A structural meta-analysis of welfare reform experiments and their impacts on children

Joseph Mullins *

December 11, 2023

Abstract

Using a model of maternal labor supply and investment in children, this paper synthesizes the findings from three separate welfare reform experiments across six sites. The proposed model maps variation in experimental design to parameters that define labor supply behavior, child care use, and the importance of money and child care arrangements in the development of child skills. The estimation procedure, which aggregates available evidence to identify the model's key causal parameters, amounts to a *structural meta-analysis*. Estimates suggest that while family resources and care arrangements play a quantitatively significant role in shaping both academic and behavioral outcomes, the predicted effects do not exhibit persistence over time. A number of counterfactuals underscore the utility of this model-based approach for understanding the mechanisms behind treatment effects and the roles played by heterogeneity and selection in shaping impacts. For example, although the experiments do not find any negative impacts on child skill development, the estimated model forecasts that these same policy reforms do entail negative effects on behavioral skills when applied to a representative sample of single mothers.

*Dept. of Economics, University of Minnesota. Email: mullinsj@umn.edu The author wishes to thank Roy Allen, Mark Colas, Anmol Bhandari, Ellen McGrattan, Jeremy Lise, Chris Flinn, and Jim Heckman for useful comments, as well as seminar participants at Wash U St Louis, UNC Chapel Hill, Minnesota, and Wisconsin. Excellent research assistance has been provided by Madison Bozich, Filip Babalievsky, Alexander Djanin and Nikhil Saxena.

1 Introduction

This paper uses an economic model of behavior to aggregate statistical information from multiple welfare reform experiments. It follows the mission statement laid out by Frisch (1933) in the very first issue of *Econometrica*: to coordinate the accumulation of empirical evidence for policy and prediction using a theoretical framework.[1] While traditional meta-analyses approach this problem by specifying the average treatment effect (ATE) as a global parameter of interest —using linear statistical models to construct a weighted average estimate of this parameter —this paper instead combines data to identify the underlying parameters of an economic model.

There are four main payoffs to taking a structural approach to meta-analysis. First, if the experimental setting is well-suited to the application of economic theory, then a model can provide an organized comparative interpretation of treatment effects across experiments. This is particularly true if treatment differences can be modeled without the need for estimating additional parameters.[2] Second, the model provides a mapping between structural parameters to a broad class of counterfactual scenarios, and can therefore be used in cases where the ATE itself does not define an explicit policy counterfactual of interest (Heckman, 1992; Heckman and Vytlacil, 2005). Third, by allowing for and estimating the latent dimensions along which individuals select into the experimental studies, the estimated model can forecast treatment impacts on broader populations of interest. Finally, an economic model allows researchers to make normative statements about how treatments (both observed and counterfactual) are valued by participants by applying the lens of revealed preference.

Using the structural approach, this paper studies the design of cash assistance and its effect on child outcomes by combining individual-level panel data from four sources: from three experimental evaluations of welfare-to-work programs in the United States, along with a representative sample of single mothers from the Survey of Income and Program Participation (SIPP). The inclusion of the SIPP is not crucial but is useful as its representative sampling frame provides a counterbalance to the highly selected experiment populations where the sampling frame is comprised of new welfare applicants and ongoing recipients.

The three experiments: Connecticut Jobs First (CTJF), the Family Transition Program (FTP) in Florida, and the Minnesota Family Investment Program (MFIP), were each conducted between 1994 and 1999 by the Manpower Demonstration Research Corporation (MDRC). Publicly available reports provide many useful site-specific details such as treatment components and sample design (Bloom et al., 2002, 2000; Miller et al., 2000; Gennetian and Miller, 2000). Each experiment restructured the conditions of cash assistance for participants with the goal of increasing labor force attachment and financial independence. In varying combinations, these experiments introduced (1) changes in benefit computation formulae; (2) mandatory employment and job search services; (3) time-limited program participation; and (4) expanded access to childcare subsidies. MDRC also collected at these sites a set of developmental indicators for school-aged children. This paper develops a dynamic model of welfare particiption, labor supply, childcare use, and child investment that provides an interpretation of each experiment's treatment effects as behavioral responses to changes in individual incentives. These changes induce variation in economic resources and childcare arrangements which in

[1] In his editorial introduction of the journal, Frisch (1933) wrote: "Statistical information is currently accumulating at an unprecedented rate. But no amount of statistical information, however complete and exact, can by itself explain economic phenomena. If we are not to get lost in the overwhelming, bewildering mass of statistical data that are now becoming available, we need the guidance and help of a powerful theoretical framework. Without this no significant interpretation and coordination of our observations will be possible."

[2] Below we define this feature as *well-articulated* variation inside the model

turn identifies the role of those inputs for skill development, as embodied by parameters of a skill formation technology.

Estimation of the model proceeds in two stages, beginning with a maximum likelihood routine that quantifies the model's behavioral primitives and allows for time-varying latent heterogeneity in labor market opportunities, along with permanent latent heterogeneity in individual preferences, labor market productivity, and labor market risk. Estimates document rich patterns of selection across all of these latent dimensions. Compared to the representative population, experimental populations at each site disproportionately select types with lower labor market productivity, as well as individuals facing temporary negative labor market shocks. The second stage of the procedure yields estimates of a technology of skill formation that maps changes in household income and childcare arrangements to behavioral and academic skill outcomes. An instrumental variables estimator uses only the variation in these inputs induced by random assignment, and offers too little precision for meaningful calculations. On the other hand, a model-implied control function approach yields much more informative estimates that provide the following lessons. A log-point increase in household resources is estimated to lead to a 1% standard deviation increase in behavioral skills. There is little evidence of any positive or negative impact of unpaid care (relative to full-time maternal care), but mean estimates suggest a loss of 6.9% of a standard deviation in behavioral skills when paid care is used instead of unpaid care. In exploring this puzzling result, allowing for heterogeneity in the ceterus paribus effects of care suggests quite meaningful differences across types. Finally, estimates suggest that all effects fade out rapidly, a finding that would likely evade empirical approaches that use persistence in skills to measure the time-varying effect of inputs.

A number of counterfactuals further demonstrate the utility of this approach to aggregating evidence and provide useful policy lessons. First, the estimated model can unbundle the many aspects of each treatment by conducting counterfactual scenarios in which treatment components are implemented in isolation. This exercise demonstrates that time limits are largely responsible for reductions in participation, while mandatory services explain most of the increase in employment relative to control groups. In the model, services can influence employment either by imposing a non-pecuniary cost on those who participate and do not work, or by increasing the rate at which job opportunities arrive while unemployed. Estimates suggest that services work through the latter channel.

A final counterfactual aims to understand the role of heterogeneity and selection in shaping treatment effects by calculating average treatment effects on key outcomes in the representative SIPP sample. A key lesson is that treatment effects do look quite different when applied to this non-selected sample of individuals and that while this population experiences welfare gains on average, there are negative consequences for the behavioral skills of children. This finding contrasts with results on the experimental populations, where the effects of the experiments have no clear sign.

This paper joins a larger body of work that uses structural models to better understand the impacts of social experiments. Todd and Wolpin (2023) provide a thorough review of the benefits of combining randomized control trials with structural modeling. The economic model of behavior provides an explicit framework with which to (1) handle the traditional bugbears of experiment design: non-compliance and substitution bias, two phenomena containing inherent information by revealed preference (Heckman et al., 1997, 2000); (2) construct ex-post improvements in policy efficiency (Todd and Wolpin, 2005; Duflo et al., 2012; Rodriguez, 2018); (3) validate models of behavior (Todd and Wolpin, 2005; Choi, 2018; Galiani et al., 2015); and (4) identify key structural parameters with broader implications (Kline and Tartari, 2016; Chan, 2017; Galiani et al., 2015). This paper differs from the bulk of the literature in two main ways. First and most obviously, the emphasis of this paper is on aggregating findings from multiple related experiments. Second, this paper will adopt the position (and demonstrate through several conceptual and quantitative examples) that the experimental variation across sites is quite important for identification and estimation of the model. It is, in the spirit of a traditional meta-analysis, focused on the task of aggregating evidence.

The rest of the paper proceeds as follows. The section below provides a more detailed discussion of this paper's methodology and relates it to existing work, while introducing the concept of "well-articulated variation", which is an important concept for how components of the treatment map into the model environment. Section 2 briefly summarizes the data. Section 3 introduces the model while Section 4 expounds on identification and explores its empirical content. Section 5 describes estimation and presents estimates of the model's primitives, as well as conducting a number of validation exercises. Section 6 uses the model for deriving particular policy lessons before Section 7 concludes.

1.1 Methodology and Related Literature

This paper is not the first to re-examine the findings of welfare-to-work evaluations in the United States, nor is it the first to demonstrate the benefits of using economic models for interpreting social experiments. A comparison with these prior studies is useful to illustrate this paper's methodological points of departure. To begin, a meta-analysis may undertake one of two related but distinct goals. The first is simply to aggregate the available evidence to identify and estimate parameters of interest with increased precision. This includes but is not limited to the average treatment effect. For example, to estimate the effect of one or more endogenous variables on a particular outcome of interest, one can use assignment to treatment across different sites as instrumental variables to provide exogenous identifying variation. In the welfare-to-work context, this approach has been used extensively to study the effect of programs on child outcomes as mediated through different combinations of inputs such as income, childcare arrangements, and education (Morris and Gennetian, 2003; Gennetian et al., 2004a,b, 2008; Crosby et al., 2010; Duncan et al., 2011).[3] The second goal is to better understand differences in treatment effects across sites by projecting them onto differences in treatment components and population characteristics. Ashworth et al. (2004) and Greenberg et al. (2005) provide two such examples in the welfare-to-work context, finding that both treatment components and population characteristics play an important role.[4]

The economic approach can also be used in pursuit of either of these two goals. It entails many advantages, but these can only be gained by additional assumptions that impose extra structure on the data, and researchers are not always in full command of when this additional structure dictates conclusions unnecessarily. This requires extra care in estimation and identification. Typically, the key causal parameters of most models can be identified even without experimental variation in treatment components. This is true whenever the treatment is *well-articulated* inside the model, which this paper defines as a model's ability to replicate treatments *without* the need for additional parameters. Todd and Wolpin (2005) provide one classic example, showing that the impacts of a conditional cash transfer can be forecast using alternative sources of variation in wages.[5] Heckman and Vytlacil (2005) note that any proposed policy that involves changes

[3]With the exception of Duncan et al. (2011) (who study only the role of income), these approaches are typically underpowered due to the weakness of random assignment as an instrument.

[4]In related work, Card et al. (2010) and Card et al. (2018) use a similar approach to study active labor market programs.

[5]Todd and Wolpin (2023) provide a number of other examples of this kind.

to endowments, prices, or constraints, can be studied in an economic model so as long there exists some source of variation in these dimensions. In the language of this paper, variation in endowments, prices, and constraints are well-articulated inside an economic model.

Three examples that are relevant to welfare-to-work will help illustrate the concept. All three treatments modeled in this paper involve some combination of (1) changes to benefit computation; (2) time-limited participation in welfare; and (3) mandatory employment services for non-working welfare participants. In the model developed below, changes in benefit formulae can be modeled as *a priori* known changes to the budget set, hence no additional parameters are required to forecast their impact. Furthermore, existing budget set variation in the data makes it possible to identify the key parameters that would determine the response to any such change. Likewise, the introduction of time limits can be modeled in a parameter-free way, since time limits alter the budget set as a known function of past decisions. Finally, employment services are an example of variation that is *not* well-articulated in the model considered here. The effect of these services are articulated in the model through additional parameters that affect payoffs and job offer arrival rates, each of which can only be estimated with randomized assignment to these services.

Unlike many prior applications of economic modeling to social experiments, this paper treats experimental variation across sites as providing additional, useful sources of variation that must be incorporated and aggregated. The estimated model should hold, in the spirit of a meta-analysis, the totality of this evidence. However, unlike traditional meta-analyses, parameters of interest are not necessarily solely identified by treatment effects due to well-articulated variation. If researchers wish to know what can be learned only from random assignment and not from other sources of variation that require more assumptions to be considered valid, then a set of estimates that exploits all of this variation jointly will not provide the answer. Three exercises in the paper aim to address this concern. The first exercise compares model estimates and forecasted treatment effects when all experimental groups are withheld from the sample to those found when all variation is used. This demonstrates the key role played by the experiments in determining parameters. Second, to address concerns about model validity, treatment groups in three separate counties are withheld from the main estimation sample to show that the model performs well in forecasting outcomes out-of-sample. Finally Section 4, in establishing identification, shows that key model parameters can be uniquely identified by restrictions that use only variation provided by random assignment, suggesting an alternative estimation approach for researchers wishing to use experimental variation in isolation.

2 Data

Upon request, MDRC provided the evaluation data for the three experiments. The core data file for each experiment contains quarterly earnings from administrative Unemployment Insurance records, monthly welfare and food stamp receipt, and basic demographic information. A supplementary follow-up survey at 3-4 years post random assignment provides information on monthly childcare expenditures, and child outcomes. Child outcomes include scales for positive, externalizing, and internalizing behaviors, parent and teacher ratings of student achievement in school, parental reports of student engagement, and indicators for whether the child has repeated a grade or been suspended. All dollar values in the analysis are deflated to year 2000 USD, and monthly benefit receipt values are aggregated to quarterly. The final sample retains only unmarried women (the overwhelming majority) who appear in the follow-up survey on children. A paid care variable is constructed from reports of monthly childcare expenditures and is equal to 1 if the respondent reports paying

Table 2.1: Summary Statistics for each data source

	CTJF	CTJF	FTP	FTP	MFIP	MFIP	MFIP	SIPP
Arm	0	1	0	1	0	1	2	0
Less than Highschool	0.39	0.405	0.424	0.476	0.302	0.355	0.349	0.273
Highschool	0.52	0.503	0.553	0.515	0.494	0.472	0.474	0.351
Some College	0.084	0.084	0.021	0.009	0.204	0.173	0.176	0.299
College	0.005	0.008	0.002	0	0	0	0	0.078
AFDC Participation	0.595	0.586	0.439	0.391	0.567	0.734	0.765	0.253
Foodstamps Participation	0.696	0.718	0.647	0.632	0.603	0.144	0.158	0.337
Mother's age	26.808	26.794	25.437	25.203	26.581	26.346	26.594	28.2
Number of Children	2.04	2.057	2.164	2.111	2.064	2.166	2.153	1.754
Employed	0.561	0.627	0.495	0.538	0.526	0.592	0.536	0.639
Earnings	582.556	590.174	361.656	423.416	616.74	595.662	491.214	1550.693
Person-Quarter Observations	12817	13140	11781	11382	7866	8118	6336	11197
Individuals	711	730	561	542	437	451	352	953

for care in the previous month from any of: Head Start, regular day care, before- or after-school care, family daycare, summer programs, or babysitters.

Data used for the SIPP are from the 1992 and 1993 panels. Initial data cleaning is accomplished using Stata code from the Center for Economic and Policy Research (CEPR, 2014). While a majority of the primary variables identifying households are consistent between panel years, additional variables from topical modules are measured differently in each year or recorded under differing variable names. The CEPR code unifies and aggregates survey questions to a set of consistent variables, in order to construct uniform extracts of each SIPP panel year. The core sample provides demographic information as well as monthly data on food stamp and welfare receipt as well as earnings, which are aggregated to quarterly frequency and deflated to year 2000 USD. A topical module on childcare expenditures offers information on whether paid care is used and on monthly expenditures. The sample is restricted to unmarried women with children, where marital status is measured at the individual's first appearance in the panel.

Table 2.1 provides summary statistics for the final sample.

3 Model

This section describes a flexible empirical model of welfare participation, labor supply, and child development.

3.1 Model Primitives

Time is discrete and indexed by t. One period in the model is equal to one quarter of a year in the data. The child's skills are malleable until they reach age 17, corresponding to $T = 68$ periods, at which point the investment problem ends. Assume that t indexes the age (in quarters) of the youngest child. When applicable, dollar units for earnings and welfare receipt will reflect monthly values. Each mother m is characterized by a permanent and unobserved type, indexed by $k(m) \in \{0, 1, ..., K\}$. In addition to this latent type, payoffs, technologies, and constraints depend on a set of observable and unobservable state variables. Observable state variables are: age of the youngest child (t), state of residence, the local unemployment rate, family size, the mother's age, the calendar year, and (if participating in a welfare-to-work experiment) treatment status, indicated by $Z \in \{0, 1, 2\}$, and application status (an ongoing recipient or new applicant). Let A_{-1} be welfare participation status from the previous period. If an individual is living in a policy environment with time limits, then they must also track their cumulative welfare receipt, given by ω. Furthermore, let $\eta \in \{0, 1, ..., K_\eta\}$ be a latent, time-varing variable that tracks labor market opportunities. Let χ be a vector that holds the value of these variables:

$$\chi = \{t, \text{StateOfResidence}, \text{FamilySize}, \text{Year}, \text{Unemp}, \text{Age}, Z, A_{-1}, \omega, \eta\}.$$

Choices and Preferences

In each period, mothers make one of 9 discrete choices, indexed by j, that correspond to combinations of participation, work, and childcare choices. Let $S_j \in \{0, 1\}$ indicate food stamp participation, $A_j \in \{0, 1\}$ indicate welfare participation, and let $P_j = S_j + A_j \in \{0, 1, 2\}$ indicate the three possible combinations of no programs (0), food stamps only (1), and food stamps plus welfare (2). Let $H_j \in \{0, 1\}$ indicate the choice to work or not work, and (if working) let $F_j \in \{0, 1\}$ indicate whether paid childcare is solicited. In addition to this discrete choice, the agent splits their net income in period t between private consumption (C) and an investment good (I).

Mothers value their private consumption, C, and their child's current stock of attributes, θ. The utility she derives in any period is given by:

$$U_{kj}(C, \theta) + \epsilon_j$$

where:

$$U_{kj}(C, \theta) = \log(C) + \alpha_{\theta,k} \log(\theta) - \alpha_{H,k} H_j - \alpha_{S,k} S_j - (\alpha_{A,k} + \alpha_P(1 - A_{-1}))A_j + \alpha_{F,k} F_j$$

and ϵ is a vector of generalized extreme value taste shocks with a nested correlation structure. These are independently and identically distributed across individuals and time periods. The nesting structure forms a decision tree with nodes corresponding to first the participation choice, then the work choice, followed by (if applicable) the childcare choice. The triple $(\sigma_1, \sigma_2, \sigma_3)$ dictates the scale of shocks in each layer of this tree. Figure 3.1 depicts the nesting structure for these taste shocks. Note the heterogeneous, nonpecuniary payoffs associated with food stamp participation, work, and paid care use. Note also that the costs of welfare participation are *state dependent*: participation is less costly if the individual was participating in the previous period.

Mothers are forward-looking, discounting the future at rate β. Thus, she values future sequences of decisions as

$$V_{kt} = \mathbb{E}_t \left\{ \sum_{s=t}^{T-1} \beta^{s-t} U_{kj(s)}(C_s, \theta_s) + \beta^{T-t} \overline{V}(\theta_T) \right\}$$

where $\mathbb{E}_t$ is her conditional expectation given information at time t and $\overline{V}$ is a terminal payoff at time T, when the child's development has concluded, equal to:

$$(1 - \beta)^{-1} \alpha_{\theta,k} \log(\theta_T).$$

Preferences U are revealed by virtue of the assumption that mothers make their participation, time use, and investment decisions in order to maximize this expected discounted present value of utilities.

Figure 3.1: Nesting structure of discrete choice taste shocks

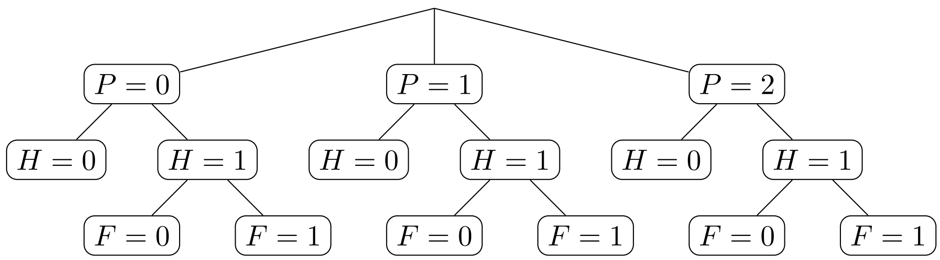

This figure depicts the nesting structure of the taste shocks ϵ, which has three layers. The terminal nodes of the tree indicate final combinations of choices, while edges leading to common nodes indicate membership to the same nest when calculating choice probabilities.

Technology and Constraints

There are two relevant technologies in the economy. First, the latent variable η indexes labor market opportunities according to the following rules. If $\eta = 0$, then this indicates that the individual has no job opportunities this period and cannot work. If $\eta > 0$, then it indexes a position in the distribution of potential earnings according to:

$$log(W_k(\chi)) = \mu_{W,k} + \beta_{W,1}\text{Unemp}_{m,t} + \beta_{W,2}\text{Age}_{m,t} + \sigma_\eta g(\eta)$$

where $g(\eta)$ indicates a position on a uniform grid on the interval $[-1, 1]$.[6].

Net income for the household in any period is then given by:

$$Y_{jk}(\chi) = y_k + H_j W_k(\chi) + \mathcal{T}(P_j, H_j W_k(\chi), \chi)$$

where y_k is (latent) non-labor income and $\mathcal{T}$ is a tax and transfer function that depends on the participation choice, on labor market earnings, on the number of children in the household, and the policy environment (which is a function of state and treatment status). Appendix C describes how this tax and transfer function is calculated.

Second, the child's attributes evolve according to a technology of skill formation with type-specific inital conditions:

$$\theta_{t+1} = \exp(\zeta) theta_t^{\delta_\theta} I^{\delta_I}, \qquad \theta_0 = \mu_{\theta,k}, \tag{3.1}$$

where ζ is an iid disturbance term. The aggregate investment good I is the outcome of an expenditure minimization problem with unit price $p_{H_j+F_j}$ that depends on the childcare arrangement for that period, which can take one of three values: (1) full-time maternal care ($H_j = 0, F_j = 0$); (2) unpaid care while the mother is working ($H_j = 1, F_j = 0$); or (3) paid care while the mother is working ($H_j = 1, F_j = 1$). The identification analysis below will establish that only the relative prices can be identified. Furthermore, only the log of these price ratios will appear in solution concepts and outcome equations. Accordingly, we define:

$$g_1 = -\log(p_1/p_0), \qquad g_2 = -\log(p_2/p_1).$$

[6]To be exact, $g(\eta) = -1 + (\eta - 1) * 2/(K_\eta - 1)$.

The term g_1 can be interpreted as the developmental effectiveness of part-time unpaid care relative to maternal care, and g_2 the effectivess of paid care relative to unpaid care. Section 4 below will clarify these interpretations.

Paid care is purchased at a price $q_k(\chi)$ given by:

$$\log(q_k(\chi)) = \mu_{q,k} + \beta_{q,1}\text{Unemp} + \beta_{q,2}\mathbf{1}\{\text{AgeYoungest} \leq 5\} + \beta_{q,3}\text{FamSize} + \gamma_{l(\chi),Z(\chi)}$$

where $l(\chi)$ indicates one of the three experimental sites and $Z(\chi)$ indicates the treatment arm within that site.

Technologies combine to produce a within-period budget constraint:

$$C + p_{H_j+F_j}I + q_k(\chi)F_j \leq Y_{kj}(\chi).$$

<u>Transitions</u>

Most state variables (such as those tracking age and year) evolve deterministically according to common sense rules. The variable η which tracks job opportunies evolves according to the following rules. Recall that when $\eta = 0$ the agent has no job offer, and these arrive for next period with probability $\lambda_{0,k}$. If a job offer arrives, it is drawn from the discrete distribution π_o which approximates a normal with mean μ_o and standard deviation σ_o.[7] When $\eta > 0$, the individual has a job offer, which can be lost next period with probability δ_k. If the job offer is not lost, another job offer η' is found with probability $\lambda_{1,k}$, with η' drawn from the same distribution π_o. This offer is accepted only if $\eta' > \eta$.

When living in a policy environment with time limits, cumulative welfare use must be tracked as:

$$\omega' = \max\{\Omega(\chi), \omega + A_j\}.$$

When an individual reaches their time limit ($\omega = \Omega(\chi)$) they are considered ineligble for welfare, implying that they receive a payment equal only to their food stamp benefit. More details on these calculations can be found in Appendix C.

3.2 Modeling the Welfare Experiments

The state variable $Z \in \{0, 1, 2\}$ indicates whether an individual has been assigned to treatment, and it effects the model environment through four channels. First, assignment to treatment ($Z > 0$) leads to a change in the calculation of welfare benefits. This is reflected in a change to the net income function Y_k through the transfer function $\mathcal{T}$. Appendix C describes the benefit formulae for all control and treatment groups. This treatment variation is well-articulated: it appears through known a priori changes to the budget set.

Second, assignment to treatment may entail mandatory employment services (indicated by $R(X, Z) = 1$). This has two potential effects. First, an additional term is added to the utility derived from choice j:

$$-\alpha_R(1 - H_j)A_j.$$

This term makes not working while participating in welfare more costly. It can be thought of as an ordeal that does not contribute to labor market outcomes. Alternatively, for individuals who are participating ($A_j = 1$)

[7] The parameterization is: $\pi_o(\eta) = [\Phi((g(\eta) - \mu_o)/\sigma_o) - \Phi((g(\eta-1) - \mu_o)/\sigma_o)]/\Phi((g(K_\eta) - \mu_o)/\sigma_o)$ with $\pi_o(1) = [\Phi((g(1) - \mu_o)/\sigma_o)/\sigma_o)]/\Phi((g(K_\eta) - \mu_o)/\sigma_o)$.

but do not have a job offer ($\eta = 0$), these services may increase the probability of finding a job for the next period to $\lambda_{k,0}(R) > \lambda\_0, k$. This is captured by a parameter λ_R, which adds to this probability according to:

$$\log\left(\frac{\lambda_{k,0}(R)}{1-\lambda_{0,k}(R)}\right) - \log\left(\frac{\lambda_{k,0}}{1-\lambda_{0,k}}\right) = \lambda_R.$$

Third, treatment may lower the cost of childcare by expanding access to subsidies. Any known subsidy to childcare is well-articulated in this model, but the expansions in the experiments do not amount to a priori known changes in prices. Hence, these effects are captured by the γ parameters in the formula for childcare costs, $q_k(\chi)$.

Fourth, treatment at two sites involves time-limited participation in welfare. This is a known a priori change to individuals' dynamic choice sets (as described in the section above) and is hence well-articulated. The remainder of this section summarizes the treatment components for each site.

CTJF Relative to the control group, a time limit of 7 quarters is introduced along with mandatory employment services if not working (a work requirement). Changes to benefit formula are such that all income is disregarded from benefit and food stamp payments, up until the poverty guidelines, at which point benefits are reduced to zero (see Appendix C).

FTP Relative to the control group, a time limit of 8 quarters is introduced along with mandatory employment services if not working (a work requirement). Changes to benefit formula are such that the first \$200 if earned income is disregarded with 50% of income thereafter (see Appendix C).

MFIP Food stamps and welfare are rolled into one cash grant, and the payment standard is increased by 20%, resulting in a substantial increase in the amount of initial income that is disregarded (see Appendix C). In treatment arm 1, participants faced mandatory employnet services if not working (a work requirement). In treatment arm 2, participants were exempt from this requirement.

3.3 Measurement Error

If either labor market earnings or childcare expenditures are positive for individual m, they are observed with additive and independent normal measurement error:

$$W^o_{m,t} = W_{k(m)}(\chi) + \xi_{W,m},\ q^o_{m,t} = q_{k(m)}(\chi) + \xi_{q,m},\ \xi_{j,m} \sim \mathcal{N}(0, \sigma^2_j$$

Likewise, skills θ are measured with normal additive error. The forthcoming estimation exercise will alternatively model academic and behavioral skills, and accordingly the measurement system is:

$$S_{l,m} = \mu_l + \lambda^l_B \log(\theta)_{B,m}) + \lambda^l_C \log(\theta)_{C,m} + \varepsilon^l_m, l \in \{1, 2, .., 8\}$$

where ε^l has variance σ^2_l and is independent across measures and observations. Section 2 discusses the 8 available measures of skills, and Section 4 clarifies the restrictions on the factor loadings (λ) that are necessary for identification.

3.4 Model Solution

The dynamic optimization problem can be phrased recursively as:

$$V_k(\chi, \theta) = \mathbb{E}_\epsilon \max_{C,I,j} \left\{U_{kj}(C, \theta) + \epsilon_j + \beta\mathbb{E}_{\theta',\chi'|\chi,j,I} V_k(\chi', \theta')\right\}$$

subject to the transition rules, technology, and constraints described in Section 3.1. The state space described by χ is, in principle, large. A number of simplifications make the problem tractable. First, Appendix A shows that the function is additively separable in $\log(\theta)$ as in Del Boca et al. (2014) and Mullins (2022):

$$V_k(\chi,\theta) = \tilde{V}_k(\chi) + \alpha_{V,t,k}\log(\theta)$$

where $\alpha_{V,t,k} = \alpha_{\theta,k} + \beta\delta_\theta\alpha_{V,t+1,k}$. The simplified dynamic problem defining $\tilde{V}$ is:

$$\tilde{V}_k(\chi,\theta) = \mathbb{E}_\epsilon \max_j \left\{\tilde{U}_{kj}(\chi) + \epsilon_j + \beta\mathbb{E}_{\chi'|\chi,j}\tilde{V}_k(\chi')\right\}$$

where $\tilde{U}$ is an indirect utility function derived from the solution to the investment problem being linear in income:

$$I = \frac{\alpha_{\theta,k}\Gamma_t}{1+\alpha_{\theta,k}\Gamma_t}\frac{Y_{kj}(\chi)}{p_{H_j+F_j}}$$

which gives:

$$\tilde{U}_{kj}(\chi) = (1+\alpha_{\theta,k}\Gamma_t)\log(Y_{jk}(\chi)) + \tilde{\alpha}_{j,k}(\chi)$$

and $\tilde{\alpha}_{j,k}$ is a choice-specific utility term given by:

$$\tilde{\alpha}_{j,k} = (-\alpha_{H,k} + \alpha_{\theta,k}\Gamma_t g_1)H_j - \alpha_{S,k}S_j - (\alpha_{A,k} + \alpha_P(1-A_{-1}))A_j + (\alpha_{F,k} + \alpha_{\theta,k}\Gamma_t g_2)F_j.$$

This expression is simply the non-pecuniary choice payoffs from static utility with two extra terms that capture how the individual values the dynamic effect that these choices have on future realizations of their child's skills. For example, the log of the inverse price ratio g_1 captures the short-run developmental benefit (or cost) of part-time unpaid care relative to full-time maternal care. The term Γ_t (derived in Appendix A) scales this short-run effect into a discounted present value payoff, while the heterogeneous term $\alpha_{\theta,k}$ scales the payoff according to how much type k values skill outcomes. The term $\alpha_{\theta,k}\Gamma_t g_2$ has an identical interpretation.

To further simplify the model solution, note that while χ has many dimensions, only three are either stochastic or endogenous to the model, these are η (job opportunities), ω (time limits), and A_{-1} (lagged welfare participation). Now consider one observation in the data, which fixes an initial age of the mother, state of residence, and sequences of policy variables, unemployment rates, and family sizes. Assuming these sequences are known to the agent the model can be solved over the quadruple (η,ω,A_{-1},k) by backward induction on t. Thus, the size of the problem is linear in sample size rather than equal to the full space of cohorts, locations and fertility sequences (which is intractably large).

4 Empirical Content and Identification

This section has three objectives. They are (1) to establish identification of the parameters of the model; (2) in doing so, demonstrate how experimental variation contributes to identification of key parameters and, (3) conversely, to show how these parameters determine treatment responses.

Let $m \in \{1,...,M\}$ index sample mothers in the data, as described in Section 2. Each observation m consists of a collection $(\boldsymbol{y}_m, \boldsymbol{X}_m, \mathcal{S}_m, Z_m)$ where: $\boldsymbol{y}_m = \{y_{m,t}\}_{t=1}^{T_m}$ is a quarterly panel of participation choices, earnings, and childcare choices and expenditures (when observed); $\boldsymbol{X}_m$ is a vector indicating the data source, state of residence, family size, education, mother's initial age, initial age of youngest child, and initial calendar year;[8] $\mathcal{S}_m$ is an 8-dimensional vector of skill measures for a focal child in the household; and $Z_m \in \{0,1,2\}$ indicates treatment status.

[8]Time-varying observables such as state unemployment rates and ages are assumed to be known functions of these initial observables

Identification of the model's key parameters proceeds in three steps. The first step relates to identification of choice probabilities, prices, and transition probabilities for η (which indexes labor market opportunities). With these in hand, additional parametric restrictions on choice probabilities that are implied by the model yield identification of preference parameters. In a final step, the model admits an outcome equation for skills that provides straightforward sources of identification for procuction parameters. Through steps (2) and (3), the identification analysis in this section also exposits the role that particular parameters play in determining treatment effects.

4.1 Identification of Choice Probabilities, Price Parameters, and Transition Parameters

Fixing the obervable components of χ, the earnings and choice outcomes adopt the structure of a hidden markov model with one permanent latent state (type k) and one time-varying latent state (η). Identification of outcome and transition probabilities in this setting may come from various sources, including (1) the presence of observables that sufficiently shift outcome probabilities (Kasahara and Shimotsu, 2009); (2) outcome probabilities being sufficiently informative (as defined by a rank condition on the emission matrix) with respect to the hidden state (Bonhomme et al., 2017); (3) or relatively weaker completeness conditions with a sufficiently long panel (Hu and Shum, 2012). Because of (1) substantial variation in state-specific policies, local unemployment rates, and randomly assigned welfare policies; (2) 9 observed choices and continuously distributed observed earnings; and (3) panels at least as long as 16 periods, identification in this setting may be achieved from any one of these three sources.

4.2 Identification of Preference Parameters

Let $p_{kj}(\chi)$ denote the probability that an agent of type k makes choice j when in state χ. The previous argument identifies these choice probabilities, $\boldsymbol{p}$, along with price and distribution parameters:

$$\Phi_1 = \{\mu_W, \beta_W, \sigma_\eta, \mu_q, \beta_q, \qquad \lambda_0, \lambda_1, \delta, \lambda_R\}.$$

Let Φ_2, Φ_3 be two sets of parameters that, in combination with the above, dictate behavior (and therefore treatment responses):

$$\Phi_2 = \{\alpha_A, \alpha_S, \alpha_H, \alpha_P\}$$

$$\Phi_3 = \{\alpha_\theta, \Gamma_t, \alpha_R, \beta, \sigma_1, \sigma_2, \sigma_3, g_1, g_2\}$$

Finally, in order to distinguish between between non-experimental and experimental sources of variation, let the state χ be decomposed into the pair (X, Z). Appendix B uses the properties of nested logit choice probabilities and the model's finite dependence (Arcidiacono and Miller, 2011) to derive a system of equations with the following structure:

$$\boldsymbol{h}_0(X, Z; \boldsymbol{p}) = \boldsymbol{\kappa}_0(\Phi_2, \Phi_3) + \boldsymbol{\kappa}_1(\Phi_3)\boldsymbol{h}_1(X, Z; \boldsymbol{p}, \Phi_1) \tag{4.1}$$

where $(\boldsymbol{h}_0, \boldsymbol{h}_1, \boldsymbol{\kappa}_0, \boldsymbol{\kappa}_1)$ are known functions. With this structure, identification of the parameters $\boldsymbol{\kappa}_0$ and $\boldsymbol{\kappa}_1$ (from which Φ_1 and Φ_2 can be inverted) is guaranteed by straightforward conditions on sufficient variation (i.e. rank independence) in $\boldsymbol{h}_1$. In principle, all variation coming from (X, Z) is valid to identify Φ_3, but the expressions also reveal that variation in Z only is sufficient. This in turn illustrates not just how the parameters in Φ_3 shape treatment effects, but also outlines an alternative minimum distance estimator if a researcher wished only to use experiment variation to identify these key parameters. The parameters in Φ_2 accordingly can be seen as "intercept" parameters that pin down levels, although this can only be said to be true for the transformation $\boldsymbol{h}_0$, not for the choice probabilities generally.

Appendix B presents each expression in the above linear system, but one example presented here will help with exposition. Define

$$\tilde{p}_{k,F}(P,X,Z) = \frac{p_k(F=1|H=1,P,\chi)}{p_k(F=0|H=1,P,\chi)}$$

as the relative choice probability of using paid care in state χ, conditional on working and conditional on welfare participation.[9] Similarly define

$$\tilde{Y}_{k,F}(X,Z) = \frac{Y_k(H=1,P=1,X,Z) - q_k(X,Z)}{Y_k(H=1,P=1,X,Z)}$$

as the cost of childcare expressed as a proportional reduction in disposable income. Fixing the participation choice and the work choice places us in the lowest layer of the nesting structure depicted in Figure 3.1, yielding:

$$\underbrace{\log(\tilde{p}_k(F,X,Z))}_{h_0} = \underbrace{\sigma_3^{-1}(\alpha_{F,k} + \alpha_{\theta,k}\Gamma_t g_2)}_{\kappa_0} + \underbrace{\sigma_3^{-1}(1+\alpha_{\theta,k}\Gamma_t)}_{\kappa_1}\underbrace{\left(\tilde{Y}_k(X,Z)\right)}_{h_1} \tag{4.2}$$

Note how this expression can be mapped into the first components of (4.1). Also note that variation Z, fixing X, is sufficient to identify the slope parameter as long as assignment to treatment varies the proportional cost of using paid care, which can be achieved either through subsidies or benefit formulae.

The other expressions that comprise (4.1) take a very similar form, and rely largely on the relationship between choice comparisons and the relative financial returns to those choices. However due to the nested logit structure and the need to difference out dynamic values, choice probabilities also appear and contribute to identification. While Appendix B explores the expressions in exact detail, the main lessons overall is that each of σ_1, σ_2, and σ_3 play a role in determining the response of individuals to changes in the financial payoffs to childcare choices, work, and participation respectively. Furthermore, the discount factor β can be related directly to the effect of time limits. In a maximum likelihood routine (which this paper employs) any variation in budget sets and choice probabilities will contribute to identification, but note that a minimum distance estimator that integrated these expressions over X would identify the parameters in a way that relies only on treatment responses.[10]

4.3 Identification of Production Parameters

Skill outcomes are measured only once in the experimental data at either three (MFIP, CTJF) or four (FTP) years after follow up. Combining the Cobb-Douglas production function with the linear investment rules yields a linear outcome equation for child skill θ at the time of measurement, t^*:

$$\log(\theta_{t^*}) = \sum_{t=1}^{t^*} \delta_\theta^{t^*-t}\delta_I\left(\log(Y_t) + g_1 H_t + g_2 F_t\right) + \delta_\theta^{t^*}\log(\theta_1)$$

where Y_t is net income, H_t indicates employment, and F_t indicates paid care use. To simplify exposition, let us represent this as:

$$\tilde{\theta}_1 = \boldsymbol{I}B(\delta,g) + \tilde{\theta}_0$$

[9] For the argument to work, any choice of P is valid

[10] Appendix B provides more details on this estimator.

where $\boldsymbol{I}$ is the vector of inputs in each time period suitably arranged and B is the vector of coefficients on each input. Let l indicate the experiment location of an observation. Random assignment implies that:

$$\mathbb{E}[\tilde{\theta}_0|l, Z] = C_l.$$

This yields a moment condition:

$$\mathbb{E}[\tilde{\theta}_1 - \boldsymbol{IB}(\delta, g) - C_l|l, Z] = 0 \tag{4.3}$$

that forms the basis of an instrumental variables estimator for the production parameters $(\delta_\theta, \delta_I, g_1, g_2)$. Identification relies on straightforward rank conditions, in particular that there are sufficiently many location-treatment combinations to move net income, work, and child care use in a rank-independent fashion. In contrast to related papers, the coefficient δ_θ is not identified by persistence in skills, but rather the relative influence of inputs at different lags. If inputs in the far past are found to be almost as important as inputs in the previous period, this implies a value of δ_θ close to 1.

Using the structure of the model implies a moment condition that capitalizes on more variation in inputs. In particular, letting $(k, \eta_1, \boldsymbol{X})$ be a triple representing type, the initial value of labor market opportunities, and observable initial conditions, the Markov structure of the model implies:

$$\mathbb{E}[\tilde{\theta}_0|\boldsymbol{I}, k, \eta_1, \boldsymbol{X}] = \mathbb{E}[\tilde{\theta}_0|k, \eta_1, \boldsymbol{X}] = \mu(k, \eta_1, \boldsymbol{X})$$

yielding a moment condition that relies on the control function μ:

$$\mathbb{E}[\tilde{\theta}_1 - \boldsymbol{I}B(\delta, g) - \mu(k, \eta_1, \boldsymbol{X})|\boldsymbol{y}, \boldsymbol{X}, Z] = 0 \tag{4.4}$$

In other words, conditional on location and the initial conditions k and η, all remaining variation in inputs (which is driven by future realizations of η and preference shocks) can be used to identify production parameters. Section 5.3 describes how these moment condtions are imeplemented in practice, using Bayesian Instrumental Variables.

5 Estimation

Estimation proceeds in two stages. Sections 3 and 4 showed parents' consideration for skill development when making their decisions is embodied by the terms $\alpha_{\theta,k}\Gamma_t$, g_1, and g_2 and that these are identified directly by age-dependent choice probabilities. Although the expression Γ_t is a known function of deeper technological parameters, it is simpler to not impose those cross-equation restrictions and simply estimate age-dependent preferences from the data.[11] Thus, a second stage estimates production parameters following estimation of all other parameters in the first stage.

5.1 Behavioral Parameters and Initial Conditions

Some additional parameters specify the distribution of initial conditions, which is necessary in order to evaluate the likelihood. The probability that an individual is of type k, conditional on observables measured at the beginning of the sample, takes a multinomial logit specification:

$$\mathbb{P}\left[k(m) = k|X_{\tau,m}\right] = \frac{\exp\left(X_{\tau,m}\beta_{\tau,k}\right)}{1 + \sum_{j=2}^{K}\exp\left(X_{\tau,m}\beta_{\tau,j}\right)}$$

[11]This does require a normalization of the scale of Γ_t, which is achieved by imposing $\Gamma_0 = 1$. In practice, it is approximated as a polynomial in age, $\Gamma_t = \exp(\beta_{\Gamma,1}t + \beta_{\Gamma,2}t^2)$.

Table 5.1: Preference Parameter Estimates

	Type-Specific Parameters					
Type	α_H	α_A	α_S	α_F	α_θ	y
$k=1$	-0.63	-0.34	-0.66	0.18	0.32	229.71
	(0.12)	(0.03)	(0.07)	(0.22)	(0.00)	(38.63)
$k=2$	-1.01	-0.46	0.21	0.85	0.28	627.79
	(0.15)	(0.03)	(0.02)	(0.28)	(0.00)	(69.32)
$k=3$	-0.97	-1.14	0.95	1.17	0.23	1105.14
	(0.15)	(0.08)	(0.08)	(0.34)	(0.00)	(122.96)
$k=4$	0.10	-1.04	1.05	-0.99	0.18	147.67
	(0.19)	(0.08)	(0.09)	(0.12)	(0.00)	(22.71)
$k=5$	-1.07	-1.82	1.73	0.20	0.17	1979.23
	(0.20)	(0.16)	(0.16)	(0.16)	(0.00)	(308.73)
	Global Parameters					
	β	σ_3	σ_2	σ_1	α_R	α_P
	0.34	1.16	0.56	0.34	0.04	1.63
	(0.05)	(0.24)	(0.04)	(0.03)	(0.02)	(0.15)

This table presents estimates of the model's preference parameters from the maximum likelihood procedure. Parentheses report standard errors.

with $\beta_{\tau,1}$ normalized to zero. The vector $X_{\tau,m}$ is a set of experiment location dummies interacted with a constant, two education dummies, household size, county dummies, and a dummy indicating if the individual is a new applicant. The initial distribution for η, $\pi_{\eta,0}$, is specified non-parametrically conditional on type, applicant status, and experimental site. For sample observations from the SIPP, η is assumed to be initially drawn from the stationary distribution.

The first stage estimates all parameters dictating preferences, prices, transitions, and initial conditions via Maximum Likelihood. An Expectation-Maximization routine produces these estimates and can tractably handle the many latent state variables. Appendix D provides more details and write the likelihood formally. For estimation, the model takes values of K (the number of types) equal to 5 and K_η (the number of grid points for η in adition to $\eta = 0$) equal to four.

5.1.1 Discussion of Estimates

Tables 5.1, 5.2, and 5.3 each respectively report estimates of parameters governing preferences, transitions, and prices. While the particular interest of this paper is in specific counterfactuals that depend on all parameters in combination, rather than any single parameter, several specific results are worth discussion. First, recall that the model permits latent heterogeneity in participation costs $(\alpha_A, \alpha_H, \alpha_S, \alpha_F)$, labor market productivity $(\mu_{W,k})$, and labor market risk $(\lambda_0, \lambda_1, \delta)$. According to these estimates, the data suggest very rich heterogeneity in each of these dimensions. Of particular note are striking differences in productivity (see the type coefficients in Table 5.3) and risk (Table 5.2). For example, estimates suggest that although

Table 5.2: Transition Parameter Estimates

	Type-Specific Parameters		
Type	λ_0	λ_1	δ
$k=1$	0.15	0.53	0.11
	(0.00)	(0.03)	(0.00)
$k=2$	0.19	0.74	0.10
	(0.01)	(0.04)	(0.00)
$k=3$	0.16	0.48	0.27
	(0.00)	(0.03)	(0.01)
$k=4$	0.06	0.66	0.02
	(0.00)	(0.03)	(0.00)
$k=5$	0.08	0.24	0.05
	(0.00)	(0.02)	(0.00)
	Global Parameters		
	μ_o	σ_o	λ_R
	-0.43	0.88	0.42
	(0.02)	(0.02)	(0.04)

This table presents estimates of the model's transition parameters from the maximum likelihood procedure. Parentheses report standard errors.

type 5 is subsantially more productive in the labor market than type 1, they find job opportunities while unemployed at half the quarterly rate, and find new job opportunities while working at half the quarterly rate. This implies that while these workers accept higher wages initially, they enjoy slower earnings growth while working. Moreover, conditional on having a job, they are half as likely to lose it compared to type 1 individuals. Similarly, while type 1 and type 3 workers are equally productive, type 3 workers are twice as likely to lose their job and therefore face more substantial wage risk.

Also of note are the parameters α_R and λ_R, which embody the two mechanisms through which mandatory services may effect behavior: work requirements could simply make not working more costly while participating in welfare, or they may assist in job search by increasing the rate at which opportunities are found. Estimates in Table 5.1 and 5.2 indicate that these services work through the latter channel, and in fact quite substantially increase the rate at which new job opportunities arrive.

Finally, note the extremely high rate of discounting implied by an estimate of β equal to 0.34 (Table 5.1). It implies an annual rate of discounting close to 99%, much higher than typically calibrated values of 2-5%. Such a high rate of discounting may be due to present bias (Chan, 2017) or lack of access to credit markets.

5.1.2 Examining Selection Across Sites

An important feature of the welfare-to-work experiments is that they selected explicitly on the application status of individuals: the sampling frame was all individuals applying or re-applying for cash assistance. While this is a natural criterion for evaluating welfare-to-work, it creates a challenge to external validity if

Table 5.3: Price Parameter Estimates

	Wages	Childcare
Type 1	5.62	5.05
	(0.02)	(0.23)
Type 2	5.92	4.83
	(0.02)	(0.27)
Type 3	5.64	4.30
	(0.02)	(0.43)
Type 4	6.27	5.33
	(0.02)	(0.23)
Type 5	6.83	5.44
	(0.02)	(0.28)
Unemployment Rate	-0.06	-0.03
	(0.00)	(0.03)
Age	0.00	-
	(0.00)	
Num. Kids	-	0.12
	-	(0.02)
Youngest ≤ 5	-	0.04
	-	(0.05)
FTP Control	-	1.35
	-	(0.17)
FTP Treat	-	1.37
	-	(0.17)
CTJF Control	-	1.29
	-	(0.17)
CTJF Treat	-	1.32
	-	(0.18)
MFIP Control	-	1.22
	-	(0.18)
MFIP Treat	-	1.33
	-	(0.18)
MFIP Incentives	-	1.36
	-	(0.18)
Measurement error (std. dev)	0.60	1.50
	(0.00)	(0.13)

This table presents estimates of the model's price parameters from the maximum likelihood procedure. Parentheses report standard errors.

Figure 5.1: Evidence of Selection in Experiments

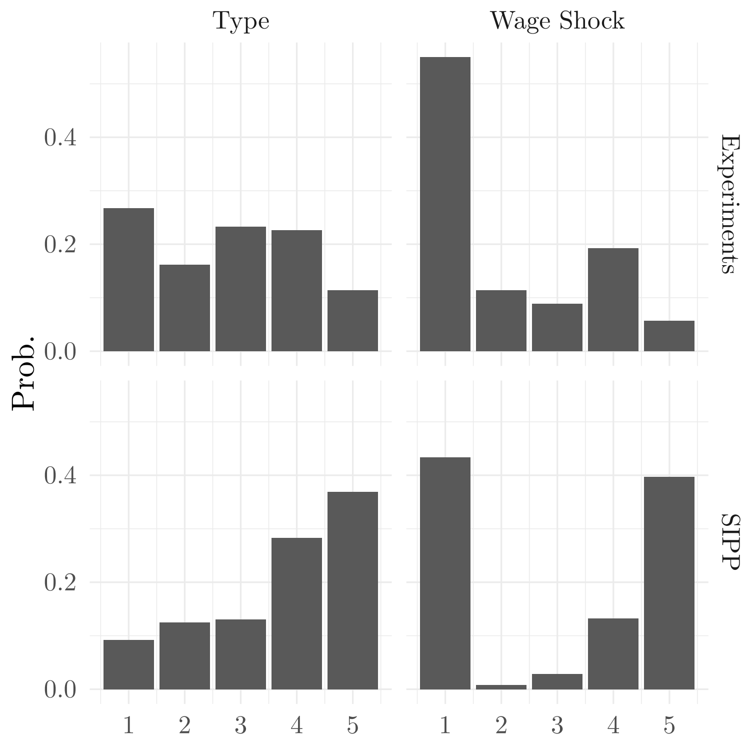

This figure depicts the frequency distribution of types (k) and initial wage shock (η) for the SIPP and for all experimental samples. These are weighted frequencies using the posterior weights over latent states at the estimates. For example: $\hat{\mathbb{P}}[\eta_1 = k_\eta|\text{Experiments}] = \sum_m q_{m1}(k_\eta)/\left(\sum_m \sum_\eta q_{m1}(\eta)\right)$ where the sum is taken over all observations m in the three experimental datasets, and the posterior weight q_{m1} is defined in Appendix D.

one wants to forecast the effect of the same policies on the population at large. Relative to this population, the experimental sampling frame is likely to oversample particular individuals at specific points in their lives. By flexibly estimating initial probabilities and by including data from the SIPP, the estimated model can speak directly to the extent of this selection. Table E.1 in Appendix E reports the estimated coefficients that dictate initial type probabilities. Due to the large dimension of $\pi_{\eta,0}$ the initial distribution of job opportunities, Figure 5.1 aggregates these joint distributions across sites to provide a simpler comparison to the SIPP sample. The Figure documents quite striking patterns of selection by both type and job opportunities. The experiments oversample low types and individuals experiencing temporarily low wage opportunities.

5.1.3 In Sample Model Fit

Figure 5.2 shows that the model fits choice probabilities and earnings outcomes very well across sites and groups of applicants. One feature of interest is that the model can replicate the "U-shape" treatment effect in employment outcomes that is common to many welfare-to-work evaluations that involve mandatory employment services. All of these studies tend to show an initial increase in employment that fades out over time (Greenberg et al., 2004). This is consistent with these services increasing the job-finding rate for those participating in welfare, leading to an initial increase in employment that fades out over time as control group

Figure 5.2: In-sample Model Fit

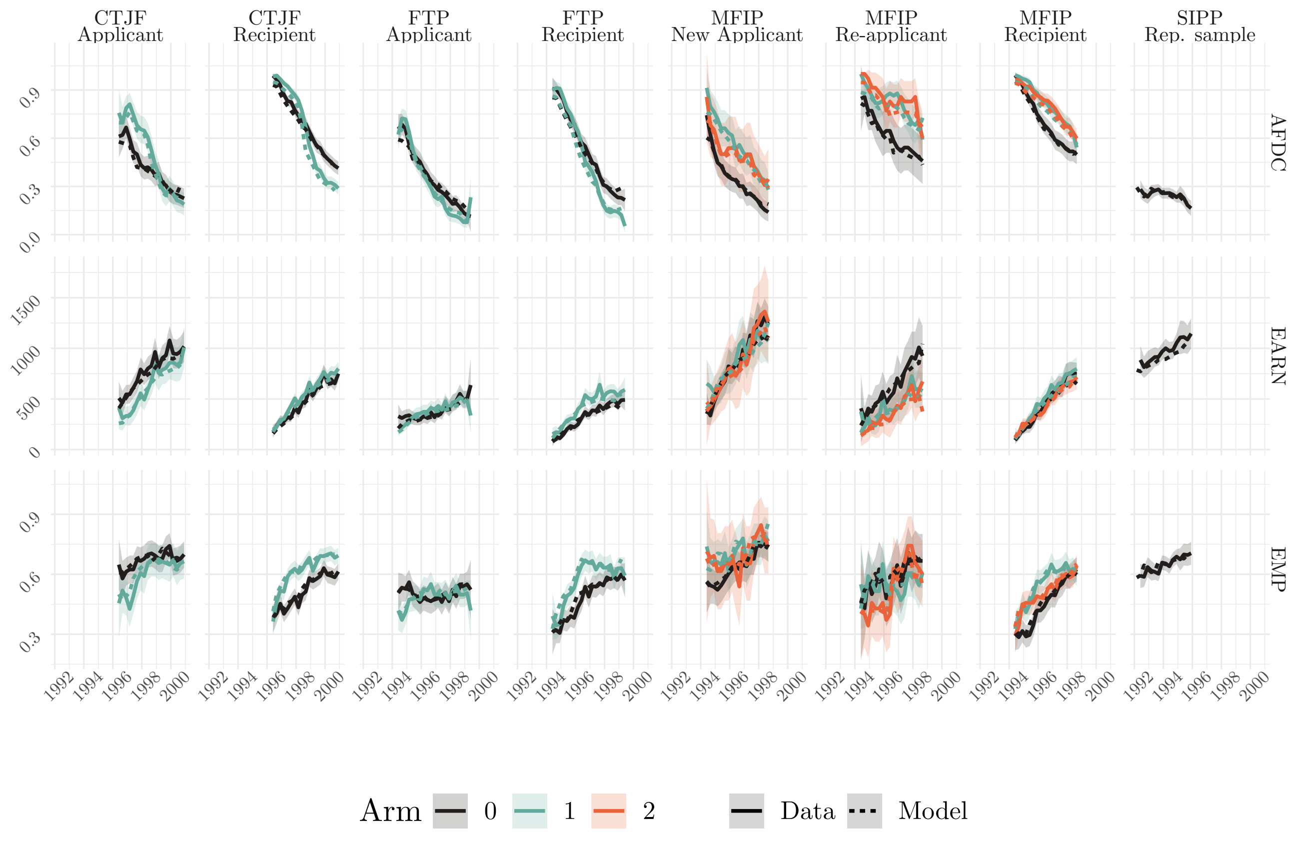

This figure shows the model's "ex-post" fit of choice probabilities and earnings. The top panel shows the fit of quarterly welfare participation (AFDC), the middle panel shows monthly earnings in year 2000 \$USD, and the bottom panel shows quarterly employment rates. Each is calculated using the posterior distribution over latent states given the estimates. For example, define $q_{m,t}(\chi, k) = \mathbb{P}[\chi, k|y_m, X_m, \hat{\Theta}]$, then predicted employment for group G in period t is equal to: $\sum_{m\in G}\sum_{k,\chi}\sum_j \mathbf{1}\{E_j = 1\}p_{kj}(\chi)q_{m,t}(\chi, k)/\left(\sum_{m\in G}\sum_{k,\chi} q_{m,t}(\chi, k)\right)$

Figure 5.3: Out-of-sample Model Fit

This figure compares average welfare participation (AFDC), earnings, and employment for new applicants in the treatment gruop in Manchester (CTJF), Anoka (MFIP), and Dakota (MFIP) counties. Ribbons indicate 95% confidence intervals for population means in each period.

members also eventually find work, albeit at a slower rate than those with access to services.

5.1.4 Out of Sample Model Fit

Following the literature (Todd and Wolpin, 2023), the estimation sample excludes members of treatment groups who are new applicants in the CTJF and MFIP experiments from Manchester, Anoka, and Dakota counties. Figure 5.3 validates the model by comparing its prediction for employment, welfare participation, and earnings for members of these three out-of-sample groups. These predictions are largely within the 95% confidence intervals of each group and match the patterns of outcomes very well.

5.2 Assessing the Contribution of Experimental Variation

The identification analysis in Section 4 argued that treatment variation is particularly crucial for identification of the discount factor β as well as treatment specific parameters α_R and λ_R which captured treatment components that are not otherwise well-articulated (i.e. replicated by equivalent variation) inside the model.

Figure 5.4: Predicted Treatment Effects for Full Sample vs Control Group Only

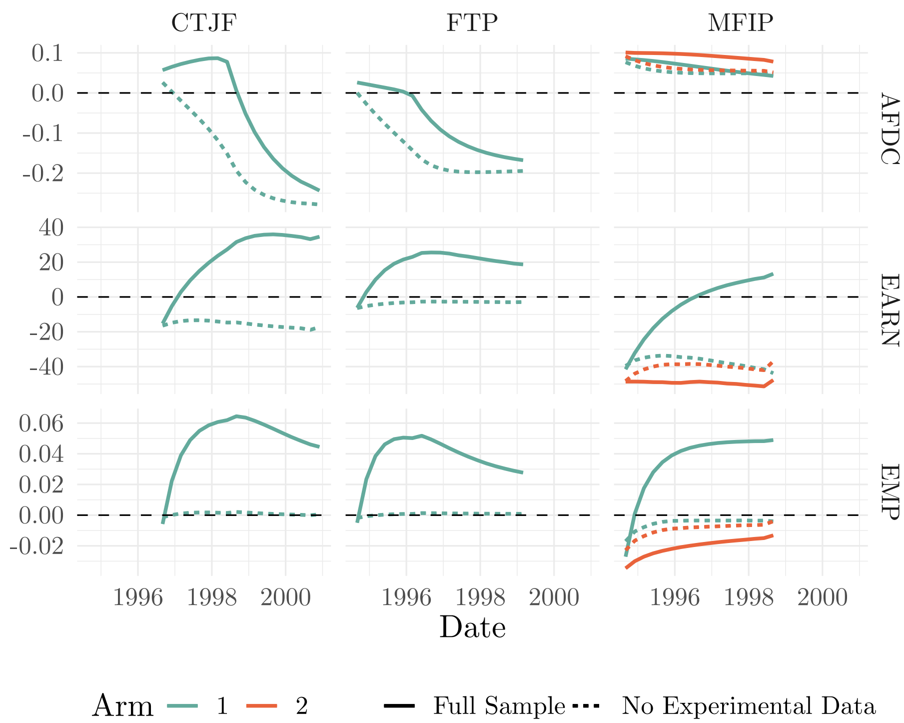

This figure compares the model's predicted treatment effects on welfare participation (AFDC), earnings, and employment for two sets of estimates. The first set is the full sample described in the paper and used elsewhere in quantitative exercises. The second set relies on estimates taken only from the control groups for the experiments and from the SIPP.

Furthermore, the analysis showed that treatment effects should contribute to the identification of the dispersion parameters $(\sigma_1, \sigma_2, \sigma_3)$ which dictate the responsiveness of individuals to changes in welfare generosity, work incentives, and childcare prices. In order to establish the important role played by experimental variation, this section presents results from an estimation exercise that uses *only* the control group and compares predictions for treatment effects as well as these key parameter estimates.

Figure 5.4 compares predicted treatment effects for each site and treatment arm from the baseline estimated model to the version estimated only on control groups. Most strikingly, there are large differences in predictions for participation in each treatment with time limits (CTJF and FTP) and for earnings and employment in treatments with work requirements (CTJF, FTP, and arm 1 of MFIP). An obvious conclusion is that for predicting treatment effects the experimental variation in the data is quite crucial. The differences in prediction for arm 2 of MFIP, which features only changes in the benefit formula, are much less pronounced, though still evident. Nonetheless, the existence of articulated variation in this dimension appears to reduce the need for experimental variation to predict policy effects.

Table 5.4 reports estimates of the parameters that the identification argued were key for responses to treatment. Echoing the large differences in treatment response to time limits, we see a large difference in β when estimated using treatment data (0.34) compared to without (0.94). Section 4 argued that this parameter was key for determining the response to time limits. Table 5.4 also reports quantitatively important differences

Table 5.4: Key Parameter Estimates for Full Sample vs Control Group Only

	β	σ_3	σ_2	σ_1
Control Group Only	0.94	1.86	0.95	0.78
	(0.00)	(0.13)	(0.03)	(0.01)
Full Sample	0.34	1.16	0.56	0.34
	(0.05)	(0.24)	(0.04)	(0.00)

This figure compares model estimates of key elasticity parameters for two sets of estimates. The first set is the full sample described in the paper and used elsewhere in quantitative exercises. The second set relies on estimates taken only from the control groups for the experiments and from the SIPP.

in the dispersion parameters σ, suggesting that the experimental variation makes an important contribution to identification of these parameters, even though there is plenty of existing variation in control group data to identify them.

5.3 Production Parameters

Adopting a scale normalization that $\mathbb{V}[\tilde{\theta}_B] = \mathbb{V}[\tilde{\theta}_C] = 1$ in the CTJF population, a minimum distance routine estimates the measurement parameters $(\lambda_B^k, \lambda_C^k, \sigma_k^2)_{k=1}^8$ by minimizing the distance between estimated variance-covariance matrices for the vector of skills $\mathcal{S}$ at each experiment site and those implied by latent variances, covariances, and measurement parameters. Table 5.5 reports the estimates from this procedure. The estimates suggest that the dedicated measurements each load meaningfully on their respective skills, while the composite measurements (engagement, grade repetition, and suspensions) appear to load meaningfully on both skills.

The estimated parameters then allow for the computation of a pair of factor scores (S_C, S_B) that are each equal to $(\tilde{\theta}_C, \tilde{\theta}_B)$ with an independent disturbance term. Making this substitution for $\tilde{\theta}$ in equation (4.3) gives:

$$\mathbb{E}[S_j - \boldsymbol{I}B(\delta, I) - C_l | Z, l] = 0.$$

Similarly, equation (4.4) becomes:

$$\mathbb{E}[S_j - \boldsymbol{I}B(\delta, g) - \mu(k, \eta_1, \boldsymbol{X}) | \boldsymbol{y}, \boldsymbol{X}, Z] = 0.$$

To reiterate the arguments made in Section 4.3, estimation with the former moment condition uses only variation in inputs driven by assignment to treatment at the different experiment locations, while the latter moment condition uses any variation in inputs conditional on initial conditions determined by $(k, \eta_1, \boldsymbol{X})$. In the model this variation comes from realizations of η (job opportunities) and ϵ (preference shocks).

A Bayesian approach to using these moment condition enables the imposition of some weak priors on the parameters $\delta_I, \delta_\theta, g_1, g_2$, which circumvents well-known issues with weak instruments. It is also useful for quantifying the information (by comparing posterior to prior) that is present in the particular source of

Table 5.5: Parameter Estimates for the Measurement of Skills

Measure	λ_B^m	λ_C^m	σ_m^2
BPI-Externalizing	-3.70	-	5.86
	(0.14)		(1.23)
BPI-Internalizing	-2.32	-	7.55
	(0.08)		(0.91)
Positive Behavior Scale	6.49	-	63.34
	(0.23)		(23.84)
School Engagement	0.16	0.96	2.03
	(0.14)	(0.14)	(0.10)
Ever Repeat Grade	0.79	-0.80	0.04
	(0.25)	(0.25)	(0.00)
Ever Suspended	-0.05	-0.05	0.08
	(0.02)	(0.02)	(0.00)
School Achievement - Parent	-	0.51	0.72
		(0.03)	(0.02)
School Achievement - Teacher	-	0.42	1.29
		(0.06)	(0.09)

This table reports minimum distance estimates of the measurement parameters introduced in Section 4.3. The minimum distance criterion uses the full set of variances and covariances with a diagonal weighting matrix containing the inverse of bootstrapped variances for each statistic. Parentheses report standard errors, which are calculated using a bootstrapping routine with 100 replacement samples.

Figure 5.5: Posterior Distributions for δ_I

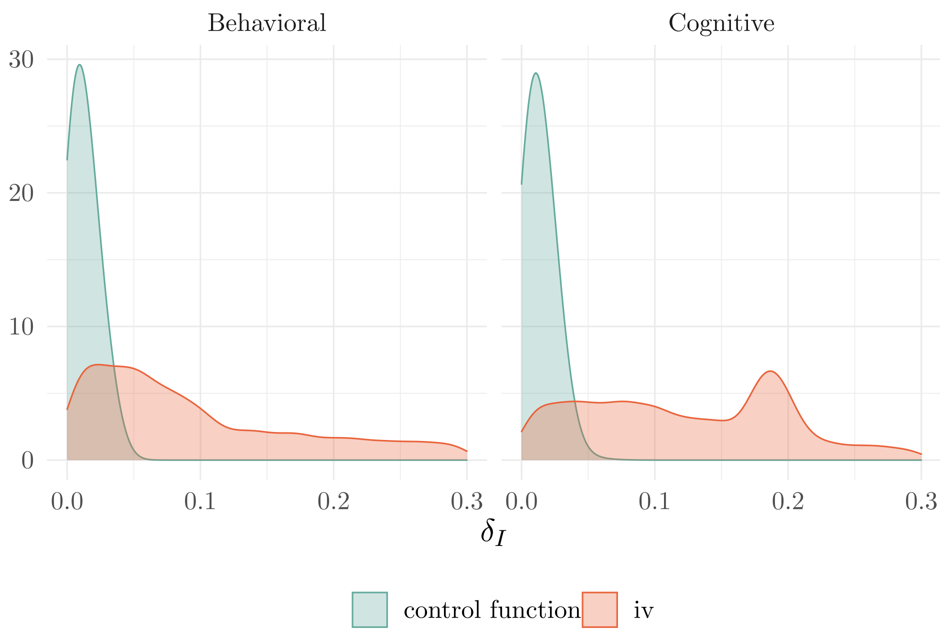

This figure shows posterior distributions for δ_I using the two approaches described in the text. Distributions are calculated from 10,000 samples drawn from the No-U-Turn Hamiltonian Monte Carlo Sampler (Hoffman et al., 2014). See Appendix D.2 for more details.

variation being used. Appendix D.2 provides more details on how to adapt these moment conditions into a Bayesian setting, on sampling techniques, and on the priors specified for key parameters.

Figure 5.5 depicts the posterior distributions for δ_I, which determines the effect of net income on future skills, for each of the two estimation strategies. The IV estimator produces larger estimate for δ_I with much more uncertainty. The averages of the posteriors for behavioral skills, for example, are 0.01 and 0.21 respectively. While this is a large discrepancy, the control function estiamate still lies comfortably within the set of likely values according to the IV posterior. It is not straightforward to benchmark either value with prior estimates in the literature, but some back of the envelope calculations are available. A value of $\delta_I = 0.2$ implies that an increase in net monthly income from \$400/month to \$480/month would result in a 3.5% of a standard deviation increase in behavioral skills over the quarter. Expressed in annual terms, this is 14% of a standard deviation for \$1000 in additional annual income. This is substantially larger than comparable numbers from Akee et al. (2018).[12] Using the control function estimate, the calculation is close 1% of a standard deviation in annual terms, and is substantially smaller.

Figures 5.6 and 5.7 show a similar pattern for g_1 and g_2, with IV estimates exhibiting larger magnitudes but also a large degree of imprecision. As with δ_I, the control function posteriors for each parameter sit comfortably inside the range of values considered reasonable by the IV posterior. The control function estimates are centered around zero, suggesting that non-maternal unpaid care has neutral developmental

[12] Combining the total effect on behavioral skills from Table 3 of that paper with the average effect of the treatment (cash rebates from Casinos) on income in Table 2, suggests around 3% of a standard deviation in behavioral skills for an additional \$1000 in annual income.

Figure 5.6: Posterior Distributions for g_1

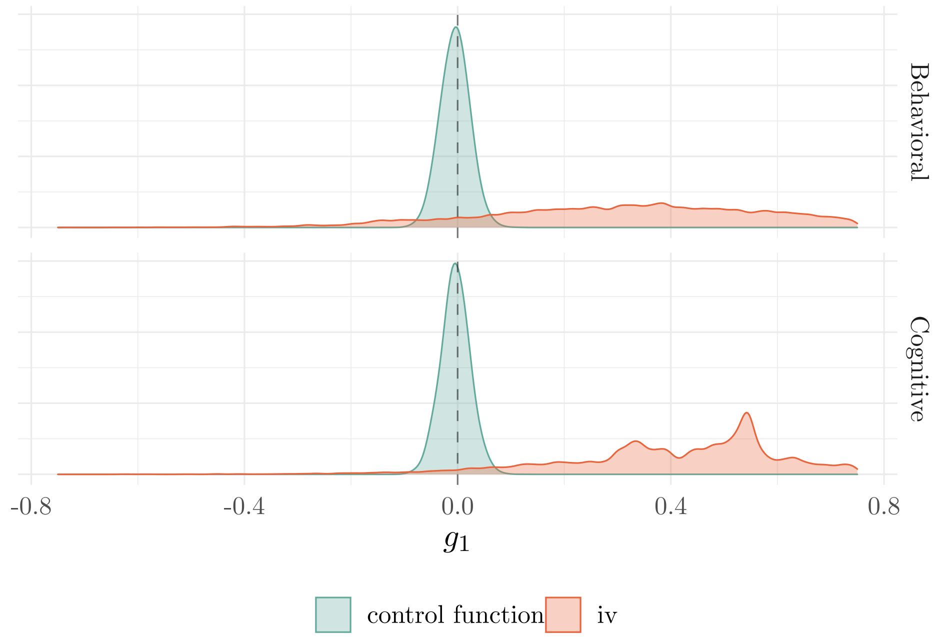

This figure shows posterior distributions for g_1 using the two approaches described in the text. Distributions are calculated from 10,000 samples drawn from the No-U-Turn Hamiltonian Monte Carlo Sampler (Hoffman et al., 2014). See Appendix D.2 for more details.

effects relative to full-time maternal care.

Most notable is the posterior for g_2 for both academic and behavioral skills, which is centered comfortably to the left of zero. The posterior mean estimate for behavioral skills is -0.069 with a 95% credible interval of $[-0.13, -0.014]$. The mean estimate implies that the choice to use paid care instead of unpaid care, all else being equal, leads to a 6.9% of a standard deviation decrease in behavioral skills. The interval bounds strongly suggests that the paid care that mothers use when working leads to worse development outcomes compared to unpaid care. In interpreting these results it is important to note that it is very difficult to adequately model the true heterogeneity in availability and quality of paid care that parents face. One interpretation of the result that $g_2 < g_1$ is that unpaid care is more likely to come from parents and other relatives and that paid care is only used when those (superior) options are not available. While the data are inadequate for thoroughly investigating this possibility, a final estimation exercise in this section will attempt to at least partially address it.

Overwhelmingly, estimates of the technology of skill formation suggest that skills are highly persistent over time (Cunha and Heckman, 2008; Cunha et al., 2010; Del Boca et al., 2014). Contrary to this evidence, the control function and IV posteriors here agree that δ_θ (which in principle determines persistence) is less than 0.5, with a posterior mean estimate of 0.22 and a 95% credible interval of $[0.01, 0.61]$. While this certainly conflicts with prior evidence, the discrepancy is likely due to the very different source of identification for this parameter. Without an initial measure of skill outcomes, there is no direct source of identification for persistence. δ_θ is instead identified by the influence of inputs in past periods relative to inputs in more recent periods. A small estimate for δ_θ suggests that measures of inputs from as many as 3 quarters prior have

Figure 5.7: Posterior Distributions for g_2

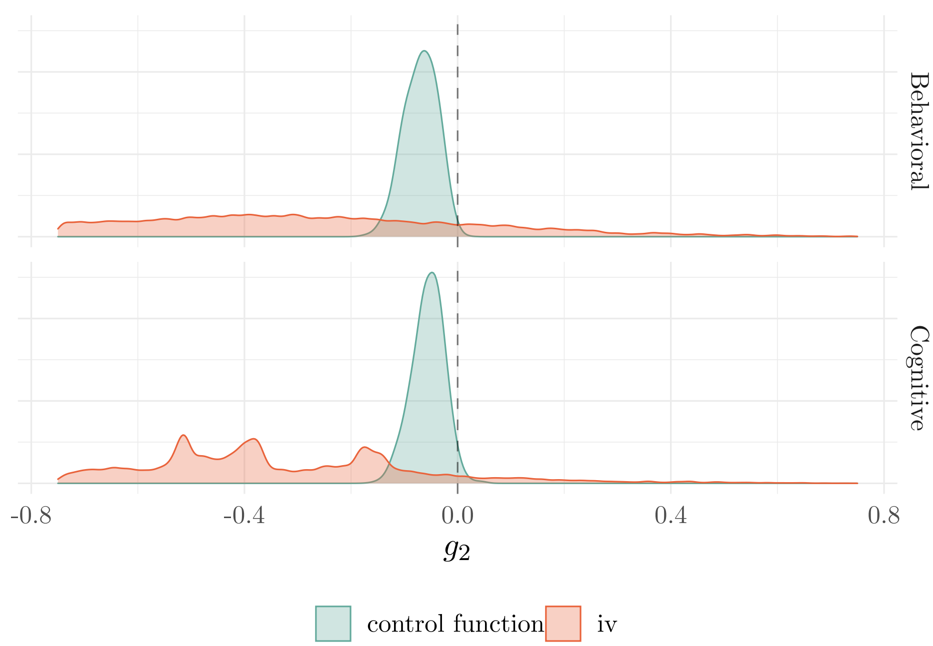

This figure shows posterior distributions for g_2 using the two approaches described in the text. Distributions are calculated from 10,000 samples drawn from the No-U-Turn Hamiltonian Monte Carlo Sampler (Hoffman et al., 2014). See Appendix D.2 for more details.

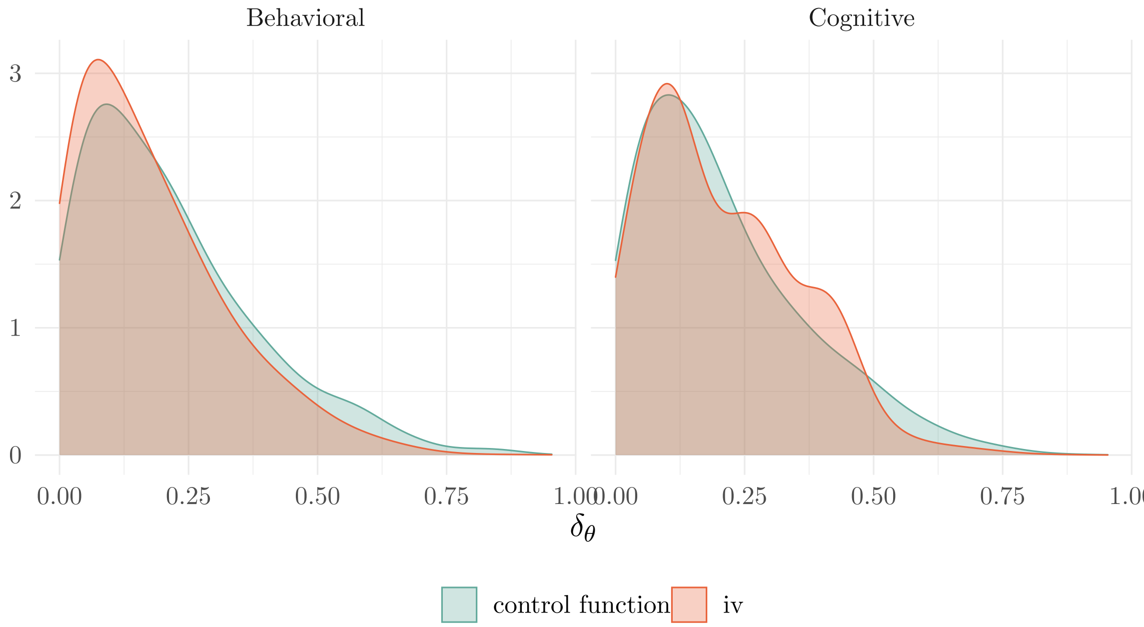

This figure shows posterior distributions for δ_θ using the two approaches described in the text. Distributions are calculated from 10,000 samples drawn from the No-U-Turn Hamiltonian Monte Carlo Sampler (Hoffman et al., 2014). See Appendix D.2 for more details.

almost no predictive power conditional on more recent periods. If variation in inputs is not sufficiently rich in the panel dimension then it may simply be difficult to effectively identify δ_θ. In light of this, it is most sensible to interpret estimates of δ_θ as a statement about how quickly the developmental effects of these inputs (and income in particular, since it is measured every period) may fade out over time. It is possible to rationalize this result with the broader literature by allowing for additional dimensions of cognitive and behavioral skills that have different input sensitivities and levels of persistence. This would relax the cross-equation restrictions between persistence in measured skills and the influence of inputs at different lags. This cannot be explored in these data, but is a promising objective for future research.

Returning to the estimates of the pair (g_1, g_2) which parameterize the ceterus paribus developmental effect of particular childcare choices, there is a concern that they are likely confounded by latent differences in the availability and quality of care options. While it is not possible to deal with this in a completely satisfactory way, it is possible to explore the mechanism by estimating type-specific values of g_1 and g_2. Figure 5.8 shows the posterior distributions using the control function approach to derive an approximate likelihood. The posteriors provide preliminary evidence that there is indeed meaningful latent heterogeneity in the quality of unpaid care relative to maternal care, and paid care relative to unpaid care. In some cases the evidence for differences across types is quite strong. For example, the posterior probability that $g_{1,2} > g_{1,5}$ is 92.6%.

Differences in the parameters across types do not appear to show any systematic ranking by productivity (as measured by $\mu_{W,k}$) or child care expenditures (as measured by $\mu_{q,k}$). For example, types (1) and (3) are equally productive in the labor market, but their posterior means for g_1 are 0.019 and -0.094. Furthermore, the posterior probability that $g_{1,1}$ lower than $g_{1,3}$ is small, at 11.3%. This suggests (perhaps unsurprisingly) that the quality and/or availability of unpaid care options varies in ways that are independent of financial resources. Similar comparisons can be made for g_2. Types (1) and (4) have similar posteriors for g_1, and type

Figure 5.8: Posterior distributions for $g_{1,k}$ and $g_{2,k}$

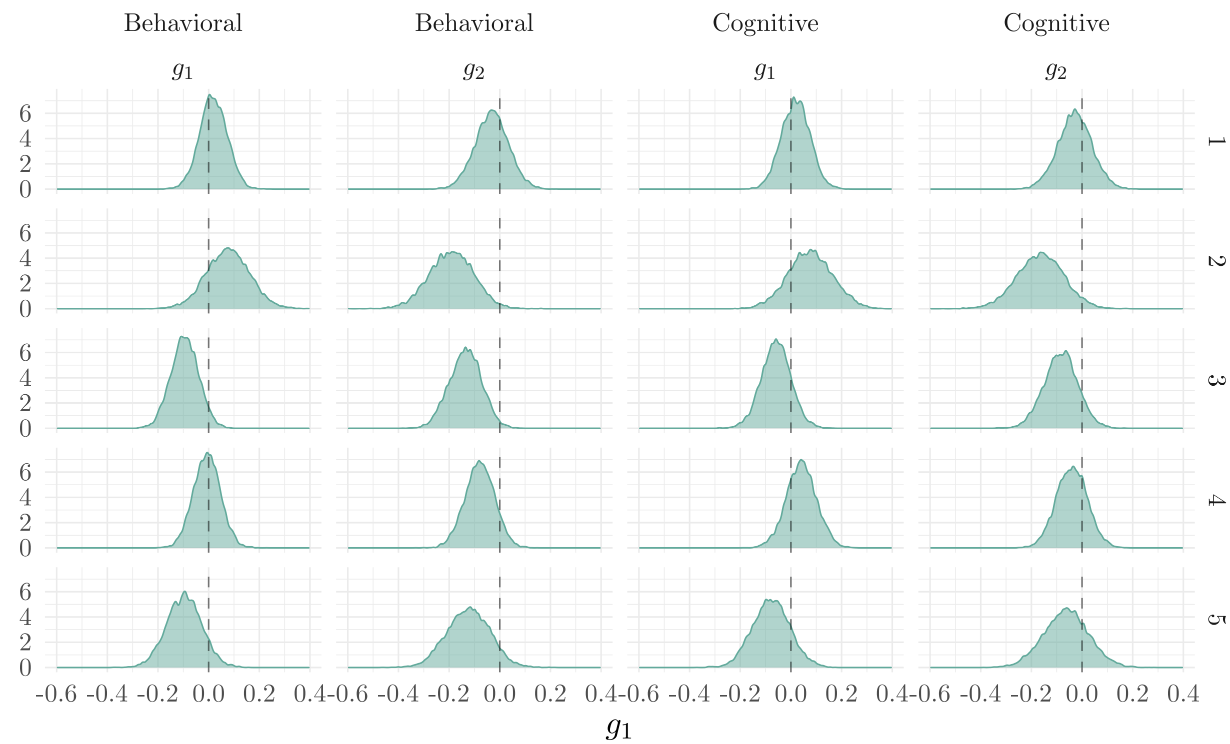

This figure shows posterior distributions for $g_{1,k}$ and $g_{2,k}$ using the control function approach described in the text. Distributions are calculated from 10,000 samples drawn from the No-U-Turn Hamiltonian Monte Carlo Sampler (Hoffman et al., 2014). See Appendix D.2 for more details.

(4) has higher average childcare expenditures, but there is very marginal evidence (a 63% posterior chance) that this care is *worse* for type (4) than for type (1). The results overall emphasize that there are important latent availability and quality choices that drive differences in the developmental costs of otherwise identical seeming care arrangements. [13]

6 Counterfactuals

The estimated model provides quantitative lessons that are not available without taking a theoretical lens to the experimental evidence. Two counterfactuals in particular demonstrate the utility of using the model to aggregate available evidence. The first counterfactual unbundles the varied components of treatment and examines the contribution of each component to the overall impact by implementing them one by one with the estimated model. The second counterfactual simulates each experiment on the representative SIPP sample, thereby overcoming the challenge of forecasting the effects of these policy reforms on the population at large when only the effects on a highly selected sample are known.

In all counterfactuals, the estimated model can likewise forecast impacts on child development outcomes and maternal welfare (as measured by equivalent variation in per-period consumption at the time of random

[13]Other work (Garcia-Vazquez, 2023; Griffen, 2019) shows promise in capturing these rich dimensions of heterogeneity using the birth cohort of the Early Childhood Longitudinal Study.

assignment). To forecast development outcomes, the counterfactual uses the posterior from the control function approach allowing for type-specific values for (g_1, g_2). With regards to welfare, Table 6.1 reports that the full treatment at each site led to welfare gains equivalent to 0.37% (FTP), 5.49% (CTJF), and 2.06% (MFIP) of consumption in each period. Of these, only the latter two have 95% credible intervals that do not contain zero. The credible intervals for forecasted impacts on cognitive and behavioral skills all contain zero, with the exception of behavioral skills in MFIP. Overall, this suggests ambiguous impacts on child skill outcomes, which reflects findings in the data across child outcomes and sites (Bloom et al., 2002, 2000; Gennetian and Miller, 2000).

6.1 Unbundling Treatment Components

Figure 6.1 and Table 6.1 report the results of the decomposition exercise for the three main components of treatment: time limits, changes to benefit formulae (incentives), and work requirements. All three components are present in the CTJF and FTP experiments, while MFIP only altered incentives and work requirements (for one arm only). Time limits in CTJF and FTP differ slightly (24 vs 21 months) but are otherwise comparable.

Time Limits

Figure 6.1 shows that time limits have a large effect on participation, with only a very small positive effect on employment. These operate with some delay, due to individuals' high discount rates: the tradeoffs associated with benefit usage only become relevant when individuals get close to exhausting their entitlement. The fact that there is only a small effect on employment when indviduals start reducing their participation suggests that income effects on employment are quite small and that, contrary to the motivation of these experiments, discouraging welfare participation does not necessarily lead to increases in employment. The model also predicts that the effects of time limits at the CTJF site are in fact larger than at FTP, underscoring the importance of initial distributions of types and labor market opportunities in determining the effect of identical treatments across sites.

Table 6.1 suggests that time limits, which reduce income and marginally increased employment, had little to no effect on skill outcomes or welfare. The negligible effect on welfare is explained by high rates of discounting: at random assignment, very little weight is placed on reductions in welfare entitlements that are at least two years in the future.

Work Requirements

This counterfactual demonstrates (see Figure 6.1) that work requirements are almost entirely responsible for the increase in employment across sites, matching very well the forecasted impact of the full treatment at each site. Although reducing participation through time limits does not have much of an effect on employment, these results indicate that increasing employment can conversely have a moderate effect on participation, which decreases by as much as 5 percentage points 5 years after random assignment.

Table 6.1 reports negative effects on skill outcomes and welfare, with credible intervals that contain zero in all but on case. Participants in MFIP appear to experience welfare losses (point estimate is a loss of 0.8% in consumption equivalent terms). This is likely from the "punitive" dimension of work requirements, embodied by the nonpecuniary cost parameter α_R.

Financial Incentives

Financial incentives lead to increases in participation and small increases in employment that fade out over time in CTJF and MFIP. When interpreting these effects, it is worth reiterating that the changes in benefit formulae are quite different across sites (see Appendix C). Each reform is designed to increase the amount of earnings that is disregarded when deducting net income from the benefit standard. In FTP, the first $200 is disregarded (up from $120) and 50% of income thereafter (up from 33 %). In CTJF, 100% of earnings are disregarded in the calculation of welfare and food stamp benefits jointly, up to the poverty guideline, before dropping to zero. In MFIP, food stamp and welfare benefits are folded into one cash grant, for which 38% (up from 33%) of income is disregarded, and a substantial amount of initial earnings is disregarded completely by increasing the benefit standard by 20%.[14]

The effect across all three sites is that a substantial number of individuals are able to benefit from combining welfare with work. However as participation rates among this group of individuals declines, the effect on those that remain (who face higher work costs and fewer employment opportunities) is closer to zero. The limited impact on employment in FTP, despite a large increase in the rate of income disregards, suggest that a full disregard on the initial portion of income may be more effective.

A (potentially) unanticipated consequence of these policies is that they lead to a more persistent *increase* in participation. Individuals who would otherwise not find it worthwhile to participate in welfare now do so because participation while working is more rewarding. At the CTJF and FTP sites, the counterfactual reveals that time limits work to counteract this effect.

Table 6.1 indicates that these changes in benefits are the main driver of welfare increases from being assigned to each treatment, yielding consumption equivalent gains that exceed those from the full treatment.

6.2 Forecasting Treatment Effects on a Representative Sample

The previous exercise offers insights into how different components of the treatment acted on the experimental populations to produce observed impacts, with some policy lessons about which components are most effective in increasing employment and reducing welfare participation. However, because the impacts are forecast exclusively on a specific population —welfare applicants and ongoing participants —it is unclear how these lessons can be extrapolated to impacts on the general population. The counterfactual in this section applies the treatment of each experiment on the SIPP sample, which is taken as a representative sample of single mothers. Relative to the experiments, the SIPP sample differs in terms of the initial distribution of both types (determined by the selection parameters β_τ and distribution of initial demographics) and job opportunties (which are assumed to be drawn from the stationary distribution).

Figure 6.2 depicts dynamic treatment effects on employment and welfare participation and Table 6.2 reports impacts on welfare and skill outcomes. Most strikingly, treatment impacts at the sites involving time limits (CTJF and FTP) are less than half the size of treatment impacts on the selected sample. The relative size of the impact is consistent with the fact that individuals in the broader population are less reliant on welfare generally. For MFIP, the effect on participation is *larger* for the representative sample, indicating that individuals in the general population who would otherwise not participate in welfare are opting in due to an increase in the financial payoff to combining welfare with work. That this effect is larger than for the selected sample speaks to the time-varying dimension of selection: the experiment oversamples individuals

[14] All income is disregarded so long as the maximum benefit is less than the payment standard minus net income.

Figure 6.1: Treatment Effects by Individual Treatment Component

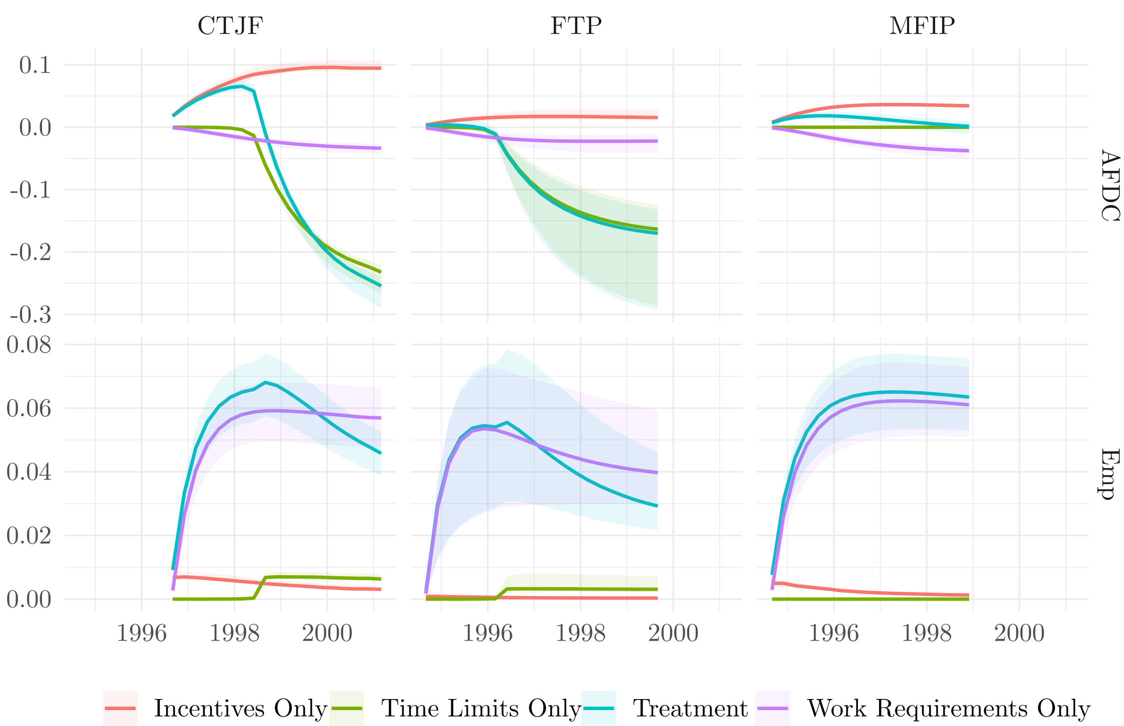

This figure depicts the model's predicted impact on welfare participation (AFDC) and employment when components of the treatment are implemented in isolation. Ribbons indicate 90% confidence intervals calculated using 75 bootstrap samples.

Table 6.1: Treatment Effects

	Treatment	Incentives Only	Work Requirements Only	Time Limits Only
		FTP		
	Treatment	Incentives Only	Work Requirements Only	Time Limits Only
Behavioral Skill	-7.27	-0.29	-7.33	-0.10
	[-18.95, 16.09]	[-1.74, 0.30]	[-18.99, 16.19]	[-0.99, 0.88]
Cognitive Skill	-3.05	0.01	-2.91	-0.38
	[-14.19, 15.37]	[-0.51, 1.33]	[-14.10, 15.58]	[-1.56, 0.50]
CEV	0.37	0.67	-0.32	-0.00
	[-18.98, 5.06]	[-31.05, 0.50]	[-0.72, 21.99]	[-0.01, 0.13]
		CTJF		
	Treatment	Incentives Only	Work Requirements Only	Time Limits Only
Behavioral Skill	-7.02	-2.72	-5.26	-0.00
	[-22.80, 1.96]	[-7.08, 0.15]	[-18.04, 2.31]	[-0.84, 0.74]
Cognitive Skill	0.02	0.28	-0.37	-0.20
	[-9.62, 10.33]	[-3.61, 3.74]	[-9.33, 8.54]	[-1.14, 0.70]
CEV	5.49	5.97	-0.59	-0.00
	[3.16, 6.25]	[3.46, 6.36]	[-1.44, 0.44]	[-0.02, -0.00]
		MFIP		
	Treatment	Incentives Only	Work Requirements Only	Time Limits Only
Behavioral Skill	-13.65	-1.68	-11.71	0.00
	[-33.46, -6.41]	[-4.06, 0.04]	[-29.25, -5.55]	[0.00, 0.00]
Cognitive Skill	-3.68	-0.01	-3.61	0.00
	[-17.56, 5.90]	[-2.05, 1.79]	[-12.31, 4.67]	[0.00, 0.00]
CEV	2.06	2.80	-0.80	0.00
	[-0.19, 2.77]	[1.38, 2.90]	[-1.53, 0.43]	[0.00, 0.00]

This table reports the model's predicted impact on behavioral skills, cognitive skills and welfare (as measured by CEV: consumption equivalent variation) when components of the treatment are implemented in isolation. Square brackets indicate 90% confidence/credible intervals calculated using 75 bootstrap samples.

Figure 6.2: Treatment Effects for the Experiment (Selected) and Representative (Non-Selected) Samples

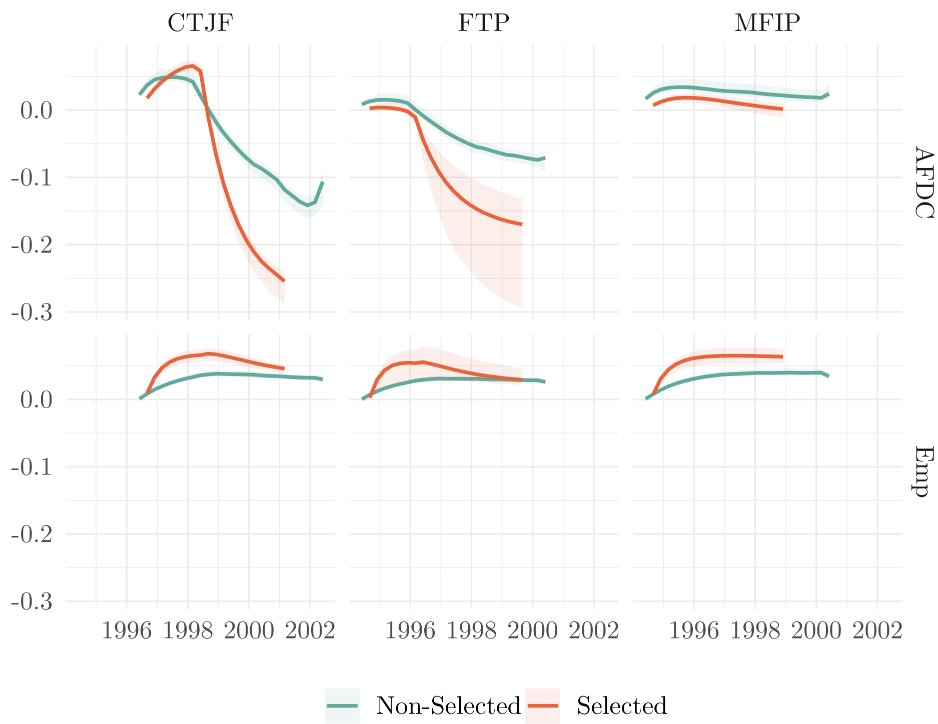

This figure depicts the model's predicted treatment effects on welfare participation (AFDC) and employment when using the estimated initial distribution at each site (selected) and when using the estimated initial distribution in the SIPP (non-selected). Ribbons indicate 90% confidence intervals calculated using 75 bootstrap samples.

who have experienced negative labor market shocks and who are otherwise less likely to be enticed by these financial incentives. Some of this selection is also due to the positive impact that services have on job arrival rates. Unemployed individuals are more likely to enroll and use these servies, leading to a persistent and economically significant increase in employment of about 4.5 percentage points. This effect is robust across sites, illustrating that work requirements introduce a permanent interplay between welfare participation and job search.

Each welfare reform leads to positive and signicant welfare gains in the general population, which the previous counterfactual exercise suggests is driven largely by increases in benefit generosity. More striking —and contrasting with the experimental evidence —is that all reforms are predicted to have negative impacts on the behavioral skills of children in the general population. The negative results here reflect the fact that in the general population, the care options used while working are more likely to have a negative impact compared to the experiment populations. Importantly, this counterfactual result helps rationalize differences between the experimental findings and those of other papers using representative samples (Bernal and Keane, 2010; Mullins, 2022; Agostinelli and Sorrenti, 2018).

Table 6.2: Treatment Effects

	FTP	CTJF	MFIP
Behavioral Skill	-10.00	-12.10	-12.25
	[-36.32, -3.09]	[-44.14, -3.94]	[-43.65, -4.52]
Cognitive Skill	-2.84	-4.03	-3.65
	[-12.38, 5.16]	[-14.97, 5.94]	[-14.20, 5.82]
CEV	1.19	2.80	2.03
	[0.83, 1.53]	[2.36, 3.19]	[1.60, 2.42]

This table reports the model's predicted impact on behavioral skills, cognitive skills and welfare (as measured by CEV: consumption equivalent variation) when using the estimated initial distribution at each site (selected) and when using the estimated initial distribution in the SIPP (non-selected). Square brackets indicate 90% confidence/credible intervals calculated using 75 bootstrap samples.

7 Conclusion

When faced with a body of experimental evidence, an economic model has proved itself useful in interpreting the data for the sake of policy lessons through counterfactuals and normative calculations. The particular example used here —welfare-to-work evaluations —is ideal because many of the treatment components are well-articulated inside the model environment, while others are intuitively parameterized. While experimental evidence is often useful for validating economic models, this paper emphasises the utility these models have for aggregating evidence, and introduces a quantitative and conceptual approach for thinking about the contribution of experimental variation to the identification and estimation of key causal parameters. The model's ability to unbundle treatments into individual components and forecast policy effects on broader populations both provided unique policy lessons that are not available in a model-free empirical analysis of treatment effects. For example, while there is no strong evidence for negative skill effects of welfare reform among the experimental participation, the estimated model has more troubling implications for that same reform on the general population.

While this paper has applied the structural approach to a meta-analysis of experiments, the methdology is certainly not limited to that environment. Very large bodies of nonexperimental evidence on important policy problems exist with little effort on the part of researchers to formally put this evidence together for policy lessons. The structural meta-analysis approach offers a promising tool for interpreting and aggregating according to the vision of Frisch (1933).

References

Agostinelli, F. and G. Sorrenti (2018): "Money vs. Time: Family Income, Maternal Labor Supply, and Child Development," *Working Paper*.

Akee, R., W. Copeland, E. J. Costello, and E. Simeonova (2018): "How Does Household Income Affect Child Personality Traits and Behaviors?" *American Economic Review*, 108, 775–827.

Arcidiacono, P. and R. A. Miller (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 79, 1823–1867.

Ashworth, K., A. Cebulla, D. Greenberg, and R. Walker (2004): "Meta-evaluation: Discovering what works best in welfare provision," *Evaluation*, 10, 193–216.

Bernal, R. and M. P. Keane (2010): "Quasi-structural estimation of a model of childcare choices and child cognitive ability production," *Journal of Econometrics*, 156, 164–189.

Bloom, D., J. J. Kemple, P. Morris, S. Scrivener, N. Verma, and R. Hendra (2000): "The Family Transition Program: Final Report on Florida's Initial Time-Limited Welfare Program," Tech. rep., Manpower Demonstration Research Corporation.

Bloom, D., S. Scrivener, C. Michalopoulos, P. Morris, R. Hendra, D. Adams-Ciardullo, and J. Walter (2002): "Jobs First Final Report on Connecticut's Welfare Reform Initiative," Tech. rep., Manpower Demonstration Research Corporation.

Bonhomme, S., K. Jochmans, and J.-M. Robin (2017): "Nonparametric estimation of non-exchangeable latent-variable models," *Journal of Econometrics*, 201, 237–248.

Card, D., J. Kluve, and A. Weber (2010): "Active Labour Market Policy Evaluations: A Meta-Analysis," *The Economic Journal*, 120, F452–F477.

——— (2018): "What works? A meta analysis of recent active labor market program evaluations," *Journal of the European Economic Association*, 16, 894–931.

Center for Economic and Policy Research (2014): "SIPP Uniform Extracts," .

Chan, M. K. (2017): "Welfare dependence and self-control: An empirical analysis," *The Review of Economic Studies*, 84, 1379–1423.

Choi, E. J. (2018): "Evaluating a structural model of labor supply and welfare participation: Evidence from state welfare reform experiments," *Working Paper*.

Crosby, D. A., C. J. Dowsett, L. A. Gennetian, and A. C. Huston (2010): "A tale of two methods: comparing regression and instrumental variables estimates of the effects of preschool child care type on the subsequent externalizing behavior of children in low-income families." *Developmental psychology*, 46, 1030.

Cunha, F., J. Heckman, and S. Schennach (2010): "Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Econometrica*, 78, 883–931.

Cunha, F. and J. J. Heckman (2008): "Formulating , Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Journal of Human Resources*, 43.

Del Boca, D., C. Flinn, and M. Wiswall (2014): "Household Choices and Child Development," *Review of Economic Studies*, 81.

Duflo, E., R. Hanna, and S. P. Ryan (2012): "Incentives Work: Getting Teachers to Come to School," *American Economic Review*, 102, 1241–78.

Duncan, G. J., P. A. Morris, and C. Rodrigues (2011): "Does money really matter? Estimating impacts of family income on young children's achievement with data from random-assignment experiments." *Developmental psychology*, 47, 1263–79.

Feenberg, D. R. and E. Coutts (1993): "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12, 189–194.

Frisch, R. (1933): "Editor's Note," *Econometrica*, 1, 1–4.

Galiani, S., A. Murphy, and J. Pantano (2015): "Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment," *American Economic Review*, 105, 3385–3415.

Garcia-Vazquez, M. (2023): "The equilibrium effects of state-mandated minimum staff-to-child ratios," *Working Paper*.

Ge, H., K. Xu, and Z. Ghahramani (2018): "Turing: a language for flexible probabilistic inference," in *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, 1682–1690.

Gennetian, L. A., D. A. Crosby, A. C. Huston, and E. D. Lowe (2004a): "Can Child Care Assistance in Welfare and Employment Programs Support the Employment of Low-Income Families?" *Journal of Policy Analysis and Management*, 23, 723–743.

Gennetian, L. A., G. Duncan, V. Knox, W. Vargas, E. Clark-Kauffman, and A. S. London (2004b): "How Welfare Policies Affect Adolescents' School Outcomes: A Synthesis of Evidence From Experimental Studies," *Journal of Research on Adolescence*, 14, 399–423.

Gennetian, L. A., K. Magnuson, and P. A. Morris (2008): "From statistical associations to causation: what developmentalists can learn from instrumental variables techniques coupled with experimental data." *Developmental psychology*, 44, 381.

Gennetian, L. A. and C. Miller (2000): "Reforming Welfare and Rewarding Work: Final Report on the Minnesota Family Investment Program Volume 2: Effects on Children," Tech. rep., Manpower Demonstration Research Corporation.

Greenberg, D., K. Ashworth, A. Cebulla, and R. Walker (2004): "Do welfare-to-work programmes work for long?" *Fiscal Studies*, 25, 27–53.

——— (2005): "When Welfare-to-Work Programs Seem to Work Well: Explaining Why Riverside and Portland Shine So Brightly," *ILR Review*, 59, 34–50.

Griffen, A. S. (2019): "Evaluating the effects of childcare policies on children's cognitive development and maternal labor supply," *Journal of Human Resources*, 54, 604–655.

Heckman, J., N. Hohmann, J. Smith, and M. Khoo (2000): "Substitution and Dropout Bias in Social Experiments: A Study of an Influential Social Experiment*," *The Quarterly Journal of Economics*, 115, 651–694.

Heckman, J. J. (1992): *Evaluating Welfare and Training Programs*, chap. Randomization and Social Policy Evaluation.

Heckman, J. J., J. Smith, and N. Clements (1997): "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *The Review of Economic Studies*, 64, 487–535.

Heckman, J. J. and E. Vytlacil (2005): "Structural Equations, Treatment Effects, and Econometric Policy Evaluation1," *Econometrica*, 73, 669–738.

Hoffman, M. D., A. Gelman, et al. (2014): "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." *J. Mach. Learn. Res.*, 15, 1593–1623.

Hu, Y. and M. Shum (2012): "Nonparametric identification of dynamic models with unobserved state variables," *Journal of Econometrics*, 171, 32–44.

Kasahara, H. and K. Shimotsu (2009): "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 77, 135–175.

Kline, P. and M. Tartari (2016): "Bounding the labor supply responses to a randomized welfare experiment: A revealed preference approach," *American Economic Review*, 106, 972–1014.

Miller, C., V. Knox, L. A. Gennetian, M. Dodoo, J. A. Hunter, and C. Redcross (2000): "Reforming Welfare and Rewarding Work: Final Report on the Minnesota Family Investment Program Volume 1: Effects on Adults," Tech. rep., Manpower Demonstration Research Corporation.

Mogensen, P. K. and A. N. Riseth (2018): "Optim: A mathematical optimization package for Julia," *Journal of Open Source Software*, 3, 615.

Morris, P. and L. Gennetian (2003): "Identifying the Effects of Income on Children's Development Using Experimental Data," *Journal of Marriage and Family*, 65, 716 – 729.

Mullins, J. (2022): "Cash Transfers in the Presence of Children's Human Capital Formation," .

Revels, J., M. Lubin, and T. Papamarkou (2016): "Forward-Mode Automatic Differentiation in Julia," *arXiv:1607.07892 [cs.MS]*.

Rodriguez, J. (2018): "Understanding the Effects of a Work-Based Welfare Policy on Child Human Capital," Working Paper.

Todd, P. E. and K. I. Wolpin (2005): "Assessing the Impact of a School Subsidy Program in Mexico : Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96.

——— (2023): "The Best of Both Worlds: Combining Randomized Controlled Trials with Structural Modeling," *Journal of Economic Literature*, 61, 41–85.

A Details of Model Solution

The additive representation can be shown by taking the form of the solution as given, and showing that the recursion preserves this relationship. Note that in the terminal period, the value function is simply $V_T(\theta) = (1-\beta)^{-1}\alpha_{\theta,k}\log(\theta)$ and so this form holds at T. We have:

$$V_{k,t}(\chi,\theta) = \mathbb{E}_\epsilon \max_{C,I,j}\left\{U_{kj}(C,\theta) + \epsilon_j + \beta\mathbb{E}_{\chi'|\chi,\theta}[\tilde{V}_{k,t+1}(\chi')] + \beta\alpha_{V,t+1,k}(\delta_I\log(I) + \delta_\theta\log(\theta))\right\}$$

where we have written θ' already in terms of I and θ. Fixing the discrete choice j fixes net income, $Y_{kj}(\chi)$ and conditional on this choice the optimal investment problem has a simple solution:

$$p_{H_j+F_j}I = \frac{\beta\delta_I\alpha_{V,t+1,k}}{1+\beta\delta_I\alpha_{V,t+1,k}}Y_{kj}(\chi)$$

which can be substituted into the problem above to get:

$$V_{k,t}(\chi,\theta) = \mathbb{E}_\epsilon \max_{j}\left\{\tilde{U}_{kj}(\chi) + \epsilon_j + \beta\mathbb{E}_{\chi'|\chi}\tilde{V}_{k,t+1}(\chi')\right\} + \alpha_{V,t,k}\log(\theta)$$

where $\alpha_{V,t,k} = \alpha_{\theta,k} + \beta\delta_\theta\alpha_{V,t+1,k}$ as in the main text and indirect utility $\tilde{U}$ takes the form in the main text. Note that $\alpha_{V,t}$ is scaled uniformly by $\alpha_{\theta,k}$ and so it can be rewritten as:

$$\alpha_{V,t,k} = \alpha_{\theta,k}\tilde{\Gamma}_t, \qquad \tilde{\Gamma}_t = 1 + \beta\delta_\theta\tilde{\Gamma}_{t+1}.$$

Accordingly, Γ_t from the main text is defined as $\Gamma_t = \delta_I\beta\tilde{\Gamma}_{t+1}$.

B Identification

This section derives the expressions that comprise the linear system in equation (4.1). The system is four dimensional, corresponding to five choice comparisons: (1) using paid care vs not, conditional on participation and work; (2) working vs not, conditional on participation; (3) participating in welfare vs not, assuming no time limit; (4) participating in welfare vs not, assuming a time limit; and (5) participating in food stamps vs no participation.

In what follows, it is useful to recall that the following fact about nested logit choice probabilites. Consider the log-odds ratio for choosing nest B_1 vs nest B_2:

$$\log\left(\mathbb{P}(B_1)/\mathbb{P}(B_2)\right) = \sigma_l^{-1}\left(\mathcal{V}_1 - \mathcal{V}_2\right)$$

where $\mathcal{V}_b$ is the inclusive value of choosing this nest and σ_l is the dispersion parameter for choices across nests in this layer, l. The inclusive value $\mathcal{V}_b$ can likewise be written as:

$$\mathcal{V}_b = \mathcal{V}_j - \sigma_{l+1}\log(\mathbb{P}(j\in B_b))$$

where $\mathcal{V}_j$ is the value of making choice j. This could be a final value if j is a single choice, or an inclusive value if j is another nest.

It will also be useful to define the concept of *finite dependence* (Arcidiacono and Miller, 2011). A model exhibits finite dependence if for any two choices there exists a sequence of future choices that returns the individual back to the same distribution of state variables. When this is the case, the formula above can be applied iteratively until the period at which the states are in the same distribution, at which point the difference in the values of the two options differences out, and we are left with a sequence of utilities and choice probabilities.

Finite dependence is simple here as the only choice affecting future states is participation. Consider the comparison of welfare participation $P_1 = 2$ vs not $P_1 = 1$ at time t. Without time limits, individuals differ at $t+1$ in terms of lagged participation, but will return to the same state at $t+2$ with $P_1' = 1$ and $P_2' = 1$. In a world with time limits, we require a two-period sequence: $(P_1' = 1, P_1'' = 1)$ and $(P_2' = 2, P_2'' = 1)$ which will return individuals to the same state at $t+3$. These sequences will be applied below when deriving ratios of participation choice probabilities.

Child Care Choices

This first equation for child care choices is derived in the main text:

$$\log(\tilde{p}_k(F, X, Z)) = \sigma_3^{-1}\left[(1 + \alpha_{\theta,k}\Gamma_t)\log\left(\tilde{Y}_{k,F}(P, X, Z)\right) + \alpha_{F,k} + \alpha_{\theta,k}\Gamma_t g_2\right] \tag{B.1}$$

and identification of the pair $\sigma_3^{-1}(1 + \alpha_{\theta,k}\Gamma_t)$ and $(\alpha_{F,k} + \alpha_{\theta,k}\Gamma_t g_2$ can be derived from variation in Z alone or from X also. For example, differencing across Z and Z' gives:

$$\log\left(\frac{\tilde{p}_{k,F}(P, X, Z')}{\tilde{p}_{k,F}(P, X, Z)}\right) = \sigma_3^{-1}(1 + \alpha_{\theta,k}\Gamma_t)\log\left(\frac{\tilde{Y}_{k,F}(P, X, Z')}{\tilde{Y}_{k,F}(P, X, Z))}\right)$$

With this slope term in hand, the age-specific term $\alpha_{F,k} + \alpha_{\theta,k}\Gamma_t$ pins down the level.

Labor Supply

As above, define

$$\tilde{p}_{k,H}(P, X, Z) = \frac{p_k(H = 1|P, \chi)}{p_k(H = 0|P, \chi)}$$

as the relative choice probability of working in state $\chi = (X, Z)$, conditional on participation P. It will be useful to also define the proportional return to working as:

$$\tilde{Y}_{k,H}(P, X, Z) = \frac{Y_k(H = 1, P, X, Z)}{Y_k(H = 0, P, X, Z)}.$$

For convenience let us also define:

$$p_{k,F}(P, X, Z) = P[F = 1|H = 1, P, X, Z]$$

With these expressions, the nested conditional choice probability for the work choice can be written as:

$$\log\left(\tilde{p}_{k,H}(P, X, Z)\right) = \sigma_2^{-1}\left(\alpha_{H,k} + \alpha_{\theta,k}\Gamma_t g_1\right) + \\ \sigma_2^{-1}\left[(1 + \alpha_{\theta,k}\Gamma_t)\log\left(\tilde{Y}_{k,H}(P, X, Z)\right) + R(X, Z)\alpha_R A - \sigma_3\log(1 - p_{k,F}(P, X, Z))\right] \tag{B.2}$$

which forms the second row of the linear system in (4.1).

Notice that three parameters now dictate the responsiveness of labor supply to different components of the treatment. The term $\sigma_2^{-1}(1+\alpha_{\theta,k}\Gamma_t)$ is an age-specific semi-elasticity of labor supply with respect to financial incentives. The term $\sigma_2^{-1}\alpha_R$ determines the response to work requirements through non-pecuniary motives, and $\sigma_2^{-1}\sigma_3$ determines the responsiveness of labor supply to childcare costs. All three terms are identified as long as there is rank-independent variation in returns to work, the existence of work requirements $(R(X, Z))$, and childcare prices. In practice this is delivered by all sources of variation in X and Z, conditional on type, but in principle one could use only variation in Z, as long as there are at least three treatments that vary returns to work, work requirements, and childcare subsidies in a rank-independent way. This paper's application has four treatments that do satisfy these reqiurements.

Adding these slope terms to the slope terms identified in the previous stage is sufficient to invert out values for σ_1, σ_2, $\alpha_{\theta,k}\Gamma_t$, $\alpha_{F,k}$ and g_2. Returning now to (B.2), the parameters g_1 and $\alpha_{H,k}$ remain to pin down the level of work choice probabilities at each age and in each state.

It may help to note here that if there exists a combination of (X, Z, Z') such that the only difference in policies is the existence of a mandatory services, this expression implies:

$$\log(\left(\frac{\tilde{p}_{k,H}(1, X, Z')}{\tilde{p}_{k,H}(1, X, Z)}\right) = \sigma_2^{-1}\alpha_R.$$

This is exactly the case for MFIP in which treatment arms (1) and (2) differ only in the imposition of work requirements and hence will be an important source of identification for α_R.

Program Participation without Time Limits

Because participation decisions dynamically influence future states, more elaborate comparisons are required. In a slight abuse of notation, let the state χ be decomposed into (X, Z, A_{-1}) so that we can indicate values of A_{-1} exclusively. Define:

$$\tilde{p}_{k,A}(X, Z, A_{-1}) = \frac{p_k[P = 2|X, Z, A_{-1}]}{p_k[P = 1|X, Z, A_{-1}]}$$

as the probability of participating in welfare relative to participating only in food stamps. The term

$$\tilde{Y}_{k,A}(X, Z) = \frac{Y_k(H = 0, P = 2, X, Z)}{Y_k(H = 0, P = 1, X, Z)}$$

is the proportional return in net income to participating in welfare assuming no work. Finally, define:

$$\tilde{u}_{k,A}(X, Z, A_{-1}) = (1 + \alpha_{\theta,k}\Gamma_t)\log(\tilde{Y}_{k,A}(X, Z)) - R(X, Z)\alpha_R - \alpha_P(1 - A_{-1}) - \sigma_2 \log\left(\frac{p_k(H = 0|P = 2, X, Z)}{p_k(H = 0|P = 1, X, Z)}\right) \quad \text{(B.3)}$$

as the difference in static payoffs from choosing $P = 2$ over $P = 1$. Putting these expressions together gives:

$$\log(\tilde{p}_{k,A}(X, Z, A_{-1})) = \sigma_1^{-1}\left[\tilde{u}_{k,A}(X, Z, A_{-1}) + \beta\mathbb{E}_{X'|X}(V(X', Z, A_{-1} = 1) - V(X', Z, A_{-1} = 0))\right].$$

A simple comparison of this log odds ratio for different values of A_{-1} identifies the application cost α_P up to scale:

$$\log\left(\frac{\tilde{p}_{k,A}(X, Z, 1)}{\tilde{p}_{k,A}(X, Z, 0)}\right) = \sigma_1^{-1}\alpha_P.$$

Then in order to difference out the dynamic values, we can exploit the model's finite dependence to get:

$$\log(\tilde{p}_{k,A}(X, Z, A_{-1})) = \sigma_1^{-1}\tilde{u}_{k,A}(X, Z, A_{-1}) - \sigma_1^{-1}\left[\beta\mathbb{E}_{X'|X}\sigma_1 \log\left(\frac{p_k(P = 1|X', Z, 1)}{p_k(P = 1|X', Z, 0)}\right)\right].$$

This forms the third row of the linear system in (4.1). Following identical logic to the previous steps, the response in program participation to benefit generosity is determined by σ_1 and a comparison between Z and Z' fixing the other variables will identify it. The fourth two is the relatively simpler comparison of food stamp participation with no participation, $p_{k,S}(X, Z) = p_k[P = 1|X, Z]/p_k[P = 0|X, Z]$:

$$\log(\tilde{p}_{k,S}(X, Z) = \sigma_1^{-1}\alpha_{S,k} + \sigma_1^{-1}(1 + \alpha_{\theta,k}\Gamma_t)\log(\tilde{Y}_{k,S}(Z, A))$$

Program Participation with Time Limits

Following the same notational convention from the previous section, we can derive an expression for $\tilde{p}_{k,A}$ in the presence of time limits. First we must further decompose the state χ into: (X, Z, A_{-1}, ω) where ω is cumulative welfare use. Then we get:

$$\log(\tilde{p}_{k,A}(X,Z,A_{-1},\omega)) = \sigma_1^{-1}\tilde{u}_{k,A}(X,Z,A_{-1}) + \sigma^{-1}\left[\beta\mathbb{E}_{X'|X}(V(X',Z,1,\omega+1) - V(X',Z,0,\omega))\right].$$

In this case, a two-period sequence of decisions places the agent back into the same state space, giving:

$$\log(\tilde{p}_{k,A}(X,Z,A_{-1},\omega)) = \sigma_1^{-1}\Bigg\{\tilde{u}_{k,A}(X,Z,A_{-1}) + \beta\mathbb{E}_{X'|X}\Bigg[(1+\alpha_{\theta,k}\Gamma_t)\log(\tilde{Y}_{k,A}(X',Z)) - \sigma_1\log\left(\frac{p_k(P=1|X',Z,1,\omega+1)}{p_k(P=2|X',Z,0,\omega)}\right) - \beta\sigma_1\mathbb{E}_{X''|X'}\log\left(\frac{p_k(P=1|X'',Z,0,\omega+1)}{p_k(P=1|X'',Z,1,\omega+1)}\right)\Bigg]\Bigg\}. \quad \text{(B.4)}$$

which forms the 6th and final row of the linear system. To see how experimental variation can help identify β using this equation, note that this expression holds also for the case where time limits do not exist. If time limits were the only difference in the policy environment between Z and Z', we would get:

$$\log\left(\frac{\tilde{p}_{k,A}(X,Z',A_{-1},\omega}{\tilde{p}_{k,A}(X,Z,A_{-1})}\right) = \beta\mathbb{E}_{X'|X}\left[\log\left(\frac{p_k(P=1|X',Z,1)}{p_k(P=2|X',Z,0)}\right) - \log\left(\frac{p_k(P=1|X',Z',1,\omega+1)}{p_k(P=2|X',Z',0,\omega)}\right)\right] + \beta^2\mathbb{E}_{X''|X}\left[\log\left(\frac{p_k(P=1|X'',Z,0)}{p_k(P=1|X'',Z,1)}\right) - \log\left(\frac{p_k(P=1|X'',Z',0,\omega+1)}{p_k(P=1|X'',Z',1,\omega+1)}\right)\right] \quad \text{(B.5)}$$

which is an expression that identifies β exlusively through the effect of time limits on participation choices.

B.1 A Minimum Distance Estimator

Having derived the system of equations that relates choice probabilities to parameters, this section briefly describes a minimum distance estimator that exclusively uses random assignment to identify Φ_3. In practice, there is one system like (4.1) for each type k and child age t. For ease of exposition, consider the estimator for just one pair (k, t) with the understanding that the approach easily generalizes by stacking such equations.

Abusing notation slightly, let Z now jointly indicate experimental site as well as treatment group. Define:

$$\boldsymbol{h}_l(Z) = \int \boldsymbol{h}_l(X, Z; \boldsymbol{p})dF(X|Z) \qquad l \in \{0, 1\}.$$

This gives a linear system:

$$\boldsymbol{h}_0(Z) = \boldsymbol{\kappa}_0(\Phi_2) + \boldsymbol{\kappa}_1(\Phi)_2\boldsymbol{h}_1(Z)$$

with a natural estimator being:

$$\widehat{\Phi_2, \Phi_1} = \arg\min(\hat{\boldsymbol{h}}_0(Z) - \boldsymbol{\kappa}_0(\Phi_2) - \boldsymbol{\kappa}_1(\Phi)_2\hat{\boldsymbol{h}}_1(Z))^T\mathbf{W}(\hat{\boldsymbol{h}}_0(Z) - \boldsymbol{\kappa}_0(\Phi_2) - \boldsymbol{\kappa}_1(\Phi)_2\hat{\boldsymbol{h}}_1(Z))$$

Which is consistent and asymptotically normal for any positive definite $\mathbf{W}$ under standard regularity conditions.

C Taxes and Transfers

Total transfers are defined as the sum of taxes, welfare (if participating) and food stamps (if participating). Let n be the numer of children, g the state in which the program takes place, y the year, and ω the current usage of time limits. E is monthly earnings.

C.1 Taxes

Taxes consist of a federal and a state computation. When earned income is sufficiently low, taxes will arrive in the form of a net payment (when income tax obligations are exceeded by the EITC). In theory, the relevant parameters to compute taxes include those that define the federal and state EITC programs, state and federal deductions and exemptions, and the marginal income tax rate with their corresponding brackets for state and federal income tax. In practice, I use the TAXSIM model of Feenberg and Coutts (1993), to approximate the tax function. Given the relevant year, state, and family size, TAXSIM computes net payments/obligations, $T^T_{mt}(e)$. This function is called for a fine grid of earnings between \$0 and \$100,000, then approximated using a polynomial spline function with year, statem and family size-specific coefficients.

C.2 Control Groups

AFDC payments follow the formula:

$$\text{AFDC}(g,n,y,E) = \max\{B(g,n,y) - (1-0.33)\max\{E-120,0\},0\}$$

where $B(g,n,y)$ is the benefit standard. Food stamp payments are:

$$\text{SNAP}(g,n,y,E) = G(g,n,y) - 0.3(0.8E + \text{AFDC} - 134)$$

C.3 MFIP

: Both arms of MFIP change benefit formulae to:

$$\text{MFIP}(g,n,y,E) = \max\{\min\{1.2(B(g,n,y)+G(g,n,y)) - (1-0.38)E, B(g,n,y)+G(g,n,y)\},0\}$$

where B and G are the benefit standards and maximum food stamp payment in Minnesota for each number of kids (n) and year (y). Hence, food stamps and welfare are disregarded at the same rate rather than the double-deduction that occurs when AFDC payments are deducted from food stamps.

C.4 CT-Jobs First

The welfare payment for CTJF also folds in food stamps and features a "cliff" in benefits at the eligibility cutoff:

$$\text{CTJF}(g,n,y,E) = \mathbf{1}\{E < PG(n,y)\}(B(g,n,y)+G(g,n,y))$$

where $PG(n,y)$ is the poverty guidelines for a family of size n in year y.

C.5 FTP

The benefit formula for FTP is:

$$\text{FTP}(g,n,y,E) = \max\{B(g,n,y) - 0.5\max\{E-200,0\},0\}$$

Food stamps are the same as the control group.

D Details of Estimation

D.1 First Stage

We collect parameters into three blocks:

$$\Theta_1 = (\alpha_A, \alpha_S, \alpha_H, \alpha_F, \alpha_\theta, \alpha_R, \alpha_P, \beta_\Gamma, \sigma, \beta, y, \mu_W, \beta_W, \sigma_\eta, \mu_q, \beta_q, \lambda_0, \lambda_1, \delta, \mu_o, \sigma_o, \lambda_R)$$

$$\Theta_2 = (\sigma_W, \sigma_q)$$

$$\Theta_3 = (\beta_\tau, \pi_{\eta,0})$$

Recall that $\boldsymbol{y}_m = \{y_{m,t}\}_{t=1}^{T_m}$ be the panel of observed outcomes for mother m in period t. Let $W_m = \{W_{m,t}\}_{t=1}^{T_m}$ be observable components of the state space for mother m in each period t. As in the main text, $\chi_{m,t}$ is the vector of state variables that determine the distribution of outcomes. For simplicity, let dependence of $\chi_{m,t}$ on η and $W_{m,t}$ be implied whenever necessary. The likelihood is:

$$\mathcal{L} = \sum_m \log \left(\sum_{k=1}^{K} \sum_{t=1}^{T_m} \sum_{\eta_t=0}^{K_\eta} \mathbb{P}[y_{m,t}|\chi_{m,t}, \Theta_1, \Theta_2] \mathbb{P}_k[\eta_{m,t+1}|\eta_{m,t}\Theta_1] \pi_{\eta,0}[\eta_1|k, X_{\tau,m}] \mathbb{P}[k|X_{\tau,m}, \Theta_3] \right)$$

The expectation maximization routine proceeds as follows, fixing a current guess of parameters, Θ^l.

E-Step Using Θ^l, construct posterior probabilities for the latent state (k, η_t) in each time period:

$$q_{mt}(\eta, k) = \mathbb{P}[k, \eta|\boldsymbol{y}_m, X_m, \Theta^l]$$

$$q_{mt}(\eta_{t+1}, \eta_t, k) = \mathbb{P}[\eta_{t+1}, \eta_t, k|\boldsymbol{y}_m, X_m, \Theta^l]$$

This problem is made tractable by exploiting the Markov structure of the outcome probabilities using the forward-back algorithm.

M-Step Using the weights q_m, update Θ^l by moving up the likelihood:

$$\mathcal{L}(\Theta) = \sum_m \sum_k \Bigg[\log(\mathbb{P}[\eta_1, k|\Theta_3]) q_{m,1}(\eta, k) + \sum_t \sum_\eta \log\left(\mathbb{P}[y_{m,t}|\chi_{m,t}, \Theta_1, \Theta_2]\right) q_{m,t}(\eta, k) + \sum_{\eta'} \mathbb{P}_k[\eta'|\eta, \Theta_1] q_{mt}(\eta', \eta, k) \Bigg] \quad \text{(D.1)}$$

This maximization step is additively separable in each of the three blocks. One iteration of the step involves:

(a) Updating Θ_1^l with 5 iterations of the LBFGS algorithm using `Optim.jl` (Mogensen and Riseth, 2018) with automatic differentiation using `ForwardDiff.jl` (Revels et al., 2016). This is completed in a number of separate block steps. Each evaluation of this piece of the likelihood requires solution of the model via backward induction.

(b) Update Θ_2^l with the new maximizer, which is given as a pair of weighted standard deviations for predicted wage and price residuals.

(c) Update β_τ by maximimizing the weighted initial type likelihood using LBFGS, and updating $\pi_{\eta,0}$ with the new maximizer which is given by a weighted frequency estimator.

This process is repeated until Θ^l and Θ^{l+1} satisfy a convergence criterion. Standard errors are estimated using the covariance of the score equation for each panel observation.

D.2 Second Stage - Production

Let X_m be a vector of location dummies and Z_m be a full set of location-treatment interaction dummies. The moment condition (4.3) implies a system of equations:

$$\boldsymbol{I}_m = X_m \Pi_1 + Z_m \Pi_2 + \eta_m$$

$$S_m = X_m A + \boldsymbol{I}_m B(\delta, g) + \xi_m$$

Where the terms η_m and ξ_m are expectation errors with unknown distribution. Their distribution is approximated as a joint normal, allowing for a likelihood to be written for the factor score S and inputs $\boldsymbol{I}$ conditional on X_m and Z_m. The priors are:

- $\delta_I \sim U[0, 2]$
- $\delta_\theta \sim U[0, 1]$
- $g_1, g_2 \sim U[-2, 2]$

To see that these priors are loose, the upper bound on δ_I implies that a log point increase in net household income would yield a 2 standard deviation increase in skills. A flat (i.e. improper) prior is specified for the parameters $(\Pi 1, \Pi_2, A)$. Estimates and figures reported in the main text are derived from 10,000 draws from a No-U-Turn Hamiltonian Monte-Carlo chain (Hoffman et al., 2014) implemented in `julia` using `Turing.jl` (Ge et al., 2018).

Rearranging the second moment condition gives:

$$S_m = \boldsymbol{I} B(\delta, g) + \sum_{k,\eta} q_{m,1}(k, \eta) \mu(k, \eta, \boldsymbol{X}) + \eta_m$$

where $q_{m,1}$ is the posterior weight over latent states defined in Appendix D.1. A lower dimensional approximation for $\mu(k, \eta, \boldsymbol{X})$ appears to suffice, with higher dimensional interaction terms seeming to make little difference to estiamates. Approximating the expectation error η_m as normal once again delivers a likelihood for S conditional on $\boldsymbol{I}, \boldsymbol{y}, \boldsymbol{X}$. The same methods and priors are used to sample from the posterior distribution.

E Tables

Table E.1: Type Selection Coefficients β_τ for each site.

	$k=2$	$k=3$	$k=4$	$k=5$
	SIPP			
Const.	0.23	1.52	1.35	1.51
	(0.57)	(0.47)	(0.44)	(0.43)
High School	0.32	-1.42	0.54	0.12
	(0.51)	(0.51)	(0.42)	(0.42)
Some College	-0.04	-2.13	0.48	1.21
	(0.56)	(0.71)	(0.45)	(0.42)
Num Kids	-0.03	-0.17	-0.34	-0.38
	(0.17)	(0.16)	(0.14)	(0.13)
	FTP			
Const.	-0.02	0.39	-2.49	-0.92
	(2.92)	(2.39)	(11.42)	(6.49)
High School	0.32	-0.37	1.78	0.57
	(0.20)	(0.19)	(0.61)	(0.28)
Some College	0.92	-0.52	2.52	1.00
	(0.70)	(1.01)	(1.08)	(0.82)
Num Kids	-0.33	-0.38	-0.49	-0.47
	(0.10)	(0.10)	(0.21)	(0.13)
New Applicant	0.62	1.01	1.03	1.28
	(0.22)	(0.20)	(0.43)	(0.27)
	CTJF			
Const.	-0.13	0.40	1.09	-1.05
	(0.74)	(0.85)	(0.53)	(0.90)
High School	0.71	-0.16	0.45	0.72
	(0.22)	(0.23)	(0.19)	(0.31)
Some College	1.07	-1.02	1.29	2.32
	(0.41)	(0.87)	(0.37)	(0.44)
Num Kids	-0.17	-0.20	-0.52	-0.15
	(0.13)	(0.14)	(0.11)	(0.17)
New Applicant	0.99	0.63	0.63	1.28
	(0.29)	(0.33)	(0.27)	(0.32)
New Haven	-0.87	-0.93	-0.92	-1.01
	(0.28)	(0.29)	(0.25)	(0.35)
	MFIP			
Const.	-2.44	0.77	0.82	-2.49
	(1.54)	(0.82)	(0.77)	(0.93)
High School	0.72	-0.60	0.30	0.99
	(0.53)	(0.31)	(0.29)	(0.43)
Some College	1.33	0.28	1.16	1.82
	(0.75)	(0.48)	(0.45)	(0.56)
Num Kids	0.29	0.01	-0.19	0.27
	(0.23)	(0.15)	(0.14)	(0.17)
New Applicant	0.86	0.14	0.32	1.74
	(0.65)	(0.44)	(0.39)	(0.44)
Re-Applicant	0.07	0.31	-0.46	0.05
	(0.71)	(0.39)	(0.39)	(0.53)
Anoka	-0.32	-0.18	0.00	0.19
	(0.73)	(0.39)	(0.35)	(0.47)
Dakota	1.74	0.99	1.12	1.71
	(0.76)	(0.64)	(0.62)	(0.66)