntary corporate pathways would need to run in parallel with regulatory frameworks such as the EU Digital Markets Act. A further question concerned channelling digital services tax revenues toward user share purchases. Greeff welcomed this as complementary but cautioned that it would replicate the tax-and-redistribute logic her proposal was designed to supersede.

Challenging (Intergenerational) Inequalities

Paul Johnson
Institute for Fiscal Studies

By the standard measures economists have used for decades, household income inequality in the UK has not worsened since 1990 and has, on some metrics, declined since 2010. Johnson takes as his starting point the divergence between this measured stability and rising public concern about inequality. He argues that the inequalities now driving political dissatisfaction are not primarily inequalities of income but of wealth, opportunity and intergenerational prospect. Examining real median gross individual earnings by birth cohort, Johnson showed that men born in the 1980s and 1990s have failed to replicate the generational income escalator enjoyed by predecessors at the same age. Pensioner incomes, by contrast, have caught up dramatically: median pensioner income after housing costs reached approximate parity with non-pensioner income around 2010, driven by the triple lock, maturing occupational pensions and housing wealth appreciation among cohorts who entered the market before prices diverged from earnings.

The central analytical move was to insist that wealth, not income, is now the primary axis of intergenerational inequality. Although wealth inequality measured conventionally has not risen dramatically, the time required to move from the 50th to the 90th percentile of the wealth distribution has roughly doubled. Inheritance has become a far larger component of lifetime resources, and work by Levell and Sturrock has demonstrated causally that children of parents who benefited from regional house price booms hold more property wealth, are more likely to work in London, and to enter the highest-paying occupations. Homeownership among 25-to-34-year-olds has fallen sharply, with the decline concentrated in the middle-income quintile — a cohort that in earlier generations would have expected to own.

Johnson then turned to policy changes that have tilted the intergenerational balance. Private-sector defined benefit pension coverage collapsed from over 40% of workers to under 10% over 30 years, with average employer contributions to defined contribution successors less than a third of their predecessors. Public spending on health, equal to education spending in 1985, exceeded twice education spending by 2025. Minimum income support above state pension age rose dramatically relative to support below it. The tax architecture compounds the asymmetry: a 27-year-old graduate faces a marginal rate of 37% including student loan repayments, against 20% for a 67-year-old on the same income.

Johnson resisted two tempting but inadequate responses. Income redistribution only partially protects against the harms he identified, because wealth and opportunity inequalities are structurally distinct and cannot be dissolved by tax-and-transfer alone. Nor can inequality be treated as separable from macro-economic management: the lack of sustained growth over the past 20 years is itself a primary driver, eroding the political space for redistributive ambition and concentrating the gains of growth among existing asset-holders. One questioner pressed on the relationship between wealth inequality and stock-market dependence. Johnson acknowledged the difficulty but returned to housing as the more tractable political terrain.

Rebalancing in the Digital Economy

Matthew Agarwala
University of Sussex

Agarwala opened by widening the frame of the conference's question. Rebalancing, he argued, is a more demanding concept than redistribution: where the latter moves resources after the fact, the former asks what is being rebalanced and through what mechanisms. He emphasised four dimensions — income and wealth, opportunity, outcomes and power — with particular attention to the last. The middle 50% of British earners now pay a lower share of their income in tax than at any point since the Second World War, which Agarwala read not as evidence of generosity but as a symptom of declining trust in what the state returns for its levies.

The affirmative case turned on digital public infrastructure — shared digital systems into which any actor, public or private, can plug. Agarwala took India's Unified Payments Interface as a proof of concept: integrating digital identity, bank account and mobile number into an interoperable system enabling instantaneous, biometrically secured transfers. UPI had reached 504 million users and processed $354 billion in monthly transactions. The Indian government transferred over $588 billion directly to beneficiaries' bank accounts, saving more than $45 billion by eliminating ghost recipients and curtailing speed money.

The second movement concerned the fiscal threat posed by AI. The effective tax rate on labour in the UK, at around 35% once income tax, national insurance and payroll levies are aggregated, is nearly double the rate on capital and algorithmic returns at approximately 19%. As AI substitutes for labour, the tax arithmetic deteriorates mechanically: Agarwala estimated that a 5-percentage-point shift from labour to capital would reduce total tax revenues by approximately 1%. This compounds with expenditure pressures from ageing populations, climate change and rising demand for redistribution. He characterised the underlying problem as constant-returns-to-scale government operating in an increasing-returns-to-scale economy.

Agarwala identified three asymmetries through which AI compounds this problem: informational, since AI systems are opaque to regulators; geographical, since profits concentrate in low-tax jurisdictions; and one of power, since a handful of firms control critical AI infrastructure. Against each, he assessed the available tax handles. Compute ranked highly as a tax handle: the semiconductor hardware on which AI runs has concentrated supply chains, physically observable inputs, and is hard to arbitrage at the point of production. Taxation of the energy and water use of data centres offered a Pigouvian rationale alongside fiscal yield. Data flows across borders were a third possible handle, though valuation remains contested. Agarwala pointed to the OECD's Pillar Two agreement as evidence that internationally coordinated minimum tax floors are achievable. Discussion pressed on the geopolitical complications of chip fabrication's concentration in Taiwan and the US. Agarwala argued that data-centre externalities provided a more tractable entry point. A participant asked whether UPI was transferable to advanced economies; Agarwala acknowledged that India's circumstances were not straightforwardly replicated but maintained that the lesson about state capacity as a precondition for rebalancing was general.

Linking the Conference: From Evidence to Rebalancing

Across two days, a common diagnosis emerged: the channels through which advantage now passes between generations are no longer well captured by the income measures that have organised post-war policy thinking. Hentall-MacCuish presented how parental time, money and educational investment compound across the life cycle in ways that conventional snapshots obscure. Sturrock showed that what the housing boom transmitted was less ownership itself than access to the geographies where opportunity is concentrated. Carneiro revealed that reinforcement of endowments emerges only when siblings compete for parental investment. Liu quantified the equilibrium feedback through which adolescent behaviour and parental responses jointly determine outcomes. Bütikofer showed that stepparents transmit roughly 60% of the residual biological-father slope, indicating that conventional persistence measures materially understate the family environment's reach. Heckman went furthest in reconstructing the measurement framework, showing that lifetime wealth measures yield IGE estimates roughly three times those obtained from wage snapshots.

The second day moved from diagnosis to instrument. Willetts traced how cohort wage stagnation, the triple lock, and the collapse of defined benefit pensions have produced a generational divide visible across earnings, taxes and wealth, proposing capital grants, deposit support and reform of the triple lock. Le Grand and Oldham developed parallel cases for capital ownership as the structural complement to income support: Le Grand for a universal £10,000 grant at 18 funded by inheritance tax, Oldham for a targeted starter capital account paired with incentivised learning. Greeff extended this ownership logic into the digital economy with Stock for Data and Creativity. Johnson argued that wealth, opportunity and intergenerational prospect have displaced income as the axes of consequential inequality. Agarwala closed the programme by examining how digital public infrastructure can expand state capacity while AI threatens to erode the tax base on which any rebalancing depends.

Taken together, the twelve talks map a coherent agenda. The empirical evidence shows that intergenerational transmission runs through time, place, wealth, family structure and information sets in ways conventional measures systematically under-record. The policy proposals share a structural intuition: where income transfers protect against poverty but do not alter the distribution of life chances, ownership of productive assets — housing, equity, human capital — can. The Share Alliance framing, articulated by Oldham, identifies starter capital accounts financed from inheritance tax as the central mechanism through which the cycle of deprivation can be broken without recourse to expanded current public spending. The wrap-up returned to Heckman's life-cycle insight: the highest returns to investment lie in early childhood, but a coherent rebalancing portfolio must address the full span from birth to retirement, and must reckon with the ambiguous effects of digital technology on the cohorts now entering adulthood. The conference closed with the recognition that the analytical foundations are in place, and that the work ahead is one of implementation — of building the political coalitions, fiscal instruments and digital infrastructure capable of carrying the diagnosis into policy.