

DISCUSSION PAPER SERIES

IZA DP No. 15436

**Time Use and the Geography of Economic  
Opportunity**

Sulagna Mookerjee  
John D. Pedersen  
David Slichter

JULY 2022

## DISCUSSION PAPER SERIES

IZA DP No. 15436

# Time Use and the Geography of Economic Opportunity

**Sulagna Mookerjee**

*SUNY*

**John D. Pedersen**

*Federal Communications Commission*

**David Slichter**

*SUNY and IZA*

JULY 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

**IZA – Institute of Labor Economics**

Schaumburg-Lippe-Straße 5–9  
53113 Bonn, Germany

Phone: +49-228-3894-0  
Email: [publications@iza.org](mailto:publications@iza.org)

[www.iza.org](http://www.iza.org)

## ABSTRACT

---

# Time Use and the Geography of Economic Opportunity\*

The community in which a child is raised has a substantial effect on their income in adulthood. To help understand what is different about communities which produce higher incomes, we document how time use differs between communities which increase vs. decrease incomes. The main differences are that, in areas which produce higher incomes, people spend more time at work, and adults spend more time with children. The data do not support some theories of what makes communities effective at producing human capital: People do not spend more time on educational activities, or on community events and institutions, in areas which increase incomes by more.

**JEL Classification:** D13, J22, J24, J62

**Keywords:** time use, intergenerational mobility

**Corresponding author:**

David Slichter

Binghamton University

4400 Vestal Pkwy E

Binghamton

NY 13902

USA

E-mail: [slichter@binghamton.edu](mailto:slichter@binghamton.edu)

---

\* The analysis and conclusions in this paper are those of the authors alone and do not necessarily reflect the views of the Federal Communications Commission. We are grateful for helpful comments from Clifford Kern, Sol Polachek, and Thomas Sinclair. All remaining errors are ours.

# 1 Introduction

It has long been believed that the community where a child is raised has a substantial impact on a child’s life course (e.g., Benabou 1993, Katz et al. 2001). Recently, Chetty and Hendren (2018a) quantified one dimension of this effect by measuring the effect of childhood community on income in adulthood (relative to living in a randomly chosen other community), finding large differences across communities.<sup>1</sup>

However, it is not known exactly *why* communities differ in their effects on the adult incomes of children raised there. A natural first step towards this is to characterize what is different between these communities. Chetty and Hendren (2018b) document that communities which increase the incomes of children who were raised there have higher standardized test scores, lower levels of segregation, and greater social capital – potentially pointing to explanations relating to the school system and social interactions. Subsequent research along these lines suggests that places which generate growth in test score performances are not better at producing incomes (Rothstein 2019, Mookerjee and Slichter 2020), and that effects on income are more correlated with historical than contemporary segregation (Andrews et al. 2017).

We contribute to the question of what makes some places increase incomes in adulthood more than others by characterizing how time use differs between these places. Differences in time use are potentially informative in this regard for two reasons. First, time use is a marker for community characteristics such as cultural values about the importance of children, work, and education. Second, the allocation of time is per se an important input into child development (e.g., Sayer et al. 2004a, Kofman and Bianchi 2012, Caetano et al. 2019).<sup>2</sup>

We use the American Time Use Survey, which provides information on how much time is spent by individuals on different activities. We observe the *quantity* of time spent but not the *quality*, as is common in the literature; however, we present a simple theoretical model in the Appendix demonstrating why one would expect differences in quality of time inputs to be reflected in the amount of time spent on each activity. For example, in communities with a comparative advantage in producing human capital through time use on social activities, we would expect to see more time spent on social activities.

We examine time use separately for three subgroups within our sample. These subgroups are (i) adults who have at least one child in the household, and are therefore likely to be caregivers to children; (ii) adults with no children in the household, who are therefore unlikely to be primary caregivers to children; and (iii) young adults (ages 15-23), to capture how the people being raised in a community spend their time.

---

<sup>1</sup>Other recent work such as Chetty et al. (2016) and Fletcher and Han (2019) also provide quantitative support for the importance of the community where a child is raised.

<sup>2</sup>Prior literature has considered that parental time spent on children might play a role in the intergenerational transmission of income inequalities, at the individual (rather than community) level (e.g., Hill and Stafford 1980, Kalil et al. 2012).

Our first empirical finding is that adults spend more time with children in communities which increase incomes in adulthood. (For the rest of the paper, we refer to communities which increase incomes in adulthood as “high-effect” communities.) This is consistent with both the views that such communities are more child-centric in the sense of placing a greater utility weight on children’s outcomes, and that time spent with family is a key input for children’s cognitive development (Caetano et al. 2019).<sup>3</sup>

Second, we find that individuals spend more time at work in high-effect communities. This finding is sensitive to controls for adults but not for the young adults cohort. This finding is consistent with the theory that communities vary in their cultural emphasis on work, and that children learn how much to emphasize their careers from observing adults (Huang et al. 2001). It is also potentially consistent with the possibility that high-effect communities are increasing children’s incomes in adulthood by placing them into healthier labor markets, rather than by increasing their human capital.

Also notable are some correlations that we fail to find. First, we examine time spent on adult care as a placebo category, and do not find that respondents in high-effect communities are any more likely to allocate time towards taking care of adult household members. This suggests that additional time on caregiving is specific to children, and not necessarily indicative of generalized values about the importance of helping others. Second, while Chetty and Hendren (2018b) find evidence that high-effect communities might have greater social capital and stronger social networks, we find no correlation between a community’s effect on incomes and time spent on social or community-building activities such as attending social gatherings, participating in community or religious organizations or events, or attending sports or entertainment events. Third, despite the higher average test scores in such communities (Chetty and Hendren 2018b), we do not find evidence that more time is spent by young adults on educational activities in high-effect communities.

The rest of the paper proceeds as follows. Section 2 describes the data. Section 3 gives the results. Section 4 concludes.

## 2 Data and methods

We obtain our measures of time use from the American Time Use Survey (ATUS). The ATUS is a Bureau of Labor Statistics Survey, conducted by the U.S. Census Bureau, on a nationally representative sample of people aged 15 and older from households that recently completed the Current Population Survey. One person from each household is randomly chosen to answer a series of questions on their time use for one day: what activities they engaged in and for how long, the location, who they were accompanied by, etc. We use survey years from 2003 to 2018.

---

<sup>3</sup>Prior papers also find that more educated parents spend more time with children (Sayer et al. 2004b, Guryan et al. 2008).

We categorize respondents into three mutually exclusive subgroups. The first is “young adults,” composed of all respondents ages 15-23, and intended to capture the time use of people who are being raised in the community. We would ideally match the age bracket used by Chetty and Hendren (2018b), who define where a child is raised by where they lived between ages 9 through 23. However, we are limited to start at age 15 because ATUS does not collect information on the time use of children under 15.

For the sample of young adults, we focus on categories of time use which might influence the respondent’s own human capital accumulation, as follows: (i) *education*, which includes time spent on classes, homework, and education-related travel, (ii) *work*, encompassing time spent actually working as well as searching, traveling, and other work-related actions (note that results are unchanged if we define work to exclusively mean time spent actually working), (iii) *watching television* as part of leisure time, such as watching sports on TV, (iv) *sports*, including participating in sports, exercising, or recreational outdoor activities, (v) *social activities* which are organized and formal, consisting of religious, social and civic actions, social services, related travel, and volunteering, and (vi) *socializing* informally, which includes time spent visiting or communicating with family, friends, and neighbors.

Our other two subgroups are adults with children in the household, and adults without children in the household. We define an adult to be a respondent age 24 or older, and adults with children are adults who report at least one person under the age of 18 residing in their household.

Among adults, our key categories of time use are (i) *child care*, which is any time spent with household children (getting them ready, helping them with schoolwork, playing with them, entertaining them, planning activities for them, time spent on grooming or healthcare, etc.) as well as any time spent with non-household children; (ii) *adult care*, which captures time spent caring for household and non-household adults and helping non-household adults with a range of activities such as running errands for them, giving them rides, helping them with technology, attending to their medical needs, etc.; (iii) *work*, encompassing time spent actually working as well as searching, traveling, and other work-related actions; and (iv) *social activities*, consisting of religious, social, and civic actions, social services, related travel, and volunteering. The goal of these measures is (a) to capture inputs directly given from caregivers to children, and (b) to proxy for attitudes, values, and behaviors common among the people children interact with.

The survey also provides standard socioeconomic and demographic information for all members of each household. Table [1](#) provides an overview of our sample from ATUS. Two variables are of particular importance for our analysis: (i) family income, which allows us to create an indicator for whether households are high or low income, and (ii) the number of children in the household, which lets us separately explore time allocation for adults with and without children in the household (which we treat as a proxy for whether they could be potential caregivers). Our definition of high and low income roughly corresponds to the median of income in the data for

each year: On average, households are denoted as high income if they earn \$60,000 and above, which corresponds to about 52% of our sample.<sup>4</sup>

We combine this data with Chetty and Hendren’s (2018b) estimates of counties’ effects on the adult incomes of children raised there. Chetty and Hendren construct estimated effects on income separately for children from families at the 25<sup>th</sup> (low) and the 75<sup>th</sup> (high) percentiles of income. They use cohorts born between 1980 and 1988, which partially overlaps with our data; the results are robust to restricting the sample to the overlapping cohorts.

For simplicity, we convert their estimates of county exposure effects into units of standard deviations above or below the average. To perform this conversion, we use their estimate of the true standard deviation of county effects, which is 0.517 for the 25<sup>th</sup> and 0.321 for the 75<sup>th</sup> percentile, in units of percent effect on income per year of exposure.<sup>5</sup> Their estimates are designed to be forecast-unbiased in the presence of estimation error, which necessitates shrinkage;<sup>6</sup> therefore, the standard deviation of their estimates is smaller than these estimates of the standard deviation of county quality. Specifically, Chetty and Hendren’s estimates are constructed using a convex combination of two estimates: (i) estimates using children who move between counties at different ages, and (ii) average outcomes for children who spend their entire childhood in one county (“permanent residents”). Chetty and Hendren argue that (i) is approximately unbiased but has high variance due to sampling error, while (ii) has low variance but is biased. In practice, most of the variation in their combined estimates comes from (ii) rather than (i).

Table 1: Descriptive statistics

	Adults		Young Adults
	Hhd child	No hhd child	
Age	40.77 (9.77)	56.63 (15.96)	18.54 (2.61)
Male	0.41 (0.50)	0.45 (0.49)	0.49 (0.50)
White	0.80 (0.42)	0.77 (0.40)	0.76 (0.42)
College grad	0.42 (0.49)	0.36 (0.48)	0.04 (0.19)
Employed	0.72 (0.45)	0.53 (0.50)	0.35 (0.47)
Observations	31,889	42,832	7,847

Notes: The table provides descriptive statistics for our sample.

For each income category, high and low, we assign the individuals in that cat-

<sup>4</sup>We are not able to divide at exactly the median because income is binned in the data.

<sup>5</sup>The mean of their measure is 0 by construction.

<sup>6</sup>See their paper for more details

egory the corresponding standard deviation measure of the effect on adult incomes for children in that income group (derived from Chetty and Hendren’s (2018b) estimates).<sup>7</sup> We then measure the difference in average time spent on each activity category between more and less effective communities by regressing time use on the measure of county effects, as follows:

$$t_{ic} = \beta_0 + \beta_1 Q_{ic} + u_{ic},$$

where  $t_{ic}$  is the time in minutes spent on an activity by individual  $i$  residing in county  $c$ , and  $Q_{ic}$  is the standard deviation measure of the effect of being raised in county  $c$  on adult income for children belonging to individual  $i$ ’s income category (high or low). We weight the regression using ATUS sampling weights. A positive estimated  $\hat{\beta}_1$  indicates that high-effect communities invest more time in that particular activity, and the magnitude is the number of additional minutes that people in a community with a 1 standard deviation above-average effect on income spend on that activity relative to people in the average community.

Because  $Q$  is estimated by Chetty and Hendren, it is important to consider how measurement error in their estimates might affect our results. Our ideal regression is

$$t_{ic} = \alpha_0 + \alpha_1 Q_{ic}^* + v_{ic},$$

where  $Q_{ic}^*$  is the true exposure effect as opposed to Chetty and Hendren’s estimate, and  $E(v_{ic}|Q_{ic}^*) = 0$ . Let  $\epsilon_{ic} := Q_{ic} - Q_{ic}^*$  be the measurement error in the estimate of  $Q^*$ .

We have

$$\begin{aligned} E(t_{ic}|Q_{ic}) &= E(\alpha_0 + \alpha_1 Q_{ic}^* + v_{ic}|Q_{ic}) \\ &= \alpha_0 + \alpha_1 E(Q_{ic}^*|Q_{ic}) + E(v_{ic}|Q_{ic}) \\ &= \alpha_0 + \alpha_1 Q_{ic} + E(v_{ic}|Q_{ic}), \end{aligned}$$

where the last equality holds because Chetty and Hendren’s measure is designed to be forecast-unbiased. It follows that our regression estimand  $\beta_1$  is equal to the ideal estimand  $\alpha_1$  if  $E(v_{ic}|Q_{ic}) = 0$ . While  $E(v_{ic}|Q_{ic}^*) = 0$ , it can still be the case that  $E(v_{ic}|Q_{ic}) \neq 0$  if  $E(v_{ic}|\epsilon_{ic}) \neq 0$ . That is, our estimates might be biased due to measurement error if the errors in Chetty and Hendren’s estimates are correlated with time usage in the same community. As discussed above, the main source of potential error in Chetty and Hendren’s estimates is that permanent residents of different communities are systematically different. Therefore, to assess whether our estimates are likely to be sensitive to measurement error, we estimate specifications using individual-level controls.

---

<sup>7</sup>For consistency with Chetty and Hendren, who give estimated effects at the 25<sup>th</sup> and 75<sup>th</sup> percentiles of the income distribution, we match our below-median income individuals to their 25<sup>th</sup> percentile estimates and our above-median income individuals to their 75<sup>th</sup> percentile estimates.



In particular, in order to rule out that these differences are simply driven by differences in socioeconomic and demographic characteristics, we estimate a specification controlling for individuals’ age, gender, race, education, and income bracket. Comparing individuals across communities who are identical in these respects eliminates the possibility that the differential time use patterns are accounted for purely by selection along these dimensions (e.g., richer or college educated people being both more likely to live in high-effect communities as well as more likely to spend time on work). With limited exceptions, our results turn out not to be sensitive to these controls.

### 3 Results

Table 2 reports the results for adults with children in the household, and Table 3 for adults without. For each sample, we report both the results from an uncontrolled specification (Panel A) and a specification including the controls described above (Panel B).

Table 2: Results for adults with children in household

	Childcare		Adult Care	Work	Soc. Act.
	Hhd	Non-hhd			
Panel (A): No Controls					
Q	2.58** (1.04)	-0.04 (0.18)	-0.04 (0.15)	9.49*** (2.86)	0.46 (0.54)
Baseline	79.99	1.53	1.41	207.29	23.45
Observations	31,889	31,889	31,889	31,889	31,889
Panel (B): With Controls					
Q	1.66* (1.00)	0.06 (0.18)	-0.00 (0.16)	3.64 (2.72)	0.56 (0.54)
Baseline	79.99	1.53	1.41	207.29	23.45
Observations	31,889	31,889	31,889	31,889	31,889

Notes: The table provides results from regressions of time spent on different activity categories by individuals aged 24 and above who have at least one child under the age of 18 residing in their household, on the effectiveness of their community at increasing incomes in adulthood. \*, \*\*, and \*\*\* denote significance at 10%, 5%, and 1% respectively. Robust standard errors in parentheses. All regressions include ATUS sampling weights.

The first set of results is for time allocated towards childcare activities.

Households in high-effect communities spend significantly more time with children overall. Each standard deviation increase in community effect on incomes is

Table 3: Results for adults with no children in household

	Childcare		Adult Care	Work	Soc. Act.
	Hhd	Non-hhd			
Panel (A): No Controls					
Q	–	0.52* (0.31)	–0.26 (0.19)	1.98 (2.67)	–0.35 (0.52)
Baseline		5.94	2.02	160.80	26.24
Observations		42,832	42,832	42,832	42,832
Panel (B): With Controls					
Q	–	0.38 (0.31)	–0.31 (0.19)	7.52*** (2.42)	–0.38 (0.52)
Baseline		5.94	2.02	160.80	26.24
Observations		42,832	42,832	42,832	42,832

*Notes:* The table provides results from regressions of time spent on different activity categories by individuals aged 24 and above who have no children residing in their household, on the effectiveness of their community at increasing incomes in adulthood. \*, \*\*, and \*\*\* denote significance at 10%, 5%, and 1% respectively. Robust standard errors in parentheses. All regressions include ATUS sampling weights.

associated with an additional 2-3 minutes per day of childcare by respondents in households with children, off of a baseline of 82 minutes per day. This is driven by additional time specifically with children from the same household, who absorb the overwhelming majority of childcare time (80 minutes per day) among such households.<sup>8</sup> Households without children also spend more time caring for children, though the amounts of time involved are small – half a minute per day per standard deviation, off of a baseline of 6 minutes per day.

We also examine time spent on taking care of adults, in the spirit of a placebo test, to see whether the larger amount of time on child care reflects a general tendency to spend time caring for others as opposed to a specific orientation to children. We find no pattern of higher time spent on adult care in high-effect communities.

The third time use category we study is work. Our results suggest that adult respondents spend more time on work in high-effect communities. The results within each household type are sensitive to the inclusion of controls, but the overall pattern suggests that adults generally spend about 5 additional minutes on work per day in counties which have one standard deviation higher effects on income. If we run the work regressions by pooling all respondents over the age of 24, regardless of whether or not there are children present, we obtain a coefficient of 4.83 (standard error of

<sup>8</sup>Results are robust to controlling for the number of children in the household.

Table 4: Results for childcare categories

	Education	Health	Other
Panel (A): No Controls			
Q	-0.48* (0.26)	0.22 (0.14)	2.83*** (0.98)
Baseline	6.83	1.82	71.34
Observations	31,889	31,889	31,889
Panel (B): With Controls			
Q	-0.35 (0.27)	0.21 (0.14)	1.80* (0.94)
Baseline	6.83	1.82	71.34
Observations	31,889	31,889	31,889

Notes: The table provides results from regressions of time spent on different childcare categories by individuals aged 24 and above who have at least one child under the age of 18 residing in their household, on the effectiveness of their community at increasing incomes in adulthood. \*, \*\*, and \*\*\* denote significance at 10%, 5%, and 1% respectively. Robust standard errors in parentheses. All regressions include ATUS sampling weights.

1.99) without controls and 5.75 (standard error of 1.84) with controls.

The last category which we explore for adults is social, civic, and religious activities. We find zero differences between high-effect and low-effect communities in these time expenditures.

Next, because we find evidence that people in high-effect communities spend more time with children, we also look in greater detail at how time with children is spent. In Table 4, we report results dividing time with children into three subcategories: time on educational activities (e.g., helping children with homework), time on health-related activities (e.g., taking children to the doctor), and all other time expenditures. The latter subcategory appears to be the main difference in time use.

Table 4 shows that, if anything, families in high-effect communities are spending less time on educational activities with their children, suggesting that high-effect communities might not be distinguished by a greater family emphasis on education. Because family time spent on educational activities is often intended to be remedial, it is also possible that this simply reflects greater satisfaction with children’s progress in school. However, the remedial motive is probably less important for children’s educational time expenditures, and we will see momentarily that the children themselves are also not spending more time on educational activities in high-effect communities.

Finally, we examine time spent by young adults, i.e. individuals aged between 15 and 23. The results are reported in Table 5 and are broadly consistent with the results among older adults.

Table 5: Results for young adults

	Educ	Work	Watch TV	Sports	Soc. Act.	Socialize
Panel (A): No Controls						
Q	-5.35 (4.54)	13.27*** (5.00)	2.26 (2.73)	-0.34 (1.57)	1.23 (2.04)	-0.19 (0.95)
Baseline	123.50	106.44	144.03	36.70	60.17	19.90
Observations	7,847	7,847	7,847	7,847	7,847	7,847
Panel (B): With Controls						
Q	-3.35 (4.24)	8.44* (4.45)	3.17 (2.72)	-0.83 (1.50)	0.38 (1.99)	0.23 (0.96)
Baseline	123.50	106.44	144.03	36.70	60.17	19.90
Observations	7,847	7,847	7,847	7,847	7,847	7,847

Notes: The table provides results from regressions of time spent on different activity categories by individuals aged between 15 and 23, on the effectiveness of their community at increasing incomes in adulthood. \*, \*\*, and \*\*\* denote significance at 10%, 5%, and 1% respectively. Robust standard errors in parentheses. All regressions include ATUS sampling weights.

The most salient difference in time usage among young adults is on work, with young adults from a one standard deviation higher-effect community spending on average 8-13 additional minutes per day working, from a baseline of 106 minutes.

We find no difference in time spent on education between high- and low-effect communities. (In fact, our point estimate is negative, i.e., time spent on education is lower in high-effect communities.)

There also seems to be little difference in time expenditure on key leisure activities. Similar to older adults, we do not find any differences in time spent on social and community activities. This is true for both structured social events and informal socializing. We also do not find differences in time spent watching TV.

## 4 Conclusion

We find that communities' effects on incomes in adulthood are correlated with how people there use their time. In particular, individuals living in areas which increase incomes tend to spend more time at work, and adults in those communities spend more time on childcare activities. This is consistent with the views that high-effect communities are distinguished by a greater cultural emphasis on work, or on children's welfare. It is also consistent with the views that high-effect communities induce higher earnings in adulthood directly as a result of additional time spent with children, or as a result of stronger labor market conditions.

However, we fail to find some correlations which would have been expected under other theories of why some communities have greater effects on income in adulthood. In particular, we do not find a positive correlation with time spent on educational activities (in general, the point estimates are negative), inconsistent with differences in cultural values about education, and inconsistent with differences in the production of human capital directly caused by time expenditure on education. Additionally, we do not find that people of any age spend additional time on social or community activities in high-effect communities, which is not consistent with the view that high-effect communities are primarily distinguished by stronger social ties and greater social capital.

## References

- [1] Benabou, Roland. (1993). Workings of a City: Location, Education, and Production. *Quarterly Journal of Economics*, 108, 619-652.
- [2] Caetano, G., Kinsler, J., Teng, H. (2019) Towards Causal Estimates of Children's Time Allocation on Skill Development. *Journal of Applied Econometrics*, 19(34), 588-605.

- [3] Chetty, R., Hendren, N. (2018a). The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects. *Quarterly Journal of Economics*, 133(3), 1107-1162.
- [4] Chetty, R., Hendren, N. (2018b). The Impacts of Neighborhoods on Intergenerational Mobility II: County Level Estimates. *Quarterly Journal of Economics*, 133(3), 1163-1228.
- [5] Chetty, R., Hendren, N., & Katz, L. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4), 855-902.
- [6] Fletcher, J., Han, J. (2019) Intergenerational Mobility in Education: Variation in Geography and Time. *Journal of Human Capital*, 13(4), 585-634.
- [7] Guryan, J., Hurst, E., & Kearney, M. (2008). Parental education and parental time with children. *Journal of Economic Perspectives*, 22(3), 23-46.
- [8] Hill, C. R., & Stafford, F. P. (1974). Allocation of time to preschool children and educational opportunity. *Journal of Human Resources*, 323-341.
- [9] Hill, C. R., & Stafford, F. P. (1980). Parental care of children: Time diary estimates of quantity, predictability, and variety. *Journal of Human Resources*, 219-239.
- [10] Huang, L., Pergamit, M., Shkolnik, J. (2001). Youth initiation into the labor market. *Monthly Labor Review*, 124.
- [11] Kalil, A., Ryan, R., & Corey, M. (2012). Diverging destinies: Maternal education and the developmental gradient in time with children. *Demography*, 49(4), 1361-1383.
- [12] Katz, L., Kling J., Liebman J. (2001). Moving To Opportunity In Boston: Early Results Of A Randomized Mobility Experiment, *Quarterly Journal of Economics*, 116(2), 607-654.
- [13] Kofman, Y., Bianchi, S. (2012). Time use of youths by immigrant and native-born parents: ATUS results. Monthly labor review / U.S. Department of Labor, Bureau of Labor Statistics, 135, 3-24.
- [14] Mookerjee, S., & Slichter, D. (2020). Test Scores, Schools, and the Geography of Economic Opportunity. Working paper.
- [15] Rothstein, J. (2019). Inequality of Educational Opportunity? Schools as Mediators of the Intergenerational Transmission of Income. *Journal of Labor Economics*, 37(S1), 85-123.

- [16] Sayer, L. C., Bianchi, S. M., & Robinson, J. P. (2004a). Are parents investing less in children? Trends in mothers' and fathers' time with children. *American Journal of Sociology*, 110, 1-43.
- [17] Sayer, L. C., Gauthier, A. H., & Furstenberg Jr, F. F. (2004b). Educational differences in parents' time with children: Cross-national variations. *Journal of Marriage and Family*, 66(5), 1152-1169.

# Appendix

## Theoretical Model

To guide our interpretation of the empirical results, we develop the following simple model of time allocation.

Suppose that families have  $m$  ways to spend their time, and choose a vector of time allocations to activities  $(t_1, \dots, t_m)$  to maximize human capital production  $y$  subject to a time constraint  $t_1 + \dots + t_m = 1$  and the constraint imposed by the human capital production function  $y \leq f(t_1, \dots, t_m)$ . Assume that  $f$  is differentiable and the technology set is convex.

Let  $t^* = (t_1^*, \dots, t_m^*)$  be the human capital-maximizing time use. By the Envelope Theorem, the marginal product  $MP_k := \frac{\partial f}{\partial t_k}$  evaluated at  $t^*$  is equal for all activities  $t_k$ .

Suppose that, for some activity  $k$ , we altered the technology of human capital production to increase  $MP_k$  at all time allocation bundles. Because the technology set is convex, every activity has a diminishing marginal product. Therefore, the new optimal input bundle  $t'$  would have a greater time use on activity  $k$  than  $t^*$  does.

This suggests that improvements in the quality of an activity are likely to result in an increase in the quantity of time spent on that activity. However, there are three caveats to this claim.

First, if one were to change the production function in a way that increases the marginal product of activity  $k$  only at levels of  $t_k < t_k^*$ , then this change to the production function would make activity  $k$  more productive without increasing time expenditure on activity  $k$ . Therefore, not all differences in activity quality will show up in quantities.

Second, time use reflects *comparative* advantages in quality, not *absolute* advantages. If one were to change the production function by multiplying all marginal products by the same scalar (e.g. by changing  $f$  to  $f' = .9 * f$ ), then one would not alter the optimal time allocation.

Third, this model assumes that families are maximizing children's human capital production. In practice, time use is also motivated by competing interests such as consumption value, institutional constraints, and social norms. To the extent that these competing motivations vary across communities, we would expect to see time use vary across communities in a way that reflects underlying values.

We can incorporate this final consideration into our framework, as follows. Suppose that, instead of maximizing human capital production, families choose a time allocation to maximize a utility function  $U = \alpha f(t_1, \dots, t_m) + (1 - \alpha)c(t_1, \dots, t_m)$ , where the function  $c$  reflects the value that families derive from other considerations such as consumption. Let  $\tilde{t}$  be the time allocation that a family would select with  $\alpha = 0$ , i.e. considering only this consumption value.

In communities where  $\alpha$  is greater, time use will be more closely aligned with  $t^*$ ; and, where  $\alpha$  is smaller, time use will be more closely aligned with  $\tilde{t}$ . That is,



time use reflects the extent of prioritizing children's human capital accumulation. Furthermore, time use differences across places might reflect differences in  $c$ : e.g., in particularly religious communities, we would expect to see people spend more time in church, regardless of its value to human capital accumulation.

Summing up, time use likely reflects a combination of (i) comparative advantages in human capital production, (ii) differences in the priority placed on child development, and (iii) differences in values related to the consumption value of activities.