Essays in Labor and Family Economics

by

Steven Lann

Bachelor of Arts in Economics, University of Florida, 2018

Master of Science in Economics, University of Pittsburgh, 2020

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Steven Lann

It was defended on

March 27th 2024

and approved by

Marla Ripoll, Department of Economics

Daniele Coen-Pirani, Department of Economics

Douglas Hanley, Department of Economics

Elizabeth Caucutt, University of Western Ontario Department of Economics

Essays in Labor and Family Economics

Steven Lann, PhD

University of Pittsburgh, 2024

This dissertation explores two topics: workplace flexibility policies and parental investment into child human capital. The first chapter studies how parents' time and goods investments into their children are affected by household income and family size. The second chapter investigates the mechanisms determining the provision of flexibility policies, as well as the welfare effects of these policies. The third chapter examines how flexibility is measured in the literature, and develops occupation-level proxy measures of flexibility for use by researchers. In Chapter 1, I introduce a model of parental investment in child human capital featuring multiple investment types, static and dynamic investment choices, and multiple children. I find evidence that children of low-income parents experience a larger quality-quantity trade-off compared to the children of high-income parents. The estimated model can replicate this and other patterns of parental investments and child outcomes. Chapter 2 provides the first economic analysis of the allocation and welfare effects of workplace flexibility policies. Using data on workplace flexibility and other non-wage amenities from the ATUS and CPS, I find that among all the amenities studied, flexibility uniquely has a significant relationship with workers' intensive labor supply and timing of labor hours. I also find evidence that employers take this labor supply endogeneity into account when choosing to offer flexibility. Next I develop a model of a labor market in which heterogeneous firms compete for workers via bundles of wages and flexibility policies. Analysis of data simulated from the calibrated model reveals a compensating wage differential of flexibility of about 7 percent. Welfare analysis shows that workplace flexibility policies strongly benefit women by reducing the gender wage gap and female unemployment, and provides a net benefit for male workers and low-productivity firms. In Chapter 3, I use the American Time Use Survey to develop four workplace flexibility policy variables that capture several important dimensions of flexibility. I develop occupation-level measures of these flexibility policies by employing a machine learning approach. These flexibility policy indices are available for 481

Census occupation codes and are shown to be highly predictive of occupation-level access to workplace flexibility policies.

Table of Contents

List of Tables

List of Figures

Preface

This dissertation is the product of not only my own hard work, but the collective guidance and support of countless people in my life. I am incredibly thankful for my committee chair, Marla Ripoll, for her patience and attentiveness for the past five years. Her guidance was critical for not only the completion of the three chapters, but also for helping me navigate the completion of the important milestones of the program. Daniele Coen-Pirani was a constant during my time here at Pitt, and as a professor and committee member he always asked the difficult questions and gave me the feedback that I needed to hear. I am grateful for him pushing me to produce better research than I thought I was capable of. Douglas Hanley provided invaluable advice on a number of technical matters, and was a constant source of great questions and insightful comments during my seminars. Elizabeth Caucutt was also an excellent external committee member.

I am also thankful for a number of my classmates. In particular, I owe much of my success in the program to my mentor David Agorastos. I have learned something from every conversation we've had, and his endless supply of wisdom, humor and positivity were instrumental to me making progress on my research during the dog days of COVID. Benjamin Matta and I became close friends playing video games together every day at the beginning of the pandemic. That friendship has endured in the years since, and he was been a constant source of levity, companionship and game theory expertise. Joung Yeob-Ha is one of my oldest friends here at Pitt, and is one of the kindest, most good-natured people I know. Others who were important to my journey are Maddie Arnsbarger, the world's best prelim study partner, and my awesome office-mates Andrew Holt and Hyung-Jin Kim.

Above all else, I owe all my successes to my loved ones. My parents, Julie and Mark Lann and their spouses, Jorge Fiallo and Kerri Lann, have provided me with immesurable love and support, and I would not be the man I am today without all four of my parents. My brother Matthew Lann is one of my greatest inspirations, and it has been the pleasure of my life for us both to be able to overcome all our challenges and graduate together. Last but certainly not least, this would not have been possible without my incredible girlfriend, Lexie

Miller. Her love and belief in me have given me the strength to persevere through the most difficult and uncertain times in my program. She is an incredible listener, and by listening to my ramblings about my research she has helped me more than she can ever realize. The completion of this dissertation marks the beginning of an incredible new chapter in our life together, and I thrilled to see what is in store for us next.

1.0 Parental Human Capital Investments and the Quality-Quantity Trade-Off

1.1 Introduction

Since the pioneering work of Becker and Tomes (1979) and Becker and Tomes (1986), parental investments in children have been recognized as a key mechanism driving the intergenerational persistence of economic status. In the last decade, papers studying intergenerational mobility have begun to incorporate the insights of Cunha and Heckman (2007) and Cunha et al. (2010) by explicitly including childhood human capital formation and parental investments of time and goods in their models. These new works reveal the extent to which disparities in child development are affected by the quantity and quality of investments provided by parents. Lee and Seshadri (2019) show that about 20% of the variance in lifetime earnings and over 50% of the variance in child human capital is explained by parental human capital and wealth before their children are born. The channel of parental time investments has been found to account for almost 40% of the observed intergenerational persistence of earnings in the United States (Yum (2018)). A sophisticated understanding of how parents choose investments in their children is thus essential to addressing important questions regarding income mobility and child development.

Del Boca et al. (2014) and Caucutt et al. (2020) characterize parents' choices of time and goods investments in the context of a model of household behavior. Up to now, this relatively nascent literature has not had much to say about how parents trade off goods and time investments at different levels of income, and there has been little treatment of parental investments in the context of heterogeneous family size. It would be desirable for models of human capital formation and parental investment to incorporate heterogeneous family size due to the importance of family size on child outcomes: While the phenomenon of the quality-quantity trade-off has long been shown in developed countries, Bagger et al. (2013) and Juhn et al. (2015) have recently established its presence in developed countries. Though family sizes are declining in the United States, these results are still relevant today; nearly 40% of families in the NLSY97 survey have three or more children.

In this paper I introduce a model of parental investments into child human capital with heterogeneous family size. Like Del Boca et al. (2014) and Caucutt et al. (2020), it is a life cycle model with a household comprised of two parents who invest time and goods into their children's human capital over a number of periods. My model has the most in common with Caucutt et al. (2020), whose model uses a similar child human capital production function and also includes an asset choice with borrowing constraints. I innovate on these models by explicitly incorporating heterogeneous family size. I am able to extend the model to an arbitrary number of children by assuming that all children in a family are identical, are all born at once, and receive the same investments. The number of children that a household has is exogenous. The initial human capital of the children is correlated with parental human capital. Family size enters into the model through CES utility over the quality and quantity of children as well as through economy of scale multipliers on the per-child investment choices in the budget constraint. Another key assumption is that the productivity of parents' time investments may be increasing in their human capital.

Examining the model analytically shows the forces behind parent's choices of investments. In order to determine the static choice of the optimal mix of input types, parents weigh the relative opportunity costs and productivities of each type of investment. The shadow price of mother's time versus father's time versus goods depends on the parents' relative wages, and is affected by family size through the relative economies of scale for each investment type. The dynamic choice of the optimal sequence of investments over time depends on relative productivities and costs of investments over time, as well as the interest rate on assets, credit access, and the degree of complementarity between present and future investments in the child human capital function. The choice of investments over time also implies a choice of the household's asset profile; households that choose to invest a lot early in childhood also choose to borrow rather than save during this time.

Next, I present novel stylized facts on the joint distributions of household income, family size, parental time and goods investments and child outcomes. First, I show evidence of heterogeneity in the quality-quantity trade-off. Using data from the NLSY97 and the PSID, I find that while low-income families face a negative relationship between family size and child outcomes, this relationship is weakened for high-income families. This is true for both child pre-labor market human capital as measured by the ASVAB in the NLSY, and child labor market outcomes as measured by lifetime income from the PSID.

I then document patterns in how parents across the income and family size distributions allocate goods and time to their children. Per-child time investments are found to not be increasing monotonically in household income; in fact, I find little variation in parental time investments across the income distribution. Total household expenditures on education goods and childcare are found to be nearly constant in family size for low-income households and increasing in family size for high-income households. In addition, I find evidence that investments exhibit economies of scale in family size; goods and time may be shared amongst the children in a household.

I estimate the model using a joint estimation procedure: The child human capital investment function is estimated by GMM, while the remaining parameters are calibrated to match the stylized facts I documented as well as other important features of parental investments. The calibrated model succeeds in generating the heterogeneity in the quality-quantity trade-off seen in the data.

Analysis of the choices and outcomes of different household types along with sensitivity analyses reveals three mechanisms driving the heterogeneity. First, parents with high human capital are found to have more productive time investments and are thus better able to take advantage of economies of scale as family size increases. Second, I find that child quality and quantity are complements in the parents' utility function, which leads to variation in the child quality-consumption margin: When family sizes are large, the complementarity between child quality and quantity encourages parents to invest heavily into their children's human capital. Higher income parents consume more and have a lower marginal utility of consumption, and are thus more willing to substitute consumption for child quality when faced with a large family size. Conversely, low-income, low-consumption parents have a high marginal utility of consumption and may be less willing or able to sacrifice consumption, resulting in a greater quality-quantity trade-off. Third, the complementarity between child quality and quantity generates heterogeneity in parents' investment choices by initial child quality. Parents give more investments to children with high initial quality compared to children with lower initial quality due to the dynamic complementarity of the child human capital function. Due to the complementarity between child quality and quantity, parents of high initial human capital children also increase their investments as family sizes increase, while investments for lower human capital children are declining with family size.

In Section 2, I introduce the model and discuss analytical results. Section 3 describes the data used to generate the stylized facts and moments used for calibration. Section 4 presents evidence of heterogeneity in the quality-quantity trade-off as well as empirical patterns in parental time and goods investments. Section 5 discusses the estimation strategy and presents the results of the estimation. In Section 6, I investigate the choices of agents in the model and discuss mechanisms. Section 7 shows the results of counterfactual exercises. Section 8 concludes.

1.2 Model

I present a life-cycle model in which two agents jointly work, raise children and accumulate assets. In the parenthood period, the agents face a complex series of trade-offs. There are both static and dynamic choices of investments in children; parents must not only determine the optimal mix of investments in any given period, but also the optimal path of investments over time. There is also a trade-off between time investments and consumption, as any time spent with children comes at the expense of labor time. Investment in children can be considered a form of saving, as the parents derive utility from the quality of their children after the parenthood period is over. This implies that there is also a choice between asset accumulation (or borrowing) and “saving” through investing in children. These choices are all affected by the parents’ human capital and number of children.

1.2.1 Description of model

The model is comprised of a married couple (denote the mother’s variables with m and the father’s with f). The agents are age 30 at time $t = 0$[1], when they are endowed with n

[1]Since all children are born at once in this model, I choose a starting age that is in between the median age for parents at first birth and last birth. In 2016 the median age of mothers at first birth was 26.9 years

children of initial quality q_0, human capital (h_m, h_f), and period 0 assets b_0. Each period in the model corresponds to one year, and the agents' life cycle proceeds as follows:

1. Parenthood, $t = 0, \ldots, T-1$: Parents invest in their children's human capital through time investments $(\tau_{m,t}, \tau_{f,t})$ and goods investments g_t. Parents work whenever they are not spending time with their kids, and also choose household consumption c_t and assets b_{t+1}. In period $T-1$, parents receive utility from the final quality and quantity of their children $V(q_T; n)$.
2. Post-parenthood, $t = T, \ldots, R$: Agents supply labor inelastically and choose household consumption c_t and assets b_{t+1}.
3. Retirement, $t = R+1, \ldots, M$: Retired agents consume from their savings and die after period M.

The two agents jointly derive utility from consumption and their children's quality[2] :

$$U_0 = \sum_{t=0}^{M} \beta^t U(c_t) + \beta^{T-1} V(q_T; n) \tag{1}$$

Utility over consumption is CRRA and the utility derived from children is CES, allowing the quality and quantity of children to be either complements or substitutes:

$$\begin{aligned} U(c_t) &= \frac{c_t^{1-\sigma}}{1-\sigma}, \\ V(q_T; n) &= \left[\nu n^{\rho} + (1-\nu) q_T^{\rho}\right]^{\frac{1}{\rho}} \end{aligned} \tag{2}$$

of age, while the median age at last birth was 31 (Carlson and Guzzo (2021),Schweizer and Guzzo (2020)). There is considerable heterogeneity in the timing of first and last births by parity and parental education, so for simplicity I choose the starting age to be 30.

[2]Parents do not spend time on leisure in this model. The inclusion of leisure would not affect the current model's most important analytical results: the within-period proportionality between investment types (equations (6) and (7)) and the intertemporal investment choice (equation 8). Leisure primarily affects the overall levels of investments, which are calibrated to match the levels of investment in the data. In addition the inclusion of leisure will introduce household bargaining into the model via differing weights on each parents' leisure utility, increasing the complexity of solving and calibrating the model. I thus abstract from leisure in order to focus on the other important trade-offs in the model.

The household budget constraint for each stage of life is given below:

$$
\begin{aligned}
&(H - \Lambda_m(n)\tau_{m,t})w_{m,t}h_m + (H - \Lambda_f(n)\tau_{f,t})w_{f,t}h_f + (1+r)b_t = c_t + b_{t+1} + n^{\phi_g}g_t, \quad t \leq T-1 \\
&\tilde{H}w_{m,t}h_m + \tilde{H}w_{f,t}h_f + (1+r)b_t = c_t + b_{t+1}, \qquad T \leq t \leq R \\
&(1+r)b_t = c_t + b_{t+1}, \qquad R+1 \leq t \leq M
\end{aligned} \tag{3}
$$

Each period, the agents receive labor income and interest on their assets if $b_t > 0$, and spend on consumption and interest on their debt if $b_t < 0$. Agents can also select future assets $b_{t+1} \geq \bar{b}$. Wages are deterministic and depend on the human capital of the agents. $H = 3200$ hours are split between work and time with children $(\tau_{m,t}, \tau_{f,t})$[3]. Parents also purchase goods g_t for their children[4].

The objects $(\Lambda_m(n), \Lambda_f(n), n^{\phi_g})$ scale up the per-child investments $(\tau_{m,t}, \tau_{f,t}, g_t)$ to total investments. The total time investment that parent i spends with their n children is given by $\Lambda_i(n)\tau_{i,t}$, with $1 \leq \Lambda_i(n) \leq n, \ i = m, f$. For instance, for $n = 2$, giving one unit of the mother's time to each child costs $1 \leq \Lambda_m(2) \leq 2$ total time. This term being less than the number of children reflects the fact that some parental time is shared between children. The derivation of these terms is discussed in Section 5. Total goods investments are given by $n^{\phi_g}g_t$ with $\phi_g \leq 1$ representing the economies of scale of the goods investments[5].

I abstract away from within-household inequality: all children are born at the same time, are identical in their initial endowments of q_0 and human capital functions, and receive the same per-child investments. This set of assumptions helps make the model more tractable while still allowing me to rationalize the observed aggregate patterns in the quality-quantity trade-off and parental investment decisions. There is also a lack of complete data on investments and child outcomes for all children in households with three or more children, which would complicate the estimation of the model with within-household inequality.

The decision to not include a fertility choice also came down to a matter of scope. The

[3]In the post-parenthood period, the agents work $\tilde{H} = 40 \times 50 = 2000$ hours per year.

[4]These goods investments include childcare expenditures, which differs from Caucutt et al. (2020) who separate goods investments from childcare expenditure in their model.

[5]The intuition for $\phi_g \leq 1$ is the same as for $1 \leq \Lambda_i(n) \leq n$, which is that some goods are shared. The $\Lambda_i(n)$ parameters can also be interpreted as reflecting the economies of scale of time investments. Due to data limitations I cannot estimate ϕ_g in the same way as I estimate the $\Lambda_i(n)$'s, so I gave them different symbolic representations to reflect this distinction.

primary aim of the model is to characterize how parental investments and child outcomes are affected by family size and parental economic status. Since I abstract from heterogeneous timing of births, the fertility choice would need to be implemented as a "period-0" decision, with all investment choices still being made afterward taking family size as given as in the current model. In order to rationalize the distribution of family size by household income seen in the data I would need to add utility shocks to households' tastes for n, the distributions which depending on the household's type. This added feature would complicate the estimation of the model while not affecting parents' investment choices. Also, Jones and Tertilt (2007) show that the relationship between income and fertility has been declining over time, which alleviates some concern about selection affecting the stylized facts I want to rationalize.

1.2.1.1 Child human capital formation

The human capital of each child evolves as follows:

$$\begin{aligned} q_{t+1} &= \mu X_t^{\theta_1} q_t^{\theta_2} \\ X_t &= \left[\alpha_{m,t}(h_m^{\psi}\tau_{m,t})^{\gamma} + \alpha_{f,t}(h_f^{\psi}\tau_{f,t})^{\gamma} + \alpha_{g,t}g_t^{\gamma}\right]^{\frac{1}{\gamma}} \end{aligned} \tag{4}$$

In addition to within-period good and time investments X_t, future child human capital depends on present human capital. Notice that this functional form gives the child human capital function the properties of dynamic complementarity ($\partial^2 q_{t+1}/\partial q_t \partial X_t > 0$) and self-productivity ($\partial q_{t+1}/\partial q_t > 0$). The different types of period-t investments will be either complements or substitutes depending on whether γ is greater than or less than zero. Another important feature of the within-period investment function are the h_i^{ψ} terms. When $\psi > 0$, parents' human capital augments their time investments; the effectiveness of parental time in improving child quality is increasing in parents' human capital.

The share parameters $\alpha_{i,t}$ are changing over time and do not necessarily sum to one. In order to represent the human capital function with share parameters that sum to one in each period, define $\tilde{\alpha}_{i,t}$ by:

$$\tilde{\alpha}_{i,t} = \frac{\alpha_{i,t}}{\sum_j \alpha_{j,t}}.$$

We can now normalize the human capital production function as follows:

$$q_{t+1} = \mu A_t \Big[\tilde{\alpha}_{m,t}(h_m^{\psi}\tau_{m,t})^{\gamma} + \tilde{\alpha}_{f,t}(h_f^{\psi}\tau_{f,t})^{\gamma} + \tilde{\alpha}_{g,t}g_t^{\gamma}\Big]^{\frac{\theta_1}{\gamma}} q_t^{\theta_2},\ A_t = \Big(\sum_j \alpha_{j,t}\Big)^{\frac{\theta_1}{\gamma}} \tag{5}$$

1.2.2 Sequential form of model

The agents solve the following sequential problem:

$$\max_{\tau_{m,t},\tau_{f,t},g_t,c_t,b_{t+1}} \sum_{t=0}^{T-1} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma} + \beta^{T-1}\big[\nu n^{\rho} + (1-\nu)q_T^{\rho}\big]^{\frac{1}{\rho}} + \sum_{t=0}^{T-1} \lambda_t(b_{t+1} - \bar{b})$$

s.t.

1a) $(H - \Lambda_m(n)\tau_{m,t})w_{m,t}h_m + (H - \Lambda_f(n)\tau_{f,t})w_{f,t}h_f + (1+r)b_t = c_t + b_{t+1} + n^{\phi_g}g_t,$

$t \leq T-1$

1b) $\tilde{H}w_{m,t}h_m + \tilde{H}w_{f,t}h_f + (1+r)b_t = c_t + b_{t+1}, \quad T \leq t \leq R$

1c) $(1+r)b_t = c_t + b_{t+1}, \quad R+1 \leq t \leq M$

2) $q_{t+1} = \mu A_t \Big[\tilde{\alpha}_{m,t}(h_m^{\psi}\tau_{m,t})^{\gamma} + \tilde{\alpha}_{f,t}(h_f^{\psi}\tau_{f,t})^{\gamma} + \tilde{\alpha}_{g,t}g_t^{\gamma}\Big]^{\frac{\theta_1}{\gamma}} q_t^{\theta_2}$

1.2.3 Analytical results

1.2.3.1 Static investment choice margins

Combining the first order conditions for the different investment types give the following results:

$$\frac{\tau_{m,t}}{\tau_{f,t}} = \left(\frac{h_f}{h_m}\right)^{\frac{\psi\gamma}{1-\gamma}} \left[\frac{\Lambda_f(n)}{\Lambda_m(n)} \frac{w_{f,t}h_f}{w_{m,t}h_m} \frac{\alpha_{m,t}}{\alpha_{f,t}}\right]^{\frac{1}{1-\gamma}} \tag{6}$$

$$\frac{\tau_{i,t}}{g_t} = h_i^{\frac{\psi\gamma}{1-\gamma}} \left[\frac{n^{\phi_g}}{\Lambda_i(n)} \frac{1}{w_{i,t}h_i} \frac{\alpha_{i,t}}{\alpha_{g,t}}\right]^{\frac{1}{1-\gamma}},\ \text{for } i \in \{m, f\} \tag{7}$$

Since each type of period t investment affects future child quality in the same way, the choice of the optimal mix of period t investments simplifies to a static choice. There are three mechanisms mediating the optimal mix of investments: the relative shadow prices and the relative productivities of each type of investments, and the elasticity of substitution between the investments. While a nonzero amount of each type of investment must be provided, investment types with a higher opportunity cost will be given in a smaller proportion than those with a lower opportunity cost. We can see from both equations that the share of each parent's time investments is decreasing in their wage. The relative shadow prices of investments also depend on family size and the relative values of the economy of scale parameters $\Lambda_m(n)$, $\Lambda_f(n)$, and ϕ_g; if for instance $\Lambda_m(n) < \Lambda_f(n)$, then mother's time exhibits greater economies of scale than father's time, and will thus have a larger share of total investments as family size increases.

Next, parents must weigh the relative productivities of the different types of investments; investment types that are more productive are given in greater proportions. Finally, the optimal investment mix depends on the elasticity of substitution between the investments. To demonstrate this, consider extreme values of γ. If γ approaches negative infinity, the ratios of investments will approach one, implying that the investments are perfect complements. On the other hand, if γ approaches one, the ratios will approach either zero or infinity, implying perfect substitutes.

1.2.3.2 Dynamic investment choice margin

Iterating the first order conditions for investments forward and combining with Euler equation for consumption allows investments to be related across time[6]:

$$\frac{\tau_{i,t+1}}{\tau_{i,t}} \leq \left[\frac{1+r}{\theta_2}\frac{w_{i,t}}{w_{i,t+1}}\frac{\tilde{\alpha}_{i,t+1}}{\tilde{\alpha}_{i,t}}\right]^{\frac{1}{1-\gamma}}, \text{ holding with equality if } \lambda_t = 0 \tag{8}$$

There are a number of forces affecting the dynamic choices of investments. Unlike the static choice of the optimal investment mix, the choice of increasing versus decreasing investments over time has implications for the household's asset choice. In the case that a household chooses to increase investments over time, they will save during early childhood, then dis-save in adolescence in order to fund larger investments and smooth consumption. The opposite will happen for a household that chooses decreasing investments: they will provide a high

[6]Since investments in a given period are proportional, I will only look at the dynamics of one type of investments; analogous results for the other investment types can be derived in a similar way.

level of investments in early childhood and will want to borrow if possible. A corollary of this result is that credit access affects parents' intertemporal choice: a binding borrowing constraint leads to a flatter profile of time investments.

Like the static investment choice, relative shadow prices and productivities across time matter. As parents' wages are assumed to be increasing over time, the opportunity cost of time investments also increases over time, encouraging high early investments and borrowing. Since wage growth is generally increasing in human capital, we should expect to see heterogeneity in the investment profiles: high human capital parents will provide more investments early in life and have a steeper declining investment profile compared to lower human capital parents. If the productivity of parental time investments are decreasing over time as reported by Del Boca et al. (2014), parents want to borrow if possible and invest more time in early childhood.

The dynamic components of the household budget constraint and child human capital function also play a role in determining the optimal sequences of investments. In particular, the interest rate of assets $(1 + r)$ and the productivity of current child human capital (θ_2) affect the choices of investments and assets. A high interest rate will encourage saving and discourage borrowing. Greater saving early in childhood incentivizes parents to invest little in their children early and more later. If the productivity of the previous period human capital is high, there are large returns to investing heavily in children early in childhood. Parents will respond with large early investments and borrowing if possible.

Note that this equation does not say anything about the overall level of investment, only the change in the level of investment over time. Family size will affect the overall level of investment though the economy of scale parameters which act as a "price" of parental time for a given number of children as well as through the parents' utility over child quality and quantity. Heterogeneity in the productivity of time investments along with the positive relationship between parent human capital and the opportunity cost of time investments will result in parent human capital also affecting the overall level of investments.

1.3 Data

1.3.1 PSID and CDS

The Child Development Supplement of the PSID (CDS) is a survey series started in 1997. The original CDS cohort consisted of a sample of children aged 0 to 12 from PSID families. Follow-up surveys for the 1997 cohort were administered in 2002 and 2007, and in 2014 a new cohort was selected.

The CDS collects detailed information on both the children and their families. A key feature of the survey is the time diary, in which parents catalog all the activities of the child for one weekday and one weekend day. These diaries detail the duration of each activity and, crucially, who else in the family participated. Using the same method as Del Boca et al. (2014) and Caucutt et al. (2020), I constructed a measurement of estimated weekly parental hours spent with each child. Since the time diaries track time spent with a particular child, the unit of measurement is time spent per-child. Some households have time diaries for multiple children, but most only track time spent with one child. This measure is split by mother and father time, and by active and passive participation in the child's activities. In this analysis I focus on the time that parent spend actively participating. Measures of weekly time investment are constructed using the method from Lee and Seshadri (2019). I use a sample of about 2,700 observations from all four waves of the CDS. I restrict the sample to only include two-parent households with four or fewer children.

The PSID began tracking household consumption of a number of goods in 1999. Among those are various forms of goods that are used to invest in child human capital. Namely, the PSID tracks childcare and education expenditures[7], as well as property tax expenditures[8]. In order to assess the joint distributions of family size, income and child expenditures, I use a panel spanning the 1999 to the 2017 waves of the PSID. In addition to the expenditure variables, the panel includes household income, number of minors in the household as well as

[7] Since I cannot observe what is spent on each child, the unit of measurement is total household expenditures.

[8] Property taxes are commonly used as a proxy for public school expenditures. It is common for parents to choose where to raise their children based on the quality of public schooling. Homes located in better school districts generally cost more than ones in worse school districts. Thus a household's level of property taxes indicates the family's implicit spending on public schooling.

the number of college-aged children that the parents have. I calculate the permanent income of each household using the method from Altonji et al. (1997), and use this as the measure of household income in the following analyses.

1.3.2 NLSY97

The 1997 National Longitudinal Survey of Youth (NLSY97) is a longitudinal survey run by the U.S. Bureau of Labor Statistics. The NLSY97 cohort is a nationally representative sample of nearly 9000 individuals born from 1980 to 1984. With 18 rounds of surveys published to date, the study provides a wealth of data about the respondents and their families, including data on household income and family size. With collaboration from the Department of Defense, the Armed Services Vocational Aptitude Battery was administered to about 80 percent of NLSY97 respondents[9]. The ASVAB is an aptitude test used by the U.S. Armed Forces for helping determine the qualification of recruits for enlistment. The scores of the NLSY97 children on the Mathematical Knowledge, Arithmetic Reasoning, Word Knowledge and Paragraph Comprehension sections are aggregated and a percentile score for each participant is reported in the 1999 survey. Also known as the Armed Forces Qualification Test (AFQT), this aggregate score is an often-used measure of pre-labor market human capital (Dickens and Flynn (2006)).

I use data from the first round of the NLSY97 as well as the ASVAB scores published in 1999. In addition to restricting the sample to respondents from households with two parents, I omit households with more than four children. Selecting only for households with married parents results in a slightly higher average parental education, income, and child ASVAB scores than the full sample. The median family in the sample has two children, but almost 40 percent of households have three or more children.

[9]The NLSY cohort was selected to test the efficacy of the new ASVAB exam because it was a nationally representative sample.

1.4 Stylized Facts

1.4.1 Heterogeneity in the quality-quantity trade-off

The following results are motivated by Juhn et al. (2015), who find preliminary evidence that the quality-quantity trade-off is affected by characteristics of parents. Using data from the NLSY79, they find that children with low-ability mothers experience a greater negative effect from the arrival of a younger sibling compared to children with high-ability mothers. In this section I present further evidence of heterogeneity in the quality-quantity trade-off. Using data from the NLSY97 and the PSID, I find that while low-income families face a negative relationship between family size and child outcomes, this relationship significantly weakens or even disappears for high-income families. This is true for both child pre-labor market human capital as measured by the ASVAB in the NLSY, and child labor market outcomes as measured by lifetime income from the PSID.

1.4.1.1 NLSY97

The analysis consists of a simple regression of the following form:

$$\text{ASVAB}_i = \beta_0 + \beta_1 \text{HH\_INC}_i + \beta_2 \text{NUMKIDS}_i + \beta_3 \text{HH\_INC}_i \times \text{NUMKIDS}_i + \beta_4 X_i + \epsilon_i \quad (9)$$

ASVAB_i is the ASVAB percentile score for respondent i. HH_INC_i is the logged 1997 household income of respondent i, and NUMKIDS_i refers to the number of children in the household. X_i is a vector of child and household characteristics, including child age and sex, age of the mother and 1997 household wealth. The coefficient of interest is β_3, the coefficient on the interaction term between household income and number of children.

Table 1 shows the estimates of how household income and family size affect child ASVAB scores. The significant positive coefficient on the interaction term implies that the quality-quantity trade-off is heterogeneous with household income; an increase in income reduces the detrimental effect of family size on child human capital. Using the coefficients from specification (2), we can quantify the comparative statics of an additional child and see that for the highest-income households the effect is in fact reversed: while an additional child decreases ASVAB performance by 0.78 percentile points for the lowest-income households, performance improves by 2.7 percentile points for the highest-income households. This suggests that there are substantial nonlinearities in the relationship between child quality and quantity for households across the income distribution.

1.4.1.2 PSID

The NLSY97 analysis gives insight into the heterogeneity of the quality-quantity trade-off, where "quality" is defined as pre-labor market human capital. I now turn my attention to the PSID to show that this result extends to child labor market outcomes. The PSID dataset is constructed similarly to Lee and Solon (2009), with a panel of child income and a snapshot of family variables from the child's teenage years.

The sample consists of the children of the first respondents of the PSID, born between 1952 and 1975. Household variables such as income and family size are collected from the years in which the children were 15 to 17 years old, and a panel of child income is collected starting from when they leave home and ending in 2017. Like with the NLSY97, I restrict the sample to two-parent households. The sample follows 685 children from 461 families, with a total of 13304 income-survey year observations. Due to a large share of the sample being from the Baby Boom generation, the median family has three children. Since nearly ten percent of the households have five children, I include these in the sample.

The specification I use is a modification of the one used in Lee and Solon (2009):

$$\begin{aligned} y_{ict} = {} & \alpha_t + \beta_1 \text{HH\_INC}_i + \beta_1 \text{HH\_INC}_i + \beta_2 \text{NUMKIDS}_i + \beta_3 \text{HH\_INC}_i \times \text{NUMKIDS}_i + \beta_4 X_{it} \\ & + \delta_1 A_i + \delta_2 A_i^2 + \delta_3 A_i^3 + \delta_4 A_i^4 + \gamma_1 (t-c-40) + \gamma_2 (t-c-40)^2 + \gamma_3 (t-c-40)^3 \\ & + \gamma_4 (t-c-40)^4 + \theta_1 \text{HH\_INC}_i (t-c-40) + \theta_2 \text{HH\_INC}_i (t-c-40)^2 \\ & + \theta_3 \text{HH\_INC}_i (t-c-40)^3 + \theta_4 \text{HH\_INC}_i (t-c-40)^4 + \epsilon_{ict} \end{aligned} \tag{10}$$

y_{ict} denotes the log family income of child i from birth cohort c in year t, A_i denotes the age of the parent, and $(t-c-40)$ is the age of the child in year t, normalized to be 0 at age 40. α_t is the year t fixed effect, and once again X_{it} is a vector of household and child controls.

The results of running specification (10) are given in Table 2. Unlike the NLSY97 results, the coefficient on parental income is not significantly different from zero. This means that the effect of parental income is fully captured by the interaction term. The significant and positive coefficient of the interaction term now carries two interpretations. The first is that intergenerational persistence of income is increasing in family size[10]. The next and more relevant interpretation is that the observed heterogeneity in the quality-quantity trade-off is still seen when we extend the definition of "quality" to include child labor market outcomes. This is evident when looking at the comparative statics of family size by income: while an extra child is associated with a 13.4% decline in child lifetime income for low-income households, this decrease is only about 2.2% for high-income households.

1.4.2 Investments by income and family size

At first glance, the comparative statics of parental investments at varying levels of income and family size seem clear. Per-child goods investments should be increasing in income and decreasing in family size. Per-child time investments should also be decreasing in family size, while the effect of wages are ambiguous, depending on the marginal utility of consumption versus child quality. If parents behaved in the way described above, we would expect to observe families of all income levels being subject to the quality-quantity trade-off. The observed heterogeneity contradicts this simple reasoning. This result instead implies that the relationship between parental investments, family size and economic status are not so clear-cut. In this section I will use data on parental goods and time investments from the PSID to document patterns in how parents allocate goods and time investments across the joint distributions of income and family size.

1.4.2.1 Time investments

As a baseline observation, I group households by income quartile and number of children and find the mean of mother and father per-child active time within each group. These

[10]This corroborates the finding of Daruich and Kozlowski (2020) that intergenerational persistence of income is affected by fertility choice

conditional means are plotted in Figure 1. Each family size/income bin is relatively large, with the smallest group having 157 observations and the largest having 452 observations. We see that for mothers there is a large decline in per-child time when family size increases from one to two children, then stays relatively flat for larger family sizes. Fathers' time seems to be much less elastic with respect to changes in household income or family size, with the spread of fathers' time being much less than that of the mothers'. An important observation is that parental time is not increasing monotonically with income; in fact, the highest-income parents seem to give less time than parents with lower incomes. In sum, this simple exercise suggests that there are some nonlinearities in the relationships between parental time investments, income and family size.

1.4.2.2 Shared versus unshared time

As stated above, an important consideration when studying how parental time inputs change with family size is the fact that siblings can share their parents' time. In order to find the degree to which parents' time is shared, I use a variable from the time diaries that denotes whether a sibling participated in a child's activities. With this variable I can further divide my previous measures of parental time into categories of shared and unshared time. In Table 3 I report the ratios of active shared time to total active time for both mothers and fathers. Due to measurement error, there is a nonzero reported amount of shared time for parents with one child. Parents with two or more children share 64% to 81% of their total time, with this percentage increasing slightly with family size. Fathers share slightly more of their time compared to mothers. The ratios of shared time are constant by parental income or education.

1.4.2.3 Goods (PSID)

Figure 2 shows how goods investments vary for households of different levels of income and number of children. The relationship between income and goods investments follows the expected pattern; higher income families have higher levels of expenditure for any number of children. There is heterogeneity in the relationship between expenditures and family size:

spending on education and childcare is roughly constant or slightly increasing in family size for low and middle-income households, but are increasing more in family size for the higher income households, in particular when going from one to two children. This implies that high-income parents may give their children fewer shared goods compared to lower-income parents.

1.5 Estimation

I employ a joint estimation procedure to estimate the parameters of the human capital investment function, $(\gamma, \alpha_{m,t}, \alpha_{f,t}, \alpha_{g,t})$, as well as the other parameters governing the agent's preferences and child human capital technology, $(\nu, \rho, \mu, \theta_1, \theta_2, \psi, \phi_g)$. The procedure is as follows: Each iteration of the estimation begins with a guess of $(\nu, \rho, \mu, \theta_1, \theta_2, \psi, \phi_g)$. Given the guess of (ψ, ϕ_g), I am able to identify $(\gamma, \alpha_{m,t}, \alpha_{f,t}, \alpha_{g,t})$ by GMM. Next, using the estimated X_t parameters I can solve and compute moments from the model in order to calibrate the rest of the parameters. The GMM and calibration methods are described below along with the derivation and estimation approach of the economy of scale parameters $\Lambda_i(n)$.

1.5.1 Estimating the human capital investment function

In order to estimate the parameters of X_t, I use a strategy similar to one employed by Caucutt et al. (2020) that makes use of the relative input results from Section 1.2:

$$\frac{\tau_{m,t}}{\tau_{f,t}} = \left(\frac{h_f}{h_m}\right)^{\frac{\psi\gamma}{1-\gamma}} \left[\frac{\Lambda_f(n)}{\Lambda_m(n)} \frac{w_{f,t}h_f}{w_{m,t}h_m} \frac{\alpha_{m,t}}{\alpha_{f,t}}\right]^{\frac{1}{1-\gamma}}$$

$$\frac{\tau_{m,t}}{n^{\phi_g} g_t} = h_m^{\frac{\psi\gamma}{1-\gamma}} \left[\frac{n^{\gamma\phi_g}}{\Lambda_m(n)} \frac{1}{w_{m,t}h_m} \frac{\alpha_{m,t}}{\alpha_{g,t}}\right]^{\frac{1}{1-\gamma}}$$

Note that I rearrange in order for the denominator of the LHS of second equation to have total goods expenditures, since only total household goods expenditures are observable in the PSID. With the exception of the $\alpha_{i,t}$ parameters, every variable in these equations is observable. Next I transform the $\alpha_{i,t}$ parameters into a form that can be estimated with the above strategy. Following the method used by Del Boca et al. (2014), I represent the $\alpha_{i,t}$ parameters as $\alpha_{i,t} = \exp(y_i + tz_i)$ for $i = \{f, g\}$. For simplicity I normalize $\alpha_{m,t} = 1$ for all t. My parameters of interest are $\Theta = (\gamma, y_f, z_f, y_g, z_g)$.

Finally, I form a new dataset containing both goods and time investments as well as parental education and family size from the PSID. This allows me to estimate Θ by generalized method of moments. The GMM estimator solves the following problem:

$$\hat{\Theta} = \arg\min_{\Theta} \mathbf{m}'(\Theta)\mathbf{W}^{-1}\mathbf{m}(\Theta), \tag{11}$$

where $\mathbf{m}(\Theta)$ are the moment conditions and $\mathbf{W}^{-1}$ is the optimal weighting matrix. My target moments are the mean the of mother/father time ratio $\frac{\tau_{m,t}}{\tau_{f,t}}$ by child age t, and the mean of the mother time/total goods ratio $\frac{\tau_{m,t}}{G_t}$ by child age t. Using data for children ages 1 to 16 gives me 32 moments to estimate 5 parameters. The optimal weighting matrix is estimated using a two-step variance-covariance estimator.

The results of the GMM estimation are presented in Table 4. I find that the substitution parameter of the investment function γ is negative, implying that mothers' time, fathers' time and goods investments are complements in the child human capital function. This result corroborates the findings of Caucutt et al. (2020), who estimate an elasticity of substitution between 0.4 and 0.5.

Figure 3 shows the sequences of the share parameters $\tilde{\alpha_{i,t}}$ implied by the estimates of (y_f, z_f, y_g, z_g). Similar to the estimated share parameters from Del Boca et al. (2014), the share of mothers' time is highest when children are young and declines steeply over time. The share of fathers' time starts lower than mothers' time and stays roughly constant over time. Unlike Del Boca et al. (2014), the initial value of the goods investments share is found to be the same as the mother's time share, but becomes the most significant input in the teenage years. This difference may be the result of differences in the investment function, in particular the addition of human capital augmenting for parental time in my model.

1.5.2 Time economy of scale derivation and estimation

Consider the time input of parent i into each child's human capital function, $\tau_{i,t}$, and the total time parent i spends with all their children, $T_{i,t}$. Parents share some time amongst all their children, $\tau_{i,t}^S$, and some time is given to each child individually, $\tau_{i,t}^U$. We can then express $\tau_{i,t}$ and $T_{i,t}$ in terms of shared and unshared time:

$$\tau_{i,t} = \tau_{i,t}^U + \tau_{i,t}^S \tag{12}$$

$$T_{i,t} = n\tau_{i,t}^U + \tau_{i,t}^S \tag{13}$$

Denote by $\phi_i(n)$ the share of $\tau_{i,t}$ that is shared:

$$\phi_i(n) = \frac{\tau_{i,t}^S}{\tau_{i,t}} = \frac{\tau_{i,t}^S}{\tau_{i,t}^U + \tau_{i,t}^S}$$

Note that this share depends on family size. We can now express $\tau_{i,t}^U$ as:

$$\tau_{i,t}^U = \tau_{i,t}^S \left(\frac{1 - \phi_i(n)}{\phi_i(n)} \right) \tag{14}$$

Finally, we can solve for $T_{i,t}$ as a function of $\tau_{i,t}$:

$$T_{i,t} = \Lambda_i(n)\tau_{i,t} \quad , \quad \Lambda_i(n) = n - n\phi_i(n) + \phi_i(n) \tag{15}$$

Table 5 shows the estimates for the $\Lambda_i(n)$ parameters, computed using the $\phi_i(n)$'s reported in Table 3. Notice that by construction the economies of scale are one for $n = 1$. These estimates imply that parental time investments have relatively high economies of scale; in order to give a unit of time to four children, parents only need to spend about 1.6 times as much total time as they would need to give that unit of time to one child. This is due to the empirical result from Section 1.4 that most of the parents' time is shared and that the proportion of shared time increases with the number of children. Due to fathers having a slightly higher proportion of shared time compared to mothers, their time is estimated to have greater economies of scale.

1.5.3 Calibration

1.5.3.1 Exogenous parameters

The CRRA parameter σ is set to a value of 1.5. I use standard values of 0.97 for the discount rate β and 0.035 for the interest rate r. The wage process of the parents is estimated using the PSID. Wage data from 1995 to 2017 were collected, along with years of education and age for both husbands and wives. These were regressed on logged wages, as shown in Table 6. For simplicity, I assume that both parents are age 30 in $t = 0$, initial assets b_0 are zero, and that $h_m = h_f = h$[11]. The parents' human capital h is a function of years of education: $h = \exp(0.0792(\text{educ} - 12))$, where 0.0792 is the estimated returns to schooling found by estimating the Mincer equation jointly for men and women in the PSID. Human capital is normalized to be one when the parents have a high school-level education. Since wages in the model are determined by h, household income and parental education are synonymous in the model. Agents in the baseline model are allowed to borrow up to 5 percent of the household's natural borrowing limit.

The distribution of initial quality q_0 depends on parental education and is estimated using Letter-Word scores in the CDS. The Letter-Word test is administered to children ages 3 to 12 in the 1997 CDS. In order to have the best proxy possible for initial ability, I consider the scores of children ages 3 to 7. The scores are age-standardized and normalized to be between 0 and 2. The mean and variance of the Letter-Word scores for each parental education group are reported in Table 7. The mean Letter-Word score is increasing with parental education and the scores are normally distributed.

1.5.3.2 Method of moments

The remaining parameters $(\nu, \rho, \mu, \theta_1, \theta_2, \psi, \phi_g)$ are calibrated using simulated method of moments (SMM). These seven parameters are targeted to eight moments from the data

[11] I make this simplifying assumption to greatly reduce the dimensionality of the problem; rather than solving the model for 16 combinations of mother and father education, I solve for four household education types. Several of those 16 education combinations have very few observations in the data due to couples' education being highly correlated. To alleviate concerns about external validity all calibration moments are computed using households whose parents have the same level of education.

concerning the allocation of time and goods from the PSID as well as child outcomes from the CDS Letter-Word tests. Since a key feature of the data is the heterogeneity in the relationship between child quality and quantity by parental economic status, my strategy in choosing moments was to target some moments for both college-educated parents and high school-educated parents.

The (ν, ρ) parameters determine the utility that parents derive from the quality and quantity of their children, and thus how investments and child quality change with family size. I thus target the ratio of total time investments between one and two-child families, for families with both college-educated parents and high school-educated parents. These moments will capture the documented facts regarding the heterogeneity in the quality-quantity trade-off.

(μ, θ_1) are the TFP of the child human capital function and the productivity of current investments X_t, respectively. These parameters govern the overall level of time and goods investments. In addition, the ψ parameter helps determine the level of investments as well as final child quality for college-educated parents compared to high school-educated parents. These three parameters are targeted to four moments: the average yearly time investments (in hours) for each parent education group and the average human capital of age-5 children from each group. I use as a proxy the 2002 Letter-Word scores of the same children in the sample described above.

The productivity of yesterday's quality θ_2 influences the dynamic choice of investments, and in particular the levels of investments over time. Thus I target the correlation between child age and time investments. ϕ_g is the economy of scale parameter for goods, and is calibrated to the ratio of goods investments between one and three and four-child families.

The estimates of the calibrated parameters are given in Table 26. The child utility substitution parameter ρ is negative, meaning that child quality and quantity are complements in the parents' utility function[12]. I find that the self-productivity of child quality θ_2 is greater than one. This suggests that not only do investments have a very high degree of dynamic complementarity, but child human capital actually appreciates over time. The

[12]Becker and Tomes (1976) shows that complementarities between child quality and quantity are not inconsistent with the existence of the quality-quantity trade-off in the data.

positive ψ parameter suggests that parental time investments are augmented substantially by parents' human capital; with the calibrated value $\psi = 0.4959$, the time investments of college-educated parents are 17% more productive than the time investments of high school-educated parents. The economy of scale parameter for goods ϕ_g is close to zero, meaning there are high economies of scale for goods investments. This is consistent with the finding from Section 1.4 that expenditures on education goods have little relationship with family size, especially for lower-income households.

1.5.4 Model fit

Figure 4 and Table 9 show the performance of the GMM estimation and the calibration with respect to their respective targeted moments. The model closely matches the input ratio moments from the PSID, in particular for the mother's time/goods ratio moments. The model obtains a close match for most parameters, with the exception of the ratios of time investments between one and two-child households. The college-educated parents decrease their time investments more than in the data, while high school-educated parents decrease their investments less than what is seen in the data.

Table 10 shows the results of the model in matching some selected untargeted moments. The model generated greater goods investments for the high school parents than what is seen in the data and fewer for the college parents. The positive relationship between parental education and final child quality is also weaker in the model compared to the data. While the child quality ratios for three or more versus one-child families are smaller for the college families in the model compared to the data, I find that this ratio is closer to one than the ratio for households with high school-educated parents. This shows that the model generates nonlinearities in agents' behavior; the relationship between child quality and quantity depends on the type of the parents. This means that the model is capable of replicating the heterogeneity in the quality-quantity trade-off seen in the data.

1.6 Model results

1.6.1 Heterogeneity in parent choices and child outcomes

Tables 11, 12 and 13 split the households in the model into groups by parental education, number of children and initial child quality and present the means of final child quality, total time investments and yearly childhood period consumption for each of the 18 groups. The final quality of the children in the model are higher on average for the college parents for any level of family size or initial quality. College parents also give more time to their children at all levels of family size or initial quality compared to high school parents, despite having already more productive investments and a greater opportunity cost of time. This implies that child quality is a normal good.

Both parental education types share similar patterns in the relationships between final quality, family size and initial quality. For both groups, final quality is declining with family size for when initial quality is low or average. When initial quality is high, there is a sharp increase in final quality between one and two children, and a smaller decline between two and three or more children. From Table 12, we can see that parents significantly increase time investments with family size for children with high initial quality. This is due to the interaction between the dynamic complementarity of the child human capital function and the complementarities between child quality and quantity. As family size increases, parents will desire greater child quality. At the same time, investments into children with higher human capital are more productive than investments into children with less human capital, causing parents of high quality children to increase investments with family size. Parents with lower human capital children will decrease investments and quality with family size because the productivity of additional investments is not large enough to outweigh the increase in the "price" of providing those investments to more children.

Looking at just the relationship between the choices of time investments and consumption and initial quality, we can see that investments are increasing and consumption is decreasing in initial quality. This is again because of the relative productivity of investments; lower initial quality reduces the incentive to invest at any level of family size, causing these parents to substitute towards their own consumption. However, for the lowest initial quality groups, we also see that final quality and time investments are decreasing less in family size for the households with more educated parents. This suggests that despite lower returns on human capital investments, the college-educated parents are more willing to maintain a higher level of investments as family sizes increase. This is due to a number of channels such as the complementarity of quality, quality being a normal good, and the higher productivity of college-educated parents' investments.

Comparing the consumption choices in Table 13 to investment and child outcomes sheds more light on the consumption/child quality margin. Mirroring the previous results, we see that consumption is decreasing in family size for the households with high initial quality children, with parents substituting consumption for highly-productive investments. While for both education groups there is a slightly U-shaped relationship between consumption and family size for average or low initial human capital children, there is a greater negative relationship for the college-educated parents compared to the high school-educated parents. This reveals a crucial mechanism driving the heterogeneity in quality-quantity trade-off. Since higher human capital parents earn more and thus consume more, the marginal utility of their consumption is lower than that of lower human capital parents. When n increases, the complementarity between child quality and quantity encourages parents to increase or maintain a high level of child quality.Higher education parents are thus more willing to substitute consumption for child quality when n increases. Lower education parents with a higher marginal utility of consumption are less willing to sacrifice their consumption, leading to a greater quality-quantity trade-off.

1.6.2 Dynamic choices and borrowing

Figure 5 shows the average choices of per-child time and goods over time for households in the model by parental education and family size. The main notable finding is that for both education levels, parents with one child give more time early and decrease investments over time more than parents with more children. This can be explained in part by looking at how agents choose assets during and after the parenthood period. Focusing on only the asset choice in the parenthood period in Figure 6, we can see that agents' optimal asset choice behavior is to borrow when kids are young, then later in childhood when wages are higher, begin to pay down their debt and begin saving for the post-parenthood period and retirement. This matches the intuition of the relationship between the slopes of investment and assets discussed in Section 1.2. With this desired behavior in mind, the heterogeneity in the slopes of time investments can be explained as follows: Parents with more children must give a larger amount of *total* investments at any given time. Because of this greater time commitment later in childhood, parents with more children are less able to take advantage of their increasing wages compared to parents with one child. Thus they borrow less in the first place, and in turn give fewer time investments early. We can see this mechanism at play in Figure 6 by how one-child parents borrow more early and begin to save earlier compared to parents with more kids.

Goods investments are found to be relatively flat over time for the high school parents and for college parents with one child. This is at odds with the data that shows goods expenditures are generally increasing with the age of children. This may be due to the interaction of two competing forces: On one hand, agents are reducing their overall level of investments over time as discussed above. On the other hand, the share of goods investments compared to time investments is increasing over time due to the estimated share parameters and increasing wages, potentially allowing goods investments to increase over time. For the aforementioned groups, the first effect cancels out the second effect; the high school parents have lower wage growth than the college parents, and as seen above the one-child, college parent households have a steeper decline in overall investments. For the college parents with more than one child, wage growth is sufficiently high and investments are not declining as steeply, allowing goods investments to increase over time as seen in the data.

1.7 Counterfactuals

1.7.1 Sensitivity analysis

In this section I evaluate how the model responds to changes in two key objects. In Table 14 I display the results of the sensitivity analysis for changes in the child utility function. First is the case where parents derive utility from the log of child quality, with family size not entering at all. We see that investments and final quality are higher in this case than in the base case. This is because the marginal utility of quality is higher in the log utility case compared to the base model when family sizes are small. However, there is a stronger quality-quantity trade-off; parents with log child utility increase their consumption with family size, while this relationship is negative on average in the base case. In the case of perfect substitutes, parents invest less and have a much steeper quality-quantity trade-off than in the base case. When utility is Cobb-Douglas between child quality and quantity, the quality-quantity trade-off is essentially eliminated for both groups of parents. Because Cobb-Douglas utility satisfies the Inada conditions, these agents will always give their children some minimally-acceptable level of quality regardless of number of children. However, after some value of quality that is close to zero the marginal utility of quality in the base case will be higher than in the Cobb-Douglas case, causing parents to give more investments and have higher human capital children overall, even in the presence of a quality-quantity trade-off.

Table 15 shows the effect of changes to the time augmenting parameter ψ [13]. The most significant observation is the negative relationship between ψ and average time investments for the college-educated parents. Rather than choosing to invest more in their children's human capital, parents with highly productive time investments choose to decrease their investments. Despite this, average child quality is increasing in ψ, meaning that high human capital parents are able to decrease time investments while still producing increasingly higher quality children. The child quality ratio between one and three or more-child families is also increasing with ψ for the college parents, meaning that higher time augmenting leads to a reduction in the quality-quantity trade-off. This is also evidenced by how the consumption-

[13] Since I normalize the human capital of high school-educated parents to be $h = 1$, ψ has no effect on these agents' choices.

child quantity correlation is decreasing with ψ. High human capital parents with large families are more willing to trade off consumption for investments when time augmenting is high.

I also report the results of removing the economies of scale of investments. We can see that not allowing parents to share their time and goods causes them to sharply decrease their investments with family size. This leads to significantly lower overall levels of investments and a stronger quality-quantity trade-off. It is evident from this experiment that the ability for parents to share investments amongst their children is a necessary feature for the model to accurately capture the complex relationships between human capital investments, family size and parental human capital.

1.7.2 Experiment: The role of credit

Recall that in the calibrated model, households are allowed to borrow up to 5% of the natural borrowing limit. In this exercise I introduce the cases of fully restricted, and fully unrestricted borrowing. The results are presented in Table 16. Most notably, we see that the access to credit has little effect on parental choices of investments; whether agents have unlimited borrowing or no borrowing at all, parents choose roughly the same levels of investments and children have the same level of final quality. The real role of borrowing is seen when comparing the consumption profiles of the no-borrowing case to the base case: When parents are not allowed to borrow, they do not change their level of investment, and in particular they invest the same early in childhood. Due to the dynamic complementarity of the child human capital function these early investments are incredibly productive, to the point that parents are willing to sacrifice a large portion of their consumption to give them. Access to credit allows parents to smooth their consumption during this early period. This result is at odds with what is seen in Abbott (2020) and Caucutt and Lochner (2020), who find that binding borrowing constraints lead to an under-investment in child human capital. However, this may be explained by differences in modeling choices regarding the specification of the human capital technology and utility from children. Most significantly, while the models in those papers feature altruistic parents, the agents in this model only care paternalistically about their children's quality.

Table 17 shows the average maximum level of borrowing for each type of household in both the base and unconstrained case. While there are small differences in the maximum borrowing for most groups, it appears that most of the agents in the base model do not have a binding constraint at any point. The groups that seem to face a binding constraint are the households with more than one child and with high initial quality. These are the groups that invest the most into their children and have the lowest consumption, so they will use the unrestricted borrowing to help increase their consumption while giving these high levels of investments.

1.7.3 Experiment: Child subsidies

In 2021, the American Rescue Plan expanded the child tax credit, providing monthly payments to parents amounting to \$3,600 per year per child under age six, and \$3,000 for all other children. This temporary policy proved to be an effective tool for reducing child poverty, with the CEA estimating that child poverty fell by nearly 3 percentage points in 2021 (Council of Economic Advisers (2023)). In order to evaluate the effects of this subsidy on parental investment choices and child human capital outcomes, I implement the expanded child tax credit in the calibrated model. For simplicity, I consider a subsidy of \$3,600 per child per year at all ages. I also consider a subsidy on goods investments of 50%.

The results of this exercise are shown in Table 18. Both policies are effective in increasing investments into child human capital: the child tax credit causes parents to provide about 8 percent more hours to each child and 9 percent more total goods per year compared to the base model. The goods subsidy increases the average household spending on total goods by about 40 percent, and parents also increase their time investments due the proportionality of investments. This results in an increase in final child human capital of about 1.4 percent and 1.5 percent for the child tax credit and goods subsidy, respectively. In the model with the child tax credit, children from households with three of more kids have about 1 percent lower human capital than only children on average. This is compared to a 2.7 percent decrease in the base model, showing how a per-child subsidy can help reduce the quality-quantity trade-off. Average household consumption and asset levels are also higher with the child tax credit compared to the base model, showing that only a portion of the cash transfer is spent on investment into children. Since the goods subsidy is explicitly an incentive to increase child human capital investments, there is not an increase in consumption due to this policy, though parents do increase their saving. In sum both policies are effective in increasing investments in children and improving child human capital, but the per-child subsidy has the added benefits of reducing the quality-quantity trade-off and increasing household consumption.

1.8 Conclusion

In this project, I develop a model of parental investments into child human capital with heterogeneous family size. The model features multiple investment types, static and dynamic investment choices, and a comprehensive specification for the child human capital function. The number of children n enter into the model through a CES utility over child quality and quantity as well as through economy of scale multipliers on the per-child investment choices in the budget constraint. I also assume that the productivity of parents' time investments are increasing in their human capital. These three features allow the model to generate interesting nonlinearites in parents' choices along the axes of parental human capital and family size.

Next, I present novel stylized facts on the joint distributions of household income, family size, parental time and goods investments and child outcomes. First, I find evidence of heterogeneity in the quality-quantity trade-off; low-income households face a negative relationship between family size and child outcomes, while high-income families do not face this trade-off. This result holds for both child pre-labor market human capital as measured in the NLSY97 and child labor market outcomes from the PSID. These observations imply that parents with different levels of income and family size choose investments for their children in a nonlinear way. Using data on parental goods and time investments from the PSID, I document patterns in how parents allocate these resources to their children. Most notably, I find that time investments are not changing monotonically with respect to income, and in fact vary little by household income. I also find evidence of investments exhibiting economies of scale.

I estimate the child human capital investment function by exploiting the proportionality of the different investment types. The remaining parameters are calibrated using moments that reflect important features of parental investments as well as the observed heterogeneity in the quality-quantity trade-off. Analysis of the choices of agents by education, family size and initial child quality reveal three key mechanisms driving the observed heterogeneity in the quality-quantity trade-off: First, high human capital parents are better able to take advantage of high productivity time investments and economies of scale of time investments compared to low human capital parents. Second, parents with higher incomes are more willing to trade off consumption for child quality when faced with an increase in family size compared to parents with lower incomes. Third, the dynamic complementarity of investments causes initial child quality to affect parents' choices of investments as family size increases. Experiments using the calibrated model show that per-child cash transfers and subsidies to goods investments are both effective policies for increasing parental investments and improving child human human capital.

There are a number of potential directions for future work, including extending the model to include schooling, as well as extending the empirical and quantitative analysis to include single-parent households. More empirical work using the Consumer Expenditure Survey along with separating goods investment from childcare in the model would result in a better treatment of monetary investment in children. As mentioned in Section 1.2, the key assumption I make to extend the model to an arbitrary n is the homogeneity of the children. Relaxing this assumption is another important direction in this line of research.

1.9 Tables and Figures

Table 1: ASVAB percentile (NLSY97)

	(1)	(2)
Log household income	8.496***	4.164*
	(2.243)	(2.274)
Number of children	-31.35***	-26.31***
	(9.543))	(9.225)
Income × number of kids	2.924***	2.584***
	(0.884)	(0.857)
Household controls	No	Yes
N	2130	2130
Adj. R^2	0.119	0.184

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Child household income (PSID)

	(1)	(2)
Log household income	0.0962	0.0714
	(0.151)	(0.302)
Number of children	-1.071**	-1.267**
	(0.527)	(0.532)
Income × number of kids	0.0963**	0.104**
	(0.0470)	(0.0472)
Controls	No	Yes
N	13304	13304
Adj. R^2	0.280	0.533

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1: Average time investments by household income quartile and family size

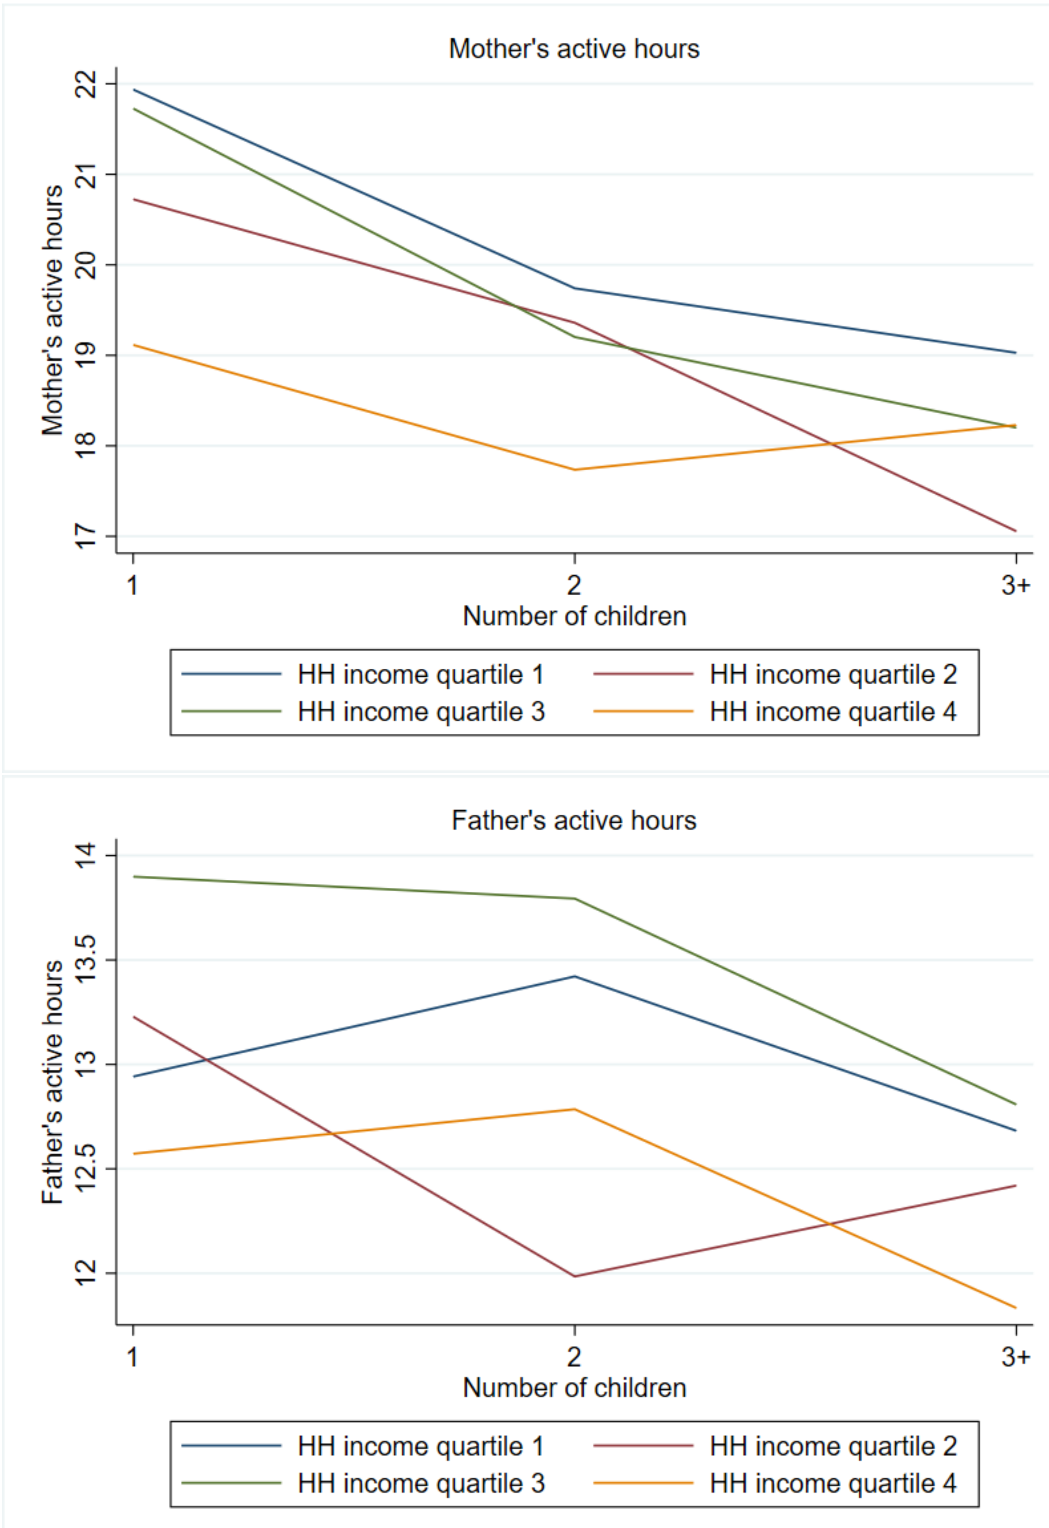

Table 3: Shared parent time ratios

Number of children	Mother shared time	Father shared time
1	0.0585	0.0628
	(0.0101)	(0.0112)
2	0.6402	0.6741
	(0.0089)	(0.0100)
3	0.7101	0.7687
	(0.0121)	(0.0126)
4	0.7725	0.8080
	(0.0222)	(0.0247)

Figure 2: Average goods investments by household income quartile and family size

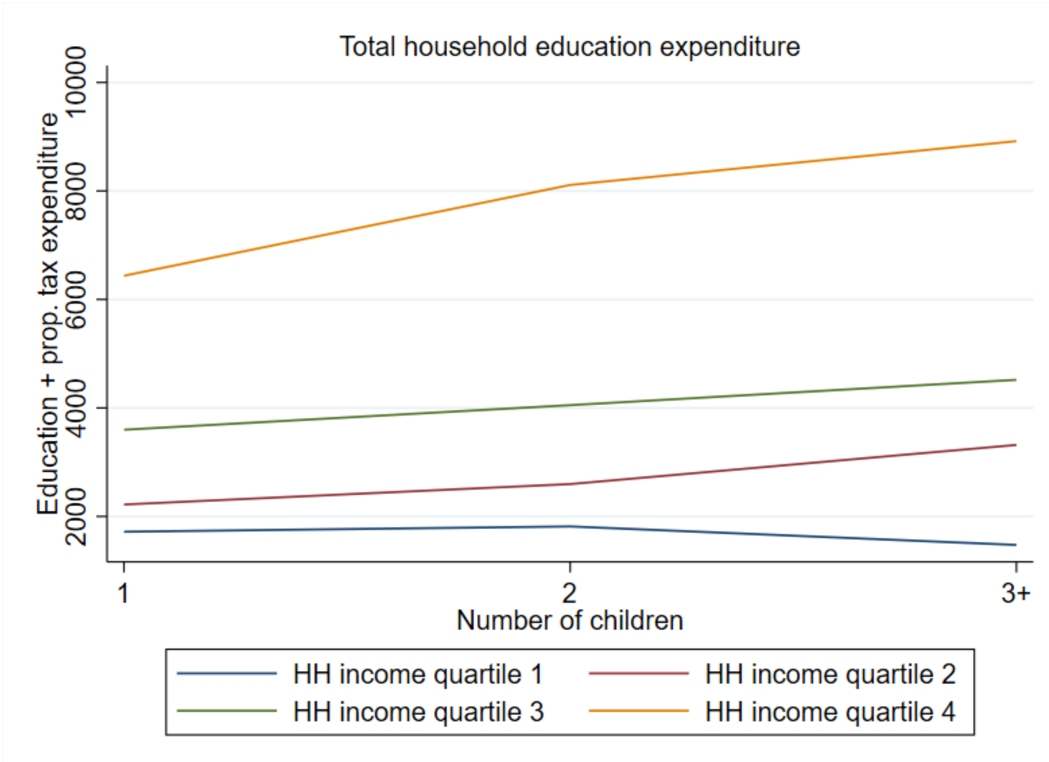

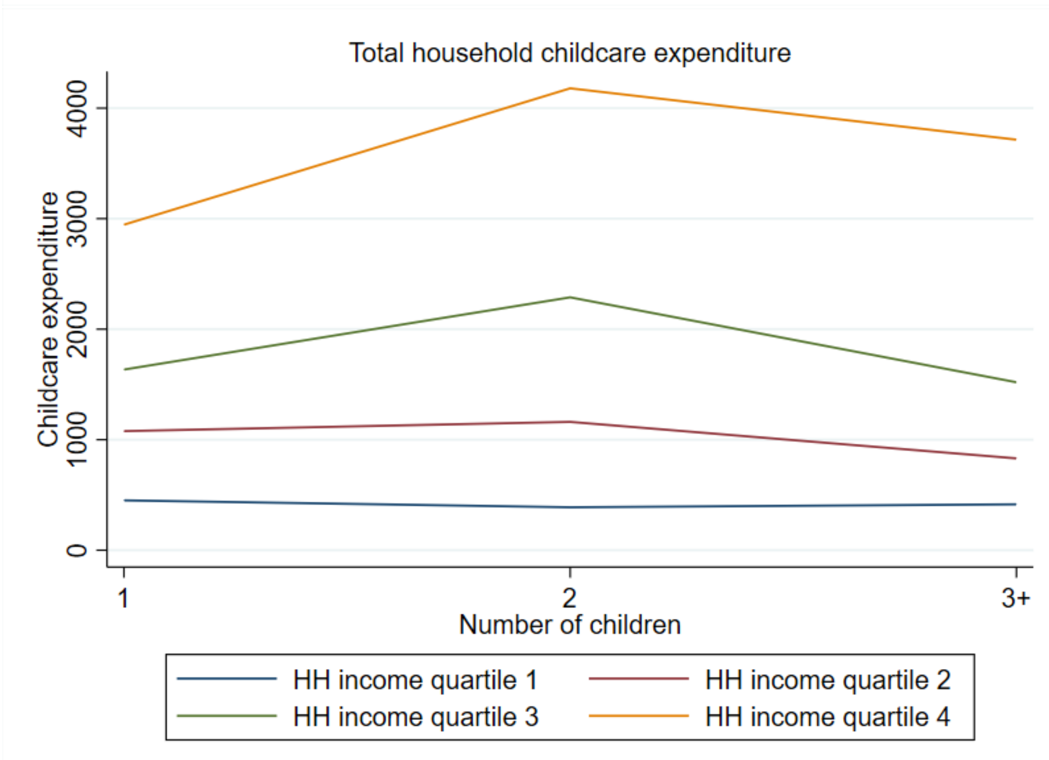

Table 4: GMM parameter estimates

Parameter	Value	Standard Error
γ	-1.476	0.0000
y_f	-1.1457	0.0005
z_f	0.0951	0.0000
y_g	0.0000	0.0023
z_g	0.1447	0.0000

Figure 3: Estimates for $\tilde{\alpha}_{i,t}$ parameters

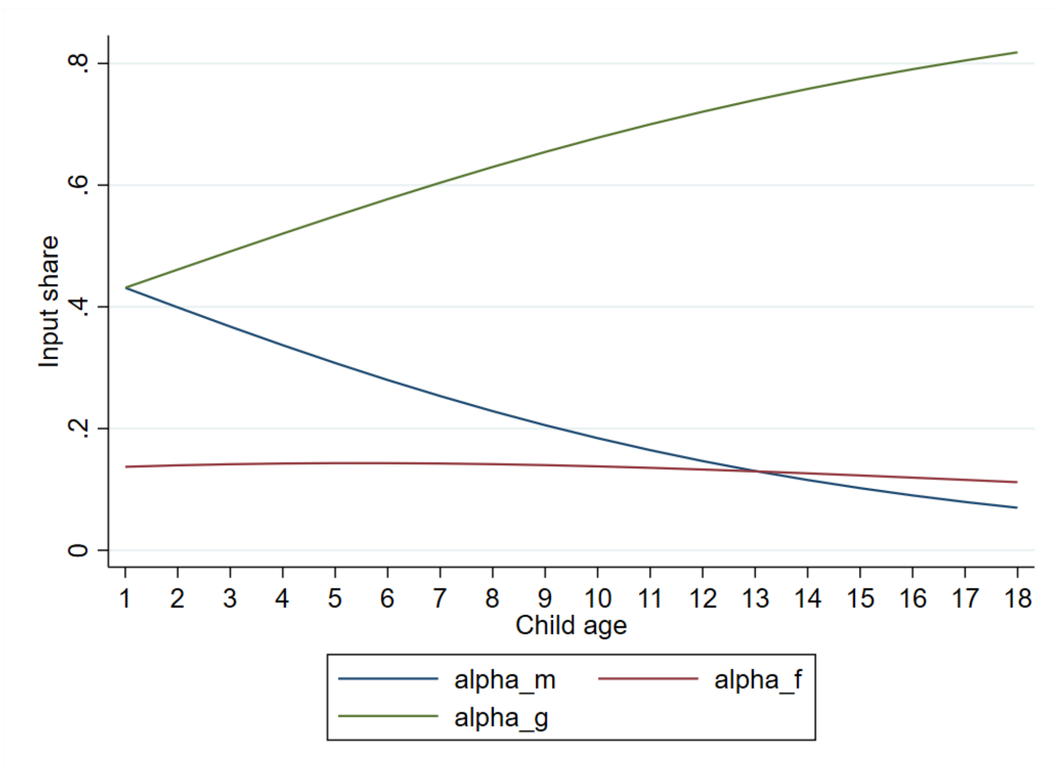

Table 5: $\Lambda_i(n)$ estimates

Number of children	$\Lambda_m(n)$	$\Lambda_f(n)$
1	1.0	1.0
	(0.000)	(0.000)
2	1.3598	1.3258
	(0.0089)	(0.0100)
3	1.5798	1.4625
	(0.0243)	(0.0251)
4	1.6925	1.5761
	(0.0668)	(0.0742)

Table 6: Logged wages

	(1) Mothers	(2) Fathers
Years of education	0.111***	0.0930***
	(0.0037)	(0.0035)
Age	0.0280***	0.0560***
	(0.0030)	(0.0030)
Age squared	-0.00019***	-0.00047***
	(0.00003)	(0.00003)
Constant	0.380***	0.263***
N	10,027	13,119
Adj. R^2	0.181	0.196

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Letter-Word score distributions

Parent education	Mean of Letter-Word score	Variance of Letter-Word score
High school dropout	0.8161	0.0551
High school	0.9234	0.0809
Some college	1.0123	0.0806
College +	1.0946	0.0843

Table 8: Calibrated parameters

Parameter	Value	Description
ν	0.7089	Child utility share parameter
ρ	-1.5552	Child utility substitution parameter
μ	0.9499	Human capital TFP parameter
θ_1	0.00506	Productivity of today's investments X_t
θ_2	1.0669	Productivity of yesterday's quality q_t
ψ	0.4959	Productivity of parents' time
ϕ_g	0.3867	Economies of scale of goods investments

Figure 4: Input ratios by child age in the data and in the model

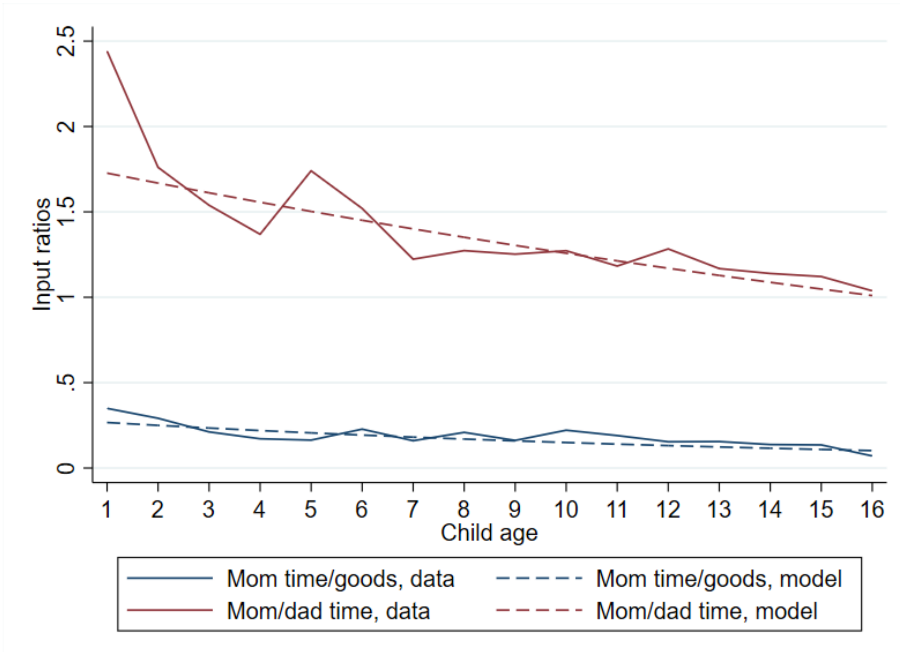

Table 9: Targeted moments

Moment	Data	Model
Avg. parental time, college-educated parents	1761.58	1804.58
Avg. parental time, HS-educated parents	1576.42	1566.78
Ratio of parental time b/t 2 and 1 children, college-educated parents	0.9827	0.9483
Ratio of parental time b/t 2 and 1 children, HS-educated parents	0.8621	0.9014
Ratio of goods b/t 3+ and 1 children	1.2741	1.2655
Parental time/child age correlation	-0.4537	-0.4511
Age 5 child quality, # std. devs from mean, college-educated parents	0.2910	0.2853
Age 5 child quality, # std. devs from mean, HS-educated parents	-0.1815	-0.1898

Table 10: Untargeted moments

Moment	Data	Model
Avg. goods, college-educated parents	9766.81	7807.66
Avg. goods, HS-educated parents	3515.54	5128.23
Parent education / child quality correlation	0.3829	0.2481
Ratio of child quality b/t 3+ and 1 children, college parents	1.0099	0.9893
Ratio of child quality b/t 3+ and 1 children, HS parents	0.9669	0.9646

Table 11: Final child quality, number of std. deviations from mean

Panel A: High school parents

Number of children	Initial child quality -1	0	1	Average
1	-1.3069	-0.0992	1.4318	-0.1589
2	-1.3457	-0.1689	1.7057	-0.1982
3+	-1.3947	-0.2035	1.5665	-0.2302
Average	-1.3531	-0.1662	1.5994	-0.2006

Panel B: College parents

Number of children	Initial child quality -1	0	1	Average
1	-1.2383	0.1350	1.5726	0.3116
2	-1.2247	0.0808	1.7628	0.3184
3+	-1.2684	0.0214	1.6676	0.2851
Average	-1.2404	0.0738	1.6949	0.3069

Column “-1”: one std. deviation or more below the mean
Column “1”: one std. deviation or more above the mean
Column “0”: within one std. deviation from the mean

Table 12: Yearly per-child time investment

Panel A: High school parents

Number of children	Initial child quality -1	0	1	Average
1	1461.4164	1944.7645	1550.6312	1778.8691
2	1134.0107	1676.7749	2082.102	1603.4833
3+	885.99648	1431.7581	1954.1477	1381.7108
Average	1124.2553	1649.4783	1922.3441	1566.7794

Panel B: College parents

Number of children	Initial child quality -1	0	1	Average
1	1482.5045	2070.9207	1774.0158	1959.3879
2	1241.6931	1822.3227	2226.352	1858.152
3+	1059.7709	1549.3439	2017.1869	1614.8446
Average	1240.5486	1790.0459	2070.4077	1804.5791

Column “-1”: one std. deviation or more below the mean
Column “1”: one std. deviation or more above the mean
Column “0”: within one std. deviation from the mean

Table 13: Average yearly consumption, childhood period only

Panel A: High school parents

Number of children	Initial child quality -1	0	1	Average
1	56,557.33	52,180.77	55,776.06	53,686.64
2	55,976.20	49,443.93	44,115.11	50,267.37
3+	57,321.73	49,710.30	42,172.92	50,375.25
Average	56,525.42	50,096.13	46,068.72	51,017.23

Panel B: College parents

Number of children	Initial child quality -1	0	1	Average
1	84,560.69	76,699.37	80,708.25	78,198.12
2	82,262.62	71,895.90	64,149.14	71,147.94
3+	82,605.86	72,475.16	62,146.40	70,977.02
Average	82,857.07	73,027.55	66,763.04	72,506.83

Column "-1": one std. deviation or more below the mean
Column "1": one std. deviation or more above the mean
Column "0": within one std. deviation from the mean

Figure 5: Time profiles of per-child time and goods investments

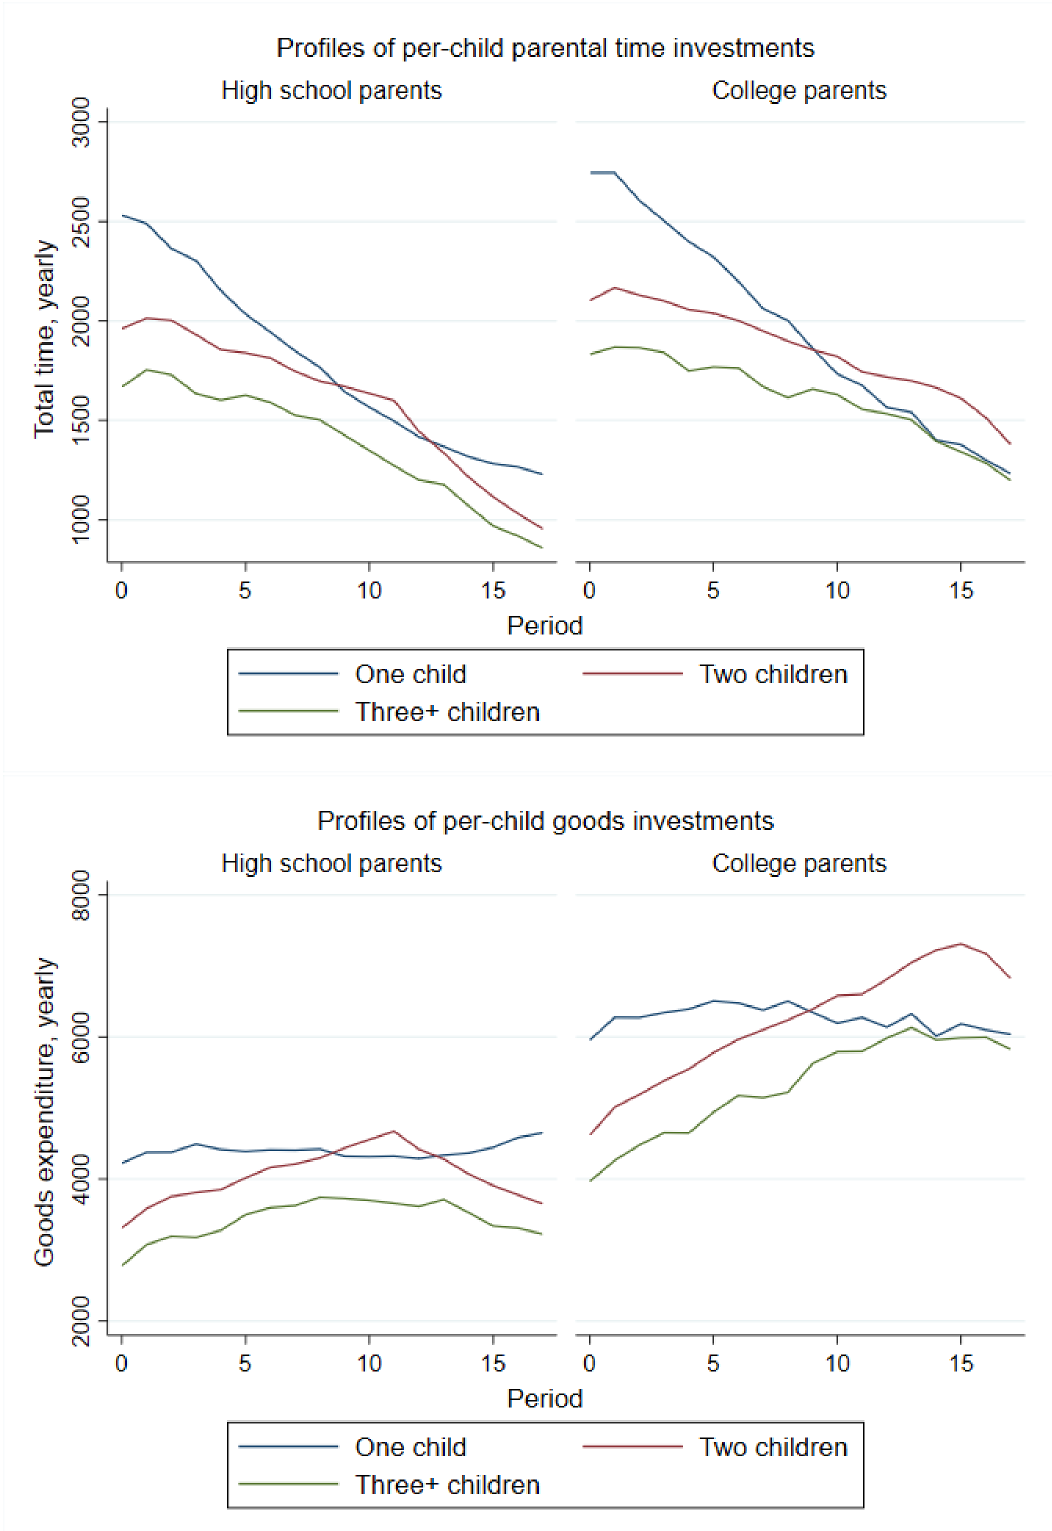

Figure 6: Time profile of asset choice during parenthood period

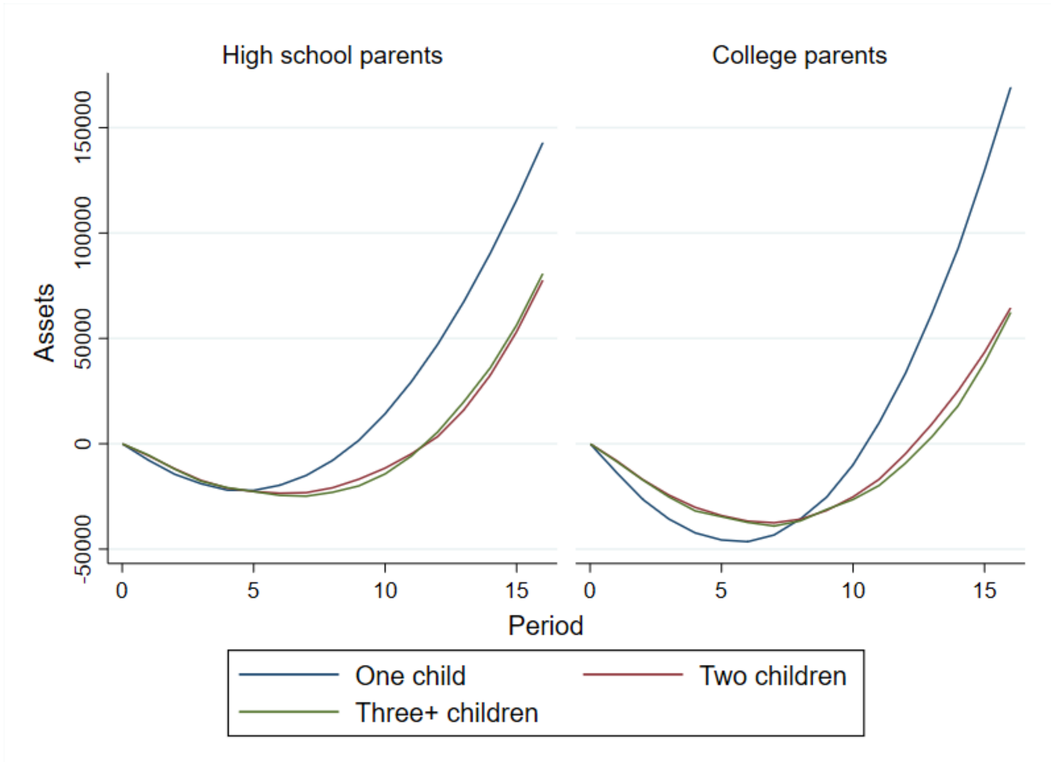

Table 14: Effect of changes to the child utility function

	Base	Log utility	Cobb-Douglas ($\rho \to 0$)	Linear ($\rho = 1$)
Avg. parental time, HS parents	1566.7794	1766.2795	1378.6597	687.1127
Avg. parental time, college parents	1804.5791	1940.8832	1662.7688	1020.3564
Ratio of child quality b/t 3+ and 1 children, HS	0.9646	0.9736	0.9892	0.9172
Ratio of child quality b/t 3+ and 1 children, college	0.9893	0.9380	0.9990	0.9355
Correlation b/t consumption and child quantity, HS	-0.1889	0.0215	-0.4463	0.0264
Correlation b/t consumption and child quantity, college	-0.2966	0.0227	-0.4784	0.0462
Avg. final child quality, HS	0.5729	0.6281	0.5611	0.5049
Avg. final child quality, college	0.7202	0.7386	0.7116	0.6597

Table 15: Effect of changes to child human capital technology

	Base	No augmenting ($\psi = 0$)	High augmenting ($\psi = 5$)	No economies of scale
Avg. parental time, HS parents	1566.7794	1566.7794	1566.7794	1045.0926
Avg. parental time, college parents	1804.5791	1808.2765	1571.0882	1206.9739
Ratio of time b/t 3+ and 1 children, HS	0.7767	0.7767	0.7767	0.3540
Ratio of time b/t 3+ and 1 children, college	0.8242	0.7928	0.9036	0.3744
Avg. final child quality, HS	0.5729	0.5729	0.5729	0.5330
Avg. final child quality, college	0.7202	0.7026	0.8773	0.6683
Ratio of child quality b/t 3+ and 1 children, HS	0.9646	0.9646	0.9646	0.8525
Ratio of child quality b/t 3+ and 1 children, college	0.9893	0.9835	1.0072	0.8679
Correlation b/t consumption and child quantity, HS	-0.1889	-0.2022	-0.2022	-0.1148
Correlation b/t consumption and child quantity, college	-0.2966	-0.2674	-0.4021	-0.1910

Table 16: Effect of credit access

	Base	No constraint	No borrowing
Avg. parental time, HS parents	1566.78	1573.48	1562.19
Avg. parental time, college parents	1804.58	1804.39	1802.67
Yearly consumption, child ages 1-6	58462.99	58617.03	52894.27
Yearly consumption, child ages 7-12	59488.15	59551.01	60669.72
Yearly consumption, child ages 13-18	60578.30	60506.66	62452.00
Avg. final child quality, HS	0.5729	0.5762	0.5723
Avg. final child quality, college	0.7202	0.7202	0.7199

Table 17: Effect of removing borrowing constraints on maximum borrowing

Panel A: High school parents, base model				Panel B: High school parents, no constraint			
Number of	Initial child quality			Number of	Initial child quality		
children	-1	0	1	children	-1	0	1
1	-21,945.01	-35,835.85	-22,295.94	1	-22,513.16	-35,900.07	-22,295.94
2	-21,430.11	-39,621.08	-61,408.46	2	-21,874.83	-39,778.25	-80,314.83
3+	-26,419.50	-39,464.50	-53,615.40	3+	-27,194.04	-39,452.49	-62,830.96
Panel C: College parents, base model				**Panel D: College parents, no constraint**			
Number of	Initial child quality			Number of	Initial child quality		
children	-1	0	1	children	-1	0	1
1	-44,842.00	-64,147.91	-50,829.70	1	-45,465.03	-64,167.45	-50,829.70
2	-30,359.02	-56,304.88	-82,493.21	2	-29,543.38	-56,619.30	-104976.55
3+	-30,690.28	-57,514.27	-83,119.57	3+	-30,585.26	-57,713.43	-102734.40

Table 18: Effect of child tax credit and investment goods subsidy

	Base model	Child tax credit, \$3,600 per child	Goods subsidy, $s = 0.5$
Mean yearly per-child time investment	1660.76	1791.12	1748.37
Mean total yearly goods investment	\$6,187.13	\$6,728.61	\$8,697.40
Mean final child quality, pct. change from base	100%	101.41%	101.51%
Pct. change final child quality, 1 and 3+ children	-2.70%	-1.12%	-2.27%
Mean childhood-period consumption	\$59,509	\$61,619	\$59,137
Mean parenthood-period assets	\$7,111	\$29,347	\$12,248

2.0 The Economics of Workplace Flexibility

2.1 Introduction

The experience of the COVID-19 pandemic brought about seismic changes to the structure of jobs, and in particular an increase in flexible work arrangements. According to the Spring 2022 American Opportunity Survey, about 58% of workers had the ability to work from home (McKinsey & Company (2022)). This is in stark contrast to the 2017-2018 American Time Use Survey, which reports that about 25% of workers participated in remote work. It was well-known even before the pandemic that the ability to choose where and when one works is a highly desirable non-wage amenity: several discrete choice and stated-preference experiments show that workers are willing to pay up to 20% of their wage for scheduling flexibility, and up to 8% of their wage for the ability to work from home (Maestas et al. (2023), Mas and Pallais (2020), Wiswall and Zafar (2018)). Despite how much workers desire these arrangements, the post-pandemic discourse around workplace flexibility often revolves around employers' efforts to return to the pre-pandemic norm of inflexible work (Peck (2023)). As the debate over the future of the workplace carries on, there is an even greater need for a detailed analysis of the costs and benefits of workplace flexibility for both workers and firms.

In this paper I contribute to this discussion by exploring the provision and effects of workplace flexibility both empirically and quantitatively. I begin by leveraging the detailed questioning in the 2017 and 2018 American Time Use Survey to construct variables to measure two dimensions of workplace flexibility: the ability to choose one's own work schedule, and the ability to work from home.

My main empirical exercise is to compare workplace flexibility to the other prominent non-wage amenities available in my data: health insurance, paid leave, and pensions. Comparing workers within occupation category, flexibility is found to be the only amenity that has an effect on workers' intensive labor supply; workers with flexibility work more hours per week and more evenly distribute their labor hours across the day compared to workers *in the same occupation* without flexibility. An analysis of the across-occupation provision of job amenities shows that flexibility is distinct in its scarcity and high across-occupation variance compared to the other amenities. The within-occupation allocation of amenities is investigated using a logistic regression technique, which shows that the prevalence of flexibility in a given occupation is less sensitive to firm size and overall amenity level compared to the other amenities. These results suggest that when deciding whether to offer workers flexibility, firms take into account not only the pecuniary cost of the policy, but also the effect of flexibility on worker labor supply and output.

In order to characterize the rich set of interactions between firms' flexibility policies and workers' job choice and labor supply, I develop and calibrate a model of a labor market in which firms compete for workers through offers of wage and flexibility policies. Similar to Sullivan and To (2014), Sullivan and To (2023) and Bonhomme and Jolivet (2009), my model features random search with take-it-or-leave-it offers of wages and amenities. I innovate on previous models by carefully modeling workers' labor supply allocation; following Cubas et al. (2022), workers can choose not only their total intensive labor supply, but also how their labor time is distributed across the day. Workers are heterogeneous in their human capital as well as their required household labor, with workers with higher housework burdens having a greater demand for flexibility. Workers have various constraints on their time which are lifted when at a job with flexibility, and will thus be more likely to accept a flexible job offer. Firms pay for offering flexibility through a pecuniary cost as well as through workers potentially moving to a less-productive work schedule. Like in Dey and Flinn (2005), firms are heterogeneous in the cost of providing the amenity, and in contrast to similar models I explicitly model firms with varying productivities.

The model is calibrated using distributions of white-collar workers from the data, and a number of counterfactual exercises are performed. Using data simulated from the model, I find that estimates of the compensating wage differential of flexibility from the ATUS data are wrong-signed due to a lack of detailed human capital controls. Controlling for human capital produces an unbiased estimate for the compensating differential of about 7 percent. Despite men and women in the model only differing by their household labor requirements, a gender wage gap of 2.4 percent is estimated. This implies that nearly 20 percent of the observed gender wage gap from the data can be attributed to gender differences in household labor hours. In an experiment in which workers provide the same labor supply regardless of flexibility status, I find that removing the labor supply endogeneity of flexibility results in flexibility being provided at a rate similar to other amenities in the data. This reinforces the finding from the data that this unique feature of workplace flexibility has a significant effect on the real-world allocation of flexibility. Decomposition exercises quantify the effect of various mechanisms on the joint distributions of flexibility and wages. Finally, by directly using the solutions to the workers' search problem I can estimate workers' valuation of flexibility. The results of this exercise show that all else equal, workers are 23 percentage points more likely to accept an offer with flexibility.

My final quantitative exercise is to compare the outcomes of workers and firms in the calibrated model to a counterfactual model without workplace flexibility. Welfare analysis shows that male workers are willing to pay up to 8 percent of their consumption to live in the economy with flexibility, while female workers are willing to pay on average about 44 percent of their consumption. This huge difference in men and women's consumption equivalence is due to the way workplace flexibility improves the labor market outcomes of women in particular: removing flexibility increases the gender wage gap from 2.4 percent to 8.1 percent, and female unemployment rises from 6 percent to 14.7 percent. Low-productivity firms are found to be the main producer-side beneficiaries of workplace flexibility, enjoying a higher per-worker profit in the calibrated model compared to the no-flexibility model due to having a relatively low cost of providing flexibility. In addition, since the ability to offer flexibility allows low-TFP firms to compete for workers on a more even footing with high-TFP firms, the share of workers employed at low-TFP firms is higher in the base model. Having a higher share of workers at high-productivity firms causes output to be higher in the economy without flexibility, but these gains from are predominately captured by the high-productivity firms in the form of higher profits. In sum, female workers are the largest beneficiary of workplace flexibility policies while male workers and low-TFP firms also enjoy a net benefit, with these welfare gains largely coming at the expense of high-productivity firm profits.

In Section 2 I describe the data used for my empirical analysis and calibration moments.

Section 3 compares workplace flexibility to other job amenities with regards to their allocation and effect on worker outcomes. Section 4 introduces the model and characterizes the labor market equilibrium. Section 5 discusses the calibration strategy and presents the results of the estimation. In Section 6 I present the results several counterfactual experiments. Section 7 quantifies the costs and benefits of flexibility for both workers and firms via a welfare analysis. Section 8 concludes.

2.2 Data

2.2.1 ATUS/CPS data

All data used in this article are sourced from IPUMS-ATUS and IPUMS-CPS . My main analysis is based on the 2017 and 2018 Leave and Job Flexibilities Module of the American Time Use Survey (ATUS). About 10,000 respondents from these two ATUS survey years were asked various questions about their access to leave, job flexibility and work schedules.

I focus on two particular types of workplace flexibility: the first is flextime, the ability for one to choose their own work schedule, and the second is the opportunity for paid work-from-home. On flextime, the Leave and Job Flexibilities Module asks the following question: "Do you have flexible work hours that allow you to vary or make changes in the times you begin and end work?" The responses to this question show a spectrum of flexibility arrangements, but for my flextime dummy variable I choose to look at the highest level of temporal flexibility, which corresponds to the response, "can frequently choose their own hours."

The questions on work-from-home are asked in sequence as follows: "As part of your job, can you work at home?" If yes, "Do you ever work at home?" If yes again, "Are you paid for the hours that you work at home, or do you just take work home from the job?" As seen by the third question, the Leave and Job Flexibilities Module makes a useful and important distinction between extra work taken home and work from home that is compensated. I choose to only count the latter as a genuine work-from-home policy. The drawback of this choice is that since each subsequent question is only asked to respondents who answered "yes" to the previous question, I only count a worker as having a job with a work-from-home policy if they actually use it. This means I am potentially ignoring a group of workers whose job has a paid work-from-home policy but who choose to not work from home.

From the Leave Module I also collect information on workers' access to paid leave. The main ATUS survey provides a comprehensive list of demographic variables such as age, education, marital status and household composition, as well as job variables, including industry and occupation codes, weekly labor supply, hourly wages and commute times. Some ATUS respondents were also administered the Annual Social and Economic Supplement (ASEC) of the Current Population Survey. The ASEC provides additional variables regarding employer characteristics, such as the size in terms of number of employees of the respondent's employer, as well as access to benefits such as health insurance and pensions. I restrict my analysis to workers between the ages of 25 and 65 who work at least 35 hours per week. The Leave Module is not given to self-employed workers, so I exclude the self-employed from my analysis. I also trim hourly wages at the 1% level. This leaves me with 3,967 observations for which I have ASEC employer data, and 7,065 observations overall. For my analysis I will focus on the ASEC-matched sample.

2.2.2 ATUS Time Diary Data

My analysis also utilizes the 2017 and 2018 ATUS time diaries. In particular, I focus on the allocation of workers' labor supply, housework and leisure throughout the day. Following Cubas et al. (2022), I collect data about total labor and housework time between the hours of 9 a.m. to 5 p.m. ("prime" working time) versus any other time of day ("off time"). My labor supply variable only counts time spent working at a respondent's main job[1]. My housework variable includes the time respondents spend on personal care, housework activities, and "caring for and helping household members".

Each respondent is administered the time diary for one day of the week, with about half

[1] The time diaries include several ancillary work activities such as waiting or socializing/leisure as part of the job, as well as time working at other jobs. I choose a strict definition of working time at the main job in order to better relate time use to the flexibility policies of workers' main job as determined from the Leave module.

of respondents being given a survey for either Saturday or Sunday. The sample is limited to diaries completed during the workweek. The time diary data is linked to ATUS/ASEC data to explore the relationship between flexibility, job characteristics and time allocation. The merged data contains 1,916 ASEC-matched observations and 3489 observations overall.

2.2.3 Summary statistics

Table 19 provides summary statistics for the ASEC-matched sample, by flexibility status. A worker is considered to have “flexibility” if their job has either flextime or paid work-from-home, or both. Workers with flexibility make up about a third of the sample, and are on average slightly younger and more likely to be male. Looking across the entire sample, they are more likely to have a college degree, be married and have kids, and on average earn a higher hourly wage and work longer hours. Flexible workers more also more likely to live in a large metropolitan statistical area (defined as one with greater than one million residents) and work for a large firm (one with at least one thousand employees).

Following Cubas et al. (2022), I compute the “prime working hours ratio” using the ATUS time diaries. This ratio measures the share of workers’ total working time that is done between the hours of 9am to 5pm. I find that flexible workers on average work about the same share of their total hours in prime time compared to inflexible workers. Workers with flexibility spend about 7 more minutes per day on housework compared to workers without workplace flexibility policies. Defining “leisure” loosely as the respondent’s waking hours not spent either working or performing housework, I find that flexible workers spend about 13 minutes less on leisure per day compared to workers who are not flexible.

Table 37 shows the intensity of flexibility policies by broad Census occupation category. There is significant variation in the provision of flexibility policies across occupations, but overall flexibility is scarce. Flexibility is most common in “white collar” office jobs such as management, business and finance, and STEM occupations. The occupation categories with the highest prevalence of flexibility is computer and mathematical science and engineering. Even among these types of occupations, flexibility is far from ubiquitous. Flextime is less common overall than paid work-from-home, but the intensity of work-from-home has a greater variation across occupations.

2.3 Workplace flexibility: a different kind of amenity

In this section I compare and contrast the following job amenities:

1. Workplace flexibility: workers who have either flextime or paid work-from-home as defined above
2. Health insurance: workers who responded "yes" to either of the ASEC questions, "the respondent was the policyholder for group health insurance that was related to current or past employment" or "the respondent's employer offers a health insurance plan to any of its employees"
3. Paid leave: workers who responded "yes" to the "Do you receive paid leave on your [current/main] job?" question in the ATUS Leave module
4. Pension: workers who answered "yes" to the IPUMS/ASEC questions "the respondent's union or employer for his or her longest job during the preceding calendar year had a pension or other retirement plan for any of the employees, and, if so, whether the respondent was included in that plan"

I will begin by comparing the effect of each amenity on a number of worker outcomes, such as hourly wages and labor supply. Next I will investigate the allocation of each amenity, both across occupation categories as well as within occupations through a logistic regression analysis. The findings presented in this section will reveal a number of interesting properties of workplace flexibility as an amenity, and will serve to motivate the use of my calibrated model.

2.3.1 Comparison of worker outcomes

In order to investigate the effect of access to various job amenities on worker outcomes, I estimate the following model:

$$y_i = \beta_0 + \beta_1 Amenity_i + \Lambda X_i + \Gamma Occ_i + \Delta Ind_i + \epsilon_i$$

The worker outcomes y_i are log hourly wages, log weekly working hours, and the prime working hours ratio for worker i. X_i is a vector of individual characteristics such as age, education, race, marital status, MSA, number of children, as well as job characteristics such as firm size. Occ_i denotes 4-digit Census occupation code fixed effects, and Ind_i denotes Census industry category fixed effects. The interpretation of the coefficients are thus the *within occupation* relationship between amenities and worker outcomes. I restrict the sample to only occupations with at least 20 observations. This reduces the sample to 3265 observations for the main ASEC-matched sample, and 1374 observations for the ASEC and time diary sample.

Results are presented in Table 21. Columns 1-4 present the regression coefficients for the model described above for each amenity, while column 5 shows the results for a regression with all of the amenities included. Panel A shows that the presence of each of the amenities is associated with a significant increase in wages compared to workers in the same occupation without those amenities. While these coefficients are not accurate estimates of the compensating wage differentials of these amenities, the similarity in the coefficients suggests that there is not a significant difference between the underlying relationship between wages and flexibility and the rest of the amenities.

The main difference between flexibility and the other amenities can be seen in Panels B and C, which measure workers' labor supply allocation. Workers with flexibility are found to work about 5% longer per week compared to workers in the same occupation without flexibility. In addition, flexibility policies are associated with a 7 percentage point decrease in the prime working hours ratio. These results strongly suggest that workplace flexibility has a significant effect on workers' day-to-day and week-to-week labor supply choices. With the exception of a small negative relationship between pensions and the prime hours ratio, the insignificant coefficients for the other amenities indicate that this effect is unique to flexibility.

One observation is that amenities 2-4 may be considered "pecuniary" amenities, where the only cost for the employer of providing the amenity is monetary[2]. Flexibility differs in that the amenity itself has an effect on workers' labor supply, and thus may affect the employer through changes to workers' labor supply allocation. Cubas et al. (2022) show that coordination between workers has an effect on firm output, and that gender differences in the allocation of labor throughout the day may exacerbate the gender wage gap. These results imply that flexibility may introduce a kind of non-pecuniary cost to the firm in which the firm's output suffers due to workers working at suboptimal times or in ways that decrease coordination.

2.3.2 Allocation of job amenities

I begin my analysis of the allocation of job amenities by comparing the intensity of the provision of the four amenities by occupation category. The results are shown in Table 22. Health insurance and paid leave are much more prevalent than flexibility in all occupation categories, and in aggregate 88% and 80% of workers have these amenities, respectively, compared to only 34% of workers having flexibility. 47% of workers have employer-provided pensions, with nearly every occupation category having a higher percentage of pensions compared to flexibility.

The bottom row of Table 22 shows the across-occupation variance of each amenity. The most striking observation is the degree of variance in the provision of flexibility across occupation categories, with the variance of flexibility being three times the size as the next-highest. This suggests that the provision of flexibility depends on the production functions of different types of occupations. On the other hand, health insurance, paid leave and pensions are provided at somewhat similar rates to workers in different occupations. This provides further evidence for the distinction I made in the previous section between pecuniary amenities and

[2]This may be debatable for paid leave, where workers leaving may have an effect on the productivity of the remaining workers. In the case of no negative spillovers from the absence of workers on leave, the cost of paid leave is then only the forgone production of the worker, which is ultimately just a pecuniary cost.

flexibility; while the decision to provide the pecuniary amenities is mainly a dollars-and-cents decision with little heterogeneity across occupations, the effect of flexibility on labor supply results in significant heterogeneity in the provision of flexibility by occupation type.

In order to investigate the within-occupation allocation of job amenities, I utilize a logistic regression framework:

$$\text{logit}(p_{Amenity_i}) = \beta_0 + \Omega AmenityIndex_i + \Lambda X_i + \Gamma Occ_i + \Delta Ind_i + \epsilon_i$$

The dependent variables in these regressions are the log odds of worker i having one of the amenities. X_i is a vector of individual and job characteristics, and occupation and industry fixed effects are included. In each regression, $AmenityIndex_i$ is the sum of the amenities that are not being used as the outcome variable. For instance, the amenity index for flexibility is the sum of workers' health insurance, paid leave and pension. These are included as independent variables in order to measure the potential within-occupation correlation between amenities.

Table 23 shows the within-occupation relationship between amenities and individual job and characteristics. Beginning with the comparison of the coefficients of the amenity indices, we can see that the relationship between flexibility and the other amenities is not significant, while the other amenities have significant positive coefficients on their amenity indices. This result may be related to the fact that the provision of pecuniary amenities depend primarily on firms' ability to pay for them; a firm being able to provide one pecuniary amenity means that might be able to afford to provide others, resulting in the co-movement between pecuniary amenities. Since flexibility has non-pecuniary costs, there will be a weaker relationship between firms' ability to pay for amenities and the provision of flexibility.

This prediction is shown to be accurate when comparing the relationship between the amenities and working at a large firm. We can see that the coefficient on large firm is significantly smaller for the provision of flexibility compared to health insurance or pensions. If one considers the size of a firm an indication of their ability to pay for non-wage amenities, then the smaller coefficient in column (1) implies that the decision for firms to provide flexibility depends on other considerations besides cost. In addition, we see that the coefficient for years of education is significantly larger for flexibility compared to the pecuniary amenities.

This suggests that firms pick and choose which workers receive flexibility, presumably with the amenity being given more frequently to workers with higher human capital. This is in contrast to health insurance and paid leave, where firms provide these benefits irrespective of within-occupation differences between workers.

2.4 Model

2.4.1 Environment

Time is infinite and discrete. There is a continuum of infinitely-lived agents with discount factor β. Each agent i has an observable type Ω drawn from $F(\Omega)$ and state variables $(h_i(\Omega), \eta_i(\Omega), \chi_i(\Omega))$. The agents' objective is to maximize the lifetime utility they receive from consumption and leisure. Each period, unemployed agents receive a wage and flexibility offer (w_j, f_j) from a firm of type $j \in \{1, 2\}$ as well as a match-specific utility shock ϵ, and must choose whether to accept or decline the offer.

There are two types of firms, denoted as type 1 and type 2. There is a continuum of firms of each type with measure one, and all firms of the same type are identical. Denote by K_1 and K_2 the set of firms of type 1 and type 2, respectively. Firms seek to maximize the expected present value of the profit they earn from a successful job match. The two types of firms have different productivities of production ($A_2 > A_1$), and all firms compete for workers through their wage and flexibility offers. Once a match occurs the firms observe the type Ω of the worker, which informs them of the worker's states, but does not observe ϵ. The firm then chooses wage and flexibility (w_k, f_k) in order to maximize their expected profit from the match.

A successful match between firm k and a worker of type Ω results in the firm providing the same (w_k, f_k) and the worker providing the same labor supply in each period, and lasts until there is an exogenous separation, which occurs at rate α. If a worker is not flexible ($f = 0$), the firms impose a minimum prime time work requirement: $n_p \geq \gamma_p$. In addition, inflexible workers have to commute to work, incurring a time cost κ. Flexible ($f = 1$) workers have no minimum prime time requirement and do not have to commute, but the firm suffers a cost $\lambda_j > 0$ for offering flexibility.

2.4.2 Workers' consumption and labor supply problem

Consider a wage and flexibility offer (w, f). Conditional on their own states and accepting the offer, workers choose consumption c and subperiod leisure (l_p, l_o) in order to maximize their flow utility $u^*(w, f; \Omega)$:

$$u^*(w, f; \Omega) = \max_{c, l_p, l_o} \left\{ \frac{\left[c^{\psi} L(l_p, l_o)^{1-\psi}\right]^{1-\sigma}}{1-\sigma} \right\} \tag{16}$$

s.t.

1) $c = w_j(n_p + n_o)$ (Budget constraint)

2) $n_p \geq \gamma_p$ if $f = 0$ (Inflexible labor constraint)

3) $0.5 = n_p + l_p + \chi\eta, \quad 0.5 = n_o + l_o + (1-\chi)\eta + \kappa 1(f = 0)$ (Time constraints)

4) $L(l_p, l_o) = \left[\nu l_p^{\rho} + (1-\nu) l_o^{\rho}\right]^{\frac{1}{\rho}}$ (Leisure aggregation)

If the worker declines the offer and is unemployed, she receives flow utility $u(\Omega)$:.

$$u(\Omega) = \frac{\left[c_u^{\psi} L(l_{u,p}, l_{u,o})^{1-\psi}\right]^{1-\sigma}}{1-\sigma} \tag{17}$$

s.t.

1) $c_u = w_u(h)$

2) $l_{u,p} = 0.5 - \chi_u \eta_u, \quad l_{u,o} = 0.5 - (1-\chi_u)\eta_u$

3) $L(l_{u,p}, l_{u,o}) = \left[\nu l_{u,p}^{\rho} + (1-\nu) l_{u,o}^{\rho}\right]^{\frac{1}{\rho}}$

2.4.3 Workers' matching problem

Suppose that each firm k has chosen type-specific wage and flexibility offers $(w_k(\Omega), f_k(\Omega))$, which the workers take as given. Denote by (W_1, F_1) and (W_2, F_2) the set of offers given by each firm of type 1 and 2. Each worker seeks to maximize their expected discounted lifetime utility. In addition to receiving the flow utility described above, employed workers also receive an i.i.d. Type 1 extreme value match quality shock in each period:

$$E\sum_{t=0}^{\infty}\Big[\beta^t u_t + \epsilon_t 1(\text{employed})\Big] \tag{18}$$

Consider a worker with types (Ω) who has drawn a match quality shock ϵ_k. The value of the worker receiving an offer from firm k of type j will be the following:

$$V_{j,k}(w_k, f_k, \epsilon_k, W_1, F_1, W_2, F_2; \Omega) = \max\Big\{V_e(w_k, f_k, \epsilon_k; \Omega), u(\Omega) + \beta V_u(\Omega)\Big\} \tag{19}$$

The value of accepting the offer, $V_e(w_k, f_k, \epsilon_k; \Omega)$, is given by:

$$V_e(w_k, f_k, \epsilon_k; \Omega) = u^*(w_k, f_k; \Omega) + \epsilon_k + \beta\Big[(1-\alpha)V_e(w_k, f_k; \Omega) + \alpha V_u(\Omega)\Big] \tag{20}$$

If the worker declines the offer, they will receive unemployment flow utility u in the current period and will be given another match in the next period. The offer will come from firm 1 with probability $\theta(h)$ and from firm 2 with probability $1-\theta(h)$, where $\theta'(h) < 0$. Denote by $F(\epsilon'_k)$ the distribution of future match quality shocks. The future value of unemployment V_u is thus:

$$\begin{aligned} V_u(\Omega) = \theta(h)\int_{K_1}\int_{-\infty}^{\infty} V_{1,k}(w_k, f_k, \epsilon'_k, , W_1, F_1, W_2, F_2; \Omega)\; dF(\epsilon'_k)\; dk\; + \\ (1-\theta(h))\int_{K_2}\int_{-\infty}^{\infty} V_{2,k}(w_k, f_k, \epsilon'_k, W_1, F_1, W_2, F_2; \Omega)\; dF(\epsilon'_k)\; dk \end{aligned} \tag{21}$$

Let $V^*_{j,k}$ be the maximum of the objective function for a worker who has received an offer from firm k of type j. Solving for $V^*_{1,k}$ and $V^*_{2,k}$ [3] for all k will give the reservation utility of a worker with types (Ω) and with offers (W_1, F_1) and (W_2, F_2):

[3]This can be solved numerically with value function iteration.

$$\bar{u}(W_1, F_1, W_2, F_2; \Omega) = \; u(\Omega) + \beta V_u^*(\Omega) \;, \tag{22}$$

$$\begin{aligned} V_u^*(\Omega) = \theta(h) \int_{K_1} \int_{-\infty}^{\infty} V_{1,k}^*(w_k, f_k, \epsilon_k', , W_1, F_1, W_2, F_2; \Omega) \; dF(\epsilon_k') \; dk \; + \\ (1-\theta(h)) \int_{K_2} \int_{-\infty}^{\infty} V_{2,k}^*(w_k, f_k, \epsilon_k', W_1, F_1, W_2, F_2; \Omega) \; dF(\epsilon_k') \; dk \end{aligned} \tag{23}$$

Knowing the reservation utility, we can compute the probability that a worker of type (Ω) will accept firm k of type j's offer of (w_k, f_k) given the other firms' offers (W_1, F_1) and (W_2, F_2):

$$p_k(w_k, f_k, W_1, F_1, W_2, F_2; \Omega) = 1\Big(V_{j,k}^*(w_k, f_k, \epsilon_k, W_1, F_1, W_2, F_2; \Omega) \geq \bar{u}(W_1, F_1, W_2, F_2; \Omega)\Big) \tag{24}$$

Denote by $\epsilon_k^*(\Omega)$ the match quality shock that will make the worker indifferent between accepting and declining k's offer:

$$\epsilon_k^*(\Omega) = \bar{u}(W_1, F_1, W_2, F_2; \Omega) - u^*(w_k, f_k; \Omega) - \beta\Big[(1-\alpha)V_{j,k}^*(w_k, f_k; \Omega) + \alpha V_u^*(\Omega)\Big] \tag{25}$$

Since the match quality shocks are drawn from a standard Type 1 extreme value distribution, we can finally compute the probability that a worker accepts the offer as

$$p_k(w_k, f_k, W_1, F_1, W_2, F_2; \Omega) = 1 - F(\epsilon_k^*(\Omega)) = 1 - \exp\{-\exp\{-\epsilon_k^*(\Omega)\}\} \tag{26}$$

2.4.4 Firms' technology and profits

There are two types of firms, low-performing firms (type 1) and high-performing firms (type 2). Each firm has one vacancy and can hire a worker who supplies prime and off-time labor (n_p, n_o). The production technology of each firm is linear in the worker's effective labor supply, which takes the following form:

$$Y_j(f) = A_j h(n_p + \gamma n_o) - \lambda_j f \tag{27}$$

The two firm types produce an identical product and have the same production function, except that high-performing firms have a larger TFP $(A_2 > A_1)$. The firms also have the ability to offer flexibility to their employees. If a worker is not flexible ($f = 0$), the firms impose a minimum prime time work requirement: $n_{p,i} \geq \gamma_p$. In addition, inflexible workers have to commute to work, incurring a time cost κ. Flexible ($f = 1$) workers have no minimum prime time requirement and do not have to commute, but the firm suffers a cost $\lambda_j > 0$ for offering flexibility.

The firms' instantaneous profit for a worker of type (Ω) who accepts offer (w, f) is thus:

$$\begin{aligned} \pi_j(w, f; \Omega) = A_j h\big[n_p(w, f; \Omega) + \gamma n_o(w, f; \Omega)\big] \\ - w\big[n_p(w, f; \Omega) + n_o(w, f; \Omega)\big] - \lambda_j f \end{aligned} \tag{28}$$

Notice that workers' time allocation depends on their states, wage and flexibility status.

2.4.5 Firm wage and flexibility choice

Consider a match between firm k and a worker of type Ω. The firms can observe workers' states (h, η, χ), but cannot observe the workers' match quality shock ϵ. Taking the other firms' offers as given, firm k chooses wage and flexibility $(w_k(W_1, F_1, W_2, F_2; \Omega), f_k(W_1, F_1, W_2, F_2; \Omega))$ in order maximize the expected present value of the profit of the match. In a successful match, firms earn the same profit in each period for the lifetime of the match. Since profits are not changing over time and separations are exogenous, maximizing the expected present value of the profit of the match is equivalent to maximizing the expected instantaneous profit.

The probability that a worker with observed type Ω accepts firm k's offer is given by Equation 26. The profit of the accepted match is given by Equation 28. Given the firms' offers (W_1, F_1, W_2, F_2), define $\Pi_k(w, f, W_1, F_1, W_2, F_2; \Omega)$ as firm k's payoff function for a choice of offer (w, f) to a worker of type Ω:

$$\Pi_k(w, f, W_1, F_1, W_2, F_2; \Omega) = p_k(w, f, W_1, F_1, W_2, F_2; \Omega) \cdot \pi_k(w, f; \Omega) \tag{29}$$

Finally, we can express firm k's wage and flexibility offer as the solution to the following expected profit maximization:

$$(w_k^*(W_1, F_1, W_2, F_2; \Omega), f_k^*(W_1, F_1, W_2, F_2; \Omega)) =_{w \in R, f \in \{0,1\}} \Pi_k(w, f, W_1, F_1, W_2, F_2; \Omega) \quad (30)$$

The firms compete for workers by considering the effect of the other firms' offers on the worker's reservation utility.

2.4.6 Across-type competition

I focus on the case of type-specific symmetric equilibrium wherein all firms of same type make the same offer. Solving the across-type competition problem thus equates to solving the problem for two representative firms of type 1 and 2.

For a worker of observed types Ω and firm type $-j$'s offer (w_{-j}, f_{-j}), define firm type j's payoff function as:

$$\Pi_j(w, f, w_{-j}, f_{-j}) = E[p_j(w, f, w_{-j}, f_{-j}; \Omega)] \times E[\pi_j(w, f; \Omega)] \quad (31)$$

and best-response function $BR_j : R \times \{0,1\} \to R \times \{0,1\}$ as:

$$BR_j(w_{-j}, f_{-j}; \Omega) = (w_j^*(w_{-j}, f_{-j}; \Omega), f_j^*(w_{-j}, f_{-j}; \Omega)), \quad (32)$$

where

$$(w_j^*(w_{-j}, f_{-j}; \Omega), f_j^*(w_{-j}, f_{-j}; \Omega)) =_{w \in R, f \in \{0,1\}} \Pi_j(w, f, w_{-j}, f_{-j}; \Omega) \quad (33)$$

Using best-response iteration, we can find the equilibrium wage and flexibility that each firm offers to a worker of observed type Ω:

$$(w_1^*(\Omega), f_1^*(\Omega)) \text{ and } (w_2^*(\Omega), f_2^*(\Omega))$$

2.4.7 Stationary distribution of workers' firm choice

For a worker of type (Ω), denote by $(p_1(w_1^*, f_1^*, w_2^*, f_2^*; \Omega), p_2(w_2^*, f_2^*, w_1^*, f_1^*; \Omega))$ the probabilities of the worker accepting firm type 1 or firm type 2 offer given the equilibrium wage and flexibility offers given above. We are interested in $p^* = [p_0^*, p_1^*, p_2^*]$, the stationary distribution of the worker's employment status. p_0^* denotes the probability that the worker is unemployed. Unemployed workers either receive with probability θ an offer from firm type 1 that is accepted with probability p_1, or receive with probability $1-\theta$ an offer from firm type 2 that is accepted with probability p_2. Employed workers transition back to unemployment with probability α. The worker's transition matrix is thus:

$$\mathbf{P} = \begin{bmatrix} 1-\theta p_1-(1-\theta)p_2 & \theta p_1 & (1-\theta)p_2 \\ \alpha & 1-\alpha & 0 \\ \alpha & 0 & 1-\alpha \end{bmatrix} \tag{34}$$

For each worker of type Ω, their stationary distribution $p^*(\Omega)$ will be the unique set of probabilities such that

$$p^*(\Omega) = p^*(\Omega)\mathbf{P} \tag{35}$$

2.4.8 Equilibrium

An equilibrium is defined by:

1. A set of equilibrium wage and flexibility offers $(W_j(\Omega), F_j(\Omega))$, $j \in \{1,2\}$
2. A set of worker consumption and leisure decision rules $c^*(w,f;\Omega)$, $l_p^*(w,f;\Omega)$, $l_o^*(w,f;\Omega)$
3. A set of match probabilities $p_k(w_k, f_k, W_1, F_1, W_2, F_2; \Omega)$

such that:

1. Workers maximize their expected discounted lifetime utility
2. Firms maximize the expected profit of each match
3. The stationary distributions of workers' firm choice are given by Equation 35
4. The resource constraint is satisfied:

$$\int_M A_{jm}h_{im}(n_{pim}+\gamma n_{pim}) - \lambda_{jm}f_m \; dm = \int_M w_m(n_{pim}+n_{pim}) \; dm + \int_M \pi_{km} \; dm,$$

where $m = \{i,j,k,\Omega\} \in M$ denotes a successful match between worker i of type Ω and firm k of type j

2.5 Calibration

All exogenous parameters and targeted moments are estimated using only "white collar workers" from the following occupation categories: management, business and financial operation, computer and mathematical science, architecture and engineering, life, physical, and social science and legal. I limit my calibration to workers from these occupations because of their pertinence to the topic of flexibility. Occupations with a low prevalence of flexibility may have some feature of their production technology that makes it impossible to offer flexibility to most workers, which is something I am not taking into account in my model. The white collar occupations where firms are plausibly able to offer most workers flexibility are the closest analog to the firms in my model.

Due to data limitations, I have no detailed information about the employers of the ATUS respondents. In order to still be able to investigate the role that firm productivity plays in the provision of flexibility policies, I construct a rough proxy for the productivity of each worker's employer: High-productivity firm $=$ 1(Over 1000 employees) $\times$ 1(MSA over 1,000,000 population). The simple intuition for this measure is as follows: First is that highly productive firms are the ones that are able to grow to a large size. Second is that due to agglomeration economies, workplaces that are located in large cities will on average have higher TFPs than those in less populous areas. The high-productivity firm described above corresponds to type-2 firms in the model, while firms with less than 1000 employees or those located in smaller MSA's are considered to be type-1 firms.

2.5.1 Exogenous parameters

The discount factor of the workers β is given a value of 0.995, with one period in the model corresponding to one month. The monthly job separation rate α is set to 0.05, which is about the average estimated monthly separation rates in the OECD (Hobijn and Sahin (2009)). The CRRA parameter σ is set to a standard value of 3.0. The minimum prime working hours for workers without flexibility, γ_p, is set to 0.328, or 26.25 hours per week. This is determined by multiplying the minimum weekly working hours of 35 hours per week for workers in my sample, by the average prime working hours ratio for workers without flexibility, which is about 75%. The commute time κ is set to the average commuting time of workers in the ATUS without flexibility of 0.055, or about 4.4 hours per week.

The unemployment wage $w_u(h)$ is equal to $0.125h$, or 25% of what the worker would earn working at a job where they are paid their marginal product h for 40 hours per week. The unemployment housework time η_u is the median housework time of unemployed respondents of the ATUS, computed to be 0.1458 or 11.66 hours per week. The prime housework ratio of unemployed workers χ_u is set to 60.50%, which is the median housework prime ratio of unemployed respondents of the ATUS.

2.5.1.1 Mapping Ω from data to model

Ω contains information on four characteristics of workers that are observed by the firm:

1. Education: High school diploma, college degree, graduate degree
2. Sex: Male or female
3. Experience (Age minus education): 5 to 9 years of experience, 10 to 19 years, and 20 or more years experience
4. Productivity: Above, below, or within one standard deviation of the mean of worker human capital

Each education/sex/experience/productivity type corresponds to a vector of model state variables $(h(\Omega), \eta(\Omega), \chi(\Omega))$, whose values are computed using the ATUS/CPS data. The housework state variables η, χ are assigned as follows: ATUS respondents are grouped by sex and experience[4], and the within-group average housework and housework ratio is used as the $(\eta(\Omega), \chi(\Omega))$. Table 24 shows the housework state variables by sex and experience. Women in the ATUS time diaries supply more household labor on average than men, and both sexes do more housework when they are younger. The percentage of housework performed during prime time ranges from 11.6% to 15.4% for men and 12.2% to 20.5% for women, with no discernible pattern in the relationship between χ and experience.

[4]Education was found to have a negligible effect on housework time allocation.

Since human capital is unobservable, I use the normalized log wages of white collar workers in the ATUS/CPS sample as an analog for human capital. For each education/experience group[5] I compute the within-group standard deviation of human capital. High-productivity workers are those with a human capital greater than one standard deviation above the group mean, and are assigned as $h(\Omega)$ the mean human capital of the high-productivity workers. The same is done for workers with human capital less than one standard deviation below the group mean, and for workers within one standard deviation of the group mean. The values used in the model are shown in Table 25.

Note that three of the four characteristics in Ω are observed by the econometrician and are included in the Mincer regressions from section 3.1. The fourth characteristic, productivity, is observed by firms but unobserved by the econometrician, and it is this unobserved variable that will allow the model to provide more accurate measurements of the compensating differential between flexibility and wages than what can be found in the data.

2.5.2 Method of moments

The remaining parameters $(A_1, A_2, \psi, \rho, \nu, \lambda_1, \lambda_2, \theta, \xi, \gamma)$ are calibrated using simulated method of moments (SMM). These ten parameters are targeted to ten moments from the data concerning the allocation of flexibility by firm type, as well as worker outcomes by firm type and flexibility status. Moments are computed for only white collar workers. A description of the calibrated parameters is provided in Table 26.

The non-pecuniary cost of flexibility is represented by γ, the productivity of off-time labor, which is assumed to be weakly less than one. This parameter is targeted to the coefficient on flexibility in the following Mincer regression:

$$\log(\text{wage}_i) = \beta_0 + \beta_1 Flex_i + \beta_2 Firm_i + \Lambda X_i + \Gamma Occ_i + \epsilon_i$$

In this equation, X_i are the characteristics in Ω that are observed by the researcher: education, sex and experience. The occupation fixed effects are also included because of the model's scope being a labor market for one white-collar occupation. This gives the coeffi-

[5] I did not group by sex because the difference in means of men and women's wages would be erroneously attributed in the model to differences in men and women's human capital.

cient on flexibility the interpretation of the within-occupation relationship between wages and flexibility, controlling for relevant observables.

A_1 is the TFP of firm type 1. Since we assume that $A_2 > A_1$, the value of this parameter is also informative of the productivity of type-2 firms. Since workers' labor supply depends on η, firms will offer different wage and flexibility bundles to workers with different η's. The TFP of the firms will help determine the bundles that are offered, as well as the relationship between flexibility offers and η. Thus I target the correlation between flexibility and η. The TFP of firm type 2, A_2, is targeted to the coefficient on firm type in the Mincer regression shown above.

The parameter ψ determines how workers value leisure relative to consumption, and thus determines the workers' total labor supply allocation. ν is the prime/off-time leisure share parameter, which affects how workers change their allocation of labor supply when the constraints on time use are lifted under flexibility. I thus jointly target the mean total weekly hours worked of workers with flexibility, as well as the mean total weekly hours worked of those without flexibility. ρ determines the degree of substitutability between prime and off time leisure, which has implications for how workers allocate their labor supply throughout the day. I thus target the prime working hours ratio for flexible workers[6].

(λ_1, λ_2) represent the pecuniary cost of flexibility for firms of type 1 and type 2, respectively. These two parameters are targeted to the share of flexible workers at low and high-productivity firms, respectively. The probability of an offer from a firm of type 2 is given by the following logistic function:

$$\theta(h) = \frac{1}{1 + e^{-\frac{h-\theta}{2\xi}}}$$

The midpoint parameter θ affects the overall level of the function at any given human capital, and is thus targeted to the share of workers at low-productivity firms. The growth rate parameter ξ determines the rate of change in the type-2 firm match probability with respect to human capital, and is targeted to the correlation between worker education and firm type.

The estimates of the calibrated parameters are given in Table 26. The leisure substi-

[6] I do not target this for inflexible workers because inflexible workers' labor supply allocation is constrained.

tution parameter ρ is positive, meaning that prime and off-time leisure are substitutes in the workers' utility function. A value for γ of 0.8691 implies labor hours in off-time are about 87% as productive as labor hours in prime-time. Type-2 firms are found to be about 67% more productive than type-1 firms, but their pecuniary flexibility costs are higher. One striking result is that the ratio of TFP to flexibility cost is approximately the same for each firm type, that is, $\frac{A_1}{\lambda_1} \approx \frac{A_2}{\lambda_2}$. This means that the pecuniary cost of flexibility as a percentage of firm output is nearly the same for all firms in the market[7].

2.5.3 Model fit

Table 27 shows the performance of the calibration with respect to the targeted moments. The model closely matches the Mincer regression coefficients on flexibility and firm type seen in the data, as well as the flexibility/housework correlation and the education/firm type correlation. The model has a higher share of workers at firms of type 1 than the data, and also has slightly higher shares of workers in both firm types with flexibility compared to the data. While the models closely matches the labor supply for workers without flexibility, it overshoots both the labor supply of flexible workers as well as the prime working hour ratio for flexible workers.

Tables 28 and 29 show the results of the model in matching some selected untargeted moments. The model succeeds in generating a negative gender wage gap despite the only difference between men and women in the model being higher household labor requirements for women. This implies that differences in household labor account for about 17% of the estimated within-occupation gender wage gap for white collar workers. The model also does a good job at matching the percentages of women and men with flexible jobs, as well as the wages and labor supply of workers by flexibility and firm type. While the prime working hours ratio for workers without flexibility is much higher in the model than in the data, the model does replicate the negative relationship between flexibility and the prime ratio that is seen in the data. The model falls short of replicating the education/flexibility correlation

[7] One major reason that has been given by firms for walking back flexibility policies are the sunk costs of expensive offices that are not used due to employees working from home. An explanation for higher-productivity firms having higher flexibility costs may be that these firms may have offices that are larger and more costly to operate compared to less productive firms

seen in the data, suggesting that the relationship between the provision of flexibility and worker human capital is quite exaggerated in the model compared to the data.

2.6 Experiments

2.6.1 The compensating wage differential of flexibility

Columns (1) and (2) of Table 30 regress log wages on flexibility status and firm type, including as controls sex, experience and education. These are the main worker characteristics that are observable in the data, and are thus included in the workers' types Ω in the model. They both find that workers with flexibility earn on average ten percent higher wages than workers without flexibility. This positive relationship is due to the lack of human capital controls in these regressions; since higher human capital workers earn higher wages and are more likely to have flexibility, an erroneous positive effect is estimated.

The model includes worker productivity in Ω, with the interpretation that there are several pieces of information such as the content of resumes and interview performance that allow firms to learn more about the human capital of workers. While firms use this information to decide optimal wage/flexibility offers, these things are unobserved by the researcher, and this lack of information about the data generating process leads to biased estimates of the compensating wage differential of flexibility in the data. This issue has historically plagued the compensating differentials literature (Brown (1980), Hwang et al. (1992), Kniesner et al. (2012)), but using a calibrated model can help provide a better estimate of the compensating wage differential of flexibility. Column (3) shows with that with the inclusion of human capital controls the sign on flexibility reverses. Workers with flexibility are now found to earn about seven percent lower wages on average compared to identical workers without flexibility.

2.6.2 The role of labor supply endogeneity

A crucial difference between flexibility and other highly-studied job amenities is the way that flexibility affects the intensive labor supply allocation of workers. Unlike amenities such as health insurance or paid leave, the decision to provide flexibility takes into account not just the pecuniary cost, but also the effect of providing flexibility on workers' labor supply. The calibrated model shows that off time labor is less productive than prime time labor, so providing flexibility has the additional cost of workers allocating more of their labor to off time. In order to determine the degree to which this additional cost affects the allocation of flexibility, I perform an experiment in which flexible workers cannot adjust their labor supply; they must use the same labor supply allocation that they would use if they did not have flexibility.

Table 31 shows the allocation of flexibility in the base model, in the no labor supply effect experiment, and the allocation of other amenities among white collar workers in the data. Removing the labor supply endogeneity of flexibility brings the allocation of flexibility in line with the other amenities: high-performance firms offer 95 percent of workers flexibility, and low-performance firms offer flexibility to 82 percent of workers. This validates the interpretation of the empirical results from Section 3 that the unique relationship between flexibility and labor supply allocation has a significant effect on the provision of flexibility. The table also shows that the relationship between human capital and amenity provision is weaker in the data for health insurance and paid leave compared to flexibility. The model is also able to replicate this feature of the data, as removing the labor supply endogeneity of flexibility weakens the relationship between worker human capital and the provision of flexibility in the calibrated model. These results imply that when offering an amenity like flexibility that will adversely affect the labor supply allocation of workers, employers will opt to discriminate and provide the amenity more readily to higher human capital workers.

2.6.3 Decomposing flexibility and wage allocations

There are three main mechanisms that determine the joint distributions of flexibility and wages in the model: competition between firms for workers, assortative matching, and the labor supply endogeneity of flexibility discussed above. In order to quantify the effects of these mechanisms, I perform a decomposition exercise by shutting down each of these features one by one. The results are shown in Table 32. Competition is removed from the model by setting the workers' discount factor β equal to zero; without the ability to consider future offers, workers accept any offer that provides a flow utility greater than that of unemployment. Firms in turn will respond by providing this minimally-acceptable offer, without needing to consider other firms' offers. Without competition firms of type 1 will provide flexibility to only 24 percent of workers, and wages are on average about 6 percent lower in absence of competition. The lower overall wages also result in a smaller compensating wage differential of flexibility.

Removing assortative matching results in an equilibrium in which type-1 firms provide flexibility more often than type-2 firms. This is due to the fact that flexibility is strongly correlated with human capital in the model. Once high-human capital workers start matching with each firm type at the same rate, the low-TFP firms will offer flexibility slightly more often due to their lower pecuniary cost of flexibility. The average difference in wages between type-1 and type-2 firms decreases from 10.43% to 0.91%, implying that assortative matching accounts for over 90% of the firm-type wage gap found in the model. As shown above, shutting down the labor supply endogeneity of flexibility increases the allocation of flexibility for both type-1 and type-2 firms. Removing this additional cost of flexibility also reduces the compensating wage differential of flexibility by about 40%. Due to the higher provision of flexibility, wages are about 3 percent lower overall compared to the base model.

2.6.4 Estimating workers' valuation of flexibility and wages

In order to estimate how much workers value different wage/flexibility offers, we can look at $(p_1(w_1, f_1, w_2, f_2; \Omega), p_2(w_2, f_2, w_1, f_1; \Omega))$, the probability that a worker of type Ω accepts an offer from each firm type given the offers of both firm types. By regressing these probabilities on the offers of each firm type, we can estimate the marginal effect of an increase in wages or a flexibility offer on the probability of accepting an offer. Table 33 shows these marginal effects. The marginal effect of flexibility and wages is the same for both firm types;

when either firms of type 1 and type 2 offer flexibility, workers are on average 23 percentage points more likely to accept their offer. Workers are 10pp less likely to accept an offer from firms of type 1 if the type-2 firms offer flexibility, and are about 15pp less likely to accept an offer from firms of type 2 if the type-1 firms offer flexibility. These negative coefficients are a consequence of the incentive that workers have to decline an offer if they believe the other firms have a better one. Since workers receive offers from firms of type 1 more often, they will more readily decline offers from type-2 firms when the type-1 offer is better because they will not have to wait as long for the better offer. A one percent increase in wages increase the probability of workers accepting an offer by about 1.1 percentage points for both firm types. Similar to flexibility, the negative effect on accepting from an increase in the other firm type's wage offer is also greater for type-2 firms.

2.7 Welfare analysis: the costs and benefits of workplace flexibility

A key result from the calibrated model are the welfare implications of workplace flexibility as an amenity offered by firms. While section 6.1 shows the costs to workers of flexibility in terms of foregone wages, we are also interested in quantifying the benefits to workers of this amenity. Also of great interest to both academics and employers is an investigation of the potential costs and benefits for firms for providing flexibility. In order to evaluate the effects of workplace flexibility on the welfare of both workers and firms, we compare their outcomes from the calibrated model to one in which there is no flexibility; in the no-flexibility economy, firms can only compete for workers through wages, and workers have no ability to avoid the restrictive time constraints from conventional work.

2.7.1 Worker welfare

Table 34 compares the welfare of workers between the base model and the no-flexibility economy. The first row of Panels A and B shows the average consumption equivalence of workers by productivity type, with the interpretation of the consumption equivalence values being the percentage change of base-model consumption that would make workers indifferent between the base model and the no-flexibility economy. The negative signs of this value mean that workers are willing to forgo consumption in order to stay in the economy with flexibility. The average consumption equivalence is decreasing with workers' productivity, which is primarily due to flexibility being more common among higher human capital workers. The second row of each panel shows the average percentage change in workers' wages from the base model to the no-flexibility economy. Wages are higher overall in the no-flexibility economy because firms no longer have a valuable amenity they can offer workers; wages must be higher than in the base model in order for workers to accept the firms' offers.

Both men and women highly desire flexibility; despite earning higher wages in the no-flexibility economy, men are willing to pay 7 to 8 percent of their base-model income in order to stay in the base model, while women willing to pay 33 to 58 percent of their income in order to have access to flexibility. There are two explanations for the dramatically higher consumption equivalence for women compared to men. First, women on average have higher household labor requirements than men, which means that time constraints from non-flexible work have a more punitive effect on women's utility from leisure compared to men, and thus women suffer more from a lack of flexibility compared to men. The household labor requirements for some women are high enough that they are not able to meet the minimum prime work hours constraint and thus cannot work unless they have a flexible job. In the no-flexibility economy these women will be chronically unemployed, as shown in Panel C. The unemployment rate for women is shown to be almost 9 percentage points higher in the no-flexibility economy compared to an increase of less than one percentage point for men.

The second reason for women's higher consumption equivalence is that women are compensated less for the loss of flexibility compared to men; while men on average earn 8 percent higher wages in the no-flexibility economy, women only earn about 4 percent higher wages. This is also shown in Panel C, which reports that the no-flexibility economy has an estimated gender wage gap that is over three times the size as the one from the base economy. This striking result demonstrates the huge role that access to workplace flexibility policies plays in mitigating the negative effects of women's higher housework labor requirements on their labor market outcomes. Taken in full, the results from this experiment tell a gendered story:

while workplace flexibility policies are highly beneficial to men, the benefits of these policies are predominantly enjoyed by women, with access to flexibility proving to be a crucial component in ensuring gender parity in the labor market.

2.7.2 Firm welfare

The analysis of the effects of workplace flexibility on firm outcomes begins with a comparison of per-worker profits between the base economy and the no-flexibility economy. Panel A of Table 35 shows the average change in per-worker profits by firm type and by worker productivity. On average, low-TFP firms earn about 4 percent lower profits in the no-flexibility economy. Since offering flexibility is less costly for low-TFP firms, these firms lose profits when they no longer have this cheap amenity to offer workers and must instead raise wages. The opposite is true for high-TFP firms, which see an average increase in profits of about 6 percent in the no-flexibility economy. This is because the high-TFP firms have a higher pecuniary cost of flexibility, and since the high-TFP firms only need to increase their wage offers by as much as the low-TFP firms does, overall they save and have higher profits.

The removal of workplace flexibility changes the way that firms compete in the labor market, with firms only able to compete for workers via wages. The effects of this change can be seen in Panel B, which shows the share of workers at high-TFP firms. The share of workers who are employed at high-TFP firms in the no-flexibility economy is 10 percentage points higher compared to the base economy. This demonstrates how the ability to offer workplace flexibility allows low-TFP firms to more effectively compete for workers with high-TFP firms. These changes in per-worker profits and the distribution of workers result in large changes to total profits by firm type: the total profits of high-TFP firms are 35 percent higher in the no-flexibility economy, while low-TFP firms collectively have 22 percent lower profits. These results show that the cost-benefit analysis of workplace flexibility for firms depends totally on firm type: the ability to offer workplace flexibility is a net benefit to low-TFP firms, while having to offer flexibility makes high-TFP firms worse off.

2.7.3 Aggregate effects

The effects of workplace flexibility on aggregate variables are shown in Table 36. The results show that total output is actually slightly higher in the no-flexibility economy, despite a higher unemployment rate. This is due to the previous finding that the share of workers at high-TFP firms is higher when there is no flexibility. Since more workers are employed at high-TFP firms, these worker-firm matches produce higher output and more than offset the foregone output of those who are pushed out of the labor force due to a lack of flexibility. This understates the difference in output, as some of the output in the base model is used to pay the pecuniary costs of providing flexibility. When looking at output net of these costs, which is equal to total labor income plus total profits, the total output in the no-flexibility economy is about 7 percent higher. In particular, total labor income is one percent higher and total profits are about 20 percent higher in the no-flexibility economy, with the labor share of income decreasing from 71.43% to 67.72%.

The comparison of the base and no-flexibility economies presents a clear set of trade-offs: the provision of workplace flexibility greatly benefits workers through the relaxing of work-related time constraints. These benefits are particularly large for women, who enjoy not only greater temporal flexibility but also a lower gender wage gap and lower unemployment compared to a counterfactual world without flexibility. Low-TFP firms also enjoy a net increase in per-worker and total profits due to flexibility being relatively cheap for them to offer while allowing them to compete for workers with high-TFP firms. On the other hand, this competition for workers via flexibility makes high-TFP firms worse off by decreasing their per-worker profit as well as reducing the number of workers they employ. This lower employment at high-TFP firms also has a deleterious effect on aggregate output, though most of these foregone gains would have gone towards high-TFP firm profits, as shown by the higher labor share of income in the base model and from Table 36.C. In sum, workplace flexibility greatly improves the outcomes of female workers along with male workers and low-TFP firms also seeing net benefits, at the expense of the profits of high-TFP firms.

2.8 Conclusion

In the midst of the current debate on the future of flexible work in corporate America, this project provides the first systematic analysis of the cost, benefits, and distribution of workplace flexibility policies. My analysis begins by utilizing the Leave and Job Flexibilities Module of the American Time Use Survey to develop two flexibility measures: flextime, or the ability to choose when one begins and ends work, and paid work-from-home. Using these measures I examine the relationship between worker outcomes and access to flexibility and other important job amenities. I find that flexibility is the only amenity to have a significant relationship with workers' intensive labor supply allocation. Logistic regressions of amenity status on observables shows that the allocation of flexibility is idiosyncratic than for other amenities; while there is significant positive comovement between the other amenities, while the provision of flexibility is unaffected by the presence of other amenities. There is also evidence of within-occupation discrimination as well as a much higher across-occupation variance of flexibility compared to other amenities. All this implies that firms are more discerning when choosing who to offer flexibility to, taking into account workers' human capital as well as the effect of flexibility on their labor supply allocation.

Next, I develop a model of random search with separation in which workers match with firms who offer them bundles of wages and flexibility policies. Firms are heterogeneous in their productivity and cost of providing flexibility, and compete in the labor market via take-it-or-leave-it wage and amenity offers. The mechanism of competition is that firms' offers affect the workers' reservation utility; making a more enticing offer will increase the probability for workers to reject a competitor's offer and wait for the more favorable match. One innovation of this model is to embed within this framework a detailed modeling of workers' time use, with workers choosing not only how much labor to supply, but *when* to supply it. Workers are heterogeneous in both their human capital and required household labor, and are thus also heterogeneous in their demand for flexibility, which relaxes various constraints on their time.

The model is calibrated using data on white-collar workers in the ATUS. Analysis of data simulated from the model shows that estimates of the compensating wage differential of flexibility using the ATUS data are wrong-signed, with the value estimated from the model to be -0.0686, a compensating differential of about 7 percent. The model also reveals that about one-fifth of the observed within-occupation gender wage gap in the data can be attributed to gender differences in required household labor. Decomposition exercises confirm that the allocation of flexibility is influenced heavily by the amenity's effect on workers' labor supply allocation. A welfare analysis comparing worker and firm outcomes between the calibrated model and a counterfactual model without flexibility reveals a clear set of costs and benefits. Access to workplace flexibility policies generate a huge benefit to female workers by substantially reducing both the gender wage gap and female unemployment. Workers of both genders additionally benefit due to the relaxed time constraints allowing both a higher labor supply as well as more leisure, with women again receiving a particularly high benefit due to their higher household labor requirements. Low-productivity firms are able to use this desirable amenity to more effectively compete for workers, and thus enjoy higher profits than in the counterfactual economy without flexibility. The main losers are high-productivity firms, who receive lower profits due to greater competition for workers from low-productivity firms as well as relatively high costs for providing flexibility.

The analyses in this paper can be improved along several dimensions with more comprehensive data. Firm-level data on wages, amenities and worker characteristics would allow for better estimates of the compensating wage differential of flexibility, and the analysis of the determinants of the provision of flexibility could be much more thorough. This would also allow for greater worker and firm heterogeniety in the model. Other directions for future improvement of the model would be the inclusion of dynamic worker states, allowing for firms to offer different combinations of flextime and/or work-from-home, and including other amenities.

2.9 Tables

Table 19: Summary statistics by flexibility status, ATUS/CPS and time diary data

	Flexibility			No flexibility		
	Mean	Std. Dev	N	Mean	Std. Dev	N
Age	42.53	7.711	1349	43.32	10.54	2618
Male	0.550	0.498	1349	0.521	0.499	2618
College degree	0.698	0.459	1349	0.399	0.490	2618
Married	0.649	0.477	1349	0.561	0.496	2618
Number of children	1.179	1.079	1349	1.012	1.086	2618
Large MSA	0.6323	0.4823	1349	0.529	0.499	2618
Large firm	0.551	0.497	1349	0.473	0.499	2618
Hourly wage	34.31	15.56	1349	23.59	12.70	2618
Hours worked/week	45.75	8.375	1349	43.61	8.024	2618
Prime working hours ratio	0.7106	0.2569	603	0.7042	0.2260	1070
Daily housework time, minutes	120.06	117.31	658	113.99	122.81	1258
Daily leisure time, minutes	420.42	195.43	654	433.71	200.34	1250

Table 20: Flexibility policies by occupation category

Occupation category	Flextime	Paid work-from-home	Flexible job	N
Management	0.3286	0.4245	0.5613	563
Business and financial operation	0.3309	0.4945	0.5927	275
Computer and mathematical science	0.2892	0.6373	0.7206	204
Architecture and engineering	0.4527	0.5135	0.6689	148
Life, physical, and social science	0.3175	0.3651	0.4762	63
Community and social service	0.2547	0.3208	0.4528	106
Legal	0.3485	0.4545	0.5606	66
Education, training, and library	0.0991	0.1381	0.1772	333
Arts, design, entertainment, sports, and media	0.3134	0.3881	0.5373	67
Healthcare practitioner and technical	0.1413	0.0967	0.1896	269
Healthcare support	0.0441	0.0294	0.0588	68
Protective service	0.0947	0.0737	0.1474	95
Food preparation and serving	0.1429	0.0220	0.1538	91
Building and grounds cleaning and maintenance	0.0536	0.0357	0.0714	112
Personal care and service	0.1176	0.1176	0.1961	51
Sales	0.2953	0.3346	0.4646	254
Office and administrative support	0.1282	0.1731	0.2457	468
Farming, fishing, and forestry	0.1154	0.0385	0.1154	26
Construction and extraction	0.0979	0.0909	0.1538	143
Installation, maintenance, and repair	0.1228	0.0614	0.1579	114
Production	0.0625	0.0110	0.0699	272
Transportation and material moving	0.0894	0.0279	0.1006	179
All occupations	0.2017	0.2475	0.3401	3967
Across-occupation variance	0.0134	0.0366	0.0462	

Table 21: Worker outcomes by access to job amenities, within occupation

Panel A: Log hourly wages					
Flexibility	0.0975***				0.0941***
	(0.0229)				(0.0224)
Paid leave		0.102***			0.0813***
		(0.0243)			(0.0240)
Employer-provided health insurance			0.147***		0.112***
			(0.0315)		(0.0314)
Employer-provided pension				0.109***	0.0980***
				(0.0180)	(0.0180)
Industry & occupation FEs	Yes	Yes	Yes	Yes	Yes
Observations	3265	3265	3265	3265	3265
Adjusted R^2	0.506	0.506	0.507	0.509	0.521
Panel B: Log weekly working hours					
Flexibility	0.0529***				0.0531***
	(0.00846)				(0.00845)
Paid leave		0.00717			0.00409
		(0.00843)			(0.00829)
Employer-provided health insurance			-0.000998		-0.00512
			(0.0123)		(0.0124)
Employer-provided pension				0.00454	0.00631
				(0.00675)	(0.00671)
Industry & occupation FEs	Yes	Yes	Yes	Yes	Yes
Observations	3265	3265	3265	3265	3265
Adjusted R^2	0.196	0.180	0.180	0.180	0.196
Panel C: Prime working hours ratio					
Flexibility	-0.0722***				-0.0739***
	(0.0187)				(0.0188)
Paid leave		0.00810			0.0148
		(0.0209)			(0.0208)
Employer-provided health insurance			-0.0165		-0.00990
			(0.0239)		(0.0240)
Employer-provided pension				-0.0298*	-0.0318**
				(0.0160)	(0.0160)
Industry & occupation FEs	Yes	Yes	Yes	Yes	Yes
Observations	1374	1374	1374	1374	1374
Adjusted R^2	0.070	0.056	0.056	0.059	0.071

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 22: Job amenities by occupation category

Occupation category	Flex Job	Health Insurance	Paid Leave	Pension	N
Management	0.5613	0.9218	0.8774	0.5151	563
Business and financial operation	0.5927	0.9309	0.8545	0.5055	275
Computer and mathematical science	0.7206	0.9412	0.9069	0.5245	204
Architecture and engineering	0.6689	0.9730	0.9122	0.5473	148
Life, physical, and social science	0.4762	0.9365	0.9524	0.5556	63
Community and social service	0.4528	0.9151	0.9151	0.5566	106
Legal	0.5606	0.8636	0.9091	0.4848	66
Education, training, and library	0.1772	0.9459	0.7748	0.6456	333
Arts, design, entertainment, sports, and media	0.5373	0.9254	0.8209	0.5373	67
Healthcare practitioner and technical	0.1896	0.9219	0.8699	0.5613	269
Healthcare support	0.0588	0.9559	0.6765	0.4118	68
Protective service	0.1474	0.9684	0.9053	0.6632	95
Food preparation and serving	0.1538	0.6813	0.4725	0.0659	91
Building and grounds cleaning and maintenance	0.0714	0.7500	0.6429	0.2589	112
Personal care and service	0.1961	0.5686	0.5294	0.2157	51
Sales	0.4646	0.8504	0.7992	0.3858	254
Office and administrative support	0.2457	0.8910	0.8291	0.4594	468
Farming, fishing, and forestry	0.1154	0.5000	0.4615	0.1538	26
Construction and extraction	0.1538	0.7203	0.4545	0.3357	143
Installation, maintenance, and repair	0.1579	0.8158	0.8158	0.3333	114
Production	0.0699	0.8419	0.7279	0.4301	272
Transportation and material moving	0.1006	0.8268	0.6536	0.4302	179
All occupations	0.3401	0.8823	0.7973	0.4737	3967
Across-occupation variance	0.0462	0.0067	0.0145	0.0136	

Table 23: Access to amenities by worker and firm characteristics, within occupation

	(1) Flexibility	(2) Health Insurance	(3) Paid Leave	(4) Pension
Amenity index	0.0509	0.911***	0.498***	0.345***
	(0.0812)	(0.103)	(0.0853)	(0.0774)
Education	0.160***	0.0326	0.0207	0.0767***
	(0.0297)	(0.0365)	(0.0293)	(0.0233)
Large Firm	0.311**	1.033***	0.111	0.698***
	(0.126)	(0.193)	(0.135)	(0.105)
Female	0.0637	-0.313	-0.0738	-0.139
	(0.137)	(0.220)	(0.158)	(0.122)
Married	0.0943	0.0224	0.180	-0.0598
	(0.136)	(0.182)	(0.140)	(0.111)
Number of Children	0.102*	-0.0305	-0.0713	0.0329
	(0.0593)	(0.0845)	(0.0631)	(0.0512)
Large MSA	-0.0789	-0.264	-0.135	-0.116
	(0.124)	(0.167)	(0.131)	(0.105)
Industry & occupation FEs	Yes	Yes	Yes	Yes
N	3077	2777	3144	3215
pseudo R^2	0.292	0.237	0.193	0.157

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 24: Housework time allocation by sex and experience

Panel A: Mean weekly housework hours			
	Experience		
	5-9 years	10-19 years	20+ years
Men	6.62	7.58	4.51
Women	10.91	10.81	7.42
Panel B: Mean housework prime ratio			
	Experience		
	5-9 years	10-19 years	20+ years
Men	12.57%	15.42%	11.64%
Women	16.58%	12.21%	20.47%

Table 25: Human capital by education, experience and productivity

		Experience								
		5-9 years			10-19 years			20+ years		
Productivity		Low	Average	High	Low	Average	High	Low	Average	High
Education	High school	0.69	0.84	1.02	0.68	0.88	1.07	0.73	0.94	1.12
	College	0.72	0.93	1.13	0.82	0.99	1.14	0.80	1.03	1.16
	Graduate	0.81	1.01	1.14	0.85	1.05	1.17	0.84	1.07	1.18

Table 26: Calibrated parameters

Parameter	Value	Description
A_1	1.3827	Type-1 firm TFP
A_2	2.3095	Type-2 firm TFP
ψ	0.3822	Consumption/leisure Cobb-Douglas parameter
ρ	0.8474	Substitution b/t prime vs off-time leisure
ν	0.2698	Leisure prime/off-time share parameter
λ_1	0.0680	Flexibility cost for type-1 firm
λ_2	0.1121	Flexibility cost for type-2 firm
θ	1.1764	$\theta(h)$ intercept parameter
ξ	0.0674	$\theta(h)$ slope parameter
γ	0.8691	Productivity of off-time labor

Table 27: Targeted moments

Moment	Data	Model
Type-1 firm worker share	62.82%	69.17%
Type-1 firm flexibility share	57.79%	61.54%
Type-2 firm flexibility share	64.45%	66.24%
Mincer regression flexibility coefficient	0.1057	0.1066
Mincer regression firm coefficient	0.0854	0.0897
Prime working hour ratio, flex workers	73.22%	83.37%
Weekly labor hours, flex workers	45.55	47.14
Weekly labor hours, inflex workers	44.17	44.37
Housework/flexibility correlation	0.0987	0.0974
Education/firm type correlation	0.1397	0.1363

Table 28: Untargeted moments

Moment	Data	Model
Mincer regression sex coefficient	-0.1402	-0.0241
Prime working hour ratio, flex workers	74.60%	87.92%
Percent of male workers with flexibility	61.00%	65.33%
Percent of female workers with flexibility	58.84%	59.81%
Education/flexibility correlation	0.1640	0.3857

Table 29: Untargeted moments: mean labor supply and relative wages by firm type and flexibility

		Data		Model		
		No Flexibility	Flexibility		No Flexibility	Flexibility
Panel A: Labor supply	Firm type-1	44.35	45.73	Firm 1	44.67	46.52
	Firm type-2	43.63	45.39	Firm 2	43.61	48.44
		No Flexibility	Flexibility		No Flexibility	Flexibility
Panel B: Wages	Firm type-1	1.00	1.153	Firm 1	1.00	1.131
	Firm type-2	1.116	1.388	Firm 2	1.099	1.243

Table 30: Mincer regressions, data and model, with and without human capital controls

	(1) Log wage, data	(2) Log wage, model	(3) Log wage, model
Flexibility	0.106***	0.107***	-0.0686***
	(0.0343)	(0.00126)	(0.000560)
High TFP firm	0.0854**	0.0897***	-0.00554***
	(0.0343)	(0.00117)	(0.000464)
Female	-0.140***	-0.0241***	-0.0368***
	(0.0345)	(0.00109)	(0.000397)
College degree	0.227***	0.0159***	0.108***
	(0.0426)	(0.00145)	(0.000556)
Graduate degree	0.370***	0.0707***	0.183***
	(0.0472)	(0.00158)	(0.000612)
Average human capital			0.281***
			(0.000612)
High human capital			0.484***
			(0.000892)
Observations	1089	46817	46817
Adjusted R^2	0.307	0.337	0.911

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 31: Effect of flexibility/labor supply interaction on provision of flexibility

	Flexibility, model	Flexibility, no labor supply effect	Flexibility, data	Health insurance, data	Paid leave, data
Percent flex, type-1 firm	61.54%	81.75%	57.79%	91.00%	87.58%
Percent flex, type-2 firm	66.24%	95.55%	64.45%	97.64%	92.08%
Education/flexibility correlation	0.3857	0.3148	0.1640	0.0689	0.0924

Table 32: Distribution of flexibility with different model elements present

Pct. flex, firm type-1	Pct. flex, firm type-2	Firm comp. diff	Avg. firm wage diff.	Pct. change in wages	Competition	Assortative matching	Labor supply endogeneity
61.53%	66.24%	-0.0686	10.43%	-	x	x	x
23.88%	84.83%	-0.0207	7.88%	-6.34%		x	x
68.53%	57.49%	-0.0868	0.91%	0.28%	x		x
81.75%	95.55%	-0.0402	11.58%	-2.94%	x	x	
23.52%	68.38%	-0.0184	0.15%	-6.46%			x
46.04%	83.81%	-0.0072	7.99%	-6.82%		x	
95.63%	97.07%	-0.0582	4.49%	-3.33%	x		
45.11%	83.03%	-0.0030	-0.13%	-6.87%			

Table 33: Effect of wage and flexibility on probability of accepting job offers

	(1) Prob accept firm type-1	(2) Prob accept firm type-2
Firm type-1 flex offer	0.234***	-0.155***
	(0.0000342)	(0.0000367)
Firm type-1 log wage offer	1.111***	-0.875***
	(0.0000720)	(0.0000772)
Firm type-2 flex offer	-0.103***	0.230***
	(0.0000342)	(0.0000367)
Firm type-2 log wage offer	-0.590***	1.161***
	(0.0000720)	(0.0000772)
Observations	270473008	270473008
Adjusted R^2	0.592	0.605

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 34: Changes in workers' welfare from base model to no-flex regime

Panel A: Men				
	Low HC	Average HC	High HC	All HC
Consumption equivalence	-7.08%	-7.51%	-8.43%	-7.58%
Difference in wages	2.47%	8.71%	10.27%	7.94%
Panel B: Women				
	Low HC	Average HC	High HC	All HC
Consumption equivalence	-33.54%	-43.55%	-57.98%	-44.34%
Difference in wages	0.0%	4.44%	4.27%	3.73%

Panel C: Changes in other worker outcomes		
	Base model	No-flex model
Gender wage gap, Mincer coefficient	-0.0241	-0.0806
Male unemployment rate	6.61%	7.42%
Female unemployment rate	6.03%	14.68%

Table 35: Changes in firm outcomes from base to no-flex economy

Panel A: Change in per-worker profits		
	Low-TFP firms	High-TFP firms
Low-productivity workers	3.42%	13.90%
Average-productivity workers	-6.61%	5.95%
High-productivity workers	-10.66%	7.33%
All workers	-4.16%	6.19%
Panel B: Change in share of workers at high-TFP firms		
	Base model	No-flex
Low-productivity workers	8.19%	16.35%
Average-productivity workers	30.06%	40.28%
High-productivity workers	58.08%	68.68%
All workers	30.83%	40.94%
Panel C: Change in total profits and profit shares		
	Low-TFP firms	High-TFP firms
Percent change in total profits	-21.77%	34.82%
Profit shares, base model	25.32%	74.68%
Profit shares, no-flex model	16.44%	83.56%

Table 36: Changes to aggregate variables from base to no-flex economy

	Base model	No-flex model
Total output, percentage of base model	100.00%	100.76%
Total output net of flexibility costs	100.00%	106.63%
Total labor income, percentage of base model	100.00%	101.08%
Total profits, percentage of base model	100.00%	120.49%
Labor share of income	71.43%	67.72%

3.0 Measuring Workplace Flexibility

3.1 Introduction

Since Claudia Goldin's seminal 2014 paper, *A Grand Gender Convergence: Its Last Chapter*, workplace flexibility has gained recognition among researchers as a job amenity with important implications for workers' labor market outcomes. The experience of the COVID-19 pandemic further increased interest in the study of flexible work arrangements, leading to a burgeoning literature on the effects of job flexibility on worker outcomes. These papers followed the lead of Goldin (2014) in the use of occupation characteristic variables as proxy measures of flexibility. In particular, Goldin selects a handful of characteristics from the Occupational Information Network (O*NET) that describe the flexibility of occupations: time pressure, contact with others, establishing interpersonal relationships, structured versus unstructured work, and freedom to make decisions[1].

Goldin describes flexibility as "a multitude of temporal matters including the number of hours, precise times, predictability and ability to schedule one's own hours." While the selected occupation characteristics succeed in describing structural features of jobs that would affect this notion of flexibility, the study of workplace flexibility would also benefit from data on the degree to which workers have access to flexible work arrangements. In this paper I address this need using worker-level data on access to flexibility policies from the American Time Use Survey. Using these data I create measures for four key flexibility policies: the ability to choose one's own hours, the ability to work from home, access to paid leave, and schedule predictability. These four measures each relate to different aspects of flexibility, and each provide workers with means to better control their work-life balance.

I conduct an exploratory analysis of the relationships between the selected O*NET characteristics, my flexibility policy measures, and "greedy" jobs. I begin by using the 2018 American Community Survey to estimate Goldin's occupation-level measure of job greedi-

[1] Some papers that use these O*NET characteristics as proxies for flexibility include Bang (2022), Lordan and Pischke (2022), Benny et al. (2021), Sobeck (2022), Schaffer and Westenberg (2019)

ness, the elasticity between earnings and weekly labor hours. These estimates are merged with the 2018 ATUS as well as a set of 100 occupation characteristic variables from the August 2018 O*NET survey. With this merged data I run logistic regressions of the selected O*NET variables and the greediness measure on each of the four flexibility policies. I find that the ability to choose when and where one works is positively correlated with greediness. I also find that contrary to Goldin's prediction of a negative relationship between flexibility and the five selected O*NET characteristics, there is a high degree of heterogeneity in the relationships between the O*NET variables and the four flexibility policies.

I devote the remainder of this paper to the development of a new set of occupation-level flexibility policy indices for future use by researchers. It is best to create occupation-level measures because they will be able to be linked to other datasets to allow researchers to use them for a variety of applications. The main challenge is how to overcome the small sample size of the ATUS, which has only about 6400 respondents from 332 of 530 Census occupations. My approach to this task involves the use of machine learning. First, I train machine learning models to predict the probability of a given worker having flexibility policies. The models are trained on a comprehensive set of individual-level variables including worker demographics such as age, sex and education, as well as job variables such as weekly labor hours, hourly wage, and salaried/hourly status. Principal component analysis is used to reduce the dimensionality of the set of 100 O*NET occupation characteristics to prevent overfitting. I use the XGBoost algorithm, which in my benchmarking tests had the best baseline predictive performance while also being one of the fastest models to execute. Since the provision of flexibility policies depends on the interactions between worker and occupation characteristics, the XGBoost algorithm's ability to flexibly model nonlinear relationships between variables makes it a powerful tool for this exercise. The models' hyperparameters and number of included O*NET principal components are tuned to maximize performance.

Next, I use the trained models to generate predicted flexibility probabilities for respondents to the 2018 ACS. I then compute the mean of these predicted probabilities for each Census and SOC occupation code to create occupation-level "flexibility propensity" measures. The interpretation of each index for a specific occupation will then be the predicted probability of an average worker in that occupation having each of the four flexibility policies.

The indices are available for 471 2018 SOC codes and 481 2018 Census occupation codes and can be found as .csv files on my personal website. The indices are highly correlated with the occupation-average provision of actual flexibility policies, and explain a significantly greater share of the variance in flexibility policies than regressions with the entire set of O*NET characteristics. Overall my indices provide a quality occupation-level measure of workers' access to flexibility policies, a complement to Goldin's approach of using O*NET occupation characteristics to measure the structural flexibility of jobs.

Section 2 describes the data used for my empirical analysis. In Section 3 I test the predictions from Goldin (2014) about the relationships between "greedy jobs", O*NET characteristics and flexibility, as well as test the value of the O*NET characteristics as proxies for flexibility. Section 4 discusses my approach to the development of new flexibility measures and describes the variables used for the prediction. Section 5 details my modeling choices for the machine learning models and reports the results of a feature importance analysis of the trained models. In Section 6 I discuss the flexibility indices, provide summary statistics and tests of their predictive ability, and list potential future improvements. Section 7 concludes.

3.2 Data

3.2.1 ATUS Leave Module

Microdata on workers' workplace flexibility policies are sourced from the American Time Use Survey's 2017 and 2018 Leave and Job Flexibilities Module, distributed by IPUMS-ATUS. The module is administered to about 10,000 respondents from the 2017 and 2018 ATUS surveys, about half of the total ATUS participants from each year. One limitation of this survey is that it is not given to self-employed workers. Thus the self-employed, who constituted about 10 percent of the labor force in 2018, are excluded from all analysis in this paper.

I use four variables that capture differing but important aspects of workplace flexibility. The first two are the ones used in Lann (2023): "flextime", which captures temporal flexibility, the ability to choose when one works, and "paid work-from-home", which measures locational flexibility, the ability to choose where one works. The flextime variable is a dummy where workers who report that they "can frequently choose their own hours." are coded as a 1. While respondents can report several levels of control over one's schedule, only the highest level of control, the response quoted above, is used to denote a temporally flexible job.

The paid work-from-home variable is based on a series of questions in which respondents are asked whether they can work from home, whether they ever work from home, and finally whether they are paid to work from home rather than simply taking work home with them. Respondents who report being paid for the work they do from home are coded as a 1. One issue with this measure pointed out in Lann (2023) is that this measure codes as a 0 workers whose employers may offer paid work-from-home if they themselves do not use it. Since the question of paid versus unpaid work-from-home is only asked to workers who report working from home, this group of paid work-from-home jobs cannot be identified.

The third flexibility variable is "paid leave", which is a dummy that is coded as a 1 if the respondent reports having paid leave at their primary job. The ATUS Leave Module also asks several follow-up questions about the availability of paid leave for personal or family illness, vacation, personal reasons, and for childcare or birth/adoption of a child. Nearly two-thirds of respondents that report having any paid leave are able to use it for at least five of the six scenarios listed above, so for simplicity I stick to the use of the dummy for any paid leave. The last flexibility variable is "predictability", which measures how far in advance a worker knows their work schedule. This is an important aspect of flexibility, as advance knowledge of one's work schedule allows for better planning and response to unforeseen circumstances, and improves work-life balance. The Leave Module asks workers how far in advance they know their work schedules, with responses ranging from less than one week to four or more weeks in advance. The predictability variable is a dummy with a 1 coded as knowing one's schedule four or more weeks in advance.

Table 37 shows the prevalence of flexibility policies by Census occupation category. Flextime and paid work-from-home are scarce overall, while there is broader access to paid leave and predictability. There is significant variation in the provision of flexibility policies across occupations, with paid work-from-home having the greatest across-occupation variance. Flexibility policies are provided most frequently in "white collar" occupations such as management, business and finance, legal, engineering and tech. Even in these types of white-collar jobs, flextime and paid work-from-home are not ubiquitous.

Like in Lann (2023), the ATUS Leave Module is linked to the main survey which contains a number of worker demographic and job characteristic variables. Also following my previous paper, the sample is restricted to workers aged 25 to 65, with wages trimmed at the 1 percent level.

3.2.2 O*NET

Following Goldin (2014), I utilize occupation characteristic data from the U.S. Department of Labor's Occupational Information Network (O*NET). I use the 23.0 release from August 2018 in order to align the occupation data with the flexibility data from the ATUS Leave Module.

The O*NET data includes 100 variables from three broad categories: work context, work activities, and education, training and experience. Work context variables describe the "physical and social factors that influence the nature of work." Included in this category are questions on required methods of communication and interpersonal interaction, physical work conditions, types of body positions (e.g. amount of sitting versus standing), and structural characteristics such as competitiveness, pace and level of responsibility. Work activities variables measure the degree to which different jobs require particular tasks. The list of tasks is comprehensive, including both physical actions and mental processes, as well as interpersonal and communication tasks and tasks involving the collection and use of data. Each task has two survey questions, one asking about the *importance* of the task for performing the job, and the other asking about what *level* of the task is required for performing the job. Schaffer and Westenberg (2019) show that these variables are highly correlated, and chose to simply add them together to use as a composite measure. I instead choose to only include one set of variables, the *importance* measures. I chose this variable due to concerns over the subjectivity of the *level* question.

The education, training and experience variables used are as follows: required level of education, related work experience, on-site or on-plant training, and on-the-job training.

Each variable is standardized to have mean zero and standard deviation one, and the 966 O*NET occupations are mapped to 527 Census occupations via a series of crosswalks. The O*NET occupation data is linked to the ATUS data at the occupation level, with 6377 observations from 332 Census occupations.

3.2.3 Job Greediness Data

In order to assess the relationship between flexibility and job greediness, I compute the same occupation-specific earnings/hours elasticity as in Goldin (2014). A higher elasticity denotes a stronger relationship between earnings and hours worked, and thus a "greedier" occupation. While Goldin used the 2009 to 2011 American Community Survey, I use the 2018 ACS. Like Goldin (2014), my sample is restricted to workers aged 25 to 64 working at least 35 hours per week, and trims workers whose annual earnings are less than 1400 hours $\times$ 0.5 $\times$ \$7.25, the federal minimum wage in 2018.

The earnings/hours elasticities are computed by estimating the following model:

$$\begin{aligned}\log(\text{earnings})_i &= \alpha + \Lambda X_i + \beta_0 \text{Female}_i + \beta_i \log(\text{Hours})_i + \Gamma \text{Occ}_i \\ &\quad + \Delta \text{Female}_i \times \text{Occ}_i + \Omega \log(\text{Hours})_i \times \text{Occ}_i + \epsilon_i\end{aligned}$$

Log annual earnings are regressed on a vector of controls X_i including an age quartic, education, race, and log weeks worked. Also included are sex, occupation fixed effects, log weekly hours worked, and the interactions between occupation and sex as well as between occupation and log hours. The occupation-specific elasticities are given by the coefficients on the occupation/log hours interaction variables. The coefficients on the occupation/sex interactions are interpreted as occupation-specific gender wage gaps.

Table 38 compares the elasticities estimated in this paper to the ones from Goldin (2014), averaged by broad occupation category. The estimates for some occupation categories are quite similar, but many categories have very different estimates. There are several explanations for this. The first is that there may have been structural changes to occupations from

2009-2011 to 2018 that may cause the greediness of various occupations to change. Second, Goldin only computes these elasticities using college graduates, while I include observations from every education level. Finally, there have been changes to the Census occupation classification since Goldin's analysis, with my data including more specific occupation categories and more occupations in each broad category.

3.3 O*NET characteristics and flexibility policies

Goldin (2014)'s analysis of the relationship between job flexibility, greediness and the gender wage gap hinges on the use of five selected O*NET characteristics that are meant to caputre structural aspects of jobs related to flexibility. The selected O*NET characteristics are as follows:

1. Time pressure: *"How often does this job require the worker to meet strict deadlines?"*
2. Contact with others: *"How much does this job require the worker to be in contact with others (face-to-face, by telephone, or otherwise) in order to perform it?"*
3. Establishing and maintaining interpersonal relationships: *"Developing constructive and cooperative working relationships with others, and maintaining them over time."*
4. Structured versus unstructured work: *"To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?"*
5. Freedom to make decisions: *"How much decision making freedom, without supervision, does the job offer?"*

Goldin (2014) posits a positive relationship between each of these characteristics and job greediness, and a negative relationship between the characteristics and flexibility. I will begin by testing the relationships between the O*NET characteristics and the earnings/hours elasticity. Next, I will use the data on flexibility policies to test the hypothesized negative relationship between the flexibility policies, the O*NET characteristics and greediness.

3.3.1 Job greediness and occupation characteristics

I explore the relationship between the selected O*NET characteristics and job greediness by regressing the earnings/hours elasticity on the characteristics at the occupation level. Table 39 shows the results of this exercise, both with the characteristics in separate regressions and then together. Looking first at the coefficients from the separate regressions, we can see that only two of the five O*NET characteristics have a significant relationship to the earnings/hours elasticity. "Time pressure" is found to have a positive relationship with greediness, but "freedom to make decisions" is found to have a negative relationship with greediness, the opposite of the predicted sign. The regression with the characteristics included together shares the same results, and the adjusted R-squared of 0.015 shows that the O*NET characteristics struggle to explain the across-occupation variance in the earnings/hours elasticity.

3.3.2 Flexibility and occupation characteristics

In order to assess the relationship between the selected O*NET characteristics and the four flexibility policies, I use a logistic regression framework:

$$\text{logit}(p_{\text{Flex}_i}) = \alpha + \Lambda X_i + \Gamma \text{Occ\_chars}_i + \epsilon_i \tag{36}$$

The dependent variables in this set of regressions are the log odds of worker i having one of the flexibility policies. X_i is a vector of individual and job characteristics, including age, sex, years of education, race, marital status and number of children, full versus part-time status, and whether they are paid hourly versus salaried. Occ_chars_i are the O*NET characteristics and earnings/hours elasticity of worker i's occupation. The O*NET characteristics and the earnings/hours elasticity are included in the same regression because of the finding in the previous section that they are largely uncorrelated. There is little concern for colinearity between these variables.

Table 40 presents the results for all occupations. "Time pressure" is found to not have a significant relationship with flextime and paid work-from-home, but has a positive relationship with paid leave and predictability. As predicted, there is a negative relationship between "contact with others" and three of the four flexibility policies. Contrary to the predicted sign, "interpersonal relationships" are found to have a positive relationship with paid work-from-home and paid leave. "Structured work" is found to be negatively related to flexibility overall, and "freedom to make decisions" has a positive relationship to flextime and a negative relationship with paid work-from-home and predictability.

The most significant result is the strong positive relationship estimated between both flextime and paid work-from-home and the earnings/hours elasticity. Lann (2023) provides evidence that across-occupation differences in the provision of flexibility are largely due to differences in the production technologies. This may provide an explanation for the observed positive relationship: the production technologies of occupations that allow for a higher provision of these policies may also feature a convex output/labor hours relationship, resulting in jobs that are both temporally and locationally flexible and "greedy". On the other hand there is a significant negative relationship between predictability and the earnings/hours elasticity, showing that greedier jobs will be harder to plan for and will have worse work-life balance. Contrary to the intuition provided by Goldin that jobs that are more demanding of workers' time will be less flexible overall, the results show that greedy jobs will allow more flexibility in some areas and less in others.

As shown in Table 37, white-collar occupation categories such as management, legal and STEM occupations are much more likely to offer flexibility policies compared to other types of occupations. It is likely then that the significant results in Table 40 reflect the fact that there are large differences in the O*NET characteristics of white-collar occupations and other occupations. Table 41 shows the results of the logistic regressions described by Equation 1 for white-collar workers only. Looking within white-collar occupations eliminates many potential differences in the jobs' production technologies, and this reduction of variance causes most of the significant relationships from Table 40 to disappear for every policy but paid leave. There is still a significant positive relationship between the earnings/hours elasticity and paid work-from-home for white-collar workers. Access to paid leave in white-collar occupations is found to be more strongly related to the O*NET characteristics than for all occupations, with the most surprising result being the strong negative relationship between paid leave and greediness that was not present in Table 40.

3.3.3 O*NET characteristics as predictors of flexibility policies

Finding useful measures of job flexibility are important because of the increasing interest in the study of flexibility since the pandemic. The ATUS Leave Module is the only source of current and publicly-available information on workers' access to flexibility policies in the United States, but the small sample size makes this data difficult to apply to other settings. It is thus desirable to find the best possible proxy measures of workplace flexibility policies that can be constructed using commonly-available data sources. While the previous results show that the selected O*NET characteristics have significant relationships with the flexibility policies, it remains to be seen whether these characteristics, or O*NET characteristics more broadly, are the best proxies available for workplace flexibility policies.

In order to find the combination of variables that are most predictive of flexibility policies, I perform a simple exercise in which I regress each policy on different combinations of individual-level variables and O*NET characteristics. I estimate linear probability models rather than logistic regressions for easier interpretation of the adjusted r-squared compared to a pseudo r-squared. Table 42 shows the results of this exercise. While the selected O*NET characteristics are able to explain a significant portion of the variance of several of the flexibility policies, the addition of the other 95 O*NET variables produces a model with much greater predictive ability. The results also show the importance of individual-level variables, with the inclusion of worker variables significantly increasing the adjusted r-squared's of the models for any combination of O*NET variables. This simple analysis understates the importance of the inclusion of both individual and O*NET variables, as these are linear models with no interactions.

The results of this exercise indicate that in order to predict workplace flexibility policies, it is important to consider both worker variables and the universe of occupation characteristic data. Lann (2023) shows that the occupation-level average access to flexibility policies is an equilibrium outcome of the labor market that depends on the structural characteristics of each occupation as well as the qualities of the workers in those occupations. Thus predicting access to flexibility policies requires not only a comprehensive set of variables, but also a model that considers the interactions between these variables.

3.4 Predicting workplace flexibility

The remainder of the paper is spend developing occupation-level flexibility policy indices for future use by researchers. My approach to the development of these indices involves the use of machine learning. First, I will use a comprehensive set of worker and O*NET job characteristics to train a machine learning model to predict the probability of a given worker having flexibility policies. Next, I will use the trained model to generate predicted flexibility policy probabilities for respondents to the 2018 ACS. Finally, I will compute the mean of these predicted probabilities for each 4-digit Census occupation code to create occupation-level "flexibility propensity" measures. The interpretation of the index for a specific occupation will then be the predicted probability of an average worker in that occupation having a given flexibility policy.

Another possible approach to developing these measures would be to to simply take the occupation-level mean of the flexibility variables in the ATUS. This method has several issues stemming from the small sample size of the ATUS Leave Module. First is that not every occupation is represented in the ATUS, giving us a measure that is only available for 332 out of 530 Census occupations. The other issue is that many occupations have vary few respondents, which will lead to an extremely noisy measure of flexibility. By leveraging the huge sample size of the ACS, my measure will be available for more Census and SOC occupations and with very narrow confidence intervals.

In this section I list and explore the predictors used in the machine learning exercise. Since I am using two different datasets (ATUS/ONET and the ACS) for this exercise, I need to be careful to only select predictors that can be found in some form in both datasets. There are some variables such as hourly/salaried status that are informative of access to the policies but are not available at the individual level in the ACS. For those variables I compute the mean at the occupation level in another representative survey and merge those averages with the ACS. Thus several "worker/job" variables included in the machine learning prediction are actually occupation-level variables.

3.4.1 Assessing the predictors of workplace flexibility: Worker characteristics

In order to develop the most accurate possible prediction of access to flexibility policies, it is not sufficient to include only occupation-level job characteristics. I also include a variety of worker variables, ranging from worker demographics to aspects of their jobs such as hourly wages. Using these individual-level variables alongside the occupation-level O*NET variables will also allow the model to take into account the interactions between worker and occupation characteristics. Across-occupation differences in worker characteristics may affect the aggregate across-occupation distribution of flexibility policies, and it is possible that the relationships between worker characteristics and job flexibility vary by occupation types. The machine learning model will be able to take these interactions into account, resulting in more accurate predictions.

The worker demographic variables used are sex, age and age squared, a dummy indicating whether the worker is white, and number of years of completed education. Two household composition variables, marital status and number of own children under age 18 in the home, are also included. Another characteristic that may be informative of access to flexibility policies is whether the worker lives in an urban or rural area. In order to measure this I created a dummy called "Large MSA" using the "msasize" variable from the ATUS. Respondents who report living in a metropolitan statistical area with a population of over 1,000,000 are given a 1 for the Large MSA variable, and a 0 otherwise.

A number of variables describing workers' jobs are also included. Lann (2023) demonstrates a positive relationship between hourly wages and weekly labor hours and access to flexibility policies, which holds both across and within occupation. I thus include hourly wage, which is computed by dividing respondents' reported yearly labor income by their usual hours worked per week times the number of weeks they worked in the last year. I also include the usual number of hours worked per week. Lann 2023 also finds that firm size is also positively related to the provision of flexibility policies. I use the same "large firm" variable as in Lann 2023, which is a 1 if the respondent's employer has over 1000 workers. Since this variable is not available in the ACS, I use the occupation-level means of the large firm variable for respondents to the Annual Social and Economic Supplement (ASEC) from

2015 to 2020.

Following the analysis from Section 3, I include the standardized earnings/hours elasticity for the workers' occupation. I also include hourly versus salaried status in the form of a dummy for which a 1 indicates being salaried. Since the ACS does not contain information about hourly versus salaried status, I use the occupation-level means from the 2017 and 2018 Current Population Survey (CPS). Another variable of interest is the broad industry category of the respondent's job. Finally, I include whether the worker has a public or private-sector employer, which is included as a dummy derived from the "worker class" variable from the ATUS and ACS.

Table 43 shows the difference in the means of the observables for workers with each flexibility policy compared to those without. For each flexibility policy, workers have significantly more education than workers without the policy. There is little difference in the ages of workers with or without flexibility, with workers with paid leave being slightly older on average. Flextime is found to be more common among male workers, and workers with paid work-from-home or predictability are more likely to be female than those without those policies. Workers with any policy are also more likely to be white. Workers with any flexibility policy are more likely to be married, and those with flextime or paid work-from-home have more children and more likely to live in a large MSA on average.

Panel C shows how various job characteristics are related to the presence of flexibility policies. Jobs with any flexibility policy pay a significantly higher hourly wage, and workers with flextime, paid work-from-home or paid leave work more hours per week compared to workers without those policies. A majority of jobs with flexibility are salaried, while most jobs without flexibility have hourly pay[2]. Private sector employers are found to be more likely to provide flextime and paid work-from-home compared to government employers, but government employees have more paid leave and predictability. Jobs with flextime, paid work-from-home or paid leave are more greedy and are more likely to be for a large firm with at least 1000 employees. Conversely, jobs with predictability are less greedy and less likely to be for a large firm.

[2]This is true when looking at either the salaried/hourly status of individual respondents or at the across-occupation mean.

Table 44 shows the provision of flexibility policies by industry category. While the variance in flexibility across industry categories is smaller than the across-occupation variance, some industries have significantly higher rates of flexibility. While some of these industries, such as financial activities and professional and business services, would be expected to have high rates of flexibility, others like manufacturing and mining, quarrying, and oil and gas extraction have higher provisions of flexibility policies than one might predict based on the low rates of flexibility in production occupations. This shows that the inclusion of the industry categories will provide the model with more information about the provision of flexibility that is not already captured by the workers' occupation category.

In addition to understanding the relationship between access to flexibility and the various worker and job variables, any predictive model will need to consider the relationship between the predictors. Including several highly-correlated variables can lead to overfitting and a loss of out-of-sample predictive accuracy. The correlations between the worker variables are shown in Table 45. The strongest correlations are between education and salaried/hourly status, education and wage, and salaried/hourly status and wage. Besides these, the correlations between most variables are relatively small, alleviating concerns about colinearity between the worker variables.

3.4.2 Assessing the predictors of workplace flexibility: O*NET characteristics

Using the detailed occupation characteristics from O*NET as predictors will improve the predictive ability of the model, but there are potential pitfalls to consider. It may not be desirable to use all 100 available O*NET variables due to concerns about overfitting. This is of particular concern in this setting, as the sample being used to train the model is relatively small at only 6377 observations. In addition, many of the characteristics are highly correlated, increasing the risk of overfitting. In order to address these concerns I use principal component analysis to reduce the dimensionality of the occupation characteristic data. Principal component analysis transforms the data into a set of 100 linear combinations (principal components) of the O*NET characteristics such that each principal component is orthogonal to the rest, and the first few components explain as much of the variance in the data as possible. By replacing the O*NET variables with the first few principal components I can avoid the pitfalls of using the full set of 100 variables while retaining most of the information.

The first 10 principal components explain about 76% of the variance in the O*NET characteristics. Table 46 shows the most important O*NET characteristics for the first 10 principal components. A brief qualitative examination of the key characteristics for each component show that they each characterize a different "type" of job. For example, the largest positive coefficients for Component 1 are for job characteristics such as use of e-mail, working with computers and letters and memos that describe "office" type jobs, while the Component 2 describes "routine" jobs that are highly structured with repetitive tasks.

In order to assess the relationship between the principal components and the flexibility policies, I regress each policy on the first 10 components. Results are displayed in Table 47. The principal components are found to be mostly significantly related to flexibility, with the signs being in directions we would intuitively expect. For instance, Component 1 has a positive relationship with each flexibility policy, which aligns with previous observations that white-collar office jobs are more likely to provide flexibility. On the other hand, we know that there are relatively low rates of flexibility in health care occupations, which explains the negative coefficient on Component 3. Overall, the principal components seem to succeed in conveying useful information from the original O*NET characteristics and should serve as good predictors of flexibility while greatly reducing the dimensionality of the model.

3.5 Machine learning

3.5.1 Model design

Since the flexibility policy measures from the ATUS are categorical variables, I will use a classifier to predict the flexibility probabilities. There are several important modeling choices that need to be made. The first is how to deal with the four distinct flexibility measures. One option is to use four separate classifiers that each predict one of the binary variables.

The other option is to combine the four dummies into one factor variable, then use a single "multiclass" classifier. While the multiclass approach would allow me to predict the probability of a worker having any combination of flexibility policies, this approach will suffer from the small sample size of the training data. There are 16 possible combinations of flexibility policies and only 6377 observations, leading to some of those combinations having very few observations. This will greatly diminish the predictive power of the multiclass model, an issue that the binary models will not have. Thus I choose to train four binary classification models to predict each of the four flexibility policy variables separately, prioritizing predictive accuracy over having a more detailed set of prediction measures.

The next important choice is which classifier to use. In order to determine which type of model I will use, I compare the predictive performance of a number of models along with the average computing time of each model. I use several measures of model performance in order to evaluate the relative strengths and weaknesses of each model. The first measure is the accuracy score, which simply measures the percentage of the test sample that was correctly predicted. While this is an important indicator of model performance, it does not tell the whole story. It is possible for a given model to have a high accuracy while not doing a good job of distinguishing between the classes. By default, classifiers will choose the class that has the highest predicted probability, so it is possible for a given model to have a high accuracy while predicting close to a 50/50 probability of flexibility. Since I care primarily about the accuracy of the predicted probabilities, I need other measures that will capture other aspects of the models' performance.

The next performance metric is the AUROC (Area Under the Receiver Operating Characteristics curve) score, which measures the degree of separability of the classes. The higher the AUROC score, the better the model is at distinguishing between the classes. The last two performance metrics, the Brier score and the log loss, are both measures of the accuracy of the predicted probabilities. The Brier score measures the difference between the predicted probability and the true class, while the log loss measures the difference between the predicted and true probability of the dependent variable. The smaller these values, the more accurate the predicted probabilities are.

In order to measure the performance of each model, I use a baseline dataset containing the individual variables as well as the 100 O*NET variables. I use cross-validation to compute the metrics for each model and flexibility policy. Table 48 shows the performance of each model, with the scores and times for each of the four outcome variables averaged together. The highest-performing models are the random forest, and gradient-boosting classifier and the XGBoost classifier. These tree-based algorithms have the advantage of being able to flexibly model nonlinear relationships between the variables. Of these, the clear best algorithms in terms of all-around performance are the gradient boosting classifier and the XGBoost classifier. These two algorithms have nearly identical scores, but XGBoost is 20 times faster to evaluate. I will thus use the binary XGBoost classifier to predict the probabilities of the four separate flexibility policies.

3.5.2 Model tuning

The next step is to tune the model to improve the predictive performance. I will be tuning the number of included features as well as the hyperparameters of the XGBoost algorithm. Hyperparameters are settings of a machine learning model that affect how a given model trains itself. For the XGBoost algorithm, which is an ensemble method that aggregates the results of a series of sequentially-built decision tree classifiers, there are a number of important parameters. Some are related to the individual trees, such as the maximum depth or the degree of random sampling of observations and features used for each tree. Others affect the overall construction and aggregation of the model, such as the number of trees to build and the learning rate. For each dataset and use case there will be combinations of these and other parameters that will improve performance over the default settings, with these being commonly found through a search over a grid of chosen values for each parameter.

It is also important to choose the number of principal components to include. As shown in Section 4, while the first few principal components explain most of the variance of the O*Net characteristics, several of the later components are more strongly correlated with flexibility than the earlier components. Despite this, it may not be best to include a large number of principal components due to the potential for overfitting. It may even be possible that using zero components and instead using the original set of O*NET variables may produce better results for a given model under a particular parameterization. In order to determine the best combination of features and parameters I use a simultaneous tuning approach: the number of principal components and hyperparameters will be tuned at the same time via a grid search. The case where the O*NET variables are used is also checked.

Table 50 shows the average performance of the models with a number of feature sets and parameterizations, with the last row showing the performance of the tuned models. The inclusion of both individual and occupation variables only gives a slight improvement in performance over using just one set of variables, and model performance is nearly identical when using either the original O*NET variables or the first 10 principal components. Tuning the models' hyperparameters and feature sets gives a similar improvement to performance as the inclusion of ONET variables improves performance, an improvement of about one percent across the four metrics.

The final step is to choose which set of tuned models to use to predict the final flexibility probabilities. I determine this qualitatively by comparing moments from the data to moments generated from the predicted flexibility probabilities. Table 49 shows the results of this exercise. Every flexibility policy is underestimated for each set of tuned models. This is largely due to differences between the ATUS and the ACS in the distribution of occupations: the ACS has a lower proportion of workers in white-collar occupations than the ATUS, with only about 21% of workers in these occupations compared to 27% in the ATUS. The tuned models that minimize the Brier score and log loss are found to greatly underestimate the probability of flextime and paid work-from-home policies, as well as generate a much higher across-occupation variance in the probability of paid leave and predictability compared to the data. For Panels B and C of Table 49, the predicted across-occupation variances of each of the flexibility policies are lower than in the data. This is expected, as 179 of the 332 occupations in the ATUS data have ten or less observations, resulting in many occupations having no one with flexibility or a very high percentage of workers with flexibility, and producing a high across-occupation variance. The median occupation in the ACS data has about 15000 observations, leading to estimates of the probability of flexibility policies that are less extreme than in the ATUS.

The results of the models that are tuned to maximize accuracy and the AUROC have very similar moments and it is difficult to determine which set of tuned models to choose. I thus include a fifth option, a set of models that are tuned to maximize a simple composite measure of model performance: (Accuracy + AUROC) / (Brier + Log loss). Panel F displays the moments for this set of models. The mean probabilities of the flexibility policies are all slightly smaller than the estimates from Panel C. The across-occupation variances of flextime and predictability are both significantly higher than for either Panels B or C, bringing these moments a bit closer in line to the data than the accuracy or AUROC-tuned models. I will thus use the tuned models with the highest composite measure for my final predictions.

3.5.3 Feature importance

While machine learning models are often less interpretable than traditional econometric models, we can still learn about the impact of each variable on the predictive performance of a model using feature importance analysis. I use a permutation importance approach to investigate which features are most important for each of the four final tuned models. Permutation importance involves randomly shuffling the values for a given feature, and computing the change in performance on a test dataset. This is done for each feature one at a time, giving a measure of the importance of each feature. This method has the advantage of being model-agnostic, meaning it can be used on any kind of model.

I compute the cross-validated permutation importances using both the accuracy and AUROC scores, then average those scores for each model. Table 51 lists the ten most important features for the four models. Looking first at the individual variables, we see that weekly hours worked and hourly wage appear in the top ten for all four flexibility policies, while the earnings/hours elasticity and private/public employer appear in three of the four. Salaried/hourly status and education are both in the top ten for flextime and paid work-from-home, and female is found to be an important of both paid work-from-home and predictability. Overall, these results show that the job variables are the most impactful individual characteristics, with most of the worker demographics and household variables not being important determinants of possession of flexibility policies.

Three of the four models were tuned to use principal components. For all three of these, we find that principal components 1, 3 and 4 are in the top ten most important features. Component 1 depends most on occupation characteristics most common in office-type jobs, which are shown to have high rates of flexibility. Component 3 describes "healthcare" jobs that a have low rates of flexibility, and component 4 relates to both being outdoors and interpersonal communication. Interestingly, component 2, which explains over 16% of the variance in the O*NET variables and describes "routine" jobs that have low rates of flexibility, is not found to be one of the most important features for predicting flexibility. The paid leave model is tuned to use the O*NET variables, and features a number of characteristics in the top ten that describe white-collar jobs, such as 'interacting with computers, level of competition, and freedom to make decisions. The results of this exercise show that both individual variables and occupation characteristics are important for the prediction of the flexibility policies.

3.6 The flexibility policy indices

The flexibility policy indices can be found in the Data section of my personal website HERE. This section contains two .csv files with the indices for each of the four flexibility policies, averaged by either 2018 Census occupation codes or 2018 SOC codes. flex_indices_occ.csv has the indices for the Census codes, while flex_indices_soc.csv has the indices averaged over the SOC codes. The indices are available for 471 SOC codes and 481 2018 Census occupation codes. They can also be used for other sets of occupation codes via crosswalking. In order to use the indices, merge these data to other datasets using the occupation code variables.

3.6.1 Summary statistics and index performance

In order to learn more about the features of the indices as well as assess their performance in measuring flexibility policies, I present a number of summary statistics and statistical tests. Table 52 reports number of correlations between the flexibility policies and indices. Panel A shows that each index is at least moderately correlated with their respective policy. Paid work-from-home and paid leave are more highly correlated with their indices than flextime and predictability. This continues a trend from Section 3, where every set of regressions explained more of the variance in paid work-from-home and paid leave compared to flextime and predictability. This suggests that there are variables missing from my data that may be particularly informative of access to these policies.

Panels B and C display the correlations between the policies and indices, respectively. The indices succeed in matching the relative size of the correlations: for instance, the largest correlation is between flextime and paid work-from-home, and the smallest is between flextime and predictability in both the data and in the indices. There is a systematically higher correlation between the indices compared to between the actual policies. This is most likely due to the fact that the flexibility probabilities were all predicted using the same set of variables and methods. Researchers should take these large correlations into consideration when designing regressions using these indices in order to avoid multicollinearity.

Next I examine the distribution of the indices across occupation categories. This is shown in Table 53, which is set up like Table 37 in order to facilitate comparison. The category-level averages of the indices are mostly within a few percentage points of the category-level averages of the actual policies, showing that the indices succeed in differentiating between occupations that have high versus low provisions of flexibility policies. One major difference is that the flexibility and paid work-from-home indices are much higher for computer and mathematical science occupations compared to the actual levels of these policies in the data. This may be due to the small sample size in the data, with only 22 observations with this type of occupation. One would intuitively expect this category to have a relatively high probability of these policies, given the high rates of flexibility for other types of similar occupations such as life, physical, and social science and engineering occupations. The indices report the high probability of flextime and paid work-from-home that one would expect from these occupations, showing that these measures are an improvement over relying on the ATUS data which suffers from a small sample size.

Last, I assess the quality of the indices as measures of access to flexibility policies by regressing the occupation-level mean of each policy on various sets of variables as well as the indices. The adjusted r-squared's for each set of regressions are reported in Table 54.

Regressions using only the indices as independent variables have significantly higher adjusted r-squared's than regressions with the entire set of O*NET characteristics and occupation-averaged individual variables. The indices explain a relatively small amount of the variance of flextime and predictability, a result that mirrors those of Table 52.A and Section 3. However, the indices explain 58 percent of the variance in paid work-from-home and 49 percent of the variance in paid leave. Overall these results indicate that the flexibility policy indices are significantly more predictive of access to flexibility policies than a linear model containing the universe of O*NET and individual variables.

3.6.2 Limitations and future improvements

Despite the success of the indices as a quality proxy for flexibility policies, there are still several possible areas for improvement. The first is that the current indices are estimated using data on flexibility and individual and occupation characteristics from 2018. This means that the measures will be less reliable the further away from 2018 one tries to use them due to structural changes in the characteristics of occupations and the determinants of of the provision of flexibility over time. Future new data on access to flexibility may allow for the training of updated models to reflect the post-pandemic changes to the allocation of flexibility.

The indices are currently only available for occupations that had a corresponding SOC occupation covered by the August 2018 ONET. As such, only 471 of the 530 SOC occupations and 481 of the 530 Census occupations found in the 2018 ACS have index values assigned. Using updated SOC codes from November 2020 release of O*NET may allow me to add more occupations and more observations for the training data. This may be a viable option as long as there weren't large changes to the occupation characteristics between these waves.

The next potential area for improvement involves a reconsideration of the definitions of my flexibility policy variables. In this paper and Lann (2023) I used somewhat strict definitions of flexibility in order to represent each policy as a binary variable. The variables are largely based on Leave Module questions that record responses on Likert scales or in other ways that measure the intensity or extensiveness of each policy. It would be possible to create new continuous variables and approach the flexibility indices as a measure of the "quality" or "intensity" of flexibility policies rather than the probability of a worker having a particular policy. This approach may result in a more informative measure, or may be complementary to the current indices.

Last, the feature importance analysis revealed that several included individual and occupation variables were not actually important for the models. A second round of feature selection could marginally improve model performance by reducing potential overfitting. This may also reduce the correlation between the indices as they may each be predicted using slightly different sets of features.

3.7 Conclusion

In this paper I contribute to the literature on workplace flexibility by introducing new measures of workplace flexibility. Using the ATUS Leave and Job Flexibilities Module, I develop four variables to measure various aspects of workplace flexibility: "flextime," the ability to choose one's hours, "paid work-from-home," the ability to choose where one works, "paid leave," and "predictability," the extent to which workers know their schedule in advance. These four variables complement Goldin (2014)'s use of O*NET characteristics to measure the structural flexibilities of jobs by providing additional information on workers' access to flexible work arrangements. An exploratory analysis shows that there are several significant relationships between the four flexibility policies, the selected O*NET characteristics, and job greediness, with some being in different directions than predicted by Goldin.

I spend most of the paper developing occupation-level measures of access to the flexibility policies. I employ a machine learning method to generate predicted probabilities of flexibility policies for workers in the ACS, then take the occupation-level means, resulting in an occupation-level "flexibility index" for each policy. This allows me to overcome the issue of the small sample size of the ATUS data. Another advantage of the machine learning approach is that the XGBoost algorithm I used can flexibly model the interactions between worker, employer and occupation characteristics, which Lann (2023) showed were important for the provision of flexibility.

The indices are available for use by researchers for future studies of workplace flexibility, and there are currently measures for 471 2018 SOC codes and 481 2018 Census occupation codes. There are several avenues for future improvement of the indices. By using the updated November 2020 O*NET release I could increase the size of the training data as well as the number of occupations that index covers. Future data such as a new ATUS Leave Module would allow me to update the indices to reflect the post-pandemic changes to the provision of flexibility. The indices will continue to be updated and improved for the foreseeable future.

3.8 Tables

Table 37: Flexibility policies by occupation category

Occupation category	Flextime	Paid WFH	Paid leave	Predictability	N
Management	0.3315	0.4257	0.8647	0.5920	902
Business and financial operation	0.3486	0.5107	0.8165	0.6208	327
Computer and mathematical science	0.3182	0.2727	0.8636	0.7727	22
Architecture and engineering	0.4253	0.4751	0.8736	0.6015	261
Life, physical, and social science	0.3441	0.3226	0.9032	0.6882	93
Community and social service	0.2649	0.3297	0.8649	0.7135	185
Legal	0.3846	0.4615	0.8803	0.6667	117
Education, training, and library	0.1424	0.1634	0.7065	0.8383	569
Arts, design, entertainment, sports, and media	0.3077	0.4530	0.7607	0.5470	117
Healthcare practitioner and technical	0.1381	0.1004	0.8180	0.6946	478
Healthcare support	0.0800	0.0533	0.6400	0.6000	75
Protective service	0.1064	0.0993	0.8582	0.6525	141
Food preparation and serving	0.1316	0.0219	0.4386	0.2675	228
Building and grounds cleaning and maintenance	0.0924	0.0252	0.5504	0.5546	238
Personal care and service	0.1386	0.0792	0.4455	0.6040	101
Sales	0.2827	0.2637	0.6751	0.3861	474
Office and administrative support	0.1413	0.1912	0.7966	0.7418	821
Farming, fishing, and forestry	0.1220	0.0244	0.4634	0.5854	41
Construction and extraction	0.1186	0.0805	0.4661	0.4195	236
Installation, maintenance, and repair	0.1105	0.0872	0.8372	0.6512	172
Production	0.0687	0.0236	0.7017	0.6202	466
Transportation and material moving	0.1022	0.0224	0.6070	0.4601	313
All occupations	0.2028	0.2183	0.7420	0.6130	6377
Across-occupation variance	0.0117	0.0271	0.0167	0.0187	

Table 38: Earnings/hours elasticity by occupation category

Occupation category	Goldin (2014)		Lann (2024)	
	Elasticity	N	Elasticity	N
Management	0.6631	26	0.7180	33
Business and financial operation	0.8872	24	0.9851	32
Computer and mathematical science	0.6988	9	0.5492	20
Architecture and engineering	0.4537	18	0.3834	24
Life, physical, and social science	0.3732	19	0.4330	26
Community and social service	0.3600	5	0.4147	6
Legal	0.7208	3	0.4640	17
Education, training, and library	0.4831	10	0.5026	16
Arts, design, entertainment, sports, and media	0.6450	18	0.4739	37
Healthcare practitioner and technical	0.3145	29	0.4472	50
Healthcare support	0.2354	5	0.4616	14
Protective service	0.6084	15	0.4520	22
Food preparation and serving	0.4962	12	0.5184	14
Building and grounds cleaning and maintenance	0.5728	6	0.5670	9
Personal care and service	0.6830	18	0.4257	27
Sales	0.8741	18	0.9278	18
Office and administrative support	0.8791	49	0.6834	57
Farming, fishing, and forestry	0.7642	6	0.5078	8
Construction and extraction	0.3130	34	0.4413	41
Installation, maintenance, and repair	0.5743	33	0.5453	33
Production	1.0406	72	0.6043	73
Transportation and material moving	0.5246	28	0.5489	41
All occupations	0.6586	469	0.5569	629

Table 39: Regressions of standardized earnings/hours elasticity on O*NET characteristics

Time pressure	0.0928*					0.0940**
	(0.0476)					(0.0477)
Contact with others		-0.00825				-0.0315
		(0.0506)				(0.0587)
Interpersonal relationships			0.0402			0.0923
			(0.0480)			(0.0594)
Structured vs unstructured				0.0434		-0.0394
				(0.0475)		(0.0789)
Freedom to make decisions					-0.0989**	-0.159**
					(0.0486)	(0.0780)
N	423	423	423	423	423	423
adj. R^2	0.007	-0.002	-0.001	-0.000	0.007	0.015

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 40: Logit regressions of flexibility measures on occupation characteristics, all occupations

	(1) Flextime	(2) Paid WFH	(3) Paid leave	(4) Predictability
Time pressure	0.0206	0.0589	0.194***	0.128***
	(0.0569)	(0.0649)	(0.0487)	(0.0488)
Contact with others	-0.167**	-0.363***	-0.0493	-0.155***
	(0.0665)	(0.0704)	(0.0567)	(0.0535)
Interpersonal relationships	0.0813	0.221**	0.283***	0.0815
	(0.0827)	(0.0874)	(0.0666)	(0.0603)
Structured vs unstructured work	-0.141	-0.748***	-0.135*	-0.250***
	(0.0905)	(0.103)	(0.0811)	(0.0759)
Freedom to make decisions	0.178**	-0.263***	-0.0941	-0.132*
	(0.0873)	(0.0956)	(0.0841)	(0.0735)
Earnings/hours elasticity	0.238***	0.363***	-0.0231	-0.213***
	(0.0452)	(0.0556)	(0.0497)	(0.0406)
Worker observables	Yes	Yes	Yes	Yes
N	6377	6377	6377	6377
pseudo R^2	0.091	0.206	0.128	0.041

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 41: Logit regressions of flexibility measures on occupation characteristics, white collar only

	(1) Flextime	(2) Paid WFH	(3) Paid leave	(4) Predictability
TimePressure	0.00754	-0.106	0.354**	0.151
	(0.126)	(0.124)	(0.168)	(0.131)
ContactWithOthers	-0.0954	-0.161	-0.500***	-0.191
	(0.118)	(0.121)	(0.157)	(0.120)
EstablishingandMaintainingInterp	-0.0632	0.0851	0.307*	-0.00564
	(0.122)	(0.117)	(0.159)	(0.115)
StructuredversusUnstructuredWork	0.111	-0.153	-0.492**	-0.131
	(0.179)	(0.187)	(0.242)	(0.180)
FreedomtoMakeDecisions	0.167	-0.212	-0.556**	-0.274
	(0.195)	(0.195)	(0.273)	(0.191)
Earnings/hours elasticity	0.0166	0.185**	-0.375***	-0.0645
	(0.0836)	(0.0867)	(0.109)	(0.0835)
Worker observables	Yes	Yes	Yes	Yes
N	1722	1722	1722	1722
pseudo R^2	0.038	0.072	0.144	0.018

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 42: Comparison of adjusted R-squared values

	Flextime	Paid WFH	Paid leave	Predictability
Individual variables only	0.087	0.208	0.133	0.080
Goldin O*NET characteristics only	0.040	0.110	0.051	0.019
Individual and Goldin O*NET variables	0.089	0.213	0.138	0.089
All O*NET characteristics	0.110	0.211	0.124	0.113
Individual and all O*NET variables	0.126	0.257	0.185	0.131

Table 43: Worker and job characteristics, by flexibility status

Panel A: Worker demographics					
	Flextime	Paid WFH	Paid leave	Predictability	Mean
Years of education	1.326***	1.894***	1.204***	0.590***	14.72
Age	-0.162	-0.109	0.539*	0.0657	43.92
Female	-0.0532***	0.0259*	-0.0204	0.121***	0.5150
White	0.0475***	0.0615***	0.0234**	0.0196*	0.8082
Panel B: Household variables					
	Flextime	Paid WFH	Paid leave	Predictability	Mean
Married	0.0842***	0.125***	0.0686***	0.0369***	0.5554
Number of children	0.122***	0.154***	-0.0512	0.0378	0.9207
Large MSA	0.0648***	0.134***	0.0210	-0.0224*	0.5438
Panel C: Occupation variables					
	Flextime	Paid WFH	Paid leave	Predictability	Mean
Hourly wage	7.786***	11.40***	7.854***	1.466***	25.9728
Weekly labor hours	1.734***	2.552***	5.832***	-0.0291	41.91
Salaried, occ. avg.	0.175***	0.254***	0.114***	0.0507***	0.4642
Private employer	0.0633***	0.0678***	-0.109***	-0.142***	0.8034
Large firm, occ. avg.	0.0216***	0.0403***	0.0651***	0.0398***	0.4785
Std. earnings/hours elasticity	0.212***	0.322***	0.0445*	-0.173***	-0.0488

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 44: Flexibility policies by industry category

Industry category	Flextime	Paid work-from-home	Paid leave	Predictability	N
Agriculture, forestry, fishing, and hunting	0.1695	0.1186	0.5424	0.5932	59
Mining, quarrying, and oil and gas extraction	0.2857	0.3571	0.8929	0.5714	28
Construction	0.1746	0.1714	0.5143	0.4444	315
Manufacturing	0.2046	0.2149	0.8069	0.6230	870
Wholesale and retail trade	0.1841	0.1580	0.7058	0.4362	690
Transportation and utilities	0.1508	0.1231	0.7877	0.6123	325
Information	0.1743	0.3394	0.8073	0.5688	109
Financial activities	0.2954	0.4114	0.7965	0.6477	457
Professional and business services	0.3210	0.4252	0.7311	0.5720	729
Educational and health services	0.1446	0.1553	0.7766	0.7530	1777
Leisure and hospitality	0.1829	0.0756	0.4902	0.3220	410
Other services	0.2800	0.2844	0.6311	0.6400	225
Public administration	0.2167	0.2063	0.9399	0.7493	383
All industries	0.2028	0.2183	0.7420	0.6130	6377
Across-industry variance	0.0037	0.0115	0.0113	0.0172	

Table 45: Correlation of worker and job characteristics

	Education	Age	Female	Children	Married	White	MSA	Hours	Salaried	Elasticity	Hourly wage	Private	Large firm
Education	1.0000												
Age	-0.0920	1.0000											
Female	0.1184	0.0129	1.0000										
Children	0.0234	-0.2986	-0.0059	1.0000									
Married	0.1038	0.0112	-0.0771	0.3336	1.0000								
White	-0.0009	-0.0371	-0.0352	0.0759	0.1333	1.0000							
MSA	0.1186	-0.0023	0.0152	-0.0282	-0.0268	-0.1223	1.0000						
Hours	0.0786	-0.0020	-0.2492	-0.0201	0.0073	0.0444	-0.0092	1.0000					
Salaried	0.5223	-0.0145	-0.0017	0.0526	0.1568	0.0640	0.1043	0.1972	1.0000				
Elasticity	-0.0992	0.0143	-0.0919	-0.0272	-0.0277	0.0114	0.0331	0.0671	0.0095	1.0000			
Hourly wage	0.4332	0.0751	-0.1144	0.0858	0.1713	0.0649	0.1903	0.0881	0.04309	0.0042	1.0000		
Private	-0.2084	-0.0481	-0.0734	0.0113	-0.0522	0.0327	0.0697	-0.0159	-0.1997	0.2603	-0.0478	1.0000	
Large firm	0.2956	-0.0301	0.0893	0.0226	0.0442	-0.0518	0.0106	0.0150	0.2584	-0.0652	0.1894	-0.2648	1.0000

Table 46: Ten most informative principal components of ONET characteristics

	Component 1	**Component 2**
Percent of variance explained	31.10%	16.53%
Five most important O*NET characteristics (highest positive coefficients)	ElectronicMail	SpendTimeMakingRepetitiveMotions
	EstablishingandMaintainingInterp	StructuredversusUnstructuredWork
	InteractingWithComputers	IndoorsEnvironmentallyControlled
	LettersandMemos	SpendTimeSitting
	CommunicatingwithPersonsOutsideO	ImportanceofRepeatingSameTasks
	Component 3	**Component 4**
Percent of variance explained	7.51%	5.71%
Five most important O*NET characteristics (highest positive coefficients)	AssistingandCaringforOthers	SellingorInfluencingOthers
	DealWithUnpleasantorAngryPeople	InanEnclosedVehicleorEquipment
	PhysicalProximity	OutdoorsUnderCover
	PerformingfororWorkingDirectlywi	OutdoorsExposedtoWeather
	DealWithPhysicallyAggressivePeop	PublicSpeaking
	Component 5	**Component 6**
Percent of variance explained	4.89%	3.07%
Five most important O*NET characteristics (highest positive coefficients)	CoachingandDevelopingOthers	ExposedtoDiseaseorInfections
	TrainingandTeachingOthers	UpdatingandUsingRelevantKnowledg
	CoordinatingtheWorkandActivities	ExposedtoRadiation
	JudgingtheQualitiesofThingsServi	FreedomtoMakeDecisions
	DevelopingandBuildingTeams	RepairingandMaintainingElectroni
	Component 7	**Component 8**
Percent of variance explained	2.33%	2.09%
Five most important O*NET characteristics (highest positive coefficients)	CommunicatingwithSupervisorsPeer	StructuredversusUnstructuredWork
	DealWithPhysicallyAggressivePeop	PerformingfororWorkingDirectlywi
	StructuredversusUnstructuredWork	PerformingAdministrativeActiviti
	PublicSpeaking	SellingorInfluencingOthers
	OutdoorsExposedtoWeather	ResolvingConflictsandNegotiating
	Component 9	**Component 10**
Percent of variance explained	1.52%	1.44%
Five most important O*NET characteristics (highest positive coefficients)	ImpactofDecisionsonCoworkersorCo	LevelofCompetition
	FreedomtoMakeDecisions	ExposedtoRadiation
	FrequencyofDecisionMaking	ImpactofDecisionsonCoworkersorCo
	ConsequenceofError	StructuredversusUnstructuredWork
	ExposedtoContaminants	WearSpecializedProtectiveorSafet

Table 47: Regressions of flexibility measures on ONET principal compoenents

	(1) Flextime	(2) Paid WFH	(3) Paid leave	(4) Predictability
Component 1	0.0147***	0.0253***	0.0172***	0.0120***
	(0.000952)	(0.000921)	(0.00104)	(0.00117)
Component 2	-0.000206	0.00189	-0.0120***	0.00101
	(0.00148)	(0.00143)	(0.00161)	(0.00182)
Component 3	-0.0216***	-0.0367***	-0.0172***	-0.0117***
	(0.00202)	(0.00195)	(0.00219)	(0.00247)
Component 4	0.0134***	0.0187***	-0.0182***	-0.0202***
	(0.00225)	(0.00218)	(0.00245)	(0.00276)
Component 5	-0.00436*	-0.00700***	-0.0119***	-0.00402
	(0.00250)	(0.00242)	(0.00272)	(0.00306)
Component 6	0.00693**	0.00238	-0.0130***	0.00910**
	(0.00352)	(0.00341)	(0.00383)	(0.00431)
Component 7	-0.0192***	-0.0133***	0.0111***	0.0445***
	(0.00396)	(0.00383)	(0.00430)	(0.00485)
Component 8	-0.00834**	-0.00941**	-0.00796*	-0.0276***
	(0.00391)	(0.00378)	(0.00425)	(0.00478)
Component 9	0.00833*	-0.000557	0.00646	0.00997*
	(0.00467)	(0.00451)	(0.00508)	(0.00571)
Component 10	0.0256***	0.0224***	-0.0268***	-0.0474***
	(0.00498)	(0.00481)	(0.00541)	(0.00609)
N	6377	6377	6377	6377
adj. R^2	0.078	0.184	0.079	0.058

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 48: Performance comparison of various classifiers

Classifiers	Accuracy	ROC AUC	Brier Score	Log Loss	Time
Logistic regression	0.6841	0.7328	0.4119	0.4983	0.76 seconds
Lasso logistic regression	0.6482	0.6807	0.4533	0.6456	7.06 seconds
Random forest	0.7550	0.7350	0.3331	0.5582	8.52 seconds
Gradient boosting	0.7683	0.7653	0.3127	0.4755	44.33 seconds
XGBoost	0.7684	0.7659	0.3118	0.4743	2.02 seconds
K-nearest-neighbors	0.7236	0.6440	0.3939	2.2196	2.52 seconds
Naive Bayes	0.6275	0.6931	0.7236	6.2730	0.28 seconds
Neural net	0.7404	0.7061	0.3833	0.6295	7.13 seconds

Table 49: Summary statistics of flex indices by model criteria

Panel A: ATUS Data				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.2028	0.2183	0.7420	0.6130
Across-occupation variance	0.0248	0.0443	0.0342	0.0382
Panel B: Highest accuracy				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.1816	0.1854	0.6944	0.5815
Across-occupation variance	0.0079	0.0292	0.0301	0.0160
Panel C: Highest AUROC				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.1807	0.1798	0.7008	0.5876
Across-occupation variance	0.0098	0.0289	0.0239	0.0157
Panel D: Lowest Brier Score				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.1096	0.1532	0.7187	0.5967
Across-occupation variance	0.0195	0.0403	0.0461	0.0583
Panel E: Lowest Log Loss				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.1038	0.1577	0.7187	0.6060
Across-occupation variance	0.0183	0.0401	0.0461	0.0511
Panel F: Highest composite measure				
	Flextime	Paid WFH	Paid leave	Predictability
Mean	0.1796	0.1795	0.6995	0.5867
Across-occupation variance	0.0109	0.0293	0.0264	0.0186

Table 50: Comparison of XGBoost models with different features and parameters

	Accuracy	AUROC	Brier Score	Log Loss
Individual variables only, XGBoost defaults	0.7663	0.7606	0.3142	0.4782
O*NET variables only, defaults	0.7484	0.7196	0.3368	0.5083
Individual and O*NET variables, defaults	0.7684	0.7660	0.3118	0.4743
Individual variables and first 10 O*NET PCs, defaults	0.7680	0.7661	0.3117	0.4741
Hyperparameter and PC tuning	0.7725	0.7776	0.3079	0.4693

Table 51: Most important features for tuned models

	Flextime	Paid WFH	Paid leave	Predictability
Number of PCs	5	11	0	5
Ten most important variables	Salaried/hourly	Hourly wage	Weekly hours	Component 4
	Weekly hours	Component 1	Hourly wage	Private employer
	Hourly wage	Education	Private employer	Weekly hours
	Greediness	Salaried/hourly	InteractingWithComputers	Component 1
	Education	Component 3	ElectronicMail	Component 3
	Component 1	Female	LevelofCompetition	Hourly wage
	Component 3	Weekly hours	AnalyzingDataorInformation	Female
	Private employer	Greediness	HandlingandMovingObjects	Component 5
	Component 4	Component 4	FreedomtoMakeDecisions	Greediness
	White	Private employer	GettingInformation	White

Table 52: Correlations between flexibility policies and flexibility indices

Panel A: Correlations between flexibility policies and indices

	Flextime	Paid WFH	Paid leave	Predictability
Correlation with index	0.5221	0.7650	0.7031	0.5397

Panel B: Correlations between flexibility policies

	Flextime	Paid WFH	Paid leave	Predictability
Flextime	1.0000			
Paid WFH	0.5267	1.0000		
Paid leave	0.1293	0.2792	1.0000	
Predictability	-0.0009	0.0584	0.2526	1.0000

Panel C: Correlations between occupation flexibility indices

	Flextime	Paid WFH	Paid leave	Predictability
Flextime	1.0000			
Paid WFH	0.9335	1.0000		
Paid leave	0.4593	0.5426	1.0000	
Predictability	0.1592	0.2806	0.6171	1.0000

Table 53: Flexibility indices by occupation category

Occupation category	Flextime index	Paid WFH index	Paid leave index	Predictability index
Management	0.3171	0.4131	0.8481	0.5949
Business and financial operation	0.3299	0.4539	0.8207	0.6041
Computer and mathematical science	0.4334	0.4207	0.8650	0.6784
Architecture and engineering	0.3559	0.4120	0.8721	0.6128
Life, physical, and social science	0.2906	0.3326	0.8532	0.6868
Community and social service	0.2062	0.2710	0.8228	0.7086
Legal	0.3749	0.4907	0.8475	0.6670
Education, training, and library	0.1422	0.1319	0.7073	0.7772
Arts, design, entertainment, sports, and media	0.2757	0.3610	0.7077	0.5690
Healthcare practitioner and technical	0.1348	0.1092	0.8018	0.6538
Healthcare support	0.0929	0.0417	0.6430	0.6187
Protective service	0.1456	0.0885	0.7835	0.6436
Food preparation and serving	0.1129	0.0210	0.3927	0.3094
Building and grounds cleaning and maintenance	0.0977	0.0278	0.5224	0.5502
Personal care and service	0.1381	0.0785	0.4283	0.5487
Sales	0.2122	0.2002	0.6446	0.4126
Office and administrative support	0.1427	0.1744	0.7539	0.6945
Farming, fishing, and forestry	0.1310	0.0260	0.3706	0.5956
Construction and extraction	0.0994	0.0538	0.5271	0.4961
Installation, maintenance, and repair	0.1140	0.0809	0.7630	0.5721
Production	0.0960	0.0317	0.6806	0.5954
Transportation and material moving	0.1106	0.0281	0.5710	0.4682
All occupations	0.1796	0.1795	0.6996	0.5867

Table 54: Comparison of adjusted R-squared values for regressions of occupation-mean flexibility policies

	Flexibility	Paid WFH	Paid leave	Predictability	N
All ONET characteristics	0.174	0.375	0.349	0.124	291
All ONET and individual variables	0.197	0.525	0.438	0.226	323
Flexibility indices	0.270	0.584	0.493	0.289	325

Bibliography

Brant Abbott. Incomplete markets and parental investments in children. Technical report, 2020. Working paper.

Joseph G. Altonji, Fumio Hayashi, and Laurence J. Kotlikoff. Parental altruism and inter vivos transfers: Theory and evidence. *Journal of Political Economy*, 105(6):1121–1166, 1997. ISSN 00223808, 1537534X. URL `http://www.jstor.org/stable/10.1086/516388`.

Jesper Bagger, Javier A. Birchenall, Hani Mansour, and Sergio Urzúa. Education, birth order, and family size. *National Bureau of Economic Research*, June 2013.

Minji Bang. Job flexibility and household labor supply: Understanding gender gaps and the child wage penalty. 01 2022.

Gary Becker and Nigel Tomes. An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy*, 87(6):1153–89, 1979. URL `https://EconPapers.repec.org/RePEc:ucp:jpolec:v:87:y:1979:i:6:p:1153-89`.

Gary Becker and Nigel Tomes. Human capital and the rise and fall of families. *Journal of Labor Economics*, 4(3):S1–39, 1986. URL `https://EconPapers.repec.org/RePEc:ucp:jlabec:v:4:y:1986:i:3:p:s1-39`.

Gary S. Becker and Nigel Tomes. Child endowments and the quantity and quality of children. *Journal of Political Economy*, 84(4):S143–S162, 1976. ISSN 00223808, 1537534X. URL `http://www.jstor.org/stable/1831106`.

Liza Benny, Sonia Bhalotra, and Manuel FernÃ¡ndez. Occupation flexibility and the graduate gender wage gap in the UK. ISER Working Paper Series 2021-05, Institute for Social and Economic Research, July 2021. URL `https://ideas.repec.org/p/ese/iserwp/2021-05.html`.

Daniela Del Boca, Christopher Flinn, and Matthew Wiswall. Household choices and child development. *The Review of Economic Studies*, 81(1 (286)):137–185, 2014. ISSN 00346527, 1467937X. URL `http://www.jstor.org/stable/43551669`.

Stéphane Bonhomme and Gregory Jolivet. The pervasive absence of compensating differentials. *Journal of Applied Econometrics*, 24(5):763–795, 2009. URL `https://EconPapers.repec.org/RePEc:jae:japmet:v:24:y:2009:i:5:p:763-795`.

Charles Brown. Equalizing differences in the labor market. *The Quarterly Journal of Economics*, 94(1):113–134, 1980. URL `https://EconPapers.repec.org/RePEc:oup:qjecon:v:94:y:1980:i:1:p:113-134`.

L. Carlson and K. B. Guzzo. Median age at last birth. family profiles, fp-21-05. bowling green, oh: National center for family marriage research. 2021. URL https://doi.org/10.25035/ncfmr/fp-21-05.

Elizabeth M. Caucutt and Lance Lochner. Early and late human capital investments, borrowing constraints, and the family. *Journal of Political Economy*, 128(3):1065–1147, 2020. doi: 10.1086/704759. URL https://doi.org/10.1086/704759.

Elizabeth M Caucutt, Lance Lochner, Joseph Mullins, and Youngmin Park. Child skill production: Accounting for parental and market-based time and goods investments. Working Paper 27838, National Bureau of Economic Research, September 2020. URL http://www.nber.org/papers/w27838.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 785–794, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939785. URL https://doi.org/10.1145/2939672.2939785.

Council of Economic Advisers. The anti-poverty and income-boosting impacts of the enhanced ctc, 2023. https://www.whitehouse.gov/cea/written-materials/2023/11/20/the-anti-poverty-and-income-boosting-impacts-of-the-enhanced-ctc/#:~:text=The%20expanded%20CTC%20would%20substantially,2.9%E2%80%935.7%20percentage%20points%20annually., Last accessed on 2024-3-30.

German Cubas, Chinhui Juhn, and Pedro Silos. Coordinated Work Schedules and the Gender Wage Gap. *The Economic Journal*, 133(651):1036–1066, 11 2022. ISSN 0013-0133. doi: 10.1093/ej/ueac086. URL https://doi.org/10.1093/ej/ueac086.

Flavio Cunha and James Heckman. The technology of skill formation. *American Economic Review*, 97(2):31–47, May 2007.

Flavio Cunha, James J. Heckman, and Susanne M. Schennach. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931, 2010.

Diego Daruich and Julian Kozlowski. Explaining Intergenerational Mobility: The Role of Fertility and Family Transfers. *Review of Economic Dynamics*, 36:220–245, April 2020.

Matthew S. Dey and Christopher Flinn. An equilibrium model of health insurance provision and wage determination. *Econometrica*, 73(2):571–627, 2005. URL https://EconPapers.repec.org/RePEc:ecm:emetrp:v:73:y:2005:i:2:p:571-627.

William T. Dickens and James R. Flynn. Black americans reduce the racial iq gap: Evidence from standardization samples. *Psychological Science*, 17(10):913–920, 2006. doi: 10.1111/j.1467-9280.2006.01802.x. URL https://doi.org/10.1111/j.1467-9280.2006.01802.x. PMID: 17100793.

Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry. Integrated public use microdata series, current population survey: Version 10.0 [dataset]. minneapolis, mn: Ipums, 2022. a. URL `https://doi.org/10.18128/D030.V10.0`.

Sarah M. Flood, Liana C. Sayer, Daniel Backman, and Annie Chen. American time use survey data extract builder: Version 3.2 [dataset]. college park, md: University of maryland and minneapolis, mn: Ipums, 2023. b. URL `https://doi.org/10.18128/D060.V3.2`.

Claudia Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, April 2014. doi: 10.1257/aer.104.4.1091. URL `https://www.aeaweb.org/articles?id=10.1257/aer.104.4.1091`.

Bart Hobijn and Aysegul Sahin. Job-finding and separation rates in the oecd. *Economics Letters*, 104(3):107–111, 2009. URL `https://EconPapers.repec.org/RePEc:eee:ecolet:v:104:y:2009:i:3:p:107-111`.

Hae-shin Hwang, W. Reed, and Carlton Hubbard. Compensating wage differentials and unobserved productivity. *Journal of Political Economy*, 100(4):835–58, 1992. URL `https://EconPapers.repec.org/RePEc:ucp:jpolec:v:100:y:1992:i:4:p:835-58`.

Larry Jones and Michele Tertilt. An economic history of fertility in the u.s.: 1826-1960. 1, 01 2007.

Chinhui Juhn, Yona Rubinstein, and C. Andrew Zuppann. The quantity-quality trade-off and the formation of cognitive and non-cognitive skills. Working Paper 21824, National Bureau of Economic Research, December 2015. URL `http://www.nber.org/papers/w21824`.

Thomas Kniesner, W Viscusi, Christopher Woock, and James Ziliak. The value of a statistical life: Evidence from panel data. *The Review of Economics and Statistics*, 94(1):74–87, 2012. URL `https://EconPapers.repec.org/RePEc:tpr:restat:v:94:y:2012:i:1:p:74-87`.

Steven Lann. The economics of workplace flexibility. Technical report, October 2023.

Chul-In Lee and Gary Solon. Trends in intergenerational income mobility. *The Review of Economics and Statistics*, 91(4):766–772, 2009. URL `https://EconPapers.repec.org/RePEc:tpr:restat:v:91:y:2009:i:4:p:766-772`.

Sang Yoon (Tim) Lee and Ananth Seshadri. On the intergenerational transmission of economic status. *Journal of Political Economy*, 127(2):855–921, 2019.

Grace Lordan and Jörn-Steffen Pischke. Does rosie like riveting? male and female occupational choices. *Economica*, 89(353):110–130, 2022. doi: https://doi.org/10.1111/ecca.12390. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/ecca.12390`.

Nicole Maestas, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger. The value of working conditions in the united states and the implications for the structure of wages. *American Economic Review*, 113(7):2007–47, July 2023. doi: 10.1257/aer.20190846. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190846.

Alexandre Mas and Amanda Pallais. Alternative work arrangements. *Annual Review of Economics*, 12(1):631–658, 2020. URL https://EconPapers.repec.org/RePEc:anr:reveco:v:12:y:2020:p:631-658.

McKinsey & Company. Americans are embracing flexible work—and they want more of it, 2022. https://www.mckinsey.com/industries/real-estate/our-insights/americans-are-embracing-flexible-work-and-they-want-more-of-it, Last accessed on 2023-10-24.

National Center for O*NET Development. O*NET OnLine. Retrieved January 23, 2024, from https://www.onetonline.org/.

Emily Peck. Companies get aggressive on return-to-office, 2023. https://www.axios.com/2023/06/13/companies-aggressive-return-to-office, Last accessed on 2023-10-24.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Steven Ruggles, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler. Ipums usa: Version 14.0 [dataset]. minneapolis, mn: Ipums, 2023. URL https://doi.org/10.18128/D010.V14.0.

David L. Schaffer and Joseph M. Westenberg. Time Flexibility, Women's Wages, and the Gender Wage Gap. *Atlantic Economic Journal*, 47(2):217–239, June 2019. doi: 10.1007/s11293-019-09622-. URL https://ideas.repec.org/a/kap/atlecj/v47y2019i2d10.1007_s11293-019-09622-9.html.

V. Schweizer and K. B. Guzzo. Age at first birth among mothers 40-44, 1990 2018. family profiles, fp-20-06. bowling green, oh: National center for family marriage research. 2020. URL https://doi.org/10.25035/ncfmr/fp-20-06.

Kristen Sobeck. Greedy jobs, labour market institutions, and the gender pay gap. *SSRN*, 2022. doi: http://dx.doi.org/10.2139/ssrn.4306651. URL https://ssrn.com/abstract=4306651.

Paul Sullivan and Ted To. Search and nonwage job characteristics. *The Journal of Human Resources*, 49(2):472–507, 2014. ISSN 0022166X. URL http://www.jstor.org/stable/23799092.

Paul Sullivan and Ted To. Job dispersion and compensating wage differentials. Technical report, Bureau of Labor Statistics, 2023.

Matthew Wiswall and Basit Zafar. Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, 133(1):457–507, 2018. URL https://EconPapers.repec.org/RePEc:oup:qjecon:v:133:y:2018:i:1:p:457-507.

Minchul Yum. Parental time investment and intergenerational mobility. CRC TR 224 Discussion Paper Series crctr224_2018_048, University of Bonn and University of Mannheim, Germany, October 2018. URL https://ideas.repec.org/p/bon/boncrc/crctr224_2018_048.html.