Work-From-Home and Intergenerational Mobility

Francesco Saverio Lenzi

Goethe University Frankfurt and European University Institute

May 17, 2026

Click here for the latest version.

Abstract

Work-from-home (WFH) technology, which has become a lasting feature of the post-pandemic economy, remains disproportionately available to high-income and college-educated workers. This paper studies how this unequal access affects intergenerational mobility through two novel mechanisms: a parental time channel and a residential sorting channel. Using U.S. microdata, I document that WFH parents devote about 23% more time, equivalent to roughly 0.15 additional hours per week, to their children's educational activities than comparable non-WFH parents. I also show that WFH families are 3.7 percentage points more likely to move to neighborhoods with higher-quality public schools, as measured by standardized test scores. To quantify the implications of these mechanisms, I develop and calibrate a spatial overlapping-generations model. The model incorporates the parental time channel through endogenous time allocation between market work, leisure, and child education, and the residential sorting channel through endogenous location choices with neighborhood spillovers. The model calibrated to U.S. data reveals a persistent intergenerational advantage: even among adults who do not themselves work from home, those with a WFH parent are 1.05 percentage points more likely to reach the top income decile, corresponding to a 12% relative increase compared to their peers. Counterfactual experiments show that more than 95% of this intergenerational advantage arises from the parental time channel rather than residential sorting. These findings suggest that WFH transmits economic advantage to children primarily through how parents invest their time at home, rather than through the neighborhood-based spillovers secured by residential sorting.

Keywords: Work-from-Home, Intergenerational Mobility, Inequality, Human capital, Child Education, Spatial Sorting

JEL Codes: C11,E00,I00,J13

Email: francescosaverio.lenzi@eui.eu

I am very grateful and deeply indebted for the insightful conversations and helpful comments received by Nicholas Bloom, Jesus Bueren, Lidia Cruces, Steven J. Davis, Georg Duernecker, Leo Kaas, Alexander Ludwig, Alexander Monge-Naranjo, Virginia Sanchez-Marcos, and Siyu Shi, and all the participants to the Macroeconomics Working Group at the EUI and the participants of the FQMG seminar in Goethe University. Financial support from the German Research Foundation grant no. 462655750 is gratefully acknowledged.

1 Introduction

The work-from-home (WFH) revolution represents one of the most significant labor market transformations in recent decades. What began as an emergency pandemic measure has solidified into a persistent feature of the modern economy. As of 2023, remote work accounts for over 28% of all full workdays in the United States (Barrero et al., 2023). A well-documented characteristic of this new regime is its unequal access: telework availability rises steeply with education and income, meaning high-wage, college-educated workers are significantly more likely to hold WFH-eligible jobs (Yasenov, 2020). These groups have historically been the most successful in transmitting economic advantage to their children (Chetty et al., 2014). This confluence raises a critical question: if the WFH technology is predominantly available to the advantaged, what are the consequences for intergenerational mobility?

This paper investigates the consequences of WFH for intergenerational mobility by identifying and quantifying two novel channels through which a parent's remote work could influence their child's long-term economic outcomes. The first is a parental time channel: by saving time previously spent commuting, WFH may enable parents to reallocate effort toward their children's education. The second is a residential sorting channel: the geographic flexibility of remote work may allow families to relocate more easily to neighborhoods with higher-quality schools.

I begin by providing direct empirical evidence for both channels using U.S. microdata. First, leveraging the American Time Use Survey (ATUS), I document that WFH is associated with a statistically significant increase of approximately 0.146 hours per week in time devoted to children's educational activities, even after controlling for a rich set of demographic and socioeconomic characteristics. In relative terms, this represents a 23% increase from the baseline for non-WFH parents. This supports the notion that saved commute time is reallocated to caregiving (Aksoy et al., 2023; Pabilonia and Vernon, 2023). Second, analyzing mobility patterns in the American Community Survey (ACS), I show that after Covid-19, WFH households with school-aged children are systematically more likely to move into neighborhoods with higher-quality public schools, as measured by standardized test scores data. I estimate a probit model that relates WFH status to the probability of moving to a better school district, while controlling for key household characteristics. The results indicate that, all else equal, WFH households are associated with a 3.7 percentage point higher probability of relocating toward a better school location compared to non-WFH households. Together, these results motivate a structural framework that integrates these empirical regularities into a theory of intergenerational transmission under remote work.

To assess the quantitative importance of these channels for intergenerational mobility, I develop and calibrate a spatial overlapping-generations model with endogenous neighborhood spillovers. Households differ in human capital and in their eligibility to work from home, and choose where to live (between a central business district and a suburb), how much to work, and how to allocate time between market work, home-based work, and child education. For WFH-eligible agents, the ability to divide labor between the office and home alters both the time available for child investment and the incentive to live farther from the CBD. Neighborhood spillovers arise endogenously from the local average of child human capital, which shapes educational productivity and feeds back into parental location choices. The model calibrated to U.S. data reveals that the influence of WFH extends to the next generation. I identify a significant intergenerational premium: even among adults who do not themselves work from home, those who had a WFH parent are 1.05 percentage points more likely to reach the top income decile, a 12% relative increase compared to their peers without WFH parents. My model also accounts for a strong *within-generation* effect: adults in WFH occupations are 8.4 percentage points more likely to reach the top income decile.

Part of my contribution lies in disentangling the sources of this intergenerational transmission. Through counterfactual experiments, I find that the parental time channel is the dominant mechanism. When I computationally shut down the ability of WFH parents to increase educational time, the intergenerational advantage collapses by over 95%. In contrast, neutralizing neighborhood spillovers has a negligible effect. This provides evidence that the intergenerational transmission of WFH advantage operates primarily through increased parental time investments. The WFH premium is transmitted not by where parents live, but by what they do with their time at home.

Related Literature. This paper connects three strands of research. First, I contribute to the literature on intergenerational mobility and spatial inequality with neighborhood externalities by introducing the work-from-home margin into a human-capital framework with endogenous location choice. The geography of childhood is a powerful determinant of adult outcomes (Chetty et al., 2014; Chetty and Hendren, 2018), a fact rationalized by models where residential sorting interacts with local spillovers. Following the seminal contributions of Benabou (1994, 1996), who formalize how peer effects and local amenities can make inequality persistent through sorting, and the early microfoundations of Durlauf et al. (1993), which highlight non-linear neighborhood dynamics, subsequent quantitative work has explored how spatial structure affects mobility. For instance, Ferreira et al. (2016) show in a Brazil-calibrated spatial model that housing costs and segmented school access endogenously shape children's skill formation and urban sorting. I build on these insights by introducing

work-from-home as a time-augmenting technology that changes parents' effective education time and mobility incentives, allowing me to decompose the time and sorting channels in U.S. data. Closest to my framework are Fogli et al. (2023) and Zheng and Graham (2022). Fogli et al. (2023) develop a quantitative overlapping-generations spatial model where household sorting, education investments and local spillovers jointly determine city-level segregation and inequality. I extend this approach by endogenizing the effective supply of parental educational time, linking labor-market technology to intergenerational transmission. Zheng and Graham (2022) instead emphasize how public-school quality and local house prices interact through property-tax finance, shaping access to education and spatial inequality, while Chyn and Daruich (2022) embed parental investments and neighborhood spillovers in an intergenerational general-equilibrium framework disciplined by exposure-based evidence. Relative to this body of work, I conceptualize the large-scale adoption of work-from-home not merely as a change in work location, but as a technology shock that re-prices parental time and loosens the geography of opportunity. By reducing the opportunity cost of time and weakening the tether of the daily commute, WFH reshapes both how advantage is produced and how it is transmitted across generations.

Second, I contribute to the growing literature on remote work and its spatial consequences. Empirical evidence from large-scale surveys documents that work-from-home (WFH) has now stabilized as a permanent feature of the post-pandemic labor market, accounting for roughly one-quarter of workdays, substantially reducing commuting, and reallocating saved time toward caregiving and home production (Barrero et al., 2021, 2023; Aksoy et al., 2022, 2023; Pabilonia and Vernon, 2023; von Gaudecker et al., 2024). Building on this evidence, macro-spatial models analyze how telework reshapes urban form, housing demand, and sorting-flattening rent gradients, reallocating population toward suburbs, and reorganizing economic activity across cities when adoption is heterogeneous across sectors (Davis et al., 2024; Delventhal and Parkhomenko, 2024; Gokan et al., 2022; Richard, 2024). These findings are consistent with the emerging *Donut Effect*, whereby remote workers relocate from central business districts toward suburban areas (Ramani and Bloom, 2021; Liu and Su, 2021). In my model, this pattern arises endogenously in partial equilibrium: WFH households are more likely to choose to reside in high-amenity neighborhoods with stronger local spillovers, as the reduced cost of distance effectively re-prices commuting. My contribution is to move beyond these contemporaneous effects and place intergenerational mobility at the center of the analysis. Rather than treating WFH as a purely labor-market or housing-market shock, I embed it within a human-capital technology featuring neighborhood spillovers, where the redistribution of both time and place induced by telework directly shapes the transmission of economic advantage across generations.

Third, I relate to the extensive literature on parental time investments and skill formation. Time-use studies document that highly educated parents devote significantly more time to cognitively stimulating childcare and that these gaps have widened over time across socioeconomic groups (Ramey and Ramey, 2009; Dotti Sani and Treas, 2016). Theoretical and empirical research on human-capital formation emphasizes dynamic complementarity in the production of skills, where early parental inputs enhance the productivity of later ones (Cunha and Heckman, 2007). Quantitative models and structural estimations link disparities in early family environments to persistent differences in adult outcomes and the effectiveness of early interventions (Caucutt and Lochner, 2020; Attanasio et al., 2020; Elango et al., 2015). I build on these insights by providing new empirical evidence on how work-from-home reshapes parental behavior. Using microdata from the American Time Use Survey (ATUS), I show that WFH parents are associated with significantly more time dedicated to educational activities with their children, consistent with a reallocation of saved commuting time toward childcare. Complementing this with evidence from the American Community Survey (ACS), I document that WFH households with school-aged children are more likely to relocate to neighborhoods with higher-quality public schools.

The remainder of the paper is organized as follows. Section 2 presents the empirical analysis based on U.S. microdata. Section 2.1 uses the American Time Use Survey (ATUS) to examine how WFH affects parental time devoted to children's education, while Section 2.2 employs the American Community Survey (ACS) to document residential moves of WFH households toward neighborhoods with higher-quality public schools. Section 3 develops a spatial overlapping-generations model that rationalizes these empirical mechanisms. Section 4 calibrates the model to pre- and post-COVID moments and conducts counterfactual experiments to isolate the channels through which WFH affects intergenerational mobility. Finally, Section 5 concludes with policy implications.

2 Empirical Evidence: Patterns in Parental Behavior and School Sorting

This section documents key descriptive relationships between work-from-home (WFH) availability, parental investment in children, and residential sorting. The objective is to identify systematic behavioral patterns that inform the mechanisms of the structural model developed in subsequent sections. I focus on two potential channels through which WFH could influence intergenerational mobility: the reallocation of parental time toward children's education and residential relocation toward neighborhoods with higher-quality schools.

It is important to note that the direction of these relationships is not obvious *ex-ante*. Regarding time use, WFH could increase educational investment by freeing up time previously spent commuting, or decrease it if the blurring of work-life boundaries leads to work crowding out family time. Similarly, the ability to live further from the workplace might lead families to prioritize school quality when moving, or it might instead lead them to prioritize other amenities like natural surroundings. The following analysis seeks to test the empirical prevalence of these competing behavioral responses and quantify the associations that motivate the model's structure.

2.1 Work-from-Home and Parental Education Time

I study how work-from-home (WFH) relates to parents' education-specific time with children using the American Time Use Survey (ATUS). The American Time Use Survey (ATUS), accessed via IPUMS (Hofferth et al., 2020) provides a record of activities in the prior 24 hours, including duration, location, and who was present. This granular, time-diary methodology offers two key advantages for my study. First, it provides an objective, real-time measure of time investment, moving beyond surveys that rely on subjective attitudinal questions or retrospective summaries of behavioral frequency.[1] Second, it enables me to isolate a precise metric of direct parental input into children's human capital formation.[2]

Sample and variables. I restrict the American Time Use Survey (ATUS) microdata to partnered, employed parents living with at least one child aged 5–15. Following Restrepo and Zeballos (2022) I exclude the year 2020 from this analysis to avoid collecting unrealistic values due to the pandemic policies and their direct influence in time use decisions; however, the main results are not influenced by this restriction. Diaries on weekends/holidays and censored work spells are excluded in the main sample[3]. I weight by ATUS person weights and all estimates are survey design-based[4]. The outcome is weekly hours[5] of education time. Following Del Boca et al. (2014), on each diary day I aggregate minutes spent in

[1] For instance, Puccioni (2015) and other studies using surveys like the ECLS-K measure parental involvement through questions about the importance of school readiness or summaries of behavioral frequency.

[2] While other studies have investigated the association between WFH and broad childcare activities (Aksoy et al., 2023; Restrepo and Zeballos, 2022; Pabilonia and Vernon, 2023; von Gaudecker et al., 2024), this paper contributes by being the first to focus specifically on education-related time investment.

[3] Diaries are excluded if the respondent was engaged in a work activity at the very start (4:00 AM) or end (4:00 AM the following day) of the diary period. Such censored work spells make it impossible to accurately measure the total duration of the work episode and would confound the analysis.

[4] In line with the complex survey design of the ATUS, all estimates use the ATUS person-level sampling weights. This ensures that the results are nationally representative of the U.S. population for the surveyed time periods.

[5] Weekly measures are computed by multiplying weekday diary by five, reflecting a standard five-day workweek. Weekend diaries are excluded from the main sample; Appendix Table 21 shows robustness to including weekend diaries.

activities that represent direct investments in children's human capital and convert to a weekly equivalent for readability. The rationale for this aggregation is to isolate time inputs that are most directly linked to the production of child human capital, as in the theory of skill formation (Cunha and Heckman, 2007, 2008). I include two core categories of activities:

1. **Direct academic support** (`030200` "Activities related to household children's education"), comprising: helping with homework (`030201`), attending school meetings or conferences (`030202`), home schooling (`030203`), and waiting associated with education (`030204`). I treat waiting as an investment input because it is a parental time cost that enables the child's educational activity.

2. **Cognitively enriching interactions** within the childcare major group `0301` ("Caring for and helping household children"): reading to or with the child (`030102`) and engaging in arts and crafts (`030104`).[6] These activities are commonly used as proximate inputs into cognitive and language development.

This combined measure captures a more comprehensive view of educational investment than homework help alone, as it includes both curriculum-based support and broader skill development, while excluding general childcare (e.g. feeding, bathing), informal play and passive supervision.

I define work-from-home (WFH) as a binary indicator for performing a substantial amount of work at home on the diary day, specifically at least two hours or 25% of a standard full-time workday. This threshold serves two purposes: first, it excludes primarily office-based workers performing minor tasks at home; second, it aligns with the intensive margin in my theoretical model where agents in teleworkable occupations allocate hours between office and home. Consequently, defining WFH as only fully remote work, as in Pabilonia and Vernon 2023, would be inappropriate for my framework, as it would represent a corner solution rather than the allocative decision I model. Because WFH is measured at the diary day, I interpret estimates as associations, not causal effects, since parents may adjust WFH on days with child-related needs.

The model controls for a comprehensive set of covariates. These include parent demographics (sex, age, race, educational attainment), household composition (number of

[6] The corrected replication code excludes activity `030107` ("Providing medical care and health services to household children"), which an earlier version of the code inadvertently included. Activity `030107` captures physical health care—not educational investment—and WFH parents are mechanically more likely to provide it since they are present at home, introducing upward bias in the measured association. Excluding it yields a cleaner proxy for education-specific investment. The WFH coefficient in the corrected code is 0.153 versus 0.146 with the erroneous inclusion; the reduction is consistent with removing a WFH-correlated confound from the outcome. See the replication package CHANGES_LOG.MD for a full discussion.

children, developmental stage of the youngest child), and socioeconomic factors (household income, metropolitan status). I also account for the spouse's characteristics (employment status, age, usual working hours) and include fixed effects for the respondent's industry, state of residence, and survey year.

Descriptive Statistics Table 1 presents survey-weighted summary statistics comparing parents who work from home (WFH) with those who do not. The sample consists of 11,045 partnered, employed parents living with children aged 5–15, with 14.6% (1,608 parents) meeting the criteria for WFH. The patterns observed are consistent with the broader literature on telework availability (Yasenov, 2020; Dingel and Neiman, 2020). WFH parents are disproportionately highly educated, 33.5% hold advanced degrees versus 21.5% of non-WFH parents, and are more likely to work in management occupations (27.4% versus 14.4%). They also have substantially higher weekly earnings ($1,739 versus $1,201) and work longer hours. In terms of time allocation, WFH parents report more time spent on both general childcare and education-specific activities. These systematic differences across observable characteristics highlight the importance of controlling for a comprehensive set of covariates in the subsequent multivariate analysis.

Table 1: Summary Statistics and Balance: WFH vs. Non-WFH Parents

	WFH		Non-WFH		Difference	
Variable	Mean	(SE)	Mean	(SE)	Difference	t-stat
Time Use (hrs/wk)						
Education Time	0.83	(0.06)	0.63	(0.02)	0.19***	3.31
Childcare Time	7.43	(0.22)	6.16	(0.10)	1.27***	5.21
Work Time	42.92	(0.39)	36.03	(0.22)	6.89***	15.46
Demographics						
Parent Age	41.19	(0.24)	39.61	(0.10)	1.58***	6.09
% Female	40.90	(1.40)	35.90	(0.60)	5.10***	3.34
Number of Children	1.95	(0.03)	1.96	(0.01)	-0.02	-0.6
Youngest Child's Age	6.34	(0.14)	6.23	(0.06)	0.12	0.80
Socioeconomic Status						
Weekly Earnings ($)	1738.75	(27.07)	1201.15	(8.81)	537.60***	18.89
% Advanced Degree	33.50	–	21.50	–	12.10***	8.75
% Bachelor's Degree	50.41	–	42.90	–	7.50***	4.82
% in Metro Area	20.78	(2.50)	22.11	(1.20)	-1.32***	-4.75
% in Management Occ.	27.40	(1.30)	14.40	(0.40)	13.00***	9.71
Unweighted Observations						
N	1608		9437			

Notes: Statistics are calculated using ATUS sampling weights to ensure national representativeness. The sample includes employed, partnered parents living with at least one child aged 5–15. Work-from-home (WFH) is defined as working at least 2 hours or 25% of a standard workday at home. Time variables are calculated as weekly equivalents from a single diary day. Differences are from linear regressions of the variable on a WFH indicator. Standard errors for proportion differences are calculated via linear probability models. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1 plots weighted means of weekly education time with 95% CIs for key subgroups.

Panel A shows a clear WFH premium of roughly one-fourth of an hour per week. Panel B reveals the familiar education gradient: parents with more schooling devote more time to children's education, regardless of work location. WFH shifts the entire distribution up: high-school, some college/BA, and advanced-degree households all show higher means when the diary includes home-based work. Panel C highlights the gender structure. Mothers devote substantially more educational time than fathers on average, and WFH raises time for both groups. The WFH–non-WFH gap is present for fathers and somewhat larger for mothers, suggesting that remote work increases involvement on both margins without eliminating the underlying gender differential. Panel D organizes means by the child's developmental stage. Educational time peaks in middle childhood (6–12), is lower for preschoolers, and falls sharply in adolescence—consistent with changing needs and autonomy. At each stage, WFH is associated with higher parental education time; the gap is largest for early and middle childhood and remains positive, though smaller, for adolescents. Overall, the figure documents a robust, subgroup-wide WFH premium in education time.

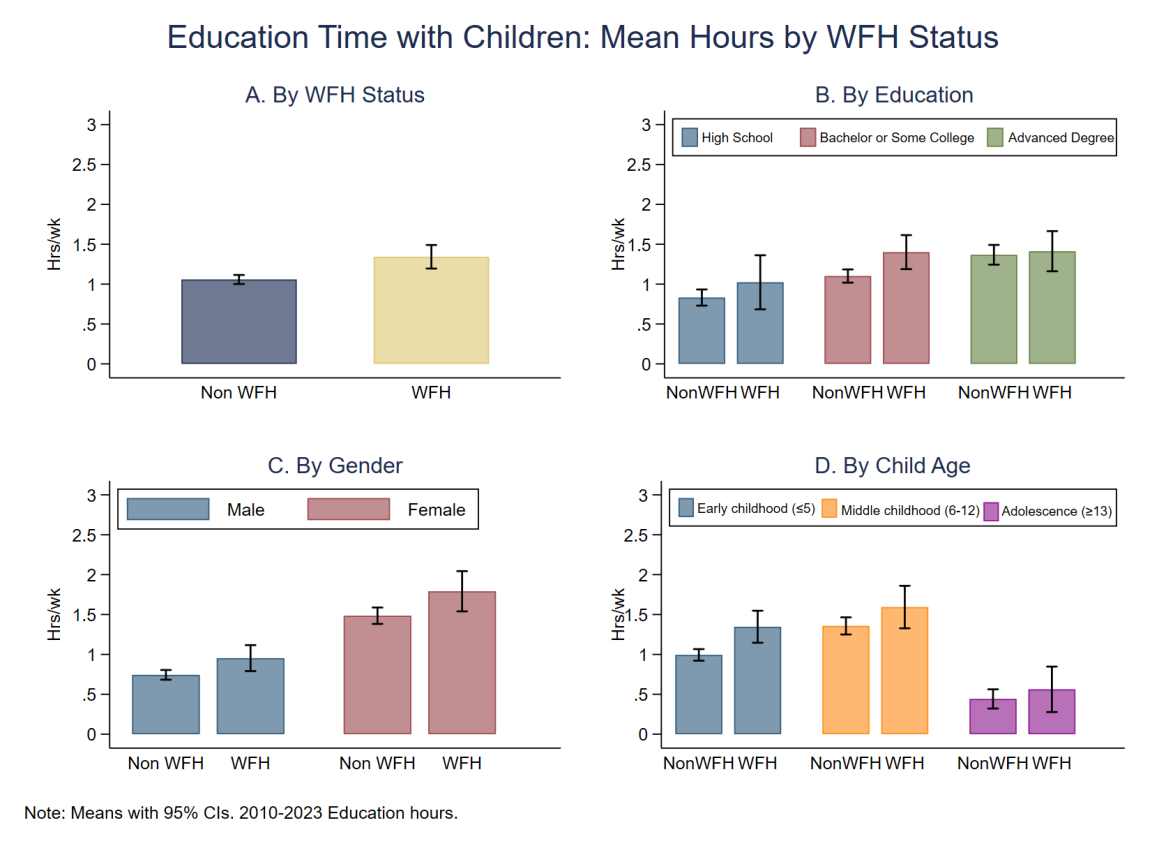


Figure 1: Education time with children (hours/week): means and 95% CIs by WFH status, education, gender, and child's age. ATUS 2010–2023 (excl. 2020).

Main Analysis To estimate the association between WFH and parental investments in children's education, I employ an OLS regression framework that controls for a comprehensive

set of observable characteristics. The baseline specification is:

$$\text{EduTime}_i = \beta_0 + \beta_1 \text{WFH}_i + \gamma' X_i + \delta_t + \lambda_s + u_i, \tag{1}$$

where EduTime_i represents weekly hours of education-specific time; WFH_i indicates whether parent i worked at least two hours or 25% of a standard workday at home on the diary day; X_i includes parent demographics (sex, education, age, race), household characteristics (number of children, youngest child's developmental stage, income, metropolitan status); and δ_t and λ_s represent year and state fixed effects. The specification also includes industry and occupation fixed effects to account for differences in job characteristics and teleworkability.[7]

Table 2 presents the estimated association between work-from-home and parental education time. Column (1) includes the full set of demographic, household, and fixed-effect controls; column (2) additionally conditions on spouse's employment status, age, and usual hours, allowing me to partial out household-level specialization patterns. The WFH coefficient is stable across both columns: 0.152 (SE = 0.063) without spousal controls and 0.153 (SE = 0.063) with them, both significant at the 5% level. Conditioning on spousal characteristics does not attenuate the WFH premium, confirming that the result is not driven by assortative mating patterns that place high-earner couples into WFH-eligible households.

The preferred estimate (column 2) implies that WFH is associated with approximately 0.153 additional hours of weekly education time, representing a 24% increase from the non-WFH baseline mean of 0.63 hours. To contextualize this effect size, we can compare it to the impact of maternal time investments found in developmental studies. For instance, Del Bono et al. (2016), using UK longitudinal data, found that a one standard deviation increase in maternal educational time raised child cognitive scores by 0.13 to 0.07 standard deviations, depending on the child's age and model specification. My estimated 24% increase in time investment due to WFH is therefore substantial, suggesting that the reorganization of work may represent a meaningful margin for increasing parental educational inputs.

The control variables reveal intuitive patterns that validate the empirical approach. Parents with bachelor's or higher degrees spend significantly more time on educational activities than those with only a high school education, consistent with human capital models of intergenerational transmission Ramey and Ramey (2009); Dotti Sani and Treas (2016). Female parents invest substantially more time than males, aligning with gender patterns in childcare specialization. The child age coefficients show a clear developmental pattern:

[7] These estimates are descriptive, as I cannot address potential self-selection into WFH occupations or unobserved factors that jointly influence both WFH arrangements and parental time allocation.

education time peaks during middle childhood (6–12 years) when academic demands increase, then declines sharply during adolescence as children gain autonomy. Each additional child is associated with increased total education time, suggesting scale effects in household educational investments.

Table 2: Work-from-Home and Parental Education Time

	(1) Weekly Education Hours	(2) Weekly Education Hours
WFH ($\geq$ 2h)	0.152**	0.153**
	(0.063)	(0.063)
Education (ref: High School)		
Bachelor/Some College	0.146***	0.120**
	(0.054)	(0.057)
Advanced Degree	0.129*	0.101
	(0.077)	(0.079)
Demographics		
Female	0.263***	0.234***
	(0.052)	(0.053)
Child Age (ref: Early $\leq$5)		
Middle (6–12)	0.252***	0.235***
	(0.047)	(0.047)
Teen (13–15)	−0.257***	−0.279***
	(0.058)	(0.060)
Household Structure		
Number of children under 18	0.150***	0.153***
	(0.025)	(0.026)
Observations	11,045	11,045
Included Controls:		
Parent age, race, income, metro status	✓	✓
Spousal characteristics (employment, age)		✓
Industry, occupation, year, and state FE	✓	✓

Notes: The dependent variable is weekly hours spent on children's educational activities (ATUS activity codes 030200–030204, 030102, 030104). WFH is defined as working at least 2 hours or 25% of a standard workday at home. All estimates use ATUS sampling weights and account for complex survey design (Taylor-linearized standard errors). Column (1) omits spouse controls; column (2) is the preferred specification. Standard errors in parentheses. *** p $<$ 0.01, ** p $<$ 0.05, * p $<$ 0.1.

The results suggest that WFH arrangements are systematically associated with greater parental involvement in children's education, even after accounting for socioeconomic status, family structure, spousal characteristics, and geographic and temporal factors. To examine whether the WFH–education link varies across population groups, I estimate interaction models by parent education, child age, and gender. The results presented in Appendix A show that the positive association between WFH and education time is concentrated among parents of school-age children (ages 6–12), when homework and structured learning activities are most relevant. The gradient by parental education is modest: the association is slightly stronger for parents with some college or a bachelor's degree but not statistically different across groups. Gender differences are small and imprecisely estimated, with both fathers and mothers showing positive point estimates.

Robustness: Alternative WFH Intensity Measures. The binary WFH indicator (at least two hours of home-based work) collapses a continuous allocation decision into a single threshold. To assess whether the main result is sensitive to this choice, I estimate three alternative specifications: a continuous measure of daily WFH hours (`wfh_hours`), the share of total work time performed at home (`wfh_share`), and a set of binary dummies at alternative thresholds (1h, 2h, 3h, 4h). Results are reported in Appendix Tables 18 and 17.

The continuous WFH hours specification implies that each additional hour of daily home-based work is associated with 0.018 additional hours per week of children's education time (SE = 0.008, $p < 0.05$). The WFH share specification—which captures the fraction of the workday performed at home on a 0–1 scale—yields a coefficient of 0.265 (SE = 0.071, $p < 0.01$), implying that a parent who shifts from fully office-based to fully home-based work would invest an additional quarter-hour per week in educational activities. The threshold robustness analysis shows coefficients of 0.192*** (SE = 0.057) at the 1-hour threshold, 0.153** (SE = 0.063) at 2 hours (the main estimate), 0.188*** (SE = 0.070) at 3 hours, and 0.197*** (SE = 0.074) at 4 hours. The consistency across thresholds confirms that the result does not depend on the specific binary cutoff.

To investigate whether the dosage–response relationship is monotone, I estimate a categorical model using four WFH intensity bins: Light (<2h), Moderate (2–4h), Substantial (4–8h), and Full-day (≥8h), with no WFH as the reference category (Appendix Table 19). The results reveal a non-monotone pattern: Substantial WFH (4–8h per day) is associated with the largest gain in education time (0.345***, SE = 0.116), followed by Light WFH (0.178**, SE = 0.071). The Moderate and Full-day categories have positive but not statistically significant coefficients. This non-monotonicity may reflect that moderate WFH provides flexibility without reducing office interactions that also support career productivity, while very long home-based days (Full-day) may crowd out marginal education activities through work pressure. The key message is that the positive WFH–education association is not confined to full-day remote work and is present at multiple intensity levels.

Weekend robustness: restricting to weekday diaries (the main sample) and expanding to include weekend diaries yield WFH coefficients of 0.153** and 0.141*, respectively (Appendix Table 21). The slight attenuation in the all-day sample reflects that WFH on weekends is less common and may capture different behavioral patterns than weekday WFH; the positive association nevertheless persists.

Extensive Margin and Intensive Margin. A natural question is whether the observed association between work-from-home (WFH) and parental education time operates through

an extensive margin (whether parents engage in any education-related activity at all) or an intensive margin—how much time they devote when such activities occur. In fact, time-diary data feature a large mass at zero, but these zeros typically reflect sampling zeros (the activity simply did not occur on the recorded day) rather than structural zeros (the activity never occurs). As emphasized by Brown and Dunn (2011), conventional models such as OLS or Tobit impose restrictive assumptions that are often inconsistent with the empirical distribution of time-use data. In particular, Tobit models treat all zeros as censored realizations from a latent normal variable, a characterization that rarely fits diary-based measures where zeroes arise from non-participation on that day rather than from censoring.

To capture both participation and intensity, I adopt a two-part modeling strategy. First, I estimate a logit model for the probability that a parent reports any positive education time,

$$\Pr(\text{EduTime}_i > 0 \mid X_i) = \Lambda\big(\alpha_0 + \alpha_1 \text{WFH}_i + \theta' X_i\big), \tag{2}$$

and report the average marginal effect (AME) of WFH in percentage points. The logit specification is flexible and requires no assumption of normality for the latent decision process[8].

For the intensive margin—conditional on engaging in educational activities—I estimate three alternative models: (i) an OLS regression restricted to diaries with positive education time, (ii) a Tobit model with left-censoring at zero, and (iii) a generalized linear model (GLM) with a Gamma family and log link. Following Brown and Dunn (2011), Hammer (2012), and Stewart (2013), the Gamma GLM provides a more appropriate specification for time-use data than the Tobit. The Gamma model assumes $\mathbb{E}[Y_i|X_i] = \exp(X_i'\beta)$ and $\text{Var}(Y_i|X_i) = \phi\,\mathbb{E}[Y_i|X_i]^2$, allowing variance to scale proportionally with the mean and accommodating the right-skew typical of time-use outcomes. Unlike the Tobit, it does not misinterpret sampling zeros as censored latent outcomes, and it maintains interpretability through the multiplicative structure implied by the log link. Formally,

$$Y_i \sim \text{Gamma}(\mu_i, \phi), \qquad \log \mu_i = \beta_0 + \beta_1 \text{WFH}_i + \gamma' X_i + \delta_t + \lambda_s, \tag{3}$$

where $\mu_i = \mathbb{E}[Y_i|X_i]$ and the WFH coefficient admits the semi-elastic interpretation $(\exp(\hat{\beta}_1) - 1) \times 100\%$ as the percentage change in the conditional mean of education time.

Results. Table 3 reports the extensive- and intensive-margin estimates. On the extensive margin (column 1), working from home increases the probability that a parent records any education-related activity by 2.0 percentage points, significant at the 10 percent level. This

[8] I employ a logit rather than a probit model to relax distributional assumptions on the unobserved decision variable. Results are unchanged if a probit specification is used.

Table 3: Extensive and Intensive Margins of WFH on Weekly Education Time

	(1) Any EduTime Logit AME (p.p.)	(2) EduTime > 0 OLS (hrs)	(3) EduTime Tobit (hrs)	(4) EduTime > 0 GLM Gamma
Panel A: Level effects				
WFH coefficient	2.0*	0.302	0.565**	0.094*
	(1.1)	(0.202)	(0.256)	(0.052)
Panel B: Relative effects				
Baseline mean (non-WFH)		3.55	0.63	3.55
Relative WFH effect	2.0 p.p.	8.5%		9.8%
Observations	10,388	2,316	11,045	2,316

Notes: Column (1) reports average marginal effects (AMEs) from a logit model for the probability of any positive education time (in percentage points; $N = 10{,}388$ due to survey design constraints). Column (2) shows OLS estimates on diaries with positive education time ($N = 2{,}316$). Column (3) reports Tobit coefficients with left-censoring at zero ($N = 11{,}045$). Column (4) reports the log-link coefficient from a Gamma GLM; the relative effect is $(\exp(0.094) - 1) \times 100 = 9.8\%$. Baseline means: $E[Y] = 0.63$ (all diaries), $E[Y > 0] = 3.55$. Appendix Figure 14 confirms superior Gamma fit. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

result is consistent with the idea that the flexibility and spatial proximity afforded by WFH facilitate at least some engagement in children's educational routines.

Turning to the intensive margin, the OLS regression estimated on the positive subsample (column 2) yields a coefficient of 0.302 hours per week, which is positive but not significant, likely due to the smaller number of diaries with nonzero education time (2,316 observations). The Tobit model (column 3), which treats all zeros as censored latent values, produces a larger and statistically significant estimate of 0.565 hours, but this specification is known to overstate effects in time-diary data because it interprets sampling zeros as censored outcomes rather than genuine non-occurrences (Stewart, 2013).

The Gamma GLM with a log link (column 4) provides the most credible intensive-margin estimate. This specification accommodates the non-negative, right-skewed nature of time-use outcomes and yields a directly interpretable multiplicative effect on the conditional mean. A comparison of the model fits, detailed in Appendix Figure 14, confirms that the Gamma distribution provides a superior fit to the empirical distribution of positive education times compared to the Normal distribution. The estimated log coefficient of 0.094 implies that, among parents who report positive education time, WFH is associated with a 9.8 percent increase in expected weekly hours (i.e., $\exp(0.094) - 1 = 0.098$). Although smaller than the Tobit estimate, the GLM effect is statistically significant and consistent in sign and magnitude with the extensive-margin results[9].

[9]See also Hammer (2012) for similar findings when comparing linear, Tobit, and Gamma models for housework time, and Stewart (2013) for a broader critique of Tobit in the ATUS context.

Across all specifications, the qualitative pattern is robust: working from home is systematically associated with greater parental investment in children's education. Robustness checks in Appendix A confirm the stability of these findings under winsorization and across leave-one-year-out re-estimations for 2010–2023, with the WFH coefficient remaining positive and statistically significant in every iteration.

Commuting Time as a Mechanism. The theoretical model predicts that WFH increases parental education time primarily by relaxing the daily time budget: hours previously spent commuting become available for household activities, including educational interactions with children. I test this channel directly using a mediation analysis in which commuting time enters as an additional control variable alongside the WFH indicator (Appendix Table 20).

Three findings support the commuting-time mechanism. First, WFH is associated with a reduction in average commuting time of 1.908 hours per week (SE = 0.110, $p < 0.001$), confirming a strong first-stage relationship. Second, conditional on commuting time, the WFH coefficient on education time falls from 0.153 (SE = 0.063) to 0.136 (SE = 0.063), implying that approximately 11% of the WFH–education premium is mediated by saved commuting time. Third, the coefficient on commuting time in the education equation is negative (-0.008^{*}, SE = 0.005), consistent with commuting crowding out educational activities. Taken together, these results suggest that the time-reallocation channel is operative, though it accounts for a minority of the total WFH premium—the majority reflects other dimensions of flexibility (reduced scheduling friction, physical proximity to children, lower cognitive load from commuting).

Main Findings. This section examines how work-from-home (WFH) arrangements relate to parental investments in children's education using ATUS time-diary data for 2010–2023 (excluding 2020).

Across specifications, the results show a clear and robust association between WFH and greater parental education time. In the full sample, the preferred OLS estimate (with all controls including spouse characteristics) implies that WFH parents spend 0.153 additional hours per week—a 24 percent increase relative to the non-WFH mean of 0.63 hours. The association is insensitive to the choice of WFH threshold: coefficients at the 1h, 3h, and 4h thresholds range from 0.188 to 0.197 and are all significant at the 1 percent level; the continuous WFH hours specification implies 0.018 additional hours of education per additional WFH hour per day. A dosage-response analysis reveals the largest gains for parents who work from home 4–8 hours per day (0.345^{***}), with positive associations at light and

full-day intensities as well.

The extensive-margin estimates reinforce this pattern: WFH increases the probability of engaging in any education-related activity by 2.0 percentage points. Conditional on reporting positive education time, the preferred Gamma GLM indicates a 9.8 percent increase in expected weekly hours. The result is robust to winsorization, to expanding the sample to include weekend diaries (0.141*, SE = 0.074), and to leave-one-year-out re-estimation. A mediation analysis attributes approximately 11 percent of the total WFH premium to saved commuting time, providing direct evidence for the time-budget mechanism in the theoretical model. Taken together, these findings suggest that the flexibility created by remote work enables parents to both participate more frequently in children's educational routines and devote more time when they do.

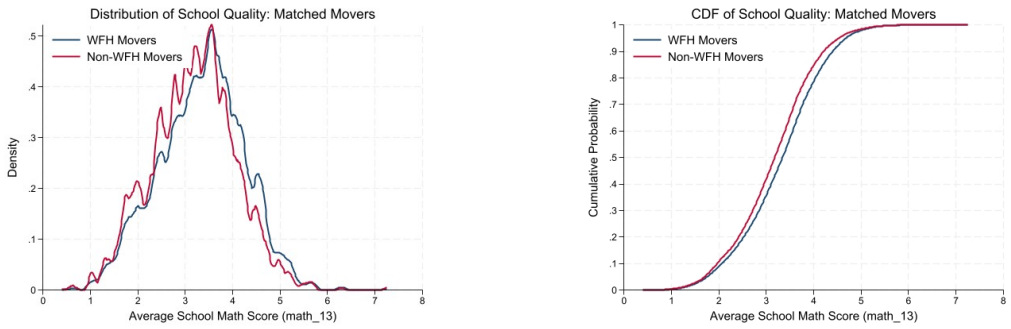
2.2 WFH and Residential Sorting into High-Quality School Districts

The geographic flexibility afforded by WFH may fundamentally alter how families choose where to live. If this flexibility allows parents to prioritize school quality over commute proximity, it could amplify educational inequality. This section evaluates whether WFH households are more likely to end up in neighborhoods with higher-performing schools.

Empirical Strategy and Data I analyze residential mobility patterns using individual-level data from the 2022–2023 American Community Survey (ACS), merged with neighborhood quality metrics from Chetty et al. (2024). For the distributional comparison, I focus on households with children who moved residences in the previous year. For the transition regressions, I use the broader sample of movers with children for whom valid origin and destination PUMA information is available. The key outcome variable is the average standardized math test score for third-grade students in 2013 at the Public Use Microdata Area (PUMA) level, denoted as `math_13`.

A simple comparison of movers would be confounded by the fact that WFH workers have higher incomes and education levels. To address this selection bias, I use propensity score matching. Specifically, I match WFH households (`WFH=1`) to observably similar non-WFH households (`WFH=0`) on education, household income, age, number of children, gender, state of residence, and local rental costs. The matching procedure, detailed in Appendix **??**, balances observables between the two groups and creates a comparable sample of movers who differ primarily in their WFH status.

Matched Destination Distributions Figure 2 compares the distribution of destination school quality for matched WFH and non-WFH movers.



(a) Kernel Density Estimates (b) Cumulative Distribution Functions

Figure 2: School Quality in Destination Neighborhoods: WFH vs. Non-WFH Movers

Notes: The figures display the distribution of average Grade 3 math scores (`math_13`) in the destination PUMAs for matched samples of WFH and non-WFH movers. The left panel shows kernel densities, and the right panel shows cumulative distribution functions.

The matched distributions show a clear rightward shift for WFH movers, indicating that they tend to relocate to PUMAs with higher school quality than observably similar non-WFH movers. This pattern is visible in both the kernel density and the CDF.

To assess whether this difference is statistically meaningful, I report a Kolmogorov-Smirnov (K-S) test in Table 4.

Table 4: Kolmogorov-Smirnov Test for Difference in School Quality Distributions

Statistic	**Value**
KS Statistic (D)	0.087
p-value	0.000
Point of Maximum Difference	3.55
Sample Size (WFH)	11,558
Sample Size (Non-WFH)	71,341

The K-S statistic of 0.087 is statistically significant, indicating that the destination-school-quality distributions of WFH and non-WFH movers differ in economically meaningful ways. This descriptive evidence is consistent with the idea that WFH flexibility facilitates residential sorting into neighborhoods with better schools.

WFH and School-Quality Transitions The distributional evidence above compares the destinations chosen by WFH and non-WFH movers. To study upgrading more directly, I next examine how WFH is associated with the *change* in school quality between origin and destination.

Empirical Strategy For movers with children and valid origin and destination PUMA information, I construct the continuous outcome

$$\texttt{school_quality_change}_i = \texttt{math_13}_i^{dest} - \texttt{math_13}_i^{orig},$$

which measures destination minus origin school quality. In the main legacy sample, I estimate an unweighted linear probability model, a logit model, and an ordered-logit model on WFH and a vector of controls including education, household income, age, age squared, number of children, gender, state fixed effects, origin school quality, origin rent, and origin local high-paying job access.

To preserve the ordinal information in the transition while being less sensitive to small changes around zero, I also estimate an ordered-logit model where the outcome is a three-category variable, `school_improvement`, indicating whether the mover transitioned to a *Worse* (0), *Similar* (1), or *Better* (2) school-quality tercile. I also estimate weighted versions of the same specification on the exact same legacy sample to isolate the role of the person weights from the role of the sample definition.

Main Results Table 5 presents the main transition results. All three unweighted estimates are positive. The linear probability model is positive and weakly significant, the logit is positive and marginally significant, and the ordered-logit is positive and statistically significant. On the same sample, the weighted versions remain positive but are somewhat attenuated, which suggests that the qualitative result is not driven by weights alone even though inference is sensitive to the choice.

Table 5: WFH and School-Quality Transitions

	(1) **LPM** *School-Quality Upgrade*	(2) **Logit** *School-Quality Upgrade*	(3) **Ordered Logit** *Worse/Similar/Better*
WFH	0.032*	0.209*	0.233**
	(0.018)	(0.109)	(0.093)
Outcome	Binary upgrade	Binary upgrade	Ordered upgrade
Model	LPM	Logit	Ordered logit
Controls	Origin set	Origin set	Origin set
State FE	Yes	Yes	Yes
Observations	6,197	6,197	6,197
R^2 / Pseudo R^2	0.310	0.299	0.260

Notes: The table reports unweighted estimates on the legacy school-age mover sample. Column (1) reports a linear probability model for the binary upgrade indicator, Column (2) reports a logit model for the same binary outcome, and Column (3) reports an ordered-logit model for terciles of school-quality change. Controls include education, state fixed effects, household income, age, age squared, number of children, gender, origin school quality, origin rent, and origin local high-paying job access. Robust standard errors are reported in parentheses. The weighted same-sample comparison yields smaller but positive coefficients: LPM 0.0205, logit 0.1682, ordered logit 0.2391. *p<0.10, **p<0.05, ***p<0.01.

Figure 3 plots the predicted probabilities from the ordered logit. The pattern is economically meaningful: relative to non-WFH movers, WFH movers are less likely to end up in the *Worse* category and more likely to end up in the *Better* category, while the probability of a *Similar* move is comparatively stable. This is exactly the type of reallocation toward better school districts suggested by the matched destination distributions.

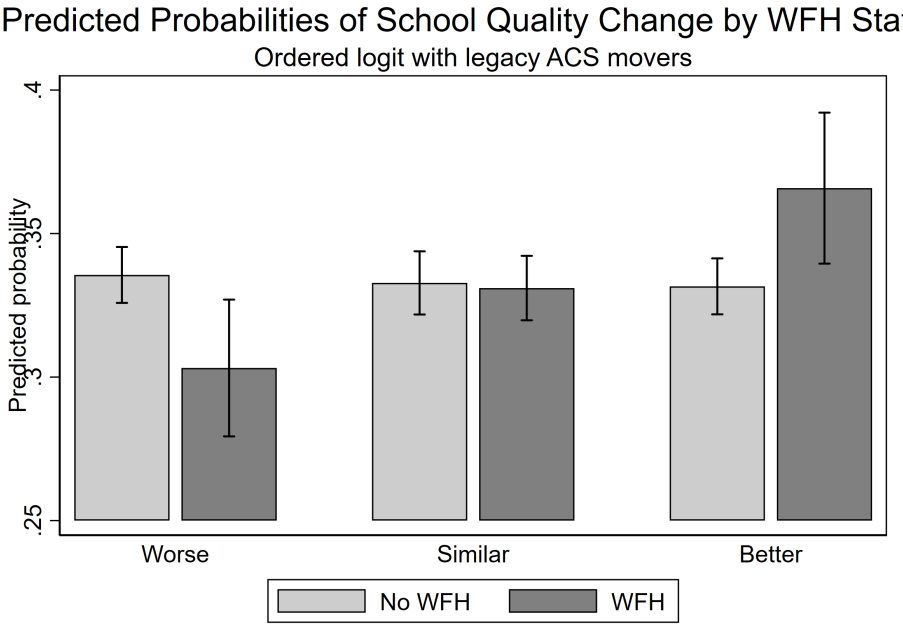


Figure 3: Predicted Probabilities of School Quality Change by WFH Status

Notes: Grouped bars show predicted probabilities for each outcome of the ordered logit in Table 5, Column (3), by WFH status. Error bars are 95% confidence intervals.

Binary Robustness Checks A useful robustness exercise is to collapse school-quality transitions into a binary indicator for moving to a better school district and to compare it with a stricter decile-improvement dummy. In the legacy same-sample comparison, both binary definitions remain positive, but the inference is noticeably more sensitive to weights than the ordered-logit specification. This reinforces the idea that the ordinal transition model preserves more of the signal in the underlying data.

Interpretation and Robustness Taken together, the ACS evidence points toward positive residential sorting by WFH households, but the strength of that conclusion depends on how the transition is measured. The matched destination distributions are robustly shifted to the right, and the legacy same-sample LPM, logit, and ordered-logit estimates all point in the same direction. The weighted and unweighted versions differ somewhat in magnitude, which suggests that weights matter for precision, but they do not overturn the positive legacy-sample pattern. This pattern suggests that the main empirical signal is a gradual reallocation toward better school districts rather than a sharp jump captured well by a zero-threshold binary outcome.

As a placebo exercise, I re-estimate the mover models on households without school-age children. Those placebo estimates remain close to zero and statistically insignificant, which is consistent with school quality being more relevant for families with children.

Discussion While the research design controls for a wide array of observables, I refrain from making strong causal claims due to potential omitted variable bias. For instance, WFH adoption might be correlated with unobserved preferences for neighborhood amenities that are themselves correlated with school quality. Even so, the revised ACS evidence now fits the paper's broader mechanism better when framed around *distributions, binary logit, and ordered-logit transitions* rather than around a single weighted logit. In that sense, the ACS section continues to support the residential-sorting channel, but the strongest version of the evidence is that WFH shifts the *distribution* of school-quality transitions in a favorable direction.

In summary, the empirical evidence indicates that WFH is associated with increased parental educational time, and the ACS results suggest that WFH also facilitates residential sorting toward better school districts when school-quality transitions are measured in a way that preserves more of the underlying information. In the next section, we build a model to understand the quantitative importance of these channels for intergenerational mobility.

3 An OLG Model of Neighborhood Sorting, WFH, and Child Education

I develop an overlapping-generations (OLG) model to formalize the empirical patterns documented in Section 2: (i) parents with work-from-home (WFH) capability devote more time to their children's education, and (ii) WFH households are more likely to sort into neighborhoods with higher school quality. The central aim of the model is to assess the quantitative contribution of each channel and to evaluate the broader implications for inequality and intergenerational mobility as WFH diffuses across the occupational structure.

The analysis proceeds in two stages. First, I analyze a partial equilibrium (PE) setting with exogenous rents to isolate the core household mechanisms. I use the simplified static framework to derive a theoretically interesting result: in the absence of exogenous forces, the model can produce a perfectly segregated economy, where the high human capital and the low human capital households live separately. Second, I embed the model in a general equilibrium (GE) framework by endogenizing housing rents. Finally, I introduce idiosyncratic location preferences to generate realistic sorting patterns for quantitative analysis.

3.1 Partial Equilibrium Analysis

I begin by analyzing the household's problem in a partial equilibrium setting. In this framework, I take key prices, i.e. housing rents (r_z), as exogenous parameters.

Economic Environment. The economy is populated by a continuum of agents of measure one who live for two periods. In the first period, an agent is a child who passively accumulates human capital and draws an innate ability a_c. In the second period, the agent becomes an adult who makes active economic decisions: choosing a residential location, allocating time between work, leisure, and child education, and consuming a final good and housing services. Furthermore, a competitive final goods sector produces output Y using only the aggregate effective labor supplied in the economy N as input ($Y = N$). Profit maximization and free entry pin the wage per unit of effective labor to the marginal product of labor, which is normalized to one ($w = 1$).

Household Types and Locations Adults are heterogeneous in three key dimensions (k, a, ι): $k > 0$ is the adult's human capital, which functions as a productivity shifter in the labor supply; $a > 0$ is the adult's innate ability, inherited from their own childhood; finally, $\iota \in \{0, 1\}$ is the occupation type, where $\iota = 1$ indicates a WFH-eligible job and $\iota = 0$ indicates

a mandatory on-site job. Occupation type (ι) is treated as an exogenous characteristic, as an analysis of the economic determinants behind occupational WFH selection is beyond the scope of this study. I acknowledge, however, that endogenizing this choice is an important avenue for future research. Adults choose to live in one of two locations, $z \in \{C, S\}$ denoting a central city (CBD) or a suburb. *Ex ante*, these locations are different only in commuting costs, $t_z > 0$, a time overhead incurred per unit of on-site work, and housing rents, $r_z > 0$. In the partial equilibrium analysis, I treat the vector of rents and commuting costs (r_C, r_S, t_C, t_S) as exogenous parameters.

Preferences and Constraints An adult's utility is derived from three distinct sources: the consumption of goods and housing, leisure, and the anticipated future human capital of their child. To model these preferences, I specify an additively separable utility function in the spirit of Chyn and Daruich (2022) and Fogli and Guerrieri (2019).

The first component, represents utility from material consumption. c and h are consumption of a numeraire good and housing services respectively. The nested Cobb-Douglas function captures the notion that adults derive utility from a composite of these two goods, with γ determining the optimal expenditure share on consumption versus housing. This composite good is then evaluated by a Constant Relative Risk Aversion (CRRA) function, where $\sigma \geq 0$ governs the household's willingness to substitute consumption across states and time; a higher σ implies a stronger preference for smooth consumption paths.

The second term, captures the utility from leisure time ℓ. The logarithmic form ensures that the marginal utility of leisure is decreasing. The parameter $\psi > 0$ scales the relative weight placed on leisure versus other utility sources.

The final term, $\beta \mathbb{E}\left[\ln k_c \mid a, z\right]$, embodies altruistic parental investment in the next generation. Parents care about their child's future human capital, k_c. The expectation operator $\mathbb{E}$ indicates that parents form rational expectations about their child's outcome based on the child's inherited innate ability a_c and the quality of the neighborhood z they are raised in. The parameter $\beta > 0$ measures the strength of this altruistic motive relative to consumption and leisure. This setup directly generates the "rug rat race" incentive: parents willingly sacrifice their own consumption and leisure to boost their child's expected human capital.

$$U = \frac{\left[c^{\gamma} h^{1-\gamma}\right]^{1-\sigma} - 1}{1-\sigma} + \psi \ln \ell + \beta \mathbb{E}\left[\ln k_c \mid a, z\right] \tag{4}$$

Labor Supply Technology An adult's effective labor supply, n, which determines their earnings, is produced by combining their human capital k with time allocated to work.

The technology for this production differs critically based on whether their occupation is eligible for remote work ($\iota = 1$). For on-site workers ($\iota = 0$), effective labor is generated through a standard function, $A^O(kn^O)^\theta$, where time spent at the office n^O is transformed into output with productivity A^O and diminishing returns $\theta \in (0, 1)$. For WFH-eligible workers, productive remote work is contingent on a key friction: adequate housing space. Only if housing consumption h meets a minimum threshold $h_{\min} > 0$ (representing a home office with the space for a desk) can they access a CES production technology. This function, allows them to combine office hours (n^O) and remote hours (n^H) as imperfect substitutes, with an elasticity of substitution $\zeta = 1/(1-\rho)$. The home arm's productivity is scaled by A^H. If this housing requirement is not met ($h < h_{\min}$), the worker is relegated to the standard on-site production technology, forfeiting the benefits of remote flexibility.

This specification, following Richard (2024), creates a crucial link between housing choices, location, and labor productivity. The requirement $h \geq h_{\min}$ means that the feasibility of productive remote work is contingent on a household's ability to consume enough housing.

$$n = \begin{cases} A^O(kn^O)^\theta, & \iota = 0, \\ \left[\left(A^O(kn^O)^\theta\right)^\rho + \left(A^H(kn^H)^\theta\right)^\rho\right]^{1/\rho}, & \iota = 1 \text{ and } h \geq h_{\min}, \\ A^O(kn^O)^\theta, & \iota = 1 \text{ and } h < h_{\min}. \end{cases} \tag{5}$$

Child Human Capital Formation The human capital of a child, k_c, is shaped by three fundamental inputs: innate ability (a_c), the quality of the neighborhood and its schools (S_z), and the time their parents invest in educational activities (e). These are combined through a Cobb-Douglas production function, as in Zheng and Graham (2022):

$$k_c = a_c S_z^\nu e^{1-\nu}, \qquad \nu \in (0, 1). \tag{6}$$

The parameter $\nu \in (0, 1)$ governs the relative importance of neighborhood spillovers versus direct parental investment.

A child's innate ability is not predetermined but is a stochastic function of their parents' ability, following a log-AR(1) process (Zheng and Graham, 2022; Fogli and Guerrieri, 2019):

$$\log a_c = \mu_a + \rho_a \log a + \varepsilon_a, \qquad \varepsilon_a \sim \mathcal{N}(0, \sigma_a^2). \tag{7}$$

Here, μ_a represents the unconditional mean of log ability, $\rho_a \in [0, 1)$ measures the persistence of ability across generations, and σ_a^2 is the variance of random shocks. This intergenerational persistence is a critical driver of inequality.

The combination of Equations (6) and (7) creates the fundamental engine for the tyranny of the advantaged. Child ability is positively correlated with parental ability ($\rho_a > 0$), the expected return on parental educational investment is in turn higher for high-ability parents. This means each marginal unit of educational time e spent with a child is anticipated to yield a higher future human capital for children of high-ability parents. Consequently, these parents face a stronger incentive to invest heavily to improve their children's competitive position.

Peer Spillovers The quality of a neighborhood's educational environment, captured by the peer spillover S_z, is not an exogenous amenity but is endogenously determined by the collective characteristics of its residents. Formally, it is defined as the average expected human capital of all children in location z (Fogli and Guerrieri, 2019):

$$S_z = \frac{1}{\mu_z} \int \mathbb{E}[k_c \mid a, z] \mathbf{1}[Z(\cdot) = z] dF(k, a, \iota). \tag{8}$$

Using the human capital production function and the process for innate ability, this definition simplifies to a closed-form expression:

$$S_z^{1-\nu} = \frac{K_a}{\mu_z} \int a^{\rho_a} e_z(k, a, \iota)^{1-\nu} \mathbf{1}[Z(\cdot) = z] dF(k, a, \iota), \tag{9}$$

where $K_a \equiv \exp(\mu_a + \sigma_a^2/2)$ is a constant scaling factor. The spillover S_z is therefore a geometric average of two key resident attributes: the innate ability of parents (via a^{ρ_a}) and the level of parental educational investment ($e^{1-\nu}$). This means neighborhoods with a higher concentration of skilled, high-investing parents naturally generate stronger positive peer effects and school quality.

Crucially, this creates a self-reinforcing cycle. Families are attracted to neighborhoods with high S_z because it boosts their own child's human capital. However, their decision to sort into a neighborhood collectively determines the value of S_z. This feedback loop can lead to persistent inequality across neighborhoods, as advantaged areas become increasingly advantaged. A key simplifying feature of the model is that, due to the log utility specification, a household's choice of educational time e is independent of the level of S_z; they take the spillover as given when making their investment decision, which avoids complex strategic interactions and makes the equilibrium analysis tractable.

Partial Equilibrium For a given distribution of adult types $F(k, a, \iota)$ and a fixed vector of location-specific parameters (r_C, r_S, t_C, t_S), the residential equilibrium of this economy

consists of a set of household policy functions for consumption $c_z(k, a, \iota)$, housing $h_z(k, a, \iota)$, labor supplies $n_z^O(k, a, \iota)$ and $n_z^H(k, a, \iota)$, educational time $e_z(k, a, \iota)$, and leisure $\ell_z(k, a, \iota)$; a location choice rule $Z(k, a, \iota) \in C, S$; and a pair of neighborhood peer spillovers (S_C, S_S) such that, for each location z:

i **Household Optimization:** Given rents r_z, commuting costs t_z, and spillovers S_z, the policy functions solve each household's utility maximization problem, subject to their budget and time constraints.

ii **Location Choice:** Each household chooses the location z that delivers the highest indirect utility, such that $Z(k, a, \iota) = \arg\max_{z \in C,S} V_z(k, a, \iota)$.

iii **Peer Spillover Consistency:** The spillover value S_z is consistent with the collective outcomes of households' location and investment decisions, satisfying Equation (9).

3.2 Residential Sorting and Comparative Statics in Partial Equilibrium

Having established the model's environment and equilibrium definition, I now characterize the resulting residential sorting patterns. In this partial equilibrium analysis, I take housing rents (r_C, r_S) and commuting costs (t_C, t_S) as exogenous, while peer spillovers (S_C, S_S) are determined endogenously through household choices.

I first derive the key household optimality conditions, which yield a convenient decomposition of the indirect utility difference between locations. This setup allows us to establish several core propositions: (i) the existence of an equilibrium with consistent peer spillovers; (ii) that sorting follows a cutoff rule in human capital, which relies on a clear single-crossing condition; (iii) that WFH eligibility unambiguously lowers the human capital cutoff for suburban residence; and (iv) that an interior sorting equilibrium exists for a non-degenerate interval of the altruism parameter β.

Household Optimization and Indirect Utility A household's problem can be simplified by first solving the intra-temporal allocation of its budget. Let total expenditure be $y = n$, where n is effective labor supply. Standard Cobb-Douglas demands yield:

$$c = \gamma y, \qquad r_z h = (1 - \gamma) y$$

The material utility component for location z can thus be expressed as a function of y and r_z:

$$U_z^{\text{mat}}(y) \;=\; \frac{\left[\chi\, y\, r_z^{-(1-\gamma)}\right]^{1-\sigma} - 1}{1 - \sigma}, \quad \chi \equiv \gamma^{\gamma}(1-\gamma)^{1-\gamma}$$

The household's problem then reduces to choosing work hours (n^O, n^H), education time e, and leisure ℓ to maximize utility, subject to the time constraint $1 = n^O + n^H + e + \ell + t_z n^O$ and the labor production technology in Eq. (5).

Let $\tau_z \equiv 1 - (1 + t_z)n^O - n^H$ denote the residual time after work and commuting for neighborhood z. The first-order conditions for an interior optimum imply a constant proportional allocation of this residual time:

$$e_z = \frac{\beta(1-\nu)}{\psi + \beta(1-\nu)}\tau_z, \qquad \ell_z = \frac{\psi}{\psi + \beta(1-\nu)}\tau_z. \tag{10}$$

This result is pivotal: any factor that increases residual time τ_z, such as WFH flexibility or a voluntary reduction in work hours, automatically increases educational investment.

The first-order conditions for work hours equate the marginal product of labor to the shadow cost of time, adjusted for commuting:

$$\lambda \frac{\partial \tilde{n}}{\partial n_z^O} = \frac{\psi + \beta(1-\nu)}{\tau_z}(1 + t_z), \qquad \lambda \frac{\partial \tilde{n}}{\partial n_z^H} = \frac{\psi + \beta(1-\nu)}{\tau_z}, \tag{11}$$

where $\lambda \equiv \frac{\partial U^{\text{mat}}}{\partial y_z}$ is the marginal utility of income. These conditions, combined with the production function, pin down optimal work hours.

The peer spillover in a location is determined by the collective investment of its residents. Given the human capital production function (Eq. (6)) and the ability process (Eq. (7)), the consistency condition (Eq. (9)) can be written as:

$$S_z^{1-\nu} = \frac{K_a}{\mu_z} \int a^{\rho_a} e_z(k, a, \iota)^{1-\nu} \mathbf{1}_{Z(k,a,\iota)=z} dF(k, a, \iota), \tag{12}$$

where $K_a \equiv \exp(\mu_a + \sigma_a^2/2)$. Households take (S_C, S_S) as given when optimizing; in equilibrium, the spillovers induced by household choices must satisfy this condition.

A household chooses the location that delivers the highest indirect utility. The utility difference between suburb and city, $\Delta V(k, a, \iota) \equiv V_S(k, a, \iota) - V_C(k, a, \iota)$, can be decomposed using the previous results:

$$\Delta V(k, a, \iota) = \underbrace{\Delta U^{\text{mat}}(k, a, \iota) + (\psi + \beta(1-\nu)) \ln \frac{\tau_S}{\tau_C}}_{\text{Adult Utility}} + \underbrace{\beta\nu \ln \frac{S_S}{S_C}}_{\text{Child Spillover Utility}} \tag{13}$$

This decomposition separates the adult's utility from consumption and leisure from the altruistic utility derived *directly* from the child's neighborhood spillover.

Existence and Properties of the Partial Equilibrium I now state the key properties of the residential equilibrium. The following regularity assumption ensures well-behaved interior solutions.

Assumption 1 (Regularity and Interiority)**.** *For each (k, a, ι, z), the household problem admits an interior solution with $n^O, n^H, e, \ell > 0$ and $h > 0$. The labor aggregator is concave, and $U_z^{mat}(y)$ is strictly increasing and concave in y.*

My first result establishes that an equilibrium with consistent peer spillovers exists.

Proposition 1 (Existence of a Residential Equilibrium with Spillovers)**.** *Under Assumption 1, there exists at least one pair of peer spillovers (S_C^*, S_S^*) that satisfies the consistency condition (Eq. (12)) given the location choices and optimal policies of households.*

Proof Sketch. Define the mapping $\mathcal{T} : (S_C, S_S) \mapsto (\tilde{S}_C, \tilde{S}_S)$ by the right-hand side of Eq. (12). A higher S_z increases the utility of location z (from the $\beta\nu \ln S_z$ term in Eq. (13)), making it more attractive. This attracts households with higher a and e, which in turn increases the resulting spillover $\tilde{S}_z$. This implies the mapping $\mathcal{T}$ is order-preserving on a compact lattice. Tarski's fixed-point theorem then guarantees the existence of at least one fixed point (S_C^*, S_S^*). (See Appendix C.1 for details). □

The core sorting result relies on a standard single-crossing condition, which I define formally.

Definition 3.1 (Single-Crossing Condition)**.** The utility difference $\Delta V(k, a, \iota)$ satisfies the Single-Crossing Condition (SCC) in human capital k if it is strictly increasing in k, i.e., $\frac{\partial \Delta V(k,a,\iota)}{\partial k} > 0$.

This condition means that as a household's human capital k increases, the suburb becomes *relatively* more attractive than the city.

Proposition 2 (Cutoff Sorting by Human Capital)**.** *Let (S_C^*, S_S^*) be a spillover fixed point. Assuming the Single-Crossing Condition holds and that neither location is dominant for all types (i.e., $\Delta V(\underline{k}, a, \iota) < 0$ and $\Delta V(\overline{k}, a, \iota) > 0$), then for each (a, ι), there exists a **unique** human capital cutoff $k^*(a, \iota)$ such that all households with $k < k^*$ choose the city and all households with $k > k^*$ choose the suburb.*

Proof. The Single-Crossing Condition ensures $\Delta V(k, a, \iota)$ is a strictly increasing function of k. Given the boundary conditions $\Delta V(\underline{k}) < 0$ and $\Delta V(\overline{k}) > 0$, the Intermediate Value Theorem guarantees the existence of a *unique* $k^* \in (\underline{k}, \overline{k})$ where $\Delta V(k^*, a, \iota) = 0$. By

monotonicity, $\Delta V < 0$ for all $k < k^*$ (city choice) and $\Delta V > 0$ for all $k > k^*$ (suburb choice). □

The SCC is not just a technical assumption, but is a theoretical feature that holds in my model due to some key trade-offs, and the separability of the utility's components. The choice of location is a trade-off between the Suburb's main cost (high commuting time, $t_S > t_C$) and the Suburb's main benefits (low housing rent, $r_S < r_C$, and potentially higher amenities/spillovers, $S_S > S_C$). The SCC holds because human capital k interacts with this trade-off in a specific way.

Lemma 1 (Primitive Conditions for Single-Crossing)**.** *Assume the suburb has higher commuting costs ($t_S > t_C > 0$) and lower rents ($r_S \leq r_C$). The Single-Crossing Condition ($\partial \Delta V / \partial k > 0$) will hold if the utility function is parameterized such that (i) housing is a normal good and (ii) the income effect on leisure and education time is dominant, leading higher-k households to work less. This dominance is satisfied if $\sigma > 1$.*

Proof Sketch. The relative attractiveness of the suburb increases with a household's productivity, k, due to the combined effect of two economic channels that both favor higher-k households. First, consider the benefit channel of housing. An increase in k raises household income y_z, and since housing is a normal good ($r_z h_z = (1-\gamma) y_z$), this leads to a greater demand for it. Consequently, the total utility gain from the suburb's cheaper rent becomes larger for more productive households. Crucially, those with higher k who purchase larger homes benefit disproportionately more from this suburban rent discount. At the same time, the cost channel of commuting also shifts in their favor. As k rises, a household's potential income increases, making residual time (leisure ℓ_z and education e_z) more valuable. Under the condition that $\sigma > 1$, the income effect from this productivity gain dominates the substitution effect, causing households to "buy" more τ_z by working fewer on-site hours (i.e., $\partial n_z^O / \partial k < 0$, see Appendix C.2). Because these higher-k households spend less time commuting, they are less exposed to the suburb's primary disadvantage: its high commute cost t_S. Thus, the main cost of suburban living becomes less burdensome as k increases. □

In summary, both effects reinforce one another: the benefits of suburban living grow while its primary cost diminishes for more productive households. This ensures that the utility differential, ΔV, shifts in favor of the suburb as k rises, establishing that $\partial \Delta V / \partial k > 0$. A formal proof is provided in Appendix C.3. I now turn to the model's central comparative static regarding WFH. This proposition shows that WFH eligibility amplifies the sorting pattern established in Lemma 1 by acting as a time-saving technology.

Proposition 3 (WFH Lowers the Suburban Cutoff)**.** *Fix rents, commuting costs, and a spillover equilibrium. Assume the suburb has a strictly higher commuting cost ($t_S > t_C$). The option to WFH ($\iota = 1$) expands the household's production frontier, and this expansion is* ***relatively more valuable*** *in the suburb than in the city. This implies:*

$$\Delta V(k, a; \iota = 1) > \Delta V(k, a; \iota = 0)$$

Consequently, the suburban cutoff human capital is lower for WFH-eligible households: $k^*(a, \iota = 1) < k^*(a, \iota = 0)$.

Proof. The proof, detailed in Appendix C.4, demonstrates that the introduction of work-from-home (WFH) eligibility lowers the productivity cutoff k^* by increasing the relative attractiveness of the suburb. The argument unfolds in three logical steps. First, WFH eligibility ($\iota = 1$) acts as a time-saving technological innovation. It allows a household to substitute "cheap-time" remote hours n^H, which cost 1 unit of time, for "expensive-time" on-site hours n^O, which cost $1 + t_z$ units of time. Through optimal substitution, a WFH-eligible household can generate the *same* effective income n as a non-eligible household but at a *strictly lower* total time cost, thereby expanding its budget set. Second, the value of this time-saving technology is location-specific. The utility gain is determined by the amount of time saved, which is maximized by replacing the most expensive on-site work. Since the suburban commute is costlier than the city's ($t_S > t_C$), the utility gain from WFH access is greater for a suburban household: $V_S(\iota = 1) - V_S(\iota = 0) > V_C(\iota = 1) - V_C(\iota = 0)$. Finally, this differential gain causes the cutoff k^* to decrease. Rearranging the inequality from the previous step shows that the net utility surplus of the suburb increases for WFH-eligible households, i.e., $\Delta V(\iota = 1) > \Delta V(\iota = 0)$. This represents an upward shift of the ΔV curve. Given the single-crossing condition (Proposition 2), where ΔV is an increasing function of k, this upward shift necessarily causes the curve to cross the $\Delta V = 0$ axis at a *lower* value of k. Therefore, I conclude that $k^*(a, \iota = 1) < k^*(a, \iota = 0)$. □

This proposition directly rationalizes the empirical finding that WFH households are more likely to sort into suburbs, which is also in line with the telework literature on the so called *donut effect* (Ramani and Bloom, 2021).

The Direct Within-Generation Effect of WFH Beyond its effect on where households live (sorting), WFH eligibility has a direct effect on household outcomes within any given location. WFH acts as a superior production technology, expanding a household's production possibility frontier over income and residual time.

Compare a non-WFH household (N, $\iota = 0$) and a WFH household (W, $\iota = 1$) of the same type (k, a) in the same location z[10]. The analysis proceeds by examining their respective production possibilities and optimal choices.

First, consider the Production Possibility Frontier (PPF). Both households produce effective income n and residual time τ. Household N has a single input, on-site work n_N^O, with a time cost of $(1+t_z)$ per unit of work. Its PPF is therefore given by $\tau_N = 1-\text{Cost}_N(n)$. In contrast, household W has access to two inputs: on-site work n_W^O (costing $1 + t_z$) and remote work n_W^H (costing 1). By optimally substituting the "cheap-time" input n_W^H for the "expensive-time" input n_W^O, household W can produce any given income level n for a strictly lower total time cost than household N.

This cost advantage leads directly to **dominance** of the WFH household's opportunity set. Since $\text{Cost}_W(n) < \text{Cost}_N(n)$ for all $n > 0$, the WFH household's PPF, defined by $\tau_W = 1 - \text{Cost}_W(n)$, strictly dominates that of the non-WFH household. The entire frontier shifts outward.

Finally, we examine the optimal choice. The two households have identical, increasing preferences over income (via U^{mat}) and residual time (via τ). Given that both are normal goods, the outward shift of the PPF guarantees that the WFH household will select an optimal bundle (n_W^*, τ_W^*) that features both higher income and higher residual time compared to the non-WFH household:

$$n_W^*(k, a, z) > n_N^*(k, a, z) \quad \text{and} \quad \tau_W^*(k, a, z) > \tau_N^*(k, a, z).$$

This technological advantage is the "direct within-generation effect" of WFH: holding location fixed, WFH-eligible households are objectively richer in both time and money. This leads them to enjoy higher consumption and invest more time in their children's education (Eq. (10)). This effect is also visible in the numerical results (Figure 6), where the WFH income curves are strictly above their non-WFH counterparts in both locations.

Finally, I establish that the model produces an interior sorting equilibrium for a plausible range of altruism.

Proposition 4 (Interior Sorting for Intermediate Altruism)**.** *Assume the single-crossing condition and that the indirect utility is continuous in β. Suppose there exist values β_0 and β_1 such that $\Delta V(\overline{k}, a, \iota; \beta_0) < 0$ (e.g., altruism is low and high-k types avoid high suburban commute costs) and $\Delta V(\underline{k}, a, \iota; \beta_1) > 0$ (e.g., altruism is high and even low-k types seek*

[10]For this analysis we assume that the relative productivities of telework hours and office hours are similar, i.e., $A_H/A_O \sim 1$. One can formally prove that there is a cutoff value for the relative TFP productivities such that WFH is not a superior technology anymore.

suburban amenities/time). Then there exists an interval (β_L, β_U) *such that for all* β *in this interval,* $\Delta V(\underline{k}, a, \iota; \beta) < 0 < \Delta V(\overline{k}, a, \iota; \beta)$*, and the unique cutoff* $k^*(a, \iota; \beta)$ *is interior.*

Proof. By continuity of ΔV in β (a consequence of the Theorem of the Maximum), the boundary values $\Delta V(\underline{k})$ and $\Delta V(\overline{k})$ are continuous functions of β. The conditions $\Delta V(\overline{k}; \beta_0) < 0$ and $\Delta V(\underline{k}; \beta_1) > 0$, combined with the Intermediate Value Theorem, guarantee the existence of thresholds β_L and β_U between which the signs at the boundaries are $\Delta V(\underline{k}) < 0$ and $\Delta V(\overline{k}) > 0$. Single-crossing (Proposition 2) then ensures a unique interior cutoff k^* for any $\beta \in (\beta_L, \beta_U)$. □

Implications for Peer Spillovers and Empirical Patterns The sorting behavior characterized by these propositions has a direct implication for neighborhood quality.

Corollary 1 (Endogenous Spillover Ranking)**.** *If higher-human-capital households sort into the suburb* ($k^*(a, \iota) < \overline{k}$ *for all* a, ι*) and educational time* $e_z(k, a, \iota)$ *is non-decreasing in* k *and* a*, then in any sorting equilibrium, the suburban peer spillover is stronger:* $S_S > S_C$*.*

Proof. The spillover S_z is the average of $a^{\rho_a} e^{1-\nu}$ among residents of z (Eq. (12)). If the suburb contains a first-order stochastically dominant distribution of k (and thus a) and households with higher k and a also invest more time e, the average of this integrand will be higher in the suburb. □

This analysis demonstrates that the partial equilibrium model successfully rationalizes the key empirical patterns.

Corollary 2 (Replicating the Empirical Patterns)**.** *The partial equilibrium model successfully rationalizes the key empirical patterns from Section 2 via three distinct, yet related, effects:*

1. ***The Direct Within-Generation Effect:*** *WFH eligibility acts as a superior production technology, allowing households to achieve both higher effective income* (n_z) *and higher residual time* (τ_z) *within any given location. This provides a direct, within-generation welfare and income gain (see Figure 6).*

2. ***The Parental Time Effect:*** *The gain in residual time* (τ_z) *from the direct effect is mechanically channeled into intergenerational investment. Per the optimal time-allocation rule (Eq. (10)), a constant fraction of this new time is allocated to educational activities* (e_z)*. This replicates the first empirical finding (Section 2.1).*

3. ***The Sorting Effect:*** *The WFH technology is disproportionately valuable in the high-commute-cost suburb* ($t_S > t_C$)*. As shown in Proposition 3, this makes the suburb*

relatively more attractive to WFH-eligible households, causing them to sort into the suburb at lower human capital cutoffs. This replicates the second empirical finding (Section 2.2), as these households move into the (endogenously) high-spillover neighborhood (Corollary 1).

Numerical Illustration The partial equilibrium model yields clear analytical predictions under the stated assumptions. To illustrate these propositions quantitatively and provide a visual intuition for the model's mechanics, I solve the household's problem numerically, with an exogenous parameterization detailed in Table 6.

Parameter	Value	Description	Source
β	0.5	Altruistic parameter	fixed
σ	2.00	CRRA risk-aversion	Econ. Literature
γ	0.76	Weight on non-durables in utility	Davis and Ortalo-Magné (2011)
ψ	1.00	Weight on leisure	Economic literature
A_O	1.00	Office productivity scale	Richard (2024)
A_H	0.82	Home-office productivity	Gibbs et al. (2023)
θ	0.82	Labor share in production	Valentinyi and Herrendorf (2008)
ρ	0.7	EOS of office and WFH labor	Delventhal and Parkhomenko (2024)
t_C	0.02	Commuting time, center	Davis et al. (2024)
t_S	0.04	Commuting time, suburb	Davis et al. (2024)
$h_{\min}$	0.1	Minimum housing for home-office productivity	Richard (2024)
ν	0.40	Elasticity of child-HC w.r.t. parental ability	De La Croix and Doepke (2003)

Table 6: Exogenously calibrated parameters for the numerical exercise.

The weight on non-durable consumption in utility, γ, comes from Davis and Ortalo-Magné (2011), reflecting the observed stability in housing expenditure shares across U.S. metropolitan areas. The relative TFP between office work A_O and home-office work A_H is informed by Gibbs et al. (2023). These values capture the relative efficiency of remote work compared to traditional office settings, based on empirical studies of productivity differentials. The labor share in production, θ, follows Valentinyi and Herrendorf (2008), who estimate sectoral factor shares from U.S. national accounts data. Commuting time parameters for the two locations, t_C and t_S, are taken from the estimates of Davis et al. (2024), reflecting empirical estimates of time costs associated with commuting from central and suburban locations, who are also using the same empirical definition of CBD and suburbs. The minimum housing requirement for home-office productivity, h_{min}, is calibrated based on Richard (2024), who links housing space to effective remote work. The intergenerational human capital elasticity parameter, ν, is drawn from De La Croix and Doepke (2003), who

estimate the relative contributions of parental ability and educational investment to child human capital formation, using a dynastic model with endogenous fertility and education choices. The parameter ρ governing the elasticity of substitution in the labor supply CES function, is taken from Delventhal and Parkhomenko (2024), who estimates this value in a spatial equilibrium model to capture how workers substitute between different types of labor (remote vs. on-site work) in response to wage differentials and commuting costs. This parameter is crucial for quantifying the responsiveness of labor allocation to changes in the economic environment, such as the rise of remote work.

I compute the optimal policies and location choices for a continuum of households heterogeneous in ability a, human capital k with occupation type ι.

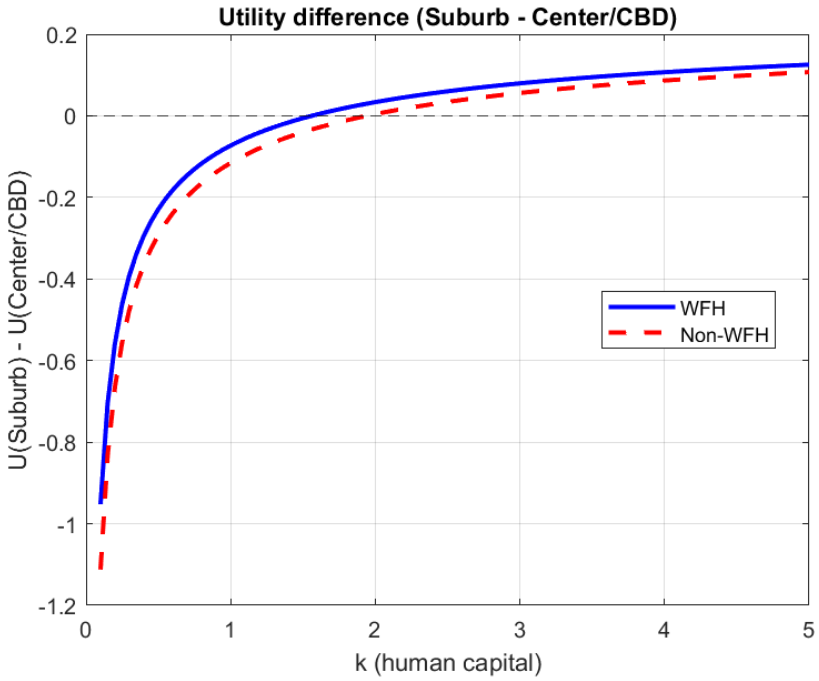


Figure 4: Single-Crossing Condition and WFH Cutoff (Propositions 2 and 3). The figure plots the utility difference between suburb and city, $\Delta V(k;\iota) = V_S(k;\iota) - V_C(k;\iota)$.

Figure 4 yields three critical results that validate the model's theoretical underpinnings. First, both curves are strictly increasing in k, numerically confirming the single-crossing condition (Lemma 1). This slope indicates that the net return to human capital is higher in the suburbs, rendering it increasingly attractive for higher-skilled households. Second, the figure clearly illustrates the sorting Effect: the WFH schedule (blue line) lies uniformly above the on-site schedule (red dashed line). This demonstrates that the WFH option amplifies the suburb's attractiveness at every skill level. Third, the unique cutoff where $\Delta V(k^*;\iota) = 0$ is lower for WFH-eligible households ($k^*_{WFH} \approx 1.56$) than for on-site workers ($k^*_{OnSite} \approx 1.96$). This downward shift is the direct result of the Sorting Effect, numerically

validating Proposition 3 and showing that WFH expands suburban sorting to a wider segment of the skill distribution.

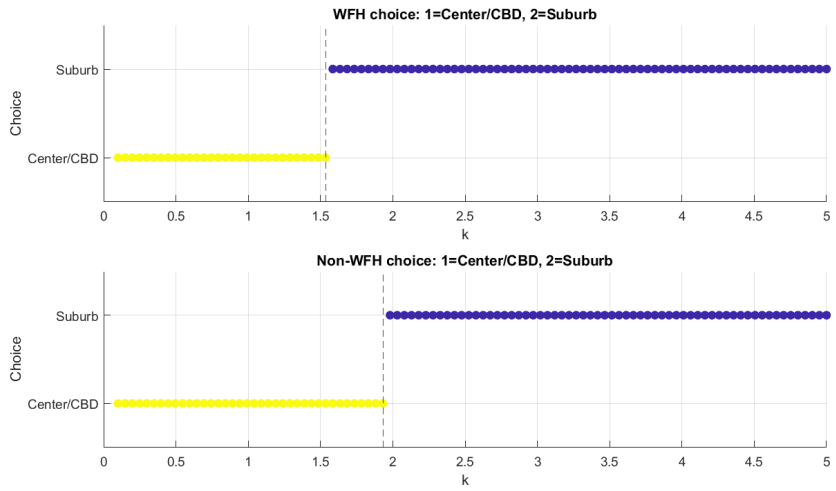


Figure 5: Equilibrium Residential Sorting by WFH Eligibility and Human Capital.

Furthermore, Figure 5 displays the discrete location choice for each household based on their human capital k. The stark vertical separation reveals a clear sorting equilibrium. The cutoff k^* is visibly lower for WFH-eligible households (top panel), providing direct, intuitive visual evidence for the Sorting Effect established in Figure 4: WFH lowers the human capital threshold for suburban residence.

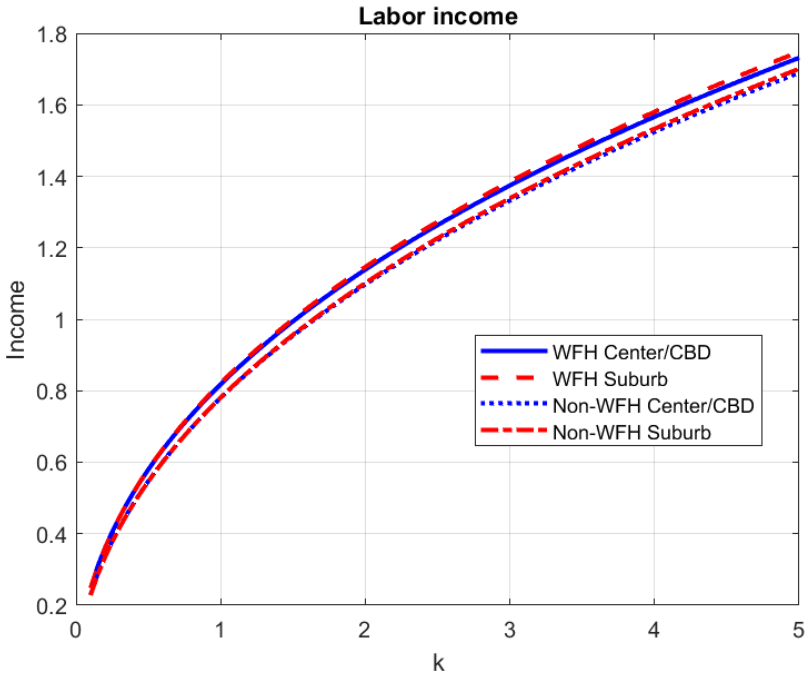


Figure 6: Labor Income by Location and WFH Status.

Figure 6 illustrates the **Direct Within-Generation Effect**, which is the mechanism

that drives the results in Figure 4. WFH acts as a superior production technology. For the same k, WFH-eligible households (solid lines) achieve strictly higher effective labor income than their on-site counterparts (dashed lines) in both locations. This is because they can substitute "cheap-time" home-work for "expensive-time" on-site work, mitigating the commute penalty ($t_S > t_C$). This technological advantage expands their production frontier, making them richer in both potential income and, subsequently, in residual time. This higher marginal return to skill in the suburbs is the economic force that drives both the upward slope of the ΔV curve and the upward shift for WFH-eligible households.

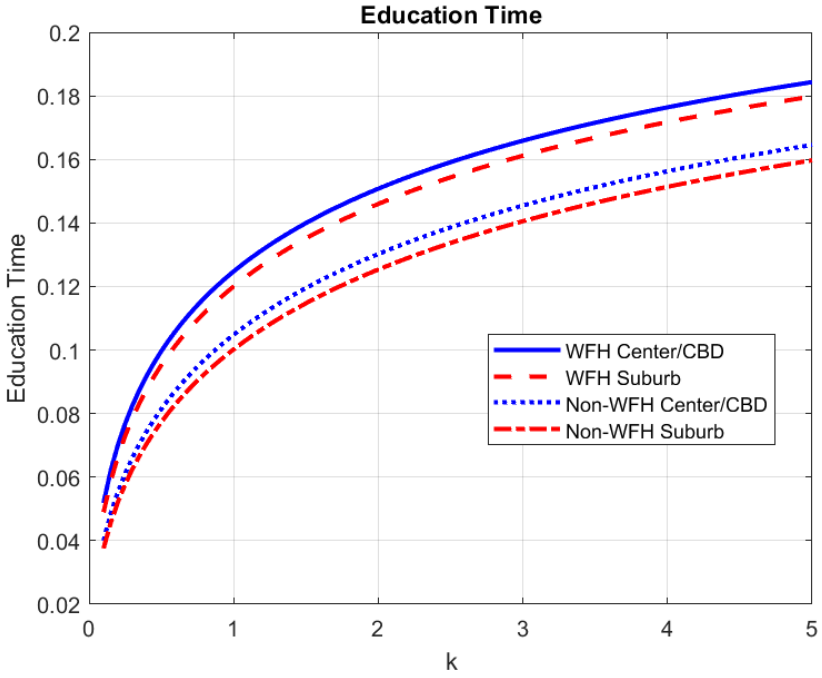


Figure 7: Educational Time Investment by Location and WFH Status.

Finally, Figure 7 shows the intergenerational consequences of these mechanisms. It clearly visualizes the *Parental Time Effect*. Within the same residential location (e.g., comparing the two "Suburb" lines), households with WFH-eligible jobs invest more time in education than their on-site counterparts. This is the mechanical consequence of the Direct Within-Generation Effect: WFH-eligible households have more residual time (τ_z), and per the optimal rule in Eq. (10), a constant proportion of this time windfall is automatically channeled into child education. Consequently, the highest level of educational investment is achieved by WFH-eligible households in the suburbs, as this group benefits synergistically from *both* the Sorting Effect (choosing the suburb) and the Parental Time Effect (the direct WFH time reallocation).

3.3 General Equilibrium and Stochastic Sorting

The partial equilibrium model yields a stark prediction: perfect sorting by household type. While theoretically elegant, this outcome is at odds with the rich heterogeneity observed in microdata, where communities are composed of families with a diverse mix of incomes, occupations, and preferences.

To bridge this gap and generate a quantitatively relevant model, I extend the baseline framework in two ways. First, I introduce a general equilibrium channel by endogenizing housing rents, which feedback into location choices. Second, I incorporate idiosyncratic location preferences, which disrupt the perfect sorting equilibrium and allow households to make different choices, thereby matching the empirical reality of imperfect sorting.

Endogenous Housing Rents In reality, attractive locations face upward pressure on housing prices. To capture this, I model a competitive housing development sector in each location z. Developers face a convex cost function for supplying housing units, H_z:

$$C_z(H_z) = \xi_z \frac{H_z^{1+\eta}}{1+\eta}, \quad \text{where} \quad \xi_z > 0, \eta > 0. \tag{14}$$

This structure implies that supplying more housing in a given location becomes progressively more expensive.

Profit maximization by these competitive developers leads to an upward-sloping housing supply curve in each location:

$$r_z = \xi_z H_z^{\eta} \quad \text{or, equivalently,} \quad H_z^S = \left(\frac{r_z}{\xi_z}\right)^{1/\eta}. \tag{15}$$

The key parameters governing this relationship are: $\xi_{\mathbf{z}}$, a location-specific cost shifter and η, the inverse elasticity of housing supply. A higher ξ_z implies that geographic or regulatory constraints (e.g., water bodies, zoning restrictions) make construction inherently more expensive, leading to higher rents for any given level of housing demand. While a larger η means supply is more inelastic (i.e., less responsive to price changes). This captures the notion that housing supply is highly rigid in the short run or in densely built areas. In this competitive market, developers earn zero economic profits in equilibrium. However, the economic rents generated by the scarce factor—land—are captured as pure profits (or Ricardian rents). I follow the standard approach by assuming these profits accrue to a class of absentee landlords who are not otherwise part of the model. This setup implies that the rental income paid by households is not recycled back into the local economy. The

formulation ensures that as more households are drawn to a desirable location z (e.g., for its WFH amenities), the resulting increase in demand for housing drives up the equilibrium rent r_z, which in turn dampens the location's attractiveness—a crucial feedback mechanism for achieving a general equilibrium.

Stochastic Location Choice To capture the heterogeneity observed in the data I must incorporate unobserved idiosyncratic factors. For this reason I introduce idiosyncratic preference shocks for each location. A household's perceived utility for location z is now given by:

$$\tilde{V}_z(k, a, \iota) = V_z(k, a, \iota) + \epsilon_z, \tag{16}$$

where ϵ_z is an i.i.d. taste shock drawn from a Type-I Extreme Value (Gumbel) distribution with scale parameter $\mu > 0$ (McFadden, 1981). $V_z(k, a, \iota)$ represents the systematic, observable utility a type-(k, a, ι) household derives from location z, based on its amenities, rents, and the household's characteristics. $\epsilon_{\mathbf{z}}$ captures all the unobserved, household-specific factors that influence its location preference (e.g., sentimental attachment to a neighborhood, the specific architecture of a home, proximity to a particular friend). The scale parameter $\mu > \mathbf{0}$ governs the variance of the idiosyncratic shocks and thus the relative importance of unobserved heterogeneity versus observed utility in the household's decision. As $\mu \to 0$, the variance of the shocks shrinks to zero. The model converges to the deterministic case, where households choose the location with the highest V_z with perfect certainty, resulting in perfect sorting. As μ increases, idiosyncratic preferences dominate the decision. Location choice becomes more random from the perspective of the model, and sorting patterns become smoother, mimicking the real-world mixing of different household types across communities.

Under this assumption, the probability that a type-(k, a, ι) household chooses location z is given by the standard logit formula:

$$P_z(k, a, \iota) = \frac{\exp\left(\mu^{-1} V_z(k, a, \iota)\right)}{\sum_{z' \in C,S} \exp\left(\mu^{-1} V_{z'}(k, a, \iota)\right)}. \tag{17}$$

This formulation completes the microfoundations of the location model. The combination of systematic utility V_z (driven by WFH amenities, rents, and wages) and stochastic preferences ϵ_z allows us to generate a quantitatively relevant equilibrium where: (i) location aggregates (e.g., the total population of each community) are determined by the systematic utility differences and housing market clearing. (ii) individual choices retain an element of randomness, ensuring that communities are composed of a heterogeneous mix of households.

3.3.1 Stochastic General Equilibrium Definition

Let the fundamental parameters of the economy be

$$\Theta = (t_C, t_S, \xi_C, \xi_S, \eta, \sigma, \gamma, \psi, \beta, \nu, \mu_a, \rho_a, \sigma_a, A^O, A^H, \theta, \rho, h_{\min}, \mu)$$

and let $F(k, a, \iota)$ be the joint distribution of adult types.

Definition 3.2 (Stochastic Residential General Equilibrium)**.** A **stochastic residential general equilibrium** is a set of: housing rents (r_C, r_S), neighborhood peer spillovers (S_C, S_S), Location choice probabilities $P_z(k, a, \iota)$ for all types, and household policy functions $x_z(k, a, \iota)$ for consumption, housing, labor, educational time, and leisure in each location, such that:

i. **Household Optimization:** Given prices (r_z) and spillovers (S_z), the policy functions $x_z(k, a, \iota)$ solve the household's utility maximization problem within each location. Location choices are described by the probabilistic choice rule in Equation (17).

ii. **Housing Market Clearing:** In each location z, the total demand for housing equals the supply provided by profit-maximizing developers:

$$H_z^D = \int P_z(k, a, \iota) \cdot h_z(k, a, \iota) dF = \left(\frac{r_z}{\xi_z}\right)^{1/\eta} = H_z^S. \tag{18}$$

iii. **Peer Spillover Consistency:** The peer spillover in each location is the probability-weighted average of the expected human capital of its children:

$$S_z = \frac{1}{\mu_z} \int P_z(k, a, \iota) \cdot \mathbb{E}[k_c \mid a, z] dF, \quad \text{where } \mu_z = \int P_z(k, a, \iota) dF. \tag{19}$$

iv. **Labor Market:** The wage for the numeraire good is pinned down at $w = 1$.

This general equilibrium definition extends the partial equilibrium case along two critical dimensions. First, housing rents (r_C, r_S) are no longer fixed but are determined endogenously by the clearing of local housing markets, which feature upward-sloping supply. Second, the deterministic location choice rule is replaced by a probabilistic one, where the probability a type chooses a location is given by Equation 17. The combination of endogenous prices and stochastic choice creates a rich framework for quantitative analysis, capturing the feedback loops between household sorting, housing prices, and neighborhood quality.

4 Calibration and Steady-State Analysis

This section details the quantitative strategy used to measure the causal effect of the work-from-home (WFH) shock on inequality and intergenerational mobility. The core idea is to use a structural model to isolate the WFH mechanism from other confounding trends.

My strategy is to first calibrate the model's structural parameters to a pre-shock steady state, which I assume is 2013, a period with low and stable WFH prevalence. I use a Simulated Method of Moments (SMM) estimator to match a set of cross-sectional and intergenerational moments from this 2013 economy. I then introduce an exogenous shock to WFH availability to match its post-covid level, while holding all the calibrated structural parameters fixed. I solve for the new, post-shock steady state and compare its outcomes to 2023 data as an untargeted validation exercise to show the model is producing meaningful results.

Data and Synthetic Metro Construction I construct a "synthetic city" to analyze spatial inequalities and work-from-home (WFH) patterns, following the methodology of Davis et al. (2024). The sample is drawn from the American Community Survey (ACS) and focuses on working-age parents with children aged 5 to 15. I pool the 30 most populous U.S. Core-Based Statistical Areas (CBSAs) and delineate each into two distinct zones: a *Central Business District (CBD)*, defined by the central county, and a *Suburb*, comprising all outlying counties within the CBSA. To maintain a monocentric structure, non-monocentric metros are excluded. This zoning rule ensures that the synthetic two-zone city is representative and comparable to the underlying microdata. Summary statistics of zone-level variables are available in Appendix **??**.

4.1 Calibration Strategy

The estimation and analysis proceed in three main stages, designed to respect causal sequencing. This approach identifies the model's stable, structural parameters (Θ) from the pre-shock 2013 environment. The subsequent simulation of the Post-Covid shock provides a clean, out-of-sample test of the model's mechanisms.

Step 1: The WFH Shock (External). The WFH state ($\iota = 1$) is not randomly assigned. Its availability is positively correlated with human capital (k), which I proxy with years of education in the data. To capture this selection, I model the probability of being WFH-

eligible using a probit model:

$$\Pr(\iota = 1 \mid k) = \Phi(\beta_0 + \beta_h k)$$

where $\Phi(\cdot)$ is the standard normal CDF. The parameters (β_0, β_h) are set externally (i.e., not estimated in the main SMM loop) for each period. They are chosen via Method of Moments (MoM) to exactly match two targets: (1) the aggregate WFH share inthe economy and (2) the correlation $\text{corr}(\iota = 1, k)$. This calibration defines the exogenous pre-COVID (2013) environment and the post-COVID shock, for which I use 2023 data, as summarized in Table 7. The algorithm used for the Method of Moments estimation is reported in Appendix **??**.

Table 7: Exogenous WFH Shock: Availability and Correlation

Period	**WFH Availability**	**corr**$(\iota = 1, k)$
Pre-COVID (2013)	3.94%	0.0700
Post-COVID (2023)	17.4%	0.1657

Notes: Data from ACS. These moments are matched externally to set the WFH availability parameters (β_0, β_h) for the 2013 and 2023 steady states.

Step 2: 2013 Steady-State Calibration (SMM). I estimate the 2013 structural parameters by Simulated Method of Moments (SMM). Let Θ collect the 12 structural parameters

$$\Theta = \{\mu_{k,\log}, \sigma_a, \sigma_{k,\log}, \rho_a, \xi_C, \xi_S, \text{commuting_multiplier}, \sigma_\varepsilon, A_H, \nu, \psi, \beta\},$$

where (ξ_C, ξ_S) are housing-supply cost shifters; commuting multiplier governs the gap between the suburb and CBD exogenous commuting time; σ_ε is location-taste dispersion; A_H is general TFP productivity for teleworked hours; ν is the neighborhood-spillover elasticity in child human capital; ψ is the weight on leisure utility; β is the parental altruism parameter; and $(\mu_{k,\log}, \sigma_{k,\log}, \rho_a, \sigma_a)$ govern the initial cross-section and transmission of k (human capital) and a (ability).

For any candidate Θ, I solve the 2013 general equilibrium (using the 2013 WFH availability from Table 7), compute the same statistic definitions as in the data, and minimize a weighted quadratic loss. Formally, I compute:

$$\min_{\Theta} \sum_{m \in \mathcal{M}_{13}} w_m \big(m_{13}^{\text{model}}(\Theta) - m_{13}^{\text{data}}\big)^2$$

with $\mathcal{M}_{13}$ including 11 moments related to economy-wide inequality (Gini, $p90/p10$, $p75/p25$), spatial sorting (suburban population share, R_S/R_C, CBD/Suburb commute ratio), time use

(ATUS education shares, WFH hours share, work hours), and intergenerational mobility (rank-rank slope, exposure gap).

Step 3: Post-COVID Shock and Validation. After the SMM calibration is complete, I validate it against a rich set of untargeted moments for the Pre-Covid (2013) economy. I then hold the calibrated structural parameters Θ fixed and conduct the counterfactual experiments. First, I introduce the 2023 WFH shock by fitting new availability parameters to exogenously match the 2023 WFH share (17.4%) and WFH-k correlation (0.1657) from Table 7. I then solve for the new, post-COVID steady state using the *original* calibrated Θ and the *new* WFH availability. Finally, to show that the steady state reached for the Post-Covid economy is not unrealistic for what concern the Pre-Covid targets, I compare the model's estimates with 2023 data.

4.2 Identification and Estimation

This subsection details how the key primitives governing human-capital formation and time allocation are pinned down, how the remaining parameters are calibrated or estimated, and how my estimates relate to the existing literature. Children's human capital is produced according to

$$k_c \;=\; a_c\, S_z^{\nu}\, e^{\,1-\nu},$$

where S_z is the zone-specific schooling/spillover index, e is parental education time, and $\nu \in (0,1)$ is the elasticity of child human capital with respect to local spillovers.

Identifying ν I exploit the well-known childhood exposure gradient in Chetty and Hendren (2018), who document that each year of childhood spent in a higher-opportunity area shifts adult outcomes toward those of permanent residents by about four percentile points. I map this statistic into the model by equating the per-period change in adult rank generated by the S-channel to the empirical exposure effect. Let ζ denote the slope of adult income rank with respect to $\ln k$ in the pre-covid steady state (this is measured internally from the baseline SS and hence provides the scale that converts human-capital gains into rank gains), and let $\Delta \ln S = \ln S_{\text{Suburb}} - \ln S_{\text{CBD}}$ and $\Delta \bar{r}^{\text{perm}}$ denote, respectively, the log-gap in S and the mean rank gap between permanent residents of the two zones. The identifying moment is

$$\zeta\, \nu\, \Delta \ln S \;=\; 0.04 \times T \times \Delta \bar{r}^{\text{perm}}, \tag{20}$$

with T set to a complete childhood ($T = 20$). Operationally, I include the SMM restriction $m_{\text{exp}}(\Theta) := \zeta\,\nu\,\Delta \ln S - 0.04T\Delta\bar{r}^{\text{perm}} = 0$ and back out ν given the measured ζ in the pre-Covid steady state. Intuitively, ζ translates model-implied human-capital gains into adult rank movements, while $\Delta \ln S$ loads the spatial gradient; the exposure slope then pins down the product $\zeta\nu$ and hence ν.

Identifying β from time-use. Conditional on ν, I target the parental education-time share using ATUS 2013. Let

$$\bar{s}_e \;=\; \frac{\mathbb{E}[e]}{\mathbb{E}[e] + \mathbb{E}[\ell] + \mathbb{E}[n]},$$

where ℓ and n denote leisure and market work. I construct $\bar{s}_e^{\text{data}}$ from ATUS microdata using education-related activity codes and choose β so that the 2013 steady state matches $\bar{s}_e^{\text{data}}$ in SMM. This pins down the trade-off between earnings/leisure and child investment in a way that is tightly connected to observed time allocation prior to Covid.

Remaining parameters. Model parameters are disciplined through a combination of external calibration, drawing on established literature, and internal calibration to match key empirical moments. Table 8 reports the full mapping. Productivity normalizations and spatial primitives are set as follows: office productivity, A_O is normalized to 1, while the telework productivity scale, A_H, is internally calibrated to match the share of work-from-home (WFH) hours in the ATUS. The resulting estimate of $A_H \approx 0.82$ aligns closely with the home-productivity factor of 0.81 reported by Gibbs et al. (2023), despite differing methodologies. The elasticity of substitution between office and home labor (ρ) is taken from the spatial-WFH literature (Delventhal and Parkhomenko, 2024). For human capital and spillovers, the spillover elasticity (ν) is identified via an exposure-based moment (Equation (20)). My estimate is higher than the ≈ 0.198 found by Zheng and Graham (2022) because the sufficient-statistic index S_z bundles school quality, peer effects, and neighborhood externalities.

The remaining parameters are disciplined by standard spatial and inequality moments. Housing supply parameters (ξ_C, ξ_S) are fit to the suburban-CBD rent ratio (ACS), while location preference shock (σ_ε) is pinned down by the suburban population share. The initial cross-sectional distribution of human capital (k) and innate ability (a) is calibrated to match the interquartile ratio, the p90/p10 income ratio, and the Gini coefficient. Finally, preference parameters for child utility (β) and leisure (ψ) are identified using time-use shares from the ATUS.

Table 8: Model Parameters: Identification and Calibration

Parameter	Value	Economic Interpretation	Identification / Source
Externally Calibrated (Literature)			
σ	2.00	Coefficient of relative risk aversion	Macro/labour literature
γ	0.76	Weight on consumption in utility	Davis and Ortalo-Magné (2011)
θ	0.82	Labor share in production	Valentinyi and Herrendorf (2008)
ρ	0.772	CES: office vs. home labor	Delventhal and Parkhomenko (2024)
A_O	1.00	Office productivity scale	Normalization
$h_{\min}$	0.10	Min. housing for home-office productivity	Richard (2024)
η	1.5	Inverse Housing Supply Elasticity	Saiz (2010)
Internally Calibrated (Model Fit)			
ξ_C, ξ_S	1.59, 1.22	Housing supply scales	Rent ratio R_S/R_C (ACS)
Commute mult.	2.00	CBD commute disutility	CBD/Suburb commute (ACS)
A_H	0.82	Home-work TFP scale	WFH hours share (ATUS)
ν	0.6000	Spillover elasticity in S_z	Exposure mapping in (20)
σ_ε	0.18	Location-preference dispersion	Suburban share (ACS)
$\mu_{k,\log}$	0.58	Mean of $\log k$	Interquartile ratio
σ_a	1.07	SD of ability a	Gini (income)
$\sigma_{k,\log}$	1.23	SD of $\log k$	p90/p10 (income)
ρ_a	0.0131	Intergen. ability persistence	Rank–rank slope
β	0.24	Weight on k_c in utility	ATUS education-time share
ψ	1.18	Weight on leisure	ATUS work/leisure shares

4.3 Pre-Covid Fit and Post-Covid Validation

This subsection evaluates the model in three steps. First, I assess the targeted fit in the pre-Covid steady state to verify that the core aggregates are disciplined. Second, I examine untargeted, disaggregated moments from the same steady state to test the internal mechanisms of spatial sorting and within-zone inequality. Finally, I validate the model's transition dynamics by comparing the shift to the post-Covid steady state against observed post-pandemic trends.

Targeted Pre-Covid Fit Table 9 reports the targeted pre-Covid fit from the SMM estimation. The model reproduces the overall shape of the 2013 income distribution closely. The interquartile ratio is essentially exact and the p90/p10 ratio deviates by only 0.03, indicating that the tails are well matched by the calibrated heterogeneity and the home–office technology. Spatial prices are also well aligned: the suburb-to-CBD rent ratio is near the ACS target and the CBD/Suburb commute ratio sits in a plausible range, suggesting the spatial budget and commuting blocks are well pinned down.

Two gaps remain and point in the same direction. The suburban population share is too low, and the share of WFH hours among WFH workers falls short of the ATUS target; raising suburban amenities and housing capacity would lift both margins, whereas the current calibration prioritizes matching inequality and prices. The rank–rank slope is below the Chetty et al. (2014) benchmark, implying slightly more mobility than in the data. The exposure moment based on Chetty and Hendren (2018) is near zero by construction and is matched tightly.

Table 9: Model Fit for Targeted Moments (Pre-Covid Steady-State)

Moment	**Data**	**Model**	**Abs. Error**
Inequality(ACS)			
Gini (income, all)	0.3588	0.3986	0.0398
p90/p10 (income, all)	7.0500	7.0758	0.0258
p75/p25 (income, all)	2.6792	2.6751	0.0041
Spatial Structure (ACS)			
Population Share, Suburb	0.7310	0.5973	0.1337
Rent Ratio (r_S/r_C)	0.8867	0.8967	0.0100
CBD/Suburb Commute Ratio	0.5100	0.5343	0.0243
Time Use & WFH (ATUS)			
Work Hours (% total time)	0.6500	0.6763	0.0263
WFH Hours Share (among WFH)	0.5500	0.4948	0.0552
Child Educ. Hrs%	0.0160	0.0174	0.0014
Intergenerational (Chetty et al., 2014; Chetty and Hendren, 2018)			
Rank-Rank Slope	0.3410	0.2652	0.0758
Exposure Gap	0.0000	0.0014	0.0014

Notes: This table reports the final fit of the SMM estimation. The model is evaluated at the estimated parameter vector. The exposure gap and the rank-rank slope estimation are based on Chetty and Hendren (2018) and Chetty et al. (2014) respectively.

Untargeted Pre-Covid Validation A stricter, untargeted test of the model's internal mechanisms appears in Table 10. Without targeting within-zone inequality or cross-zone income ratios, the model correctly generates richer suburbs than the CBD and reproduces the ordering and magnitudes of within-zone dispersion. The suburb-to-CBD mean income ratio is very close to the data, which validates the sorting mechanism over amenities, commuting,

and housing. The CBD is more unequal than the suburb on every metric, as in the data; interquartile ratios are particularly tight, while Ginis and the Theil index are somewhat high, mirroring the slightly elevated aggregate Gini in the targeted table.

Conditional on WFH status, dispersion is more mixed: group-specific Gini coefficients are higher while p90/p10 ratios are lower than in the data, suggesting residual within-occupation dispersion that is stronger than needed relative to tail separation.

Table 10: Pre-Covid Untargeted Spatial and Inequality Moments (Model vs Data)

Untargeted Moment (2013)	Data	Model	Δ (M−D)
A. Spatial Income Sorting			
Mean income ratio (S/C)	1.0660	1.0530	−0.0130
B. Within-Zone Inequality (Gini)			
Gini (income, CBD)	0.3541	0.3708	0.0167
Gini (income, Suburb)	0.3421	0.3664	0.0243
C. Within-Zone Inequality (p90/p10)			
p90/p10 ratio (CBD)	6.9440	6.6660	−0.2780
p90/p10 ratio (Suburb)	6.7450	6.5740	−0.1710
D. Within-Zone Inequality (p75/p25)			
p75/p25 ratio (CBD)	2.6450	2.5970	−0.0480
p75/p25 ratio (Suburb)	2.6280	2.5690	−0.0590
E. WFH Inequality			
Gini (income, WFH)	0.3260	0.3708	0.0440
Gini (income, NonWFH)	0.3421	0.3855	0.0430
p90/p10 ratio (WFH)	6.3380	6.1650	−0.1730
p90/p10 ratio (NonWFH)	7.0000	6.5360	−0.4640
F. General Entropy			
Theil–T index	0.2070	0.2460	0.0390

Visual validations complement these tables. Figure 8 compares the model's pre-Covid quintile transition matrix to the Chetty et al. (2014) reference. The diagonal is strong in both, with the thickest mass in the tails, and the off-diagonal leakage is of the right order across all parent quintiles. The model is more mobile in the middle of the distribution, consistent with the underprediction of the rank–rank slope in Table 9.

Post-Covid Validation To validate the transition logic, Table 11 compares model pre- and post-Covid steady states and overlays 2023 data when available. The model produces a large increase in WFH intensity—roughly a fourfold rise in the WFH-hours share—which is in the right order of magnitude relative to post-pandemic evidence. Figure 9 shows that the model also captures the steepening gradient of WFH adoption across income quintiles, which is crucial for understanding the distributional impact of the shock.

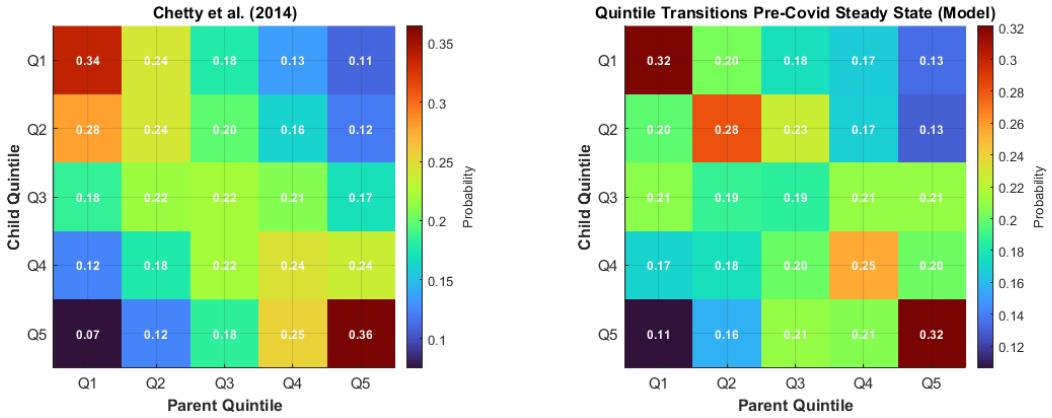


Figure 8: Intergenerational quintile transitions: Estimates of Chetty et al. (2014) versus model pre-Covid.

Inequality aggregates remain broadly flat across steady states, consistent with the view that the pandemic did not immediately overturn the income distribution. Spatial adjustment moves in the expected qualitative direction—a slightly higher CBD population share and a higher suburb-to-CBD rent ratio—although the empirical rent differential widens more than in the model. The rank–rank slope falls modestly, indicating slightly higher mobility in the post-Covid steady state; given the baseline slope is already somewhat low, I interpret this change cautiously.

Table 11: Model vs Data: Pre- to Post-Covid Comparison

Moment	**Model Pre-Covid**	**Model Post-Covid**	**Data 2023**	**Pre vs Post Change**
WFH-hours share	0.017	0.076	0.100	+343%
Gini coefficient	0.386	0.386	0.360	+0.1%
p90/p10 ratio	6.54	6.52	7.07	-0.3%
CBD population share	0.402	0.404	0.316	+0.5%
Rent ratio (Suburb/CBD)	0.856	0.861	1.160	+0.6%
Spillover ratio (S/C)	1.05	1.07	—	+1.9%
Rank–rank slope	0.243	0.239	—	-1.8%

Notes: The table reports model steady states before and after the Covid shock and compares to available 2023 aggregates where possible. The last column shows the model's relative change from pre- to post-Covid.

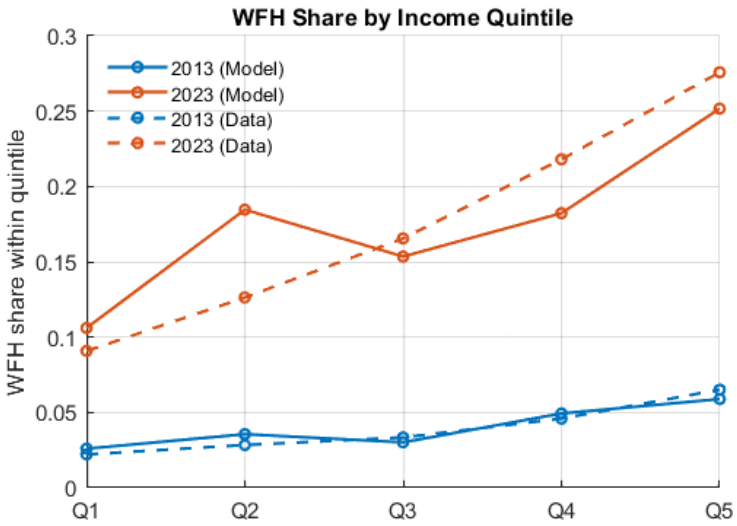


Figure 9: WFH share by income quintile in 2013 and 2023: model and data.

Summary of Model Performance In sum, the model is well-disciplined on the core features relevant for this study's mechanisms. It matches the overall shape of inequality, replicates the geography of income sorting without direct targeting, and generates a qualitatively realistic response to the WFH shock. In the quantitative exercises that follow, I therefore emphasize and interpret results through objects that the model reproduces most credibly: distributional tail probabilities, spatial differences, and intergenerational transitions across quantiles.

4.4 Intergenerational Mobility: a Tale of Two Tiers

The post-Covid steady state delivers only a modest change in aggregate mobility when viewed through economy-wide summary measures, yet this apparent stability hides sharp heterogeneity by work modality and family background. Table **??** documents that the rank–rank slope moves little across steady states, while transition probabilities change only slightly in the aggregate. These averages, however, bundle together groups facing different opportunity sets and different childhood investments.

A useful way to see this is to separate (i) a *within-generation* tier driven by the child's own access to WFH and (ii) an *intergenerational* tier driven by the parent's WFH status, conditioning on the child's own work modality. Panel (A) of Figure 10 shows a clear first-order shift in the income distribution between WFH and non-WFH children: WFH children are more likely to reach the upper tail. In the post-Covid steady state, the probability of being in the top decile is 0.142 for WFH children versus 0.087 for non-WFH children, an absolute gain of 5.54 percentage points (a relative lift of about 64%). At the top-quintile threshold,

the corresponding probabilities are 0.271 versus 0.182, an 8.92 percentage point gain (about 49%). These gaps summarize the model's **direct (within-generation) WFH effect**: conditional on skill and location choices implied by equilibrium, WFH expands the feasible set over time and earnings and shifts children upward in the adult income distribution.

The second tier is intergenerational. Conditioning on *child non-WFH*, having a WFH parent raises top-tail attainment: Pr(Top10) increases from 0.084 to 0.099 (a 1.54 percentage point gain), and Pr(Top20) increases from 0.178 to 0.200 (a 2.24 percentage point gain). Conditioning instead on *child WFH*, parental WFH still matters: the top-decile probability increases from 0.139 to 0.159 (a 1.98 percentage point gain), and the top-quintile probability rises from 0.266 to 0.295 (a 2.95 percentage point gain). Panel (B) of Figure 10 visualizes the parental effect among child non-WFH, while Panel (C) shows the full four-way decomposition.

While these unconditional tail probabilities are informative, they also conflate causal mechanisms with selection on parental human capital. To address this directly, Figure 12 reports *conditional tail probabilities by parent human-capital bin*. Two patterns stand out. First, top-tail attainment rises steeply with parent k for both child types, indicating that composition is quantitatively important for the unconditional gaps. Second, *within a parent-k bin*, parental WFH is associated with a positive tail advantage for both child non-WFH and child WFH, but the magnitude varies across the distribution: the conditional gaps are small in the bottom bins and become economically meaningful in the upper bins (e.g. in the top parent-k bins, the conditional Top10 and Top20 gaps are on the order of 1–3 percentage points). This conditional perspective is the appropriate object for interpreting intergenerational "advantage" in the model, because it holds fixed the main driver of baseline earnings capacity.

Table 12 summarizes the core tail objects used throughout the paper: the within-generation WFH tier (Panel A), and the parental WFH tier conditioning on the child's own modality (Panel B). These are the key "paper numbers" that connect the model's equilibrium allocation to intergenerational outcomes.

Table 12: Probability of Top-Tail Success by WFH Status (Post-Covid Steady State)

	Probability		Absolute	Relative
	WFH Group	**Non-WFH Group**	**Difference (pp)**	**Difference (%)**
Panel A: Direct Effect (Child's WFH Status)				
P(Child in Top 10%)	0.1421	0.0867	+5.54	+63.9%
P(Child in Top 20%)	0.2708	0.1816	+8.92	+49.1%
Panel B: Intergenerational Effect (Parent's WFH Status)				
Children without WFH:				
P(Top 10%)	0.0994	0.0840	+1.54	+18.3%
P(Top 20%)	0.2002	0.1777	+2.24	+12.7%
Children with WFH:				
P(Top 10%)	0.1585	0.1387	+1.98	+14.3%
P(Top 20%)	0.2952	0.2657	+2.95	+11.1%

Notes: Values are computed in the post-Covid steady state. Panel A compares children by their own WFH status. Panel B conditions on the child's modality and compares outcomes by parental WFH status.

The rank-based view delivers a consistent message. Figure 11 plots expected child rank as a function of parent rank, separately for child non-WFH and child WFH, and within each panel compares parental WFH versus parental non-WFH. In both child groups, the parental-WFH curve lies above the parental-non-WFH curve throughout the parent-rank distribution, implying a broadly positive intergenerational premium associated with parental WFH.

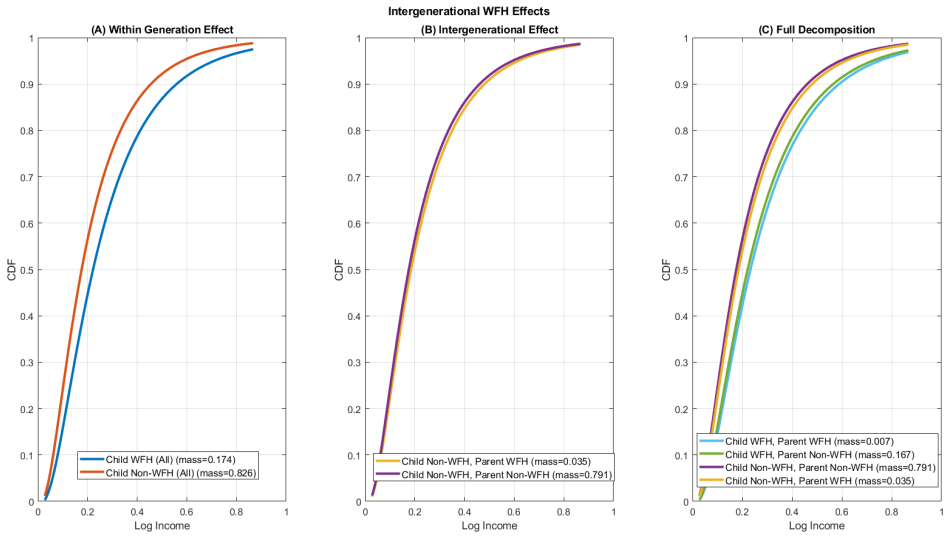


Figure 10: Within-generation and intergenerational WFH effects on child income distributions. Panel (A) compares WFH and non-WFH children. Panel (B) compares non-WFH children by parental WFH status. Panel (C) shows all four groups jointly.

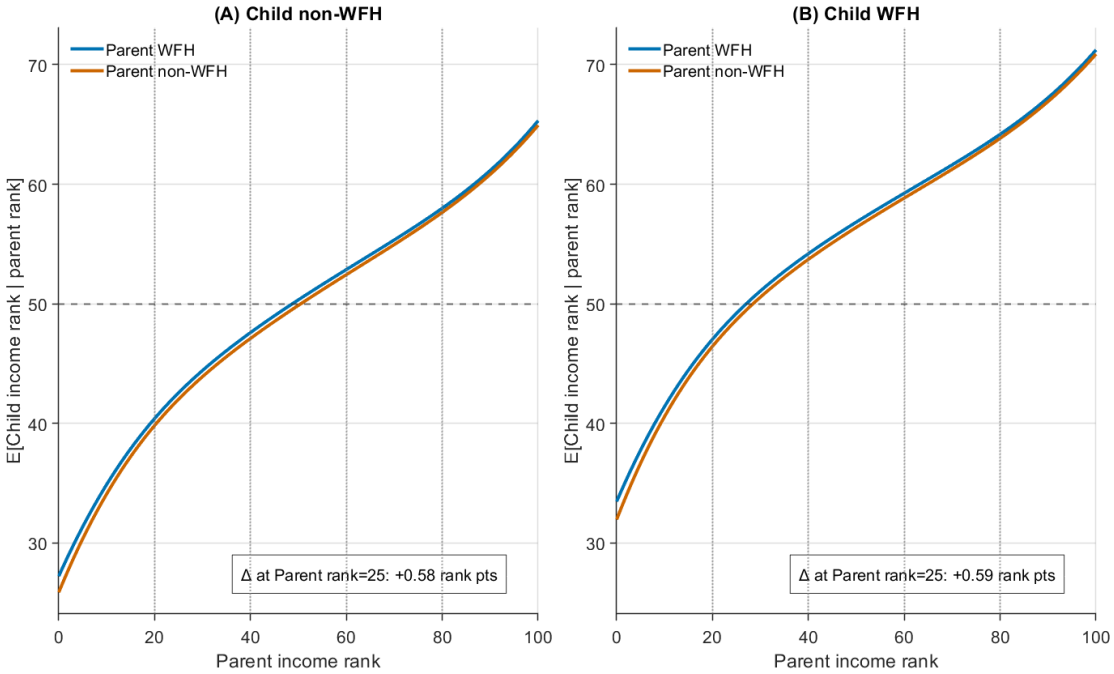


Figure 11: Rank–rank profiles by child modality. Each panel plots $\mathbb{E}$[child rank | parent rank] and compares parental WFH versus parental non-WFH.

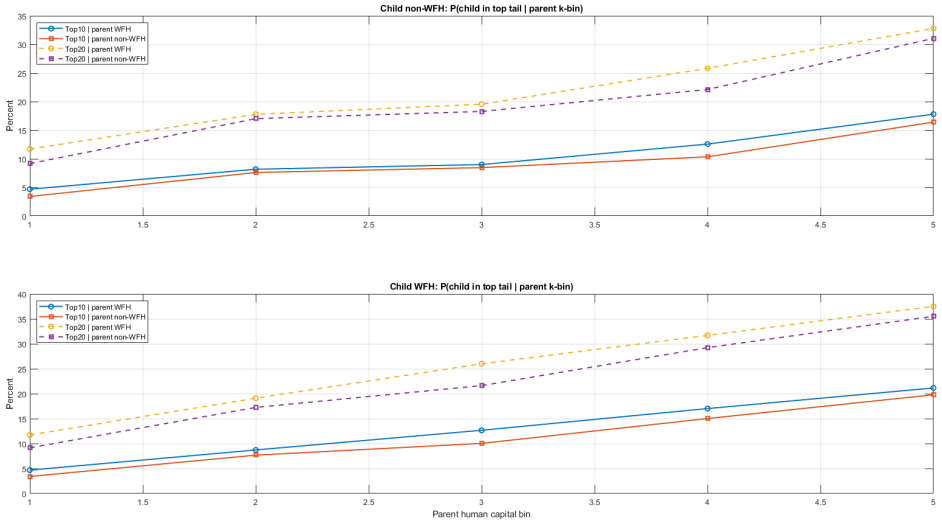


Figure 12: Conditional tail probabilities by parent human-capital bin. The figure reports Pr(Top10) and Pr(Top20) within parent-k bins, separately for child non-WFH and child WFH, and compares parental WFH versus parental non-WFH within bin.

Finally, the spatial objects are informative about the *sorting* margin, but in this calibration they do not mechanically reinforce the child-level WFH advantage. In the post-Covid steady state, the probability of residing in the "high-S" zone is 0.317 for WFH children versus 0.389 for non-WFH children (Table 13), a negative gap of 7.16 percentage points. Conditional on parental status, the high-S residence probabilities are similar (0.385 for WFH

parents versus 0.376 for non-WFH parents). Thus, in this two-zone calibration, spatial allocation does not account for the intergenerational gaps documented above; instead, the most salient patterns operate within zones and within parent-k strata.

Table 13: Spatial sorting and segregation (post-Covid steady state)

Group / metric	Value	Definition
Residence in the high-spillover zone ($z = \text{High } S_z$)		
Child WFH	0.317	$\Pr(z = \text{High} \mid \text{child WFH})$
Child non-WFH	0.389	$\Pr(z = \text{High} \mid \text{child non-WFH})$
Given parent WFH	0.385	$\Pr(z = \text{High} \mid \text{parent WFH})$
Given parent non-WFH	0.376	$\Pr(z = \text{High} \mid \text{parent non-WFH})$
Segregation		
Dissimilarity (WFH vs non-WFH)	0.072	Two-group dissimilarity across zones

Notes: Values are computed on the child distribution in the post-Covid steady state. The "high-S zone" corresponds to the zone with the larger equilibrium spillover index.

Why, then, do aggregate mobility statistics barely move? The answer is composition and offsetting margins. The post-Covid economy features a much larger WFH share, and WFH children enjoy substantially higher top-tail success. At the same time, conditional transition patterns vary across parent-k strata, and the reweighting of mass across child/parent modality groups can leave rank-based aggregates nearly unchanged. The conditional objects in Figure 12 make this explicit: the economy-wide tail numbers reflect both within-bin gaps and where the post-Covid mass sits across parent-k bins.

There are reasons to treat the magnitude of the within-generation effect with caution. In the model, earnings are tightly linked to skill, location, and the home–office technology, so WFH can operate as a relatively clean time-to-output margin. In reality, idiosyncratic earnings risk, firm-side constraints, task indivisibilities, and hours frictions may dampen the translation of saved commuting time into earnings. These forces would compress the within-generation wedge without overturning its sign.

4.5 Counterfactuals: The Source of the Parental Advantage

The analysis in the previous section isolated a clear intergenerational force: a distinct, positive outcome gap for children of WFH-eligible parents, even for those children who do not work from home themselves. This raises a critical question: what is the source of this advantage? The model proposes two primary mechanisms: first, the parental time channel, whereby WFH parents reallocate saved commute time toward their children's education , and second, the residential sorting channel, whereby WFH flexibility allows families to move to higher-quality neighborhoods with stronger peer spillovers.

To disentangle these two forces, I conduct two counterfactual experiments starting from the Post-Covid steady state. First, I neutralize the parental time channel. Second, I neutralize the neighborhood spillover channel.

In the first experiment, which I call the "No Education Channel" counterfactual, I computationally shut down the time investment differential. I force WFH-eligible parents to choose the exact same amount of education time as their non-WFH counterparts who share the same dyad of human capital and ability (a, k). This isolates the sorting channel, leaving it intact while removing any extra time investment. The result, shown in Figure 13, is striking and immediate. The clear stochastic dominance that defined the parental WFH advantage in Figure 10 Panel (B) completely collapses. The two CDFs—for children of WFH and non-WFH parents—are now nearly indistinguishable, lying almost perfectly on top of one another.

The quantitative results in Table 14 confirm this visual collapse. In the baseline Post-Covid steady state, non-WFH children with WFH parents had a +1.05 percentage point advantage in reaching the top income decile over their peers with non-WFH parents. In the "No Education Channel" counterfactual, this gap shrinks to just +0.05 percentage points, a reduction of over 95%. The advantage in reaching the top quintile similarly evaporates.

In the second experiment, the "No Spillover Channel" counterfactual, I do the opposite. I leave the parental time channel fully active but neutralize all benefits from sorting. I achieve this by equalizing the neighborhood spillovers in both the CBD and the Suburb, setting them to the average level from the Pre-Covid steady state. As shown in Table 14, this has almost no effect on the intergenerational gap. The advantage for top-decile attainment remains at +1.01 percentage points, nearly identical to the +1.05 percentage point gap in the baseline Post-Covid steady state.

The decomposition, therefore, points unambiguously to one mechanism. The intergenerational advantage of WFH is not driven by where WFH parents live, but by what they do while they are at home. The ability to reallocate time toward child education is the dominant force. The residential sorting channel, while a prominent feature of the WFH discussion, appears to play a secondary and minor role in explaining the transmission of advantage to the next generation in this framework.

Table 14: Decomposing the Intergenerational Effect (Counterfactuals)

	P(Child in Top 10%) *among non-WFH Children*		
	(1) Baseline Post-Covid	(2) Counterfactual 1: No Education Channel	(3) Counterfactual 2: No Spillover Channel
Child of WFH Parent	0.0965	0.0865	0.0958
Child of Non-WFH Parent	0.0860	0.0860	0.0857
Intergenerational Gap (pp)	**+1.05 pp**	**+0.05 pp**	**+1.01 pp**

Notes: This table isolates the intergenerational effect seen in Panel B of Table 12. Column (1) shows the baseline gap. Column (2) shows the gap when the parental time channel is shut down. Column (3) shows the gap when the neighborhood spillover channel is shut down.

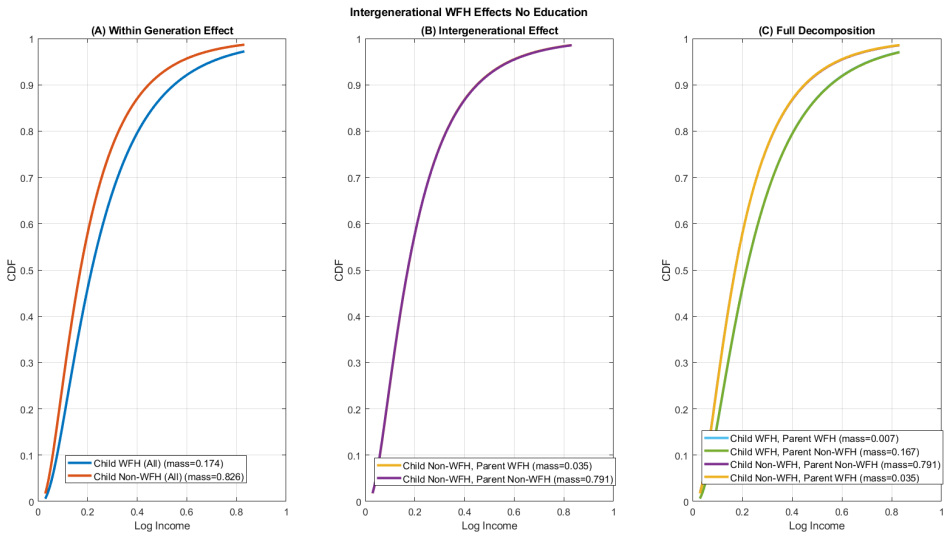


Figure 13: Counterfactual 1: No Education Time Channel

Notes: The graph shows the child income rank CDFs in a counterfactual where WFH parents cannot invest more time in education than non-WFH parents. The stochastic dominance seen in the baseline (Figure 10, Panel B) is eliminated.

The stark result that the neighborhood sorting channel plays a negligible role must be interpreted within the specific spatial structure of the model. My calibration, based on a two-zone city (CBD and Suburb), results in equilibrium peer spillovers that are relatively similar[11]. This limited spatial heterogeneity necessarily mutes the potential impact of residential choice. In reality, metropolitan areas are characterized by a fine-grained mosaic of neighborhoods with vastly different school qualities and peer effects. Consequently, the counterfactual analysis likely provides a lower-bound estimate for the importance of the sorting channel. In a model with more spatially differentiated neighborhoods, the ability to use WFH flexibility to access exclusive, high-spillover districts would almost certainly play

[11]For instance, for the Post Covid steady state, the Spillovers were [0.615, 0.657] for CBD and Suburb respectively.

a larger role in transmitting advantage. My finding that the time-investment channel is dominant is therefore most accurately stated as: within a metropolitan structure where the available residential options offer relatively comparable spillovers, what parents do at home is a more powerful margin of adjustment than where they choose to live.

Sensitivity. The suite of sensitivity analyses conducted in Appendices **??**, **??**, and **??** confirms the robustness of the central finding: the intergenerational transmission of the WFH advantage operates predominantly through the parental education time channel, not through residential sorting or neighborhood spillovers. Varying the education elasticity (ν) in the child human capital production function (Appendix **??**) shows that the magnitude of the intergenerational gap is directly tied to the productivity of parental time investments. When parental time is highly productive ($\nu = 0.2$), the gap is large (+3.23 pp), but it shrinks as the return to education time diminishes, becoming negligible ($\nu = 0.9$, +0.18 pp). Crucially, in the "No Education" counterfactual, this gap collapses to zero for all values of ν, proving that differential time investment is the essential mechanism. Furthermore, artificially imposing large neighborhood spillover contrasts (Appendix **??**) reveals that even substantial geographic inequality (up to $\pm 30\%$) has a minimal effect on the gap, which only increases meaningfully under an extreme and unrealistic $\pm 40\%$ contrast. Finally, varying the WFH productivity parameter (A_H) in Appendix **??** demonstrates that the intergenerational advantage, while moderately amplified by higher remote work efficiency, remains positive across a wide range of values and is still entirely eliminated when the education channel is shut down. Collectively, these tests provide compelling evidence that the reallocation of parental time, not spatial sorting, is the primary engine of intergenerational advantage in the age of remote work.

5 Conclusions

This paper examines how the rise of work from home (WFH), a labor market transformation disproportionately concentrated among high-income and highly educated workers, affects intergenerational inequality in the United States. Using multiple microdata sources and a quantitative spatial model, I identify two main mechanisms through which parental WFH influences children's long-run outcomes: a parental time channel and a residential sorting channel.

Empirically, I begin by examining how WFH alters parental education time. Using ATUS time diaries and controlling for a rich set of demographic, occupational, and geographic

factors, I find that WFH parents devote 0.146 additional hours per week to children's education, a 23 percent increase relative to the 0.63 hour baseline for non-WFH parents. This pattern appears on both margins: WFH parents are 1.9 percentage points more likely to engage in any educational activity (extensive margin) and, conditional on engaging, spend 8.4 percent more time (intensive margin). These findings are robust across alternative specifications and consistent with the interpretation that reduced commuting time or increased schedule flexibility allows parents to reallocate time toward cognitively enriching interactions at home.

Next, I examine residential mobility using the American Community Survey (ACS) combined with standardized school test scores. Kolmogorov–Smirnov tests reveal a systematic rightward shift in the distribution of school quality among WFH families with school-aged children, indicating sorting into better school districts after the diffusion of remote work. A probit specification confirms this pattern quantitatively: conditional on demographics, education, and income, WFH households are 3.7 percentage points more likely to relocate to a neighborhood with higher-quality schools relative to non-WFH households.

To interpret these empirical findings and assess their long-run implications, I develop a spatial overlapping-generations model that embeds both mechanisms. In a partial equilibrium version, I characterize analytically the theoretical forces driving location and time allocation choices, showing that WFH acts as a time-augmenting technology that lowers the shadow cost of parental investment and weakens the spatial constraint of the commute. I then extend the model to a quantitative version calibrated to U.S. data on income distribution, telework prevalence, and time-use patterns.

The calibrated model uncovers a novel and quantitatively important intergenerational channel. Even among adults who do not work from home, those with a WFH parent are 1.05 percentage points more likely to reach the top income decile and 1.9 percentage points more likely to reach the top quintile, corresponding to relative increases of 12 and 10 percent, respectively. The model also generates a sizeable within-generation effect: WFH-eligible adults are about 8.4 percentage points more likely to reach the top income decile than their non-WFH counterparts. This estimate should, however, be interpreted with caution. In the model, WFH eligibility is treated as exogenous and wages are deterministic, so agents cannot self-select into remote work or face idiosyncratic earnings shocks. As a result, the within-generation advantage likely represents an upper bound on the true effect of WFH.

Finally, I conduct counterfactual experiments to isolate the mechanisms behind the intergenerational transmission of WFH advantage. When I eliminate the additional educational time available to WFH parents, the intergenerational advantage collapses by over 95 percent,

while neutralizing neighborhood spillovers has a negligible effect. These results suggest that WFH transmits advantage primarily through the reallocation of parental time toward human capital formation, rather than through residential sorting or neighborhood effects. The diffusion of WFH thus reshapes the geography of opportunity not only by changing where parents live, but more fundamentally by changing how they use their time at home.

Future Research This analysis isolates two mechanisms that appear first-order—parental time and neighborhood spillovers—while acknowledging that complementary forces likely operate as well. First, the model focuses on parental time and neighborhood spillovers as the primary mechanisms for human capital transmission. However, other intergenerational channels may be at play. For example, WFH might alter the visibility of work in the home, influencing children's "work ethic" or career aspirations. It could also change household consumption patterns, such as the ability to afford private tutoring or other educational resources, which are not explicitly modeled here. Second, a key simplification in this model is that WFH eligibility is an exogenous state correlated with human capital, following the approach in related literature (Davis et al., 2024; Richard, 2024). In reality, households and firms make complex, endogenous choices about remote work. Parents with a high preference for child-rearing may be more likely to select into WFH-eligible jobs. Endogenizing this choice would be a valuable extension to disentangle how much of the "parental time effect" is a causal result of WFH flexibility versus a selection effect of parents who prioritize time investment. Third, the model generates a very large direct (within-generation) effect of WFH. This is partly because the model translates saved commuting time directly into productive labor hours. As noted in Section 4, real-world earnings are subject to idiosyncratic risk, firm-side heterogeneity, and hours constraints that would likely dampen this effect. Incorporating these frictions would provide a more conservative estimate of the within-generation gains.

Policy Implications The findings have important policy implications. If the WFH-induced advantage operates primarily through parental time investment, as my experiment suggests, it means that inequality of opportunity could widen as advantaged parents gain a new, powerful tool to invest in their children. This "rug rat race" may accelerate, leaving children from non-WFH (and often less-advantaged) households further behind.

Policies aimed at mitigating this divergence could focus on compensating for this time gap. For instance, public investment in high-quality extra tutoring hours or enriched after-school programs for children from disadvantaged backgrounds could help level the playing field. Furthermore, while the sorting channel was found to be secondary in this framework, policies that reduce school quality gaps between neighborhoods remain crucial for addressing

the broader landscape of spatial inequality.

In conclusion, the WFH revolution is not just a spatial or labor market shock; it is an intergenerational one. By reallocating the crucial resource of parental time, it threatens to widen existing gaps in human capital accumulation, solidifying advantage at the top of the income distribution.

References

Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M., and Zarate, P. (2022). Working from home around the world. *Brookings Papers on Economic Activity*, 2022(2):281–360.

Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M., and Zarate, P. (2023). Time savings when working from home. In *AEA Papers and Proceedings*, volume 113, pages 597–603. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Attanasio, O., Meghir, C., and Nix, E. (2020). Human capital development and parental investment in india. *The review of economic studies*, 87(6):2511–2541.

Barrero, J. M., Bloom, N., and Davis, S. J. (2021). Why working from home will stick. Technical report, National Bureau of Economic Research.

Barrero, J. M., Bloom, N., and Davis, S. J. (2023). The evolution of work from home. *Journal of Economic Perspectives*, 37(4):23–49.

Benabou, R. (1994). Human capital, inequality, and growth: A local perspective. *European Economic Review*, 38(3-4):817–826.

Benabou, R. (1996). Equity and efficiency in human capital investment: the local connection. *The Review of Economic Studies*, 63(2):237–264.

Brown, J. E. and Dunn, P. K. (2011). Comparisons of tobit, linear, and poisson-gamma regression models: An application of time use data. *Sociological Methods & Research*, 40(3):511–535.

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

Chetty, R., Friedman, J. N., and Stepner, M. (2024). The economic impacts of covid-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics*, 139(2):829–889.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The quarterly journal of economics*, 133(3):1107–1162.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity. *The geography of*.

Chyn, E. and Daruich, D. (2022). An equilibrium analysis of the effects of neighborhood-based interventions on children. Technical report, National Bureau of Economic Research.

Cunha, F. and Heckman, J. (2007). The technology of skill formation. *American economic review*, 97(2):31–47.

Cunha, F. and Heckman, J. J. (2008). Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of human resources*, 43(4):738–782.

Davis, M. A., Ghent, A. C., and Gregory, J. (2024). The work-from-home technology boon and its consequences. *Review of Economic Studies*, 91(6):3362–3401.

Davis, M. A. and Ortalo-Magné, F. (2011). Household expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261.

De La Croix, D. and Doepke, M. (2003). Inequality and growth: why differential fertility matters. *American Economic Review*, 93(4):1091–1113.

Del Boca, D., Flinn, C., and Wiswall, M. (2014). Household choices and child development. *Review of Economic Studies*, 81(1):137–185.

Del Bono, E., Francesconi, M., Kelly, Y., and Sacker, A. (2016). Early maternal time investment and early child outcomes. *The Economic Journal*, 126(596):F96–F135.

Delventhal, M. and Parkhomenko, A. (2024). Spatial implications of telecommuting. *Available at SSRN 3746555*.

Dingel, J. I. and Neiman, B. (2020). How many jobs can be done at home? *Journal of public economics*, 189:104235.

Dotti Sani, G. M. and Treas, J. (2016). Educational gradients in parents' child-care time across countries, 1965–2012. *Journal of marriage and Family*, 78(4):1083–1096.

Durlauf, S. N. et al. (1993). *Neighborhood feedbacks, endogenous stratification, and income inequality*. Social Systems Research Institute, University of Wisconsin.

Elango, S., García, J. L., Heckman, J. J., and Hojman, A. (2015). Early childhood education. In *Economics of Means-Tested Transfer Programs in the United States, Volume 2*, pages 235–297. University of Chicago Press.

Ferreira, P. C., Monge-Naranjo, A., and Torres de Mello Pereira, L. (2016). Of cities and slums. *FRB St. Louis Working Paper*, (2016-22).

Fogli, A. and Guerrieri, V. (2019). The end of the american dream? inequality and segregation in us cities. Technical report, National Bureau of Economic Research.

Fogli, A., Guerrieri, V., Ponder, M., and Prato, M. (2023). Scaling up the american dream: A dynamic analysis. Technical report, Working paper.

Gibbs, M., Mengel, F., and Siemroth, C. (2023). Work from home and productivity: Evidence from personnel and analytics data on information technology professionals. *Journal of Political Economy Microeconomics*, 1(1):7–41.

Gokan, T., Kichko, S., Matheson, J. A., and Thisse, J.-F. (2022). How the rise of teleworking will reshape labor markets and cities. Technical report, CESifo Working Paper.

Hammer, B. (2012). *Statistical models for time use data: an application to housework and childcare activities using the Austrian time use surveys from 2008 and 1992*. na.

Hofferth, S. L., Flood, S. M., Sobek, M., and Backman, D. (2020). American time use survey data extract builder: Version 2.8. *(No Title)*.

Leuven, E. and Sianesi, B. (2018). Psmatch2: Stata module to perform full mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing.

Liu, S. and Su, Y. (2021). The impact of the covid-19 pandemic on the demand for density: Evidence from the us housing market. *Economics letters*, 207:110010.

McFadden, D. (1981). Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications*, 198272.

Pabilonia, S. W. and Vernon, V. (2023). Who is doing the chores and childcare in dual-earner couples during the covid-19 era of working from home? *Review of Economics of the Household*, 21(2):519–565.

Puccioni, J. (2015). Parents' conceptions of school readiness, transition practices, and children's academic achievement trajectories. *The Journal of Educational Research*, 108(2):130–147.

Ramani, A. and Bloom, N. (2021). The donut effect of covid-19 on cities. Technical report, National Bureau of Economic Research.

Ramey, G. and Ramey, V. A. (2009). The rug rat race. Technical report, National Bureau of Economic Research.

Restrepo, B. J. and Zeballos, E. (2022). Work from home and daily time allocations: evidence from the coronavirus pandemic. *Review of Economics of the Household*, 20(3):735–758.

Richard, M. (2024). The spatial and distributive implications of working-from-home: A general equilibrium model. Technical report, Tech. rep., Stanford Institute for Economic Policy Research.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.

Stewart, J. (2013). Tobit or not tobit? *Journal of economic and social measurement*, 38(3):263–290.

Valentinyi, A. and Herrendorf, B. (2008). Measuring factor income shares at the sectoral level. *Review of Economic Dynamics*, 11(4):820–835.

von Gaudecker, H.-M., Holler, R., Simon, L., and Zimpelmann, C. (2024). Can work from home help balance the parental division of labor? Technical report, ECONtribute Discussion Paper.

Yasenov, V. (2020). Who can work from home? Technical report, IZA Discussion Papers.

Zheng, A. and Graham, J. (2022). Public education inequality and intergenerational mobility. *American Economic Journal: Macroeconomics*, 14(3):250–282.

A Heterogeneity, Outlier Robustness, and LOYO

A.1 Heterogeneous Effects by Parent Education, Child Age, and Gender

Figure 14 compares the empirical distribution of positive education diary values to fitted Normal and Gamma distributions. The Gamma provides a substantially better fit to the right-skewed distribution of positive education times, motivating its use in the GLM intensive-margin specification (Section 2.1).

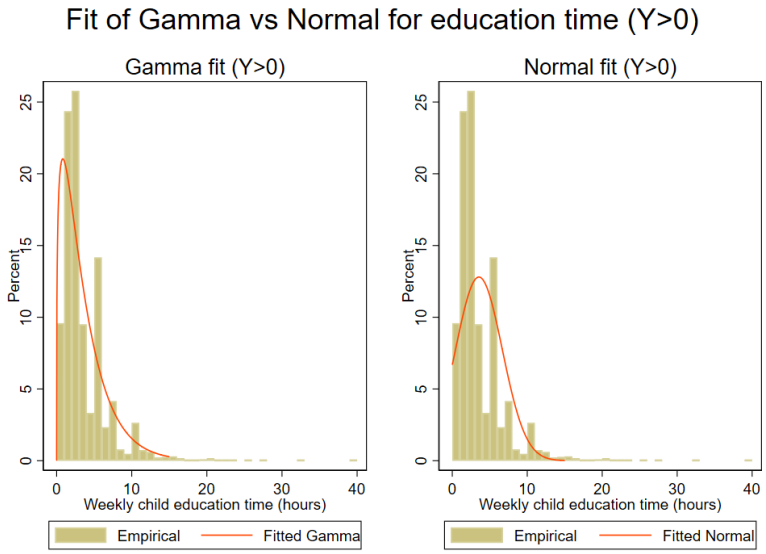


Figure 14: Empirical distribution of positive education diary values vs. fitted Normal and Gamma distributions. ATUS 2010–2023 (excl. 2020), conditional on positive education time.

Figure 15 plots estimated WFH coefficients from interactions of WFH with parent education level. The positive WFH–education association is present across all education groups. The association is slightly larger for parents with some college or a bachelor's degree relative to high school graduates, though the differences are not statistically significant. Advanced degree holders show a positive but smaller point estimate, consistent with the hypothesis that highly educated parents already dedicate substantial education time and face diminishing returns to additional flexibility.

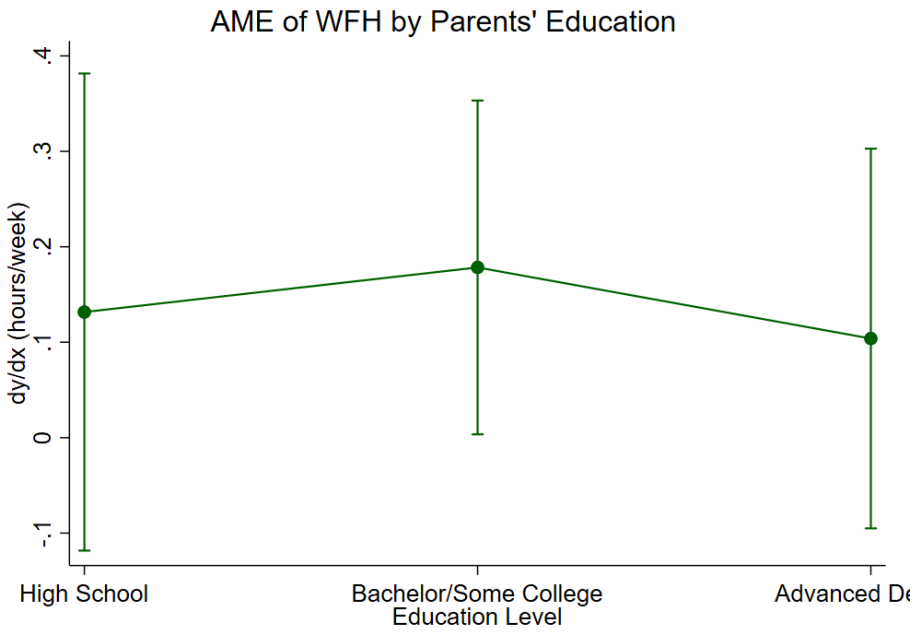


Figure 15: Heterogeneous WFH effects on education time by parent education level. Estimates from OLS interaction model. ATUS 2010–2023 (excl. 2020). Survey-weighted with 95% confidence intervals.

Figure 16 plots the WFH coefficient by child age group. The association is most pronounced for parents of school-age children (6–12 years), precisely the range when homework help, tutoring, and structured cognitive activities are most relevant. For parents of preschoolers (5 and under) the point estimate is positive but smaller; for teenagers (13–15) the WFH coefficient remains positive but with wider confidence intervals. This age gradient is consistent with the paper's theoretical mechanism: the educational value of parental time input is highest during the period of intensive skill formation associated with primary schooling.

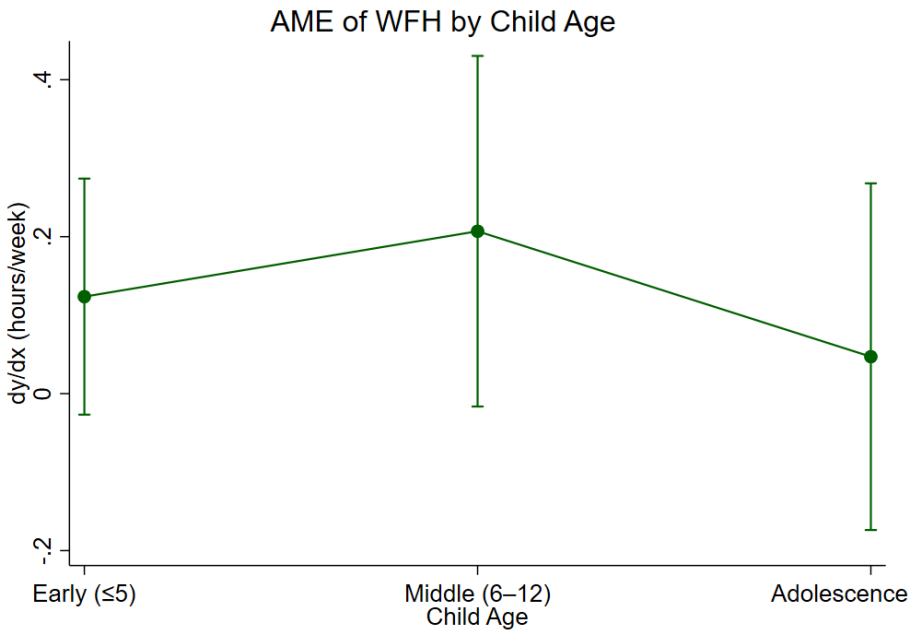


Figure 16: Heterogeneous WFH effects on education time by child age group. Estimates from OLS interaction model. Survey-weighted with 95% confidence intervals.

Figure 17 plots the WFH coefficient separately for fathers and mothers. Both groups show a positive association with WFH. The point estimate is slightly larger for mothers, consistent with the aggregate gender gap in education time documented in Table 1. The confidence intervals overlap substantially, and a formal interaction test does not reject equality of the WFH effects across genders.

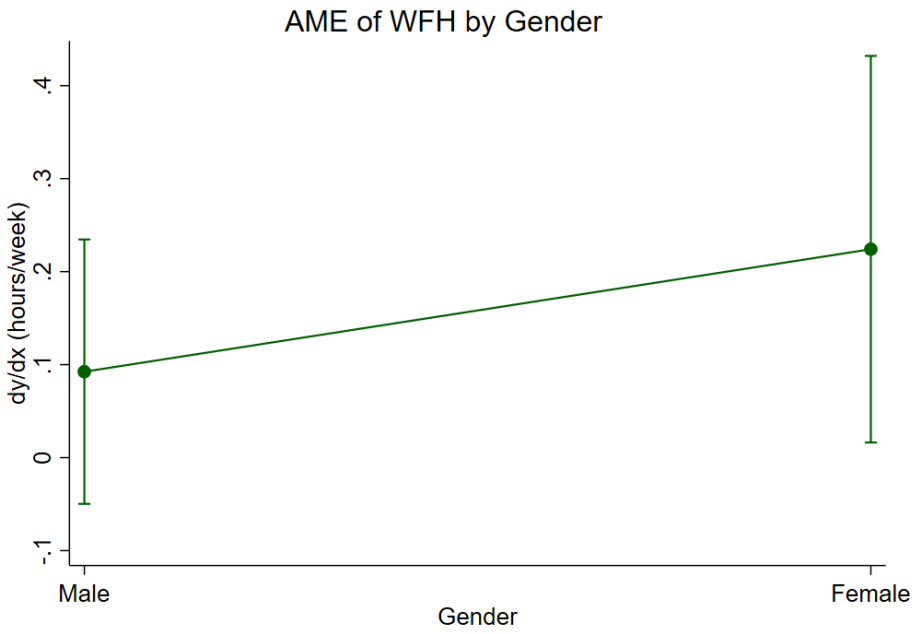


Figure 17: Heterogeneous WFH effects on education time by parent gender. Survey-weighted with 95% confidence intervals.

In summary, the heterogeneity analysis shows that the positive association between WFH and parental education time is most pronounced for parents of school-age children, modestly larger for middle-educated households, and broadly similar across genders. These patterns reinforce the interpretation that WFH facilitates reallocation of time toward cognitively intensive child investments when such activities are most relevant, rather than representing a uniform shift across all demographic groups.

A.2 Winsorization and Standardized Outcomes

I assess sensitivity to outliers and scaling by (i) winsorizing the dependent variable at the 1st and 99th percentiles and re-estimating (1), and (ii) standardizing the dependent variable using survey-weighted mean and standard deviation to obtain a "standard-units" outcome. The WFH coefficient remains positive and statistically significant under both transformations, with only modest attenuation after trimming extreme values. This confirms that the main result is not an artifact of a small number of very long education episodes (Table 15).

A.3 Leave-One-Year-Out (LOYO) Robustness

To verify that no single cohort year drives the result, I re-estimate (1) excluding one year at a time. The estimated WFH coefficient is positive and statistically significant in every

Table 15: Robustness to Outliers and Scaling

	Outlier Robustness		
	(1) Main OLS	(2) Winsorized (1%)	(3) Standardized
Work from Home (≥2h)	0.153**	0.141**	0.080**
	(0.063)	(0.056)	(0.033)
Observations	11,045	11,045	11,045

Notes: *** p < 0.01, ** p < 0.05, * p < 0.1. Winsorized specification trims top and bottom 1% of education time distribution. Standardized specification uses z-scores of education time. All specifications include full set of controls from Table 2. Survey-weighted estimates (ATUS wt06); Taylor-linearized standard errors.

LOYO run; magnitudes vary in a narrow band around the main estimate (Table 16). The design-based standard errors and p-values are stable across omissions, which is reassuring given the exclusion of 2020 from the analytical sample.

Table 16: Leave-One-Year-Out (LOYO) Robustness: WFH Effect on Weekly Education Time

Omitted year	WFH coef (hrs/wk)	SE	p-value
2010	0.145**	(0.064)	0.024
2011	0.175***	(0.064)	0.007
2012	0.146**	(0.065)	0.024
2013	0.151**	(0.065)	0.019
2014	0.161**	(0.065)	0.014
2015	0.167**	(0.065)	0.011
2016	0.180***	(0.065)	0.005
2017	0.151**	(0.065)	0.019
2018	0.110*	(0.059)	0.062
2019	0.142**	(0.063)	0.024
2021	0.158**	(0.065)	0.015
2022	0.162**	(0.068)	0.016
2023	0.159**	(0.068)	0.019

Notes: Survey-weighted (wt06) ATUS; Taylor-linearized SEs.
Controls: parent sex, education, age, #children, child age group, income, metro, race, industry FE, state FE, spouse employment/age/education, and year FE. One year omitted each row.
$*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Note on LOYO values: The values in Table 16 are derived from the corrected replication code (`02_main_analysis.do`). Each row re-runs the full survey-weighted regression from equation (1) with one year dropped. The corrected activity coding (excluding `030107`) applies to all iterations. The LOYO band (0.110–0.180) is consistent with the full-sample estimate of 0.153, confirming temporal stability.

B Additional Robustness: WFH Intensity, Mechanism, and Weekend

B.1 Robustness to WFH Binary Threshold

The main specification defines WFH as working at least 2 hours from home. Table 17 shows that the positive association between WFH and parental education time is insensitive to this choice. At the 1-hour threshold the coefficient increases to 0.192*** (SE = 0.057), possibly because this broader definition includes more households with flexible schedules but fewer hours of true home-based work; at higher thresholds (3h, 4h) the estimates range between 0.188*** and 0.197***. All four specifications are significant at the 5% level or better, and the point estimates cluster around 0.15–0.20 hours per week.

Table 17: Robustness to WFH Threshold Definition

	(1)	(2)	(3)	(4)
	Weekly Child Education Hours (hrs/wk)			
	≥1h	≥2h (main)	≥3h	≥4h
WFH ≥ 1h	0.192***			
	(0.057)			
WFH ≥ 2h (main)		0.153**		
		(0.063)		
WFH ≥ 3h			0.188***	
			(0.070)	
WFH ≥ 4h				0.197***
				(0.074)
Observations	11,045	11,045	11,045	11,045
Full controls	✓	✓	✓	✓

Columns differ only in WFH binary threshold (minimum diary-day WFH minutes).
All specifications include controls from Table 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B.2 Continuous WFH Measures

Table 18 compares the binary WFH dummy to two continuous WFH intensity measures: daily WFH hours (`wfh_hours`) and WFH share of total work time (`wfh_share`). All three specifications tell a consistent story. The continuous WFH hours coefficient implies that each additional hour of diary-day work from home is associated with 0.018 additional hours per week of children's education (SE = 0.008, $p < 0.05$). The WFH share coefficient of 0.265 (SE = 0.071, $p < 0.01$) implies that a shift from fully office-based to fully home-based work would increase education time by about 15 minutes per week, comparable in magnitude to

the binary estimate. The slight reduction in sample size for column (3) (N = 9,755 vs. 11,045) reflects that the WFH share is undefined for non-working diary days.

Table 18: WFH and Education Time: Continuous vs. Binary Measures

	(1)	(2)	(3)
	Weekly Child Education Hours (hrs/wk)		
	Binary	WFH Hours	WFH Share
WFH dummy (≥2h)	0.153**		
	(0.063)		
WFH hours (daily)		0.018**	
		(0.008)	
WFH share (0–1)			0.265***
			(0.071)
Observations	11,045	11,045	9,755
Full controls	✓	✓	✓

Notes: Col. (1) binary WFH (≥2h), col. (2) daily WFH hours (continuous), col. (3) WFH share of total work time. All specifications include the full set of controls from Table 2.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Survey-weighted (ATUS wt06); Taylor-linearized SEs.

B.3 Dosage-Response: WFH Intensity and Education Time

Table 19 and Figure 18 present a non-parametric dosage-response analysis using four WFH intensity bins. The reference category is no WFH (0 diary-day WFH hours).

Table 19: Dosage-Response: WFH Intensity and Education Time

WFH Intensity (ref: No WFH)	Education Hrs/Wk
	(1)
Light WFH (<2h)	0.178**
	(0.071)
Moderate WFH (2–4h)	0.064
	(0.097)
Substantial WFH (4–8h)	0.345***
	(0.116)
Full-day WFH (≥8h)	0.123
	(0.082)
Observations	11,045
Full controls	✓

Notes: Omitted category: No WFH (≥0h). All controls as in Table 2.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The dosage-response is non-monotone. The largest and most precisely estimated gain accrues to parents who work from home 4–8 hours per day ("Substantial WFH": 0.345***,

SE = 0.116). Light WFH (<2h) also yields a significant positive effect (0.178**, SE = 0.071). The Moderate (2–4h) and Full-day (≥8h) categories have positive but statistically insignificant point estimates. The non-monotonicity is consistent with diminishing returns at very high WFH intensity: parents working a full day from home may face greater work pressure that partially crowds out educational time, while parents with 4–8 hours of home-based work benefit from the time-budget flexibility without the same work intensity.

[Figure placeholder: dosage-response plot]
`fig_dosage_response.png`
Copy from Replication_Package/Output/Figures/

Figure 18: WFH dosage-response: estimated association between WFH intensity and parental education time (hours/week) with 95% confidence intervals. Reference category: No WFH. Controls as in Table 2.

B.4 Commuting Time as a Mechanism

Table 20 presents the commuting mediation analysis described in Section 2.1. Column (1) replicates the baseline WFH coefficient (0.153**, SE = 0.063). Column (2) adds weekly commuting time as a control; the WFH coefficient falls to 0.136 (SE = 0.063), with the reduction implying that 11% of the WFH premium is mediated by saved commuting time. Column (3) shows the first stage: WFH reduces commuting by 1.908 hours per week (SE = 0.110, $p < 0.001$).

Table 20: Commuting Time as Mechanism: Mediation Analysis

	(1)	(2)	(3)
	Outcome: Education time (hrs/wk)		Outcome: Commute (hrs/wk)
	(1) Education (no commute)	(2) Education (+ commute)	(3) Commute (1st stage)
Work from Home (≥2h)	0.153**	0.136**	−1.908***
	(0.063)	(0.063)	(0.110)
Commute time (hrs/wk)		−0.008*	
		(0.005)	
Observations	11,045	11,045	11,045

Notes: Col. (1) baseline WFH → education. Col. (2) adds weekly commute as control.
Col. (3) WFH → commute (first stage). Mediation fraction = (0.153 − 0.136)/0.153 = 11%.
All columns include full controls from Table 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The 11% mediation share implies that the large majority of the WFH–education premium (≈ 89%) operates through channels other than commute time savings alone. These additional channels likely include reduced scheduling friction (no fixed commute departure time), the physical proximity of WFH parents to children during after-school hours, and

reduced cognitive/physical fatigue relative to in-office work. The mediation result is directionally consistent with the theoretical model, which predicts that WFH increases parental education time by relaxing the time budget; it further reveals that the time-budget channel captures only part of the total effect.

B.5 Weekend Robustness

The main analysis uses weekday diaries only (scaled to weekly equivalents by multiplying by five), following standard practice in the time-use literature for measuring work-related associations. Table 21 shows that expanding the sample to include weekend diaries (scaled by seven, with a weekend indicator added as a control) produces a qualitatively identical but slightly smaller WFH coefficient (0.141*, SE = 0.074). The slight attenuation is expected: WFH on weekends is less common and may reflect different work patterns (e.g., catching up on tasks rather than a regular WFH arrangement), which could dilute the estimated association. The positive sign and near-5% significance confirm that the result is not confined to weekdays.

Table 21: Robustness: Weekday vs. All-Day Sample

	(1)	(2)
	Education Hours per Week	
	(1) Weekdays only (×5)	(2) All days (×7)
Work from Home (≥2h)	0.153**	0.141*
	(0.063)	(0.074)
Observations	11,045	22,144

Col. (2) includes weekend diaries with a weekend indicator.
All controls from Table 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C Proofs for Section 3.1

C.1 Proof of Proposition 1 (Existence of Equilibrium)

The proof relies on Tarski's (1955) fixed-point theorem.

1. **Define the State Space:** Let the space of possible peer spillovers be a compact lattice $S = [S_{\min}, S_{\max}]^2$, where $S_{\min} > 0$ and $S_{\max} < \infty$ are the minimum and maximum possible spillovers generated by any sorting configuration.

2. **Define the Mapping $\mathcal{T}$:** Define a mapping $\mathcal{T} : S \mapsto S$ that takes a *guessed* pair of

spillovers $\vec{S} = (S_C, S_S)$ and returns the *implied* pair $\vec{\tilde{S}} = (\tilde{S}_C, \tilde{S}_S)$ generated by the resulting household decisions.

(a) Given $\vec{S}$, solve all households' optimization problems to find their policy functions and indirect utilities $V_z(k, a, \iota; \vec{S})$ for $z \in \{C, S\}$.

(b) Determine the location choice $Z(k, a, \iota; \vec{S}) = \arg\max_z V_z$ for all types.

(c) Calculate the implied spillovers $(\tilde{S}_C, \tilde{S}_S)$ using the consistency condition (Eq. (12)), which integrates the choices $e_z(\cdot)$ over the set of households who choose each location:

$$\tilde{S}_z^{1-\nu} = \frac{K_a}{\mu_z} \int a^{\rho_a} e_z(k, a, \iota)^{1-\nu} \mathbf{1}_{Z(\cdot)=z} \, dF(k, a, \iota)$$

where $\mu_z = \int \mathbf{1}_{Z(\cdot)=z} \, dF$.

3. **Show $\mathcal{T}$ is Order-Preserving (Monotone):** I must show that if $\vec{S}' \geq \vec{S}$ (i.e., $S'_C \geq S_C$ and $S'_S \geq S_S$), then $\mathcal{T}(\vec{S}') \geq \mathcal{T}(\vec{S})$.

   - Consider an increase in the guessed suburban spillover, $S'_S > S_S$, holding S_C constant.
   - From Eq. (13), an increase in S_S raises the utility difference ΔV for all households, as $\beta\nu > 0$.
   - This causes the cutoff $k^*(a, \iota)$ to *decrease* (or stay the same). A set of "marginal" households with $k \in [k^*(S'_S), k^*(S_S)]$ now switches from C to S.
   - Assuming e_z is non-decreasing in k and a, these marginal households have (on average) higher a and e than the initial city residents and are similar to the inframarginal suburban residents.
   - This new sorting *increases* the average $a^{\rho_a} e^{1-\nu}$ in the suburb, thus $\tilde{S}'_S \geq \tilde{S}_S$. It also *decreases* (or leaves unchanged) the average quality in the city, so $\tilde{S}'_C \leq \tilde{S}_C$.
   - A symmetric argument holds for an increase in S_C. The mapping $\mathcal{T}$ is order-preserving.

4. **Conclusion:** $\mathcal{T}$ is an order-preserving mapping on a compact lattice S. By Tarski's fixed-point theorem, there exists at least one fixed point $\vec{S}^* = \mathcal{T}(\vec{S}^*)$, which is a residential equilibrium.

C.2 Proof of Dominant Income Effect (for Lemma 1)

I show that if $\sigma > 1$, the income effect dominates the substitution effect, and optimal on-site work n_z^O is decreasing in productivity k.

For simplicity, consider a non-WFH worker ($\iota = 0$) in location z. The household's problem is to choose n_z^O to maximize utility, yielding the FOC (Eq. (11)):

$$\lambda \cdot \frac{\partial n}{\partial n_z^O} = \frac{\psi + \beta(1-\nu)}{\tau_z} \cdot (1 + t_z)$$

Let $\Psi \equiv \psi + \beta(1-\nu)$ and $T_z \equiv 1 + t_z$. Substitute the components:

- $\lambda = \frac{\partial U^{\text{mat}}}{\partial y_z} = [\chi y_z r_z^{-(1-\gamma)}]^{-\sigma} \cdot \chi r_z^{-(1-\gamma)} = \chi^{1-\sigma} r_z^{-(1-\gamma)(1-\sigma)} (y_z)^{-\sigma}$
- $y_z = n = A^O (k n_z^O)^\theta$
- $\frac{\partial n}{\partial n_z^O} = A^O \theta k^\theta (n_z^O)^{\theta-1}$
- $\tau_z = 1 - T_z n_z^O$

Plugging these in, the FOC is an equation $F(n_z^O, k) = 0$:

$$F(n_z^O, k) = \underbrace{C_z \cdot \left(A^O (k n_z^O)^\theta\right)^{-\sigma} \cdot \left(A^O \theta k^\theta (n_z^O)^{\theta-1}\right)}_{LHS(n_z^O, k)} - \underbrace{\frac{\Psi \cdot T_z}{1 - T_z n_z^O}}_{RHS(n_z^O)} = 0$$

where $C_z = \chi^{1-\sigma} r_z^{-(1-\gamma)(1-\sigma)}$ is a constant w.r.t. k and n_z^O. Simplifying the LHS:

$$LHS(n_z^O, k) = C_z (A^O)^{1-\sigma} \theta \cdot k^{\theta(1-\sigma)} \cdot (n_z^O)^{\theta(1-\sigma)-1}$$

By the Implicit Function Theorem, $\frac{dn_z^O}{dk} = -\frac{\partial F / \partial k}{\partial F / \partial n_z^O}$. I analyze the signs of the partial derivatives.

1. **Partial w.r.t. k:** The RHS is not a function of k.

$$\frac{\partial F}{\partial k} = \frac{\partial LHS}{\partial k} = \left[C_z (A^O)^{1-\sigma} \theta \cdot (\cdot)\right] \cdot \theta(1-\sigma) k^{\theta(1-\sigma)-1}$$

   Since all terms in brackets are positive, the sign is determined by $\theta(1-\sigma)$. If $\sigma > 1$, then $(1-\sigma) < 0$, and $\frac{\partial F}{\partial k} < 0$.

2. **Partial w.r.t. n_z^O:**

$$\frac{\partial F}{\partial n_z^O} = \frac{\partial LHS}{\partial n_z^O} - \frac{\partial RHS}{\partial n_z^O}$$

   The first term is $\frac{\partial LHS}{\partial n_z^O} = [\dots] \cdot (\theta(1-\sigma) - 1)(n_z^O)^{\theta(1-\sigma)-2}$. Since $\theta \in (0,1)$ and $\sigma > 1$, the exponent $(\theta(1-\sigma) - 1)$ is strictly negative. Thus, $\frac{\partial LHS}{\partial n_z^O} < 0$.

   The second term is $\frac{\partial RHS}{\partial n_z^O} = \frac{\Psi (T_z)^2}{(1 - T_z n_z^O)^2} > 0$.

   Therefore, $\frac{\partial F}{\partial n_z^O} = (\text{Negative}) - (\text{Positive}) < 0$.

Combining these results:

$$\frac{dn_z^O}{dk} = -\frac{\partial F/\partial k}{\partial F/\partial n_z^O} = -\frac{(\text{Negative})}{(\text{Negative})} = -(\text{Positive}) < 0$$

Thus, for $\sigma > 1$, optimal on-site work n_z^O is a decreasing function of productivity k.

C.3 Proof of Lemma 1 (Primitive Conditions for Single-Crossing)

I must prove that $\frac{\partial \Delta V(k,a,\iota)}{\partial k} > 0$. I analyze the derivative of the utility difference in Eq. (13), holding S_S and S_C fixed (as they are taken as given by the household):

$$\frac{\partial \Delta V}{\partial k} = \left[\frac{\partial V_S}{\partial k} - \frac{\partial V_C}{\partial k}\right]$$

Using the Envelope Theorem, the derivative of the indirect utility V_z w.r.t. the parameter k is the partial derivative of the Lagrangian w.r.t. k, holding the choice variables constant.

$$\frac{\partial V_z}{\partial k} = \frac{\partial \mathcal{L}}{\partial k} = \lambda \frac{\partial n_z}{\partial k}$$

where $\lambda = \frac{\partial U^{\text{mat}}}{\partial y_z}$ is the marginal utility of income, and $\frac{\partial n_z}{\partial k}$ is the marginal product of human capital, evaluated at the optimal choices. For a non-WFH worker: $n_z = A^O(kn_z^O)^\theta$, so $\frac{\partial n_z}{\partial k} = A^O\theta k^{\theta-1}(n_z^O)^\theta$. The derivative of the utility difference is:

$$\frac{\partial \Delta V}{\partial k} = \lambda_S \cdot \frac{\partial n_S}{\partial k} - \lambda_C \cdot \frac{\partial n_C}{\partial k}$$

Substitute $\lambda_z \propto (y_z)^{-\sigma}$ and n_z:

$$\frac{\partial \Delta V}{\partial k} \propto (y_S)^{-\sigma} \cdot (k^{\theta-1}(n_S^O)^\theta) - (y_C)^{-\sigma} \cdot (k^{\theta-1}(n_C^O)^\theta)$$

Substitute $y_z = A^O(kn_z^O)^\theta$:

$$\frac{\partial \Delta V}{\partial k} \propto (n_S^O)^{\theta(1-\sigma)} k^{\theta(1-\sigma)} \cdot k^{\theta-1} - (n_C^O)^{\theta(1-\sigma)} k^{\theta(1-\sigma)} \cdot k^{\theta-1}$$

$$\frac{\partial \Delta V}{\partial k} \propto k^{\theta-\sigma\theta-1} \cdot \left[(n_S^O)^{\theta(1-\sigma)} - (n_C^O)^{\theta(1-\sigma)}\right]$$

A simpler proof relies on the logic in the main text's proof sketch. The household's utility is $V_z(k) = U^{\text{mat}}(y_z(k), r_z) + \Psi \ln(\tau_z(k)) + \beta\nu \ln S_z$.

$$\frac{\partial \Delta V}{\partial k} = \left[\frac{\partial U_S^{\text{mat}}}{\partial k} - \frac{\partial U_C^{\text{mat}}}{\partial k}\right] + \Psi\left[\frac{\partial \ln \tau_S}{\partial k} - \frac{\partial \ln \tau_C}{\partial k}\right]$$

1. **Residual Time Channel:** $\frac{\partial \ln \tau_z}{\partial k} = \frac{1}{\tau_z}\frac{\partial \tau_z}{\partial k} = \frac{-T_z}{\tau_z}\frac{dn_z^O}{dk}$. From Appendix C.2, $\frac{dn_z^O}{dk} < 0$. Therefore, $\frac{\partial \ln \tau_z}{\partial k} > 0$. This gain is larger in the suburb because $T_S > T_C$. The gain in residual time τ_z from higher k is greater in the suburb. Thus, the Residual Time Channel is positive.

2. **Material Utility Channel:** $\frac{\partial U_z^{\text{mat}}}{\partial k} = \frac{\partial U^{\text{mat}}}{\partial y_z}\frac{dn_z}{dk}$. As noted in the text, y_z is increasing in k, so h_z is increasing in k. The utility gain from lower rent $r_S < r_C$ is increasing in k, which is a positive component. This is offset by $\lambda_S < \lambda_C$ due to $y_S > y_C$.

The combination of the strictly positive Residual Time Channel and the positive housing-benefit component of the Material Utility Channel ensures that $\frac{\partial \Delta V}{\partial k} > 0$.

C.4 Proof of Proposition 3 (WFH Lowers the Suburban Cutoff)

The proof is a comparative static on the equilibrium sorting pattern from Proposition 2.

1. **Value of WFH:** Let $V_z(k, a, \iota)$ be the indirect utility in location z for type (k, a, ι). WFH eligibility ($\iota = 1$) provides access to a superior labor production technology (Eq. (5)) compared to an on-site-only household ($\iota = 0$). This technology allows the household to substitute remote hours n^H for on-site hours n^O.

2. **Location-Specific Value:** The primary benefit of this substitution is saving on the time cost of commuting, which is $(1 + t_z)n^O$. Because this time cost is strictly higher in the suburb ($t_S > t_C$), the utility gain from accessing the WFH technology is larger in the suburb.

$$V_S(k, a, \iota = 1) - V_S(k, a, \iota = 0) > V_C(k, a, \iota = 1) - V_C(k, a, \iota = 0)$$

3. **Impact on Cutoff:** Rearranging this inequality gives:

$$V_S(k, a, \iota = 1) - V_C(k, a, \iota = 1) > V_S(k, a, \iota = 0) - V_C(k, a, \iota = 0)$$

   This is exactly $\Delta V(k, a; \iota = 1) > \Delta V(k, a; \iota = 0)$. For any given k, the utility difference in favor of the suburb is higher for WFH-eligible households.

4. From Proposition 2, I know that $\Delta V(k)$ is a strictly increasing function of k for both $\iota = 0$ and $\iota = 1$. An upward shift of this entire increasing function (from the $\iota = 0$ curve to the $\iota = 1$ curve) necessarily means it crosses the $\Delta V = 0$ axis at a *lower* value of k.

5. Therefore, $k^*(a, \iota = 1) < k^*(a, \iota = 0)$.