

The Best of Both Worlds: Combining RCTs with Structural Modeling

Petra E. Todd¹ and Kenneth I Wolpin²

January, 2020

¹University of Pennsylvania

²Rice University and University of Pennsylvania

Contents

1	Introduction	1
2	Early related literature	7
2.1	Early studies of the reliability of models to forecast decision-making	7
2.2	Studies of the reliability of nonexperimental evaluation estimators	8
3	Example of how to use the structural approach to perform an ex ante evaluation and to analyze the effects of alternative policy designs	10
3.0.1	A Simple Model of Welfare Participation:	10
4	Model Validation	17
5	Applications	20
5.1	Conditional cash transfer programs	20
5.1.1	RCT studies	20
5.1.2	Quasi-experimental studies	26
5.2	Welfare programs	27
5.2.1	RCT Studies	27
5.2.2	Quasi-experimental studies	31
5.3	Early childhood programs	32
5.3.1	RCT studies	32
5.3.2	Quasi-experimental studies	35
5.4	Relocation/migration subsidies	36
5.4.1	RCT studies	36
5.5	Other programs	39
5.5.1	RCT studies	39
5.5.2	Quasi-experimental studies	42
6	Evaluating effects of programs with spillover or general equilibrium effects	44
6.1	RCT studies	44
6.2	Quasi-experimental studies	47
7	Conclusion	48

Abstract

There is a long-standing debate about the extent to which economic theory should inform econometric modeling and estimation. This debate is particularly evident in the program/policy evaluation literature, where reduced form, experimental or quasi-experimental, and structural modeling approaches are viewed as rival methodologies. Proponents of the reduced form approach are often highly critical of the assumptions invoked in structural applications. Proponents of the structural approach point to the limitation of the reduced form approach in not being able to inform about program impacts prior to implementation or about the costs and benefits of program designs that deviate from the one that was implemented. In this paper, we argue that there is an emerging view of a natural synergy between these two approaches, melding them to exploit the advantages and ameliorate the disadvantages of each. We provide examples of how data from randomized controlled trials, the exemplar of reduced form practitioners, can be used to enhance the credibility of structural estimation. We also illustrate how the structural approach usefully complements experimental analyses by enabling evaluation of counterfactual policies/programs. Lastly, we survey many recent studies that combine these methodologies in various ways across different subfields within economics.

1 Introduction

The use of modern econometrics and computational methods in the practice of empirical research in economics has stimulated much debate. The history of this debate, spanning many decades, is exemplified by the titles of the following influential books: *Measurement Without Theory* (Koopmans, 1947), *Specification searches: Ad hoc inference with nonexperimental data* (Leamer, 1978), and *Mostly Harmless Econometrics: An Empiricist's Companion* (Angrist and Pischke, 2008). More recently, attention has focused on the choice of empirical methodologies for conducting research in policy/program evaluation. The distinguishing feature of alternative approaches to evaluation is the extent to which economic theory informs econometric modeling and estimation.

Popular terminology identifies one evaluation approach as "reduced form," which often invokes the notion of an "experiment" in that there is an identifiable group that is subject to the program or policy, a treatment group, and another group that is not, a comparison or control group. Reduced form analyses that are based on an explicit randomization, a randomized controlled trial (RCT) or on a (so-called) natural experiment are deemed to be experimental. Analyses not based on an explicit or natural randomization are deemed quasi-experimental (for example, the use of difference-in-difference, matching or regression discontinuity methodologies). The aim of a reduced form evaluation is to estimate the impacts of existing programs or policies. A second evaluation approach is popularly termed "structural," which generally consists of a fully specified behavioral model, usually though not necessarily parametric. The aim of the structural approach is to evaluate existing policies and to perform counterfactual program/policy experiments, including the evaluation of new hypothetical policies.¹

Reduced form and structural approaches have long been considered to be rival methodolo-

¹A critical feature of the structural approach is structural invariance, which is required for *ex ante* evaluation, that is, evaluation of a program that has not yet been implemented or is an untried modification of an existing program. *Ex post* evaluations within the structural approach make use of program variation, for example, observations on individuals before and after a program is implemented, and specify how the structure is altered due to the program. Structural invariance is not relevant for reduced form *ex post* analysis, which usually does not specify the mechanisms through which the program affects outcomes.

gies for conducting empirical economics research (e.g. Heckman, 2001, Angrist and Pishke, 2010). Proponents of the reduced form approach are often highly critical of the assumptions invoked in structural applications, whereas proponents of the structural approach argue that reduced form analyses are severely limited in their scope, not being able to inform about program impacts prior to implementation or about the costs and benefits of program designs that deviate from an existing program. In this paper, we argue the merits of an emerging view, that there is a natural synergy between experimental and structural approaches. A new literature over the last two decades, reviewed in this essay, combines these two approaches, exploiting the advantages and ameliorating the disadvantages of each.

Lalonde (1986), a precursor to the more recent literature, was among the first to exploit the synergy between experiments and structural estimation. Lalonde's insight was to use an RCT, in this case the National Supported Work (NSW) Demonstration training program, to test the validity of alternative nonexperimental estimators of program impacts. He found that different nonexperimental estimators yielded different impact estimates and, furthermore, that the estimates deviated substantially from the experimental benchmarks. Heckman and Hotz (1989) developed preprogram exogeneity tests that were useful in narrowing the range of estimates, but a wide range of estimates remained even after applying these tests.²

The perceived failure of nonexperimental methods to reproduce experimental results added to a prior literature critiquing the value of tightly connecting economic theory to estimation.³ Taken together, this cumulative body of work helped spur a movement that rejected the use of the structural approach, based on formal economic modeling, in favor of a reduced form, purely statistical, approach.⁴ Angrist and Pishke (2010) declared that a "credibility revolution" took place with researchers increasingly relying on experimental and

²Heckman, Ichimura and Todd (1997) argued that one reason that the estimators Lalonde (1989) considered did not perform well was that his data were not rich enough and that the econometric models perform better with better data.

³Results based on the estimation of demand systems, at least as far back as Stone (1954), generated a large literature questioning the empirical value of the neoclassical model of demand (see, for example, Blaug (1980)).

⁴The reduced form approach is sometimes referred to as "causal modeling," even though it eschews the modeling of mechanisms.

quasi-experimental research designs that were more credible than the structural approach.

The goal of the reduced form approach is to estimate the effects of existing treatments or policies on outcomes without having to specify a model of the mechanisms through which the treatment effect occurs. The approach requires data on a treated group and on an untreated comparison group. The main threat to validity in comparing the outcomes of the two groups is nonrandom treatment selection. To obtain a reliable estimate of the treatment effect requires either random assignment to treatment or some other exogenous assignment rule, such as a lottery, generating an instrument for treatment assignment.⁵ If done well, an RCT provides an unbiased estimate of the mean effect of treatment on the treated under minimal assumptions. RCTs can also be used to examine treatment impact heterogeneity in a straightforward way when the sample sizes are sufficient to permit subgroup analyses. However, as noted in prior research, RCTs also have significant limitations.⁶

In the context of this essay, the most relevant limitation of RCTs is their limited scope. RCTs are often costly, which makes it infeasible to extensively vary the treatment design or the length of treatment exposure within the experiment. Most often there is only a single treatment, as for example, in the Mexican PROGRESA program studied by Todd and Wolpin (2006) and by Attanasio et. al. (2011) and the Indian teacher incentive program studied by Duflo et. al. (2012). Researchers may be interested in the potential impacts and costs of a range of hypothetical programs with different design parameters, particularly if interest centers around designing a program that achieves some optimality criteria for a given budget. RCTs provide information on the particular design that was implemented and are typically uninformative about the potential costs and benefits from alternative program designs.

In the structural approach, the researcher specifies and estimates a formal economic behavioral model.⁷ The model structures that researchers use in empirical work vary according

⁵Heckman and Urzua (2009) provides a critical assessment of the role of instrumental variables in answering relevant economic questions.

⁶Leamer (1982) provides an early discussion of the interpretation of experimental results. Deaton (2010) critically reviews the role of field experiments in development economics. See Imbens (2009) for a response to both Deaton (2010) and Heckman and Urzua (2009).

⁷The theoretical basis for these models span both neoclassical and behavioral economics. The surveys by Keane, Todd and Wolpin (2011), primarily of the former, and DellaVigna

to the policy issue being addressed and include static or dynamic models as well as partial equilibrium or general equilibrium frameworks. The key limitation of the structural approach is that the estimation almost always relies on additional atheoretic assumptions about functional forms and error distributions, usually chosen for computational convenience. Even more fundamentally, researchers may disagree on the appropriate behavioral framework.⁸ These issues raise what is perhaps the most vexing problem in empirical research, namely model validation and selection.

Depending on the model specification, structural methods can be used for (i) analyzing the effect of extending the program to different subpopulations, (ii) simulating program impacts, costs and take-up rates under alternative program designs, (iii) analyzing program impacts over a time horizon that exceeds the length of time observed in the data and (iv) analyzing program impacts in the presence of spillover or general equilibrium effects. In some contexts, the structural method can also be used for purposes of *ex-ante* evaluation, that is, to predict the effects of a program intervention prior to its implementation. This makes it possible to study the potential impacts and costs of alternative program designs prior to implementing them.

This paper builds on a previous JEL survey by Heckman (2010) that described ways of “building bridges” and finding a “middle ground” between structural modeling and reduced form program evaluation approaches. In that paper, Heckman makes explicit the economics implicit in local average treatment effect (LATE) evaluation approaches and he proposes methods for moving beyond LATE to identify and estimate parameters of greater policy relevance. Drawing on a theorem of Vytlacil (2002) that shows that the LATE model of Imbens and Angrist (1994) is equivalent to a nonparametric version of the generalized Roy model, Heckman (2010) provides an economic interpretation of LATE within the Roy model framework. He surveys methods developed in Heckman and Vytlacil (2005), Heckman, Urzua and Vytlacil (2006), Cunha, Heckman and Navarro (2007), and Carniero, Hansen and Heckman (2003) for generalizing and extending LATE analysis for two-outcome and multiple-outcome models, including ordered and unordered choice models and he introduces policy relevant

(2018), of the latter, provide a number of examples.

⁸An example would be the choice of a unitary, collective or non-cooperative model of household decision making.

treatment effects (PRTE). Heckman (2010) emphasizes the value of placing the policy questions foremost and asking how the questions can be answered with statistics, rather than focusing on what parameters can be easily obtained with statistics and then asking if they happen to be policy relevant.

Along these lines, this paper takes the policy questions at the stages of designing, implementing and refining a program to be the central focus and develops and estimates behavioral models to be able to address such questions. The models we describe typically specify in greater detail than in Heckman's papers the theoretical mechanisms that determine outcomes and choices as well as program components. Imposing additional structure and functional form assumptions carries a risk of model misspecification, but it also provides the framework needed to carry out ex ante policy evaluation, to analyze changes in the design of program components, and to accommodate possible general equilibrium effects. We survey a variety of approaches in the recent literature for combining data from RCTs with structural modeling to increase the credibility of inference from such models. These approaches are illustrated in papers that span a number of subfields within economics.⁹

There are two ways that RCT data can be used to enhance the credibility of structural methods. The first is for purposes of model validation and selection, using either the treatment group or the control group as a "holdout" sample in performing out-of-sample model fit tests. Such a strategy mitigates the impact of data mining that is inherent in the formulation of structurally estimated empirical models and that limits the applicability of standard model selection criteria.¹⁰ When estimation is feasible without treatment variation (see below), the estimated model can be used to forecast the choices and outcomes of the holdout sample, either the treatment group or the control group, and the forecasts compared to the actual holdout sample data. If the model forecasts are "sufficiently" accurate, then the model is deemed to fit well and to be potentially useful for other purposes, such as analyzing the effect of varying parameters of the program design. A second way that researchers can use RCT data is to base estimation on both the treatment and control group. In this case,

⁹A parallel literature exploits (presumed exogenous) policy regime shifts in a manner similar to RCTs. Although our main focus is on RCTs, we present examples from that literature as well.

¹⁰See Keane and Wolpin (2007) and Schorfheide and Wolpin (2012, 2016).

variation induced by the treatment provides an additional, and sometimes necessary, source of variation for identifying and estimating model parameters and improving precision. These two approaches of using the RCT data can be combined, that is, researchers can first use either the control group or treatment group data as a holdout sample and then afterwards reestimate the model using both groups.

A requirement for combining these approaches is that the experimental data go beyond measurement of treatment status and outcomes. Successful empirical implementation of behavioral models requires that the key variables governing the decision-making described by the model be measured. For example, as part of the PROGRESA experiment in Mexico, the government collected extensive survey data from the families in both treatment and control villages, which allowed researchers to implement reduced form modeling strategies (including RCT, regression discontinuity and matching estimators) as well as to specify and structurally estimate rich models of family behavior that allow for counterfactual program analysis.

The paper develops as follows. Section two describes two earlier strands of literature that laid the groundwork for forecasting policy effects using behavioral models and for evaluating the models' performance against experimental benchmarks. Section three illustrates how and when structural models can be used for *ex ante* evaluation, discusses both nonparametric and parametric approaches, and considers the value of incorporating RCT data. Section four describes alternative approaches to assessing model validity. Section five surveys many recent studies across different subfields within economics that combine RCT/quasi-experimental and structural modeling approaches in different ways. Section six focusses on a smaller set of recent papers that develop models to account for spillover effects or general equilibrium effects in evaluating policy effects. Section seven concludes.

2 Early related literature

2.1 Early studies of the reliability of models to forecast decision-making

The problem of forecasting the effects of hypothetical social programs is part of the more general problem of studying the effects of policy changes prior to their implementation that was described by Marschak (1953) as one of the most challenging problems facing empirical economists.¹¹ In practice, in the early discrete choice literature, researchers used random utility models (RUMs) to predict the demand for a new good prior to its being introduced into the choice set.¹² Both theoretical and empirical criteria were applied to evaluate the performance of the models. Empirically, the model's performance could sometimes be assessed by comparing the model's predictions about demands for good with the *ex post* realized demand. In one of the earliest applications of this idea, McFadden (1977) used a RUM to forecast the demand for the San Francisco BART subway system prior to its being built and then checked the forecast's accuracy against the actual subway demand data. A later study by Lumsdaine, Stock and Wise (1992) studied the performance of alternative models at forecasting the impact of a pension bonus program on older workers' retirement. The authors first estimated the models using data gathered prior to the bonus program and then compared the models' forecasts to actual data on workers' departures.

There are a few early examples of empirical studies that compared the forecasting performance of economic models in estimating treatment effects to those obtained from randomized experiments. For example, Moffit (1979) used a labor supply model to forecast the effects of the Gary Negative Income Tax Experiment, which provided wage subsidies and income guarantees to low income people. Wise (1985) developed and estimated a housing demand model and used it to forecast the effects of a subsidy program. He compared his models' forecasts to the subsidy effects observed under a randomized experiment.

¹¹See Heckman (2001).

¹²Much of the initial empirical research was aimed at predicting the demand for transportation modes.

2.2 Studies of the reliability of nonexperimental evaluation estimators

As previously noted, there has been a long-standing debate in the literature over whether social programs can be reliably evaluated without a randomized experiment. Several of the early papers were in the context of evaluating job training programs. Lalonde's (1986) influential paper compared the performance of some standard econometric estimators against RCT benchmarks using data from the National Supported Work (NSW) experiment.¹³ The evaluation estimators he considered included cross-section, difference-in-difference and control function regression estimators applied to treatment data from NSW and comparison group samples drawn from the Panel Survey of Income Dynamics (PSID) and from the Current Population Survey (CPS). He found that the impact estimates he obtained differed across estimators and that the resulting range of estimates was too wide to be useful. Heckman and Hotz (1986) developed pre-program exogeneity tests that could be applied to rule out particular estimators. The approach they suggested is to estimate treatment effects using preprogram data when the program effects are known to be zero. Deviation from zero is taken as indicative of the estimator being biased.

Dehejia and Wahba (1999, 2002) also analyzed the NSW data, applying a class of estimators based on propensity-score matching.¹⁴ They found small biases and argued that matching estimators are more reliable than traditional econometric methods in reproducing the RCT results. However, Smith and Todd (2005), in a reanalysis of the data, found the Dehejia and Wahba (1999, 2002) results to be highly sensitive to their sample selection criteria.¹⁵

Heckman, Ichimura and Todd (1997) and Heckman, Ichimura, Smith and Todd (1998) applied matching estimators to data from the JTPA (Job Training Partnership Act) exper-

¹³The NSW program provided job training to unemployed, urban disadvantaged populations.

¹⁴These estimators were introduced in the statistics literature by Rosenbaum and Rubin (1983). Traditional propensity-score matching methods pair each program participant with a single nonparticipant, where pairs are chosen based on the degree of similarity in the estimated probabilities of participating in the program (the propensity scores).

¹⁵Most estimators, including the standard regression estimators considered by Lalonde (1986), exhibit small biases in the data subsamples that DW used for their analysis.

iment. They show that data quality is crucial to the performance of the estimator. The estimators were found to perform well in replicating RCT results only when they were applied to comparison group data satisfying the following criteria: (i) the same data sources (i.e., the same surveys or the same type of administrative data or both) are used for participants and nonparticipants, (ii) participants and nonparticipants reside in the same local labor markets, and (iii) the data contain a rich set of variables relevant to modeling the program participation decision. If the comparison group data fails to satisfy these criteria, the performance of the estimators diminishes greatly.

Glewwe, Kremer, Moulin, Zitzewitz (2004) estimated effects of introducing flip-charts in schools using both an RCT and a difference-in-difference approach. The RCT indicated the experimental treatment effect to be essentially zero in magnitude and precisely estimated, but the difference-in-difference estimator did not replicate the RCT results.

A recent study by Griffen and Todd (2017) compared experimental Head Start impact study estimates to nonexperimental estimates obtained using comparison group data from the ECLS-B. They applied both conventional regression evaluation estimators and matching estimators. Some of the estimators closely reproduced the experimental results. The difference-in-differences matching estimator exhibited the best overall performance in terms of low bias values and in capturing the pattern of statistically significant treatment effects.

In summary, the question of whether nonexperimental estimators offer a viable alternative to RCTs is still a matter of some debate. However, much evidence has been accumulated to provide guidance as to when a nonexperimental approach is likely to be successful. Having high quality survey data and a comparison group that is highly comparable to the treated group are important to any reliable estimation strategy.

The goal of this literature on nonexperimental estimators has largely been to estimate the effect of an existing program on a treated group. The frameworks developed do not specify the mechanisms through which the treatment effect occurs and in most cases are not suitable for studying the effects of modifying a program's design.

3 Example of how to use the structural approach to perform an *ex ante* evaluation and to analyze the effects of alternative policy designs

In this section, we illustrate with an example how the structural approach can be used for purposes of *ex ante* evaluation and for studying effects of alternative policy designs.

3.0.1 A Simple Model of Welfare Participation:

Nonparametric *ex ante* Evaluation We consider two states of the world, the current state where there is no welfare program and a hypothetical state with a welfare program. In the hypothetical state, there is a welfare benefit, $b(y_i, n_i) \geq 0$, offered to unmarried women with children who do not work; the benefit level depends on the woman's (denoted by i) nonearned income y_i and on the number of children n_i . In either state of the world, the woman decides whether to work or not. If she works $L_i = 0$, and if not $L_i = 1$. The woman's utility function, which is assumed not to depend on the state of the world (although see below), is given by

$$U_i = U(C_i, L_i; \epsilon_i) \tag{1}$$

where C_i is woman i 's consumption and ϵ_i shifts a woman's marginal utility of leisure relative to consumption. In the current state, a woman faces the budget constraint

$$C_i = y_i + w_i(1 - L_i), \tag{2}$$

where w_i is the woman's wage if she chooses to work. In the hypothetical state, the budget constraint reflects the additional potential income from the welfare program, namely

$$C_i = y_i + w_i(1 - L_i) + b(y_i, n_i)L_i. \tag{3}$$

A woman works if $U_i(L_i = 0|y_i, w_i, b(y_i, n_i), \epsilon_i) \geq U_i(L_i = 1|y_i, w_i, b(y_i, n_i), \epsilon_i)$, where $b(y_i, n_i) = 0$ in the current state.¹⁶ If the program is offered, the take-up rate depends on the number of eligible women (women for whom $b(y_i, n_i) > 0$) who choose not to work. The

¹⁶Todd and Wolpin (2008) and Wolpin (2013) consider a setting where the choice is continuous hours of work.

model implies that a woman who chooses not to work without the welfare program and who is eligible for the program is always better off choosing to take up welfare. We later consider the consequences for *ex ante* evaluation in an augmented model where there may be a stigma effect of taking welfare.

The basis for the nonparametric estimator stems from the simple insight that the budget constraint in the hypothetical state can be rewritten as

$$\begin{aligned} C_i &= (y_i + b(y_i, n_i)) + (w_i - b(y_i, n_i))(1 - L_i) \\ &= \tilde{y}_i + \tilde{w}_i(1 - L_i), \end{aligned} \tag{4}$$

Comparing equation 2 to 4, it can be seen that the form of the budget constraint is identical for both states of the world, with and without the welfare program. Under the assumption that the unobservable preference shifter (ϵ_i) is statistically independent of all observables, the implication of this observation is that given data in the no-welfare state, the effect of the hypothetical program can be estimated by comparing the employment status of women who have n_i children, nonearned income y_i and wage offer w_i to women also with n_i children, but with nonearned income $\tilde{y}_i = y_i + b(y_i, n_i)$ and wage offer $\tilde{w}_i = w_i - b(y_i, n_i)$.

Todd and Wolpin (2008) develop a matching estimator that can be used to recover the effect of the program for the situation where the program can be represented as a different parameterization of the existing budget constraint. Letting $H_i(y_i, w_i, b(y_i, n_i), \epsilon_i) = 1 - L_i(y_i, w_i, b(y_i, n_i), \epsilon_i)$, the matching estimator of the policy impact on the employment rate, based on a sample of J women, is

$$\hat{\Delta} = \frac{1}{n} \sum_{\substack{j=1 \\ j, i \in S_P}}^J \hat{E}[H_i | y_i = y_j + b(y_j, n_j), w_i = w_j - b(y_j, n_j)] - [H_j(y_j, w_j, n_j)], \tag{5}$$

where S_p is the region of overlapping support. For each woman, $j = 1, \dots, J$, in the sample with observed tuple (y_j, w_j, n_j) , we average the employment rate over all women with observed tuple $(y_i + b(y_i, n_i), w_i - b(y_i, n_i), n_i)$ and subtract the actual employment status of woman j .¹⁷ The impact of the program is the average of these differences over all J women in the

¹⁷It is actually not necessary to match on the number of children, but only on the combination of non-earned income and number of children that leads to a given welfare benefit.

sample.¹⁸

The matching estimator can be used to analyze the impact of a menu of policies by altering the benefit schedule. The only qualification to the estimation, given the model, is that the sample needs to be large enough for the matching to be credible.¹⁹ Given a menu of alternative program designs, a policy maker can choose a design to satisfy a particular social welfare function subject to any cost constraints.

Parametric *ex ante* Evaluation Extensions of the model that support nonparametric estimation are limited, because it is not always possible to represent programs in terms of the budget constraint in the no-program state.²⁰ Most researchers therefore adopt parametric models. Before considering an explicit case where nonparametric estimation is infeasible, it is useful to work through the estimation of the parametric model with the hypothetical welfare program considered above. The following structure establishes the conventional baseline parametric model:

$$\begin{aligned}
 U_i &= C_i + \alpha_i L_i + \lambda C_i L_i, \\
 \alpha_i &= x_i \beta + \epsilon_i, \\
 C_i &= (y_i + b(y_i, n_i)) + (w_i - b(y_i, n_i))(1 - L_i), \\
 w_i &= z_i \gamma + \eta_i,
 \end{aligned} \tag{6}$$

Matching on number of children would be necessary if fertility directly affects the work decision without welfare, for example, if the marginal utility of work depended on the number of children.

¹⁸The estimator can be modified to control for relevant conditioning variables by exact matching on those variables. Matches can only be performed for women whose y_j, w_j, n_j and associated $\tilde{y}_i, \tilde{w}_i, n_i$ values both lie in the support of y, w, n . TW (2008) demonstrate how this matching estimator can be implemented using kernel density functions for the matching.

¹⁹With sufficient sample size, it would be possible to also match women in terms of demographic variables such as age, race/ethnicity/education that could affect their preferences.

²⁰Wolpin (2013) discusses the viability of nonparametric *ex ante* evaluation under a variety of extensions of similar models, including allowing for partial observability of wages, fixed costs of work, childcare costs, kinked budget constraints, endogenous fertility and life cycle dynamics.

where, in addition to the terms previously defined, x_i is a vector of observed preference shifters, and z_i is a vector of observed and η_i unobserved determinants of wage offers. The wage function is specified to allow for the fact that only accepted wages are generally observed.²¹ The employment decision is determined by a comparison of the alternative-specific utilities, $U_i(L_i = 0)$ if the woman works and $U_i(L_i = 1)$ if the woman does not work:

$$\begin{aligned} U_i(L_i = 0) &= y_i + z_i\gamma + \eta_i, \\ U_i(L_i = 1) &= (1 + \lambda)(y_i + b(y_i, n_i)) + x_i\beta + \epsilon_i. \end{aligned} \quad (7)$$

The latent variable function, the difference in utilities, $U_i(L_i = 0) - U_i(L_i = 1)$, is thus

$$\begin{aligned} v_i^*(x_i, w_i, \eta_i, \epsilon_i) &= -\lambda(y_i + b(y_i, n_i)) + (z_i\gamma - b(y_i, n_i)) - x_i\beta + \eta_i - \epsilon_i \\ &= \xi_i^* + \xi_i \end{aligned} \quad (8)$$

where $\xi_i = \eta_i - \epsilon_i$, $\xi_i^* = -\lambda(y_i + b(y_i, n_i)) + (z_i\gamma - b(y_i, n_i)) - x_i\beta$.

To perform the *ex ante* analysis, set $b(y_i, n_i) = 0$. In that case, $\xi_i^* = -\lambda y_i + z_i\gamma - x_i\beta$, and the likelihood function for a sample of I women in the no-welfare state is

$$\mathcal{L}(\theta; x_i, z_i) = \prod_{i=1}^{I} \Pr(L_i = 0, w_i | x_i, z_i, y_i)^{1-L_i} \Pr(L_i = 1 | x_i, z_i, y_i)^{L_i}, \quad (9)$$

where θ is the parameter vector to be estimated, $\Pr(L_i = 0, w_i | x_i, z_i, y_i) = \Pr(\xi_i \geq -\xi_{it}^*(x_i, z_i, y_i) | \eta_i = w_i - z_i\gamma) f(\eta_i = w_i - z_i\gamma)$ with $f(\cdot)$ the density of η_i , and $\Pr(L_i = 1 | x_i, z_i, y_i) = \Pr(\xi_{it} < -\xi_{it}^*(x_i, z_i, y_i))$.²²

To complete the parameterization, assume that ϵ and η are joint normal with variance-covariance matrix, $\Lambda = \begin{pmatrix} \sigma_\epsilon^2 & \cdot \\ \sigma_{\epsilon\eta} & \sigma_\eta^2 \end{pmatrix}$. The parameters of the model to be estimated include $\beta, \gamma, \lambda, \sigma_\epsilon^2, \sigma_\eta^2$, and $\sigma_{\epsilon\eta}$. As is well known, joint normality is sufficient to identify the wage

²¹Todd and Wolpin (2008) and Wolpin (2013) show that a distributional assumption is required to perform an *ex ante* evaluation when wages are partially observed. Although the wage offer function can be estimated without distributional assumptions, the constant in the wage offer function, which is necessary for the *ex ante* evaluation, cannot be separately identified (see Heckman (1990), Wolpin (2013)).

²²The number of children only enters the model through the welfare schedule. Allowing for either a preference or cost of children, and assuming fertility is not a choice, does not change the conclusions from the analysis. As shown in Wolpin (2013), nonparametric *ex ante* evaluation is not feasible if fertility is a choice.

parameters (γ and σ_η^2) as well as $\frac{(\sigma_\eta^2 - \sigma_{\epsilon\eta})}{\sigma_\epsilon^2}$ (Heckman 1979). With the exclusion restriction that there is a variable in x that is not in z , identification doesn't have to rely solely on the distributional assumption. The data on work choices identify $\beta/\sigma_\xi, \gamma/\sigma_\xi$ and λ/σ_ξ . To identify σ_ξ , note that there are three possible types of variables that appear in the likelihood function, variables that appear only in z , that is, only in the wage function, variables that appear only in x , that is, only in the utility function, and variables that appear in both z and x . Having identified the parameters of the wage function (the γ 's), the identification of σ_ξ (and thus also $\sigma_{\epsilon\eta}$) requires the existence of at least one variable that appears only in the wage equation, a variable in z and not in x . With that exclusion restriction, all of the elements of ξ_i^* are identified.

The identification argument is independent of the existence of the welfare program. That is, the model parameters can be identified from data either with or without the program in place. With parameter estimates in hand, the *ex ante* impact of the welfare program on employment, $\Pr(L_i = 0|x_i, z_i, n_i, y_i, b(y_i, n_i)) - \Pr(L_i = 0|x_i, z_i, n_i, y_i, b(y_i, n_i) = 0)$, can be obtained for various welfare benefit schedules $b(y_i, n_i)$.

To understand the contribution of the parametric model, note that the hypothetical program considered above excluded working women from eligibility. Suppose, more realistically, that the program allows working women to receive welfare benefits, but that women who work are subject to reduced benefits that depend on their earnings. Specifically, assume that there is a benefit reduction (tax) rate that is proportional to earnings and that net benefits are given by $b(n_i, y_i) - \tau(n_i)w_i \geq 0$, where $\tau(n_i)$, the tax rate on earnings, depends on the number of children. The budget constraint in this case is

$$\begin{aligned}
C_i &= y_i + w_i(1 - L_i) + (b(y_i, n_i) - \tau(n_i)w_i)L_i, \\
&= (y_i + b(y_i, n_i)) + (w(1 + \tau(n_i)) - b)(1 - L_i) - \tau(n_i)w_i, \\
&= \tilde{y}_i + \tilde{w}_i(1 - L_i) - \tau(n_i)w_i.
\end{aligned} \tag{10}$$

Clearly, the form of the budget constraint no longer conforms to the case without the welfare program. Nonparametric estimation of the *ex ante* program effect using the previously described matching estimator is infeasible.

On the other hand, the parametric model parameters can be estimated in the absence of

any data on the welfare program and the model can be used to assess the policy effects of the welfare program with the earnings tax. The latent index governing labor supply decisions is given by

$$\begin{aligned} v_i^*(x_i, w_i, \eta_i, \epsilon_i) &= -\lambda y_i - (1 + \lambda)b(y_i, n_i) + z_i\gamma((1 + \lambda)\tau(n_i) + 1) - x_i\beta + ((1 + \lambda)\tau(n_i) + 1)\eta_i \\ &= \xi_i^* + \xi_i \end{aligned}$$

Using the estimated model, a policy maker can be provided with a menu of program options that vary the benefit schedule and tax rate. An ex-ante evaluation of their effects on employment, take-up rates and costs can be performed.

Most of the literature we will review adopts parametric models, either static or dynamic, of individuals' decision making processes. In the context of the previously described model, one dynamic extension would be to allow the wage offer function to depend on prior work experience. Additional extensions might include additional choices, such as schooling, fertility and marriage.²³ Assuming the woman maximizes discounted expected lifetime utility and future realizations of preferences and wage offers are unknown, the decision problem involves solving a discrete choice dynamic programming (DCDP) problem. There are now a number of survey articles that provide detailed discussions of available methods for estimating DCDP models.(See Keane et. al. (2011)) The focus of this essay is rather on how to synthesize the RCT and structural modeling approaches to best advantage.

Incorporating an RCT Suppose a government is contemplating the introduction of a welfare program. To better understand the program's impact on female employment, the government decides to do an RCT. Given the cost of conducting an RCT, the government chooses only one benefit schedule, $b(n_i, y_i)$ and sets $\tau(n_i) = 0$. The sampling frame includes all unmarried women with at least one child, independent of their employment status. Women are randomized into two groups, one of which is offered the program, the treatment group, and one of which is not, the control group. In addition, the government collects data on the women's wage histories, unearned income, fertility, marital status and employment.

²³For example, see Keane and Wolpin (2010).

The experimental impact estimates showed significantly lower employment rate after one year for women in the treatment group.

After completing the RCT, the government makes the data available to researchers. Given that the treatment effect has already been calculated (including for subgroups based on observable characteristics collected in the survey, e.g., race, education, employment histories, etc.), some researchers decide the data offer nothing more to study. They advise the government to do additional RCTs to study the impact of varying the benefit schedule. Other researchers begin work on developing estimable models for the purpose of evaluating variations in the program’s design.

The latter researchers have decisions to make with regard to choice of the model and estimation sample. Model selection is done through a process by which a researcher tries to improve the fit of a model during a model building phase by altering the model structure and re-assessing within-sample model fit. This process is sometimes referred to as data mining and it carries with it the dangers of models being overparameterized to fit the data. Given this process, many models with different structures will fit the data equally well. Standard errors are also incorrect if they do not account for the iterative model selection process. An alternative to within-sample fit statistics is to use a hold-out sample and look at an out-of-sample fit criterion.²⁴

To make the decision about whether to withhold some of the data in estimation, the researcher should have a model in mind. To see why, consider the previous model of welfare participation decisions augmented to include a direct effect of welfare participation on utility, that is, a stigma effect associated with program take-up. Specifically, let the utility function be given by

$$U_i = C_i + \alpha_i L_i + \lambda C_i L_i - \varphi_i P_i, \tag{12}$$

where $P_i = 1$ indicates that the woman takes up the program (conditional on eligibility), $P_i = 0$ if she does not and $\varphi_i = \bar{\varphi} + \omega_i$ is the woman’s disutility of participating in the welfare program (stigma). To make the point most clearly, assume that the program only applies to

²⁴From a Bayesian perspective, one would never hold out data; the marginal likelihood carries with it a penalty for models with more parameters (see Schorfheide and Wolpin (2012) for a discussion of these issues).

non-working women. In that case, the budget constraint is

$$C_i = y_i + w_i(1 - L_i) + b(y_i, n_i)P_i \quad (13)$$

The choice set for an eligible woman is now, work, $L_i = 0$, not work and take up the program, $L_i = 1$ and $P_i = 1$, or not work and not take up the program, $L_i = 1$ and $P_i = 0$. The alternative-specific utilities are:

$$U_i(L_i = 0) = y_i + z_i\gamma + \eta_i, \quad (14)$$

$$U_i(L_i = 1, P_i = 1) = 7(1 + \lambda)(y_i + b(y_i, n_i)) - \bar{\varphi} - \omega_i + x_i\beta + \epsilon_i, \quad (15)$$

$$U_i(L_i = 1, P_i = 0) = (1 + \lambda)y_i + x_i\beta + \epsilon_i. \quad (16)$$

As can be seen, the stigma effect is identified from the proportion of women who are eligible for the welfare program, but choose not to take it up.²⁵ It is clear that the stigma effect cannot be identified using only control group data and that estimating the model only using control group data will not generate accurate forecasts of program effects without good *a priori* evidence on $\bar{\varphi}$.²⁶

If a researcher commits to holding out either the treatment or the control group, all data mining in terms of model development must be based only on the subsample used for estimation. If all the data are used for estimation, then out-of-sample validation is eschewed. As the review of this literature demonstrates, there does not seem to be a consensus, certainly on models, but also on the choice of estimation sample.

4 Model Validation

As illustrated above, a major benefit of a structural modeling approach is that it allows for *ex ante* evaluation of policy interventions as well as consideration of alternative policy

²⁵A woman will not take up welfare, $P_i = 0$, if $\omega_i \geq (1 + \lambda)b(y_i, n_i) - \bar{\varphi}$, and will take it up otherwise.

²⁶However, evidence for the existence of stigma based on "eligible" women not taking up the program relies on there not being significant measurement error in the data used to infer eligibility. Women classified as eligible may be observed not to take up the program because they are in fact not eligible, which could rationalize a model in which there is no stigma. It would be possible to estimate the classification error only using the treatment group data, in which case the control group could serve as a holdout sample.

designs and eligibility criteria. However, models typically rely on extra-theoretic modeling and distributional assumptions, so model validation is an important concern.

There are primarily three different approaches that researchers take to assess model validity. The first is to check robustness to alternative modeling assumptions, which was a suggestion of Leamer (1983). This requires estimating many different versions of the model and comparing the results obtained, which can be computationally intensive.

A second traditional way of considering model validity is to examine within-sample fit. Once the model parameters are estimated, including the parameters of the distributions of any unobservables, it is possible to use the estimated model to simulate the choices and outcomes for any group of individuals. To examine the within-sample model fit, one compares the actual choices and outcomes observed in the data to that simulated under the model. Formal tests of within-sample fit can be conducted (for example, Pearson chi-square test).²⁷ Such tests, however, are biased towards not rejecting the model when the researcher has engaged in data mining.

A third way of evaluating a model's validity is to use a holdout sample. Under this approach, the model is estimated based on a data subsample and then used to predict the behavior of the part of the sample that was held out. In the case of an RCT, the use of a holdout sample as a validation tool has strong intuitive appeal. The RCT alters the structure of the decision problem faced by the agents in the treatment group and simultaneously ensures that distribution of unobservables is the "same" across treatment and control groups. Depending on whether there is observable variation that can mimic the structural change (as in altering the budget constraint in the welfare example above), it may be possible to recover the model parameters using only the control (or treatment) group. To be able to accurately forecast the reaction of agents to the treatment based only on data from the control sample would seem to be a non-trivial test of the model. Forecasting the treatment effect from the control sample would also appear to provide a basis for selecting among (or combining) models.

To our knowledge, Schorfheide and Wolpin (2016) is the only paper to go beyond the

²⁷In the context of structural estimation, it is formally necessary to adjust degrees of freedom of the test for estimated parameters. See Heckman (1984) and Andrews (1988).

intuitive argument and provide a formal justification for the use of a holdout sample. Their approach is to cast the problem of model selection as a principal-agent problem. A policy maker, the principal, would like to predict the effects of a treatment at varying treatment levels. The data are available to the policy maker from an RCT that has been conducted for a single treatment level. To assess the impact of alternative treatments, the policy maker engages two modelers, the agents, each of whom estimates their preferred structural model and provides measures of predictive fit. Modelers are rewarded in terms of model fit. SW consider two data venues available to the policy maker. In the first, the no-holdout venue, the modelers have access to the full sample of observations and are evaluated based on the marginal likelihood function they report, which, in a Bayesian framework, is used to update model probabilities. Because the modelers have access to the full sample, there is an incentive to modify their model specifications and thus overstate the marginal likelihood values. SW refer to this behavior as data mining. More specifically, data mining takes the form of data-based modifications of the prior distributions used to obtain posteriors. In the second, the holdout venue, on the other hand, the modelers have access only to a subset of observations and are asked by the policy maker to predict features of the sample that is held out for model evaluation. Data mining creates a trade-off between providing the full sample, which would otherwise be optimal for prediction, and withholding data. SW provide a qualitative characterization of the behavior of the modelers under the two venues based on analytical derivations and use a numerical example to illustrate how the size and the composition (in terms of observations from the control and treatment groups) of the holdout sample affects the risk of the policy maker. Their numerical example shows that it is possible for the holdout venue to dominate the no-holdout venue because of the data mining that occurs if the modelers have access to the full sample. The lowest level of risk in their example is attained by holding back 50% of the sample (where the control and treatment sample are of equal size) and providing the modelers only with data either from the control or from the treatment group.

In what follows, we review papers from the new literature that combines structural modeling with data from RCTs or from quasi-experiments. Some of the studies use all the data in estimation and some exclude either the treatment or control group for use as a holdout

sample. The estimated models are used for various purposes. Most studies use the model estimates to evaluate the effects of policies that deviate in some ways from the policy that was implemented, as described in the previous example. However, some papers are also concerned with spillover effects from treated individuals onto untreated individuals or with general equilibrium effects arising from demand and supply side market responses and they develop to account for these possibilities.

5 Applications

5.1 Conditional cash transfer programs

One class of programs that has been studied using the structural approach is conditional cash transfer (CCT) programs, particularly in the area of education. We first describe two dynamic models that were developed and estimated to study the effects of the Progresa CCT program in rural areas of Mexico on schooling, labor supply and fertility outcomes. Then we describe a simpler static model that was also used to study impacts of CCT programs in Mexico and Ecuador on school and child work choices. Third, we describe a model that was developed to study teacher attendance decisions in India and to analyze the effect of a teacher attendance subsidy and bonus program. Each of these three studies exploits data from an RCT in different ways to estimate and validate the structural model and then uses the model to perform a range of counterfactual experiments. Lastly, we describe a study of the Progresa CCT program that uses quasi-experimental data from urban areas to study food demand.

5.1.1 RCT studies

Effects of the Progresa program on schooling and fertility outcomes In 1997, the Mexican government introduced a conditional cash transfer (CCT) program in rural areas that provided a subsidy to families for each child that met a school attendance goal. The initial program, called PROGRESA was afterwards extended to urban areas (and renamed Oportunidades and later Prospera). Similar programs have been adopted in numerous other countries (for example, in Bangladesh, Brazil, Colombia, Guatemala, Malawi, Nicaragua, and Pakistan).

To evaluate the initial program, the Mexican government conducted a randomized social experiment, in which 506 rural villages were randomly assigned to either participate in the program or serve as controls. Randomization, under ideal conditions, allows mean program impacts to be assessed through simple comparisons of outcomes for treatments and controls. The program was effective in increasing school attendance; treatment effects, measured as the difference in average attendance rates of children in the treatment and control villages one year after the program, ranged from 5 to 15 percentage points depending on age and sex (Behrman et. al. 2005, Schultz 2004).

An important limitation of large scale social experiments, such as PROGRESA, is that it is often prohibitively costly to vary the experimental treatments in a way that permits evaluation of a variety of policies of interest. In the PROGRESA experiment, all eligible treatment group households faced the same subsidy schedule, so it is not possible to evaluate the effects of alternative subsidy schemes through simple treatment-control comparisons. In addition, because the experiment lasted only two years, one cannot directly assess the long term impacts of the program on completed schooling.

Todd and Wolpin (TW) and Attanasio et. al. (AMS) analyze the impact of the PROGRESA program on school attendance via the estimation of a DCDP model of decision-making about children's schooling. They use their model estimates to compare the effects of the existing subsidy program to the effects of various alternative program designs. Both papers adopt the DCDP approach, use data derived from the same source and perform similar counterfactual exercises; however, the models used differ non-trivially in their structure.

We first provide a general description of the PROGRESA data and then describe the two models, their different approaches to using the data, and their empirical findings. A baseline survey was conducted in October 1997 of all households in both the treatment and control villages prior to the implementation of the program. The experiment began in the 1998/99 school year and continued for two years.²⁸ The program (which included a child health component as well) provided benefits that, on average, amounted to about 25% of family income. The school attendance subsidy component amounted to about 75% of total

²⁸Within the treatment villages, only households that satisfied an eligibility criterion based on a "marginality" index were provided with the subsidy.

payments. The subsidy began at grade 3 and increased with each additional completed year of schooling to offset the increased opportunity cost of attending school as children become older. The subsidy level was the same for girls and boys up to grade 6, but was larger for girls in grades 7 to 9.

In the TW model, each year a married couple decides on whether each of their children between the ages of 6 and 15 will attend school, remain at home or, for those age 12 to 15, work in the labor market (mutually exclusive). They also decide whether the wife will become pregnant (while fecund). The couple receives utility in each period from their current stock of children, their children's current years of schooling, their school attendance, and from any children at home. There is also a utility cost to attending school (grades 7-9) that depends on the distance from the village to a school. Households differ in their preferences for the choice variables according to their discrete "type" and households have time-varying preference shocks. The household's income includes the parent's income and the wage income of the children who work. A child's wage (offer) depends on the child's age and sex, the distance to the nearest city, household type and unobserved shocks. Model parameters are estimated by simulated maximum likelihood.

The AMS model also includes the binary choice of school or work (excluding the "at home" option), but, unlike the TW model, assumes that each child's utility is maximized independently of that of the parents or of other children. The school/work decision is made at each age from 6 to 17, at which time there is a terminal payoff that depends on the number of years of schooling completed. The child receives a wage offer in each period that is village/education/age-specific. If the child rejects the wage offer and instead attends school, the child receives a utility payoff (positive or negative) that depends on observable preference shifters (parental background, the child's age and the state of residence), the number of years of past attendance, on observable variables that affect the cost of attending primary or secondary school (distance to a secondary school) on a child's unobserved discrete preference "type" and on a time-varying preference shock that is assumed to be distributed as extreme value.

AMS putatively allow for a direct effect of the program on school attendance utility,

either a "feel good" effect from participating in the program or a "stigma" effect.²⁹ As noted in our welfare example presented above, the possibility that there may be an intrinsic value of participation in the program *per se* would require that both the treatment and control households are used in estimation. AMS use both groups in estimation. In contrast, TW, hold out treated households from the estimation sample, using these households for external validation of the model.

TW compare the predicted effects of the PROGRESA program on completed schooling, as implemented, with that of alternative programs. Model simulations of households from the time of marriage until the last born child reaches age 16 show that the average years of completed schooling in the absence of the program would be 6.29 for girls and 6.42 for boys and that 19.8 percent of girls and 22.8 percent of boys would have completed the 9th grade . The model predicts an increase in completed schooling of about one-half year for both boys and girls, or 26.0 percent of the maximal potential increase for girls and 28.9 percent for boys.³⁰

As noted, the PROGRESA subsidy schedule rewards school attendance starting at grade 3. However, attendance in grades 3 -5 is almost universal, making the subsidy at early grade levels essentially an income transfer. TW calculated that the per-family cost of the program could be held roughly constant if the subsidy in grades 3-5 were eliminated and the subsidy in grades 6-9 were increased by about 45 percent. Under the modified plan, the proportion of girls completing ninth grade increases by 3.4 percentage points and proportion of boys by 3.8 percent, although there was a small decline in the proportion of children who complete at least sixth grade. TW also use the model to evaluate alternative hypothetical programs, such as a bonus for completing 9th grade, a school building program that decreases the distances that study need to travel to attend school and a family unconditional income transfer program.

²⁹As Wolpin (2013) points out, because AMS do not fully specify the constrained optimization problem, it turns out that their model is observationally equivalent to one in which there is no direct program effect on utility. Thus, estimation of the partial equilibrium decision model did not require both the treatment and control groups.

³⁰Interestingly, this estimate corresponds closely with that obtained by Behrman et. al. (2005) and Schultz (2004) using non-structural approaches.

AMS perform two counterfactuals. As in TW: (i) they simulate the impact of eliminating the subsidy to primary school and redistributing the savings to increase the subsidies at later grades and (ii) they simulate the impact of building schools. Like TW, they find the effect of the first counterfactual to be large. However, the metric used by AMS is not directly comparable to that of TW. They find that the budget-neutral effect of eliminating the subsidy at younger ages increases school attendance rates by as much as 100 percent at age 15-16. Also consistent with TW, AMS find a large effect of building schools on school attendance of older children as in their modified program. The TW and AMS findings are, perhaps, surprisingly robust given the quite significant differences between the model structures and estimation samples.

Effects of CCTs on schooling and work in Mexico and Ecuador

A study by Leite, Narayan and Skoufias (2015) uses microsimulation methods to perform an *ex ante* evaluation of conditional cash transfer program impacts on school enrollment and child working. The model they use is based on an earlier model originally developed in Bourguignon, Ferreira and Leite (2003) in studying effects of the Bolsa Escola conditional cash transfer program in Brazil. The model is a static discrete choice random utility model where the options are for each child to not attend school, to combine/schooling and working, or to only attend school. The model assumes that decisions to send a child to school are independent of parents' working decisions, and that decisions about multiple siblings are made independently and that family composition is exogenous. The utility depends on family income, inclusive of child wages and any program transfers associated with the alternative school/work choice combinations. The model incorporates a "means-test" to approximate program eligibility.

Estimation of the model does not require panel data and is much less demanding in terms of computational complexity than the TW and AMS models described above. Nevertheless, when Leite et. al. (2015) compare the *ex ante* predictions from the model to experimental benchmark estimates from the Mexican Progresa experiment and the Bono de Desarrollo program in Ecuador, they find that the model produces reliable forecasts.³¹

³¹Under the Bono de Desarrollo program, beneficiary households receive grants of \$15 per month under the conditions that children of ages 6-16 years are regularly enrolled in

Effects of teacher attendance subsidies in India

Duflo, Hanna and Ryan (2012) analyze the impact of financial incentives and attendance monitoring of teachers in rural India on teacher absenteeism. In September 2003, an NGO implemented a randomized controlled trial in which 60 of 120 schools were randomly assigned to a treatment group in which teacher monthly salaries were determined by a non-linear function of the days in the month they were in attendance. Attendance in the treatment group was monitored by requiring a photograph be taken of the teacher and students at the beginning and end of each school day using a camera with tamper-proof time and date functions. The salary structure consisted of a flat payment for attending 20 days in the month, a 5 percent bonus payment for each day above 20 (about 3 days per-month on average) and a 5 percent penalty for each day below 20 (up to 10 days missed). Teachers in the control group schools faced the same flat payment, with neither a bonus for additional days above 20 or a penalty for days fewer than 20. Attendance was monitored through random checks once a month and control group teachers were reminded that they could be fired for excessive absences.

Duflo et. al. (2012) specify a finite horizon DCDP model of the teacher's daily decision about attendance. They estimate different specifications, including i.i.d. preference shocks to absence days, observed and unobserved heterogeneity in preferences, and serially correlated preference shocks. As seen in Table 1, they estimate the model on the treatment group and use the control group to select and validate the model specification. Their choice of the treatment group for the estimation is based on the RCT design, namely that there are essentially two treatments, the financial incentive and the use of the camera for monitoring absences. Estimating the model on the control group alone would, if there is the additional cost of monitoring from the use of the camera technology, lead to an understatement of the counterfactual impact of the program based only on implementing the financial incentive. However, validating the model based on estimates using the treatment group and predicting attendance of the control group is not affected by the combined treatment. The validation

school with an attendance rate of at least 80% per month and children of ages 0-5 years make scheduled visits to health centers. Coverage reached one million households (5 million people). Schady and Araujo (2008) present experiment impact estimates showing positive effects on school enrollment and negative effects on child work.

exercise identified two specifications as being similar in their out-of-sample performance.

As noted, the RCT included only one form of financial incentives. However, it is possible, given model estimates, to calculate the optimal incentive scheme, that is, the financial incentive structure that produces the same absentee rate at least cost. When they use the model to find the optimal incentive structure, they find that the optimal structure saves 22 percent of the average cost associated with the incentive structure that was implemented in the experiment.

5.1.2 Quasi-experimental studies

Effects of the Progresa program on food demand

Angelucci and Attanasio (2013) analyze the effect of the Progresa cash transfer program on food demand in urban areas of Mexico. Their data come from a quasi-experiment that made the program available to households in some localities but not in others. The authors use two different evaluation estimators - a propensity score matching estimator and a difference-in-difference estimator - applied to longitudinal data from households in treated localities and matched control localities. The matching estimates show that the program led to an increase in the food expenditure share and an increase in high protein food consumption.

One of the goals of the paper is to assess whether a standard Engel curve model could be used to do an *ex ante* prediction of the treatment impacts. The Engel curves relate food expenditure shares and high protein food expenditure shares to total expenditure. They estimate Engel curve demand models using data on control households collected at times before and after the program and on treatment households collected prior to the program. The parameter estimates indicate the food is a necessity and high protein foods are a luxury.

The impact estimates based on the matching and difference-in-difference estimators showed the treatment group increased their expenditure share on food, which is inconsistent with the decline predicted by the estimated Engel curve. For high-protein food, the quasi-experimental evidence and the estimated Engel curve both predict an increase, but the treatment effect estimate is larger in magnitude than that predicted by the Engel curve. When the Engel curve is estimated on the treatment group before and after the program, the parameter estimates change substantially, which the authors interpret as additional evidence

of model misspecification.

The authors hypothesize that the Engel curve which assumed to represent consumption demand of a unitary household, does not account for the fact that the Progresa cash transfers were given to women and that decision-making within the household may instead be the result of a bargaining process. When they reestimate the Engel curve only on the subset of single female-headed households, they find that the parameter estimates using treatment group data from before and after the program are stable.

5.2 Welfare programs

5.2.1 RCT Studies

Effects of cash transfers in India

Alatas et. al. (2016) analyze the effects of a cash transfer program in Indonesia called PHK, specifically, how the mechanism that is used to enroll people affects program take-up rates and impacts. PHK enrolled 2.4 million households, each receiving \$130 per year for 6 years, with eligibility determined on the basis of an asset test. The authors carried out an RCT that varied the enrollment process across villages. In 200 treated villages, they introduced a “self-targeting” scheme, whereby households had to travel to apply for the program at a registration site and to take an asset test to determine eligibility. The RCT also randomly varied the application costs across treated households by varying the distance needed to travel to the registration site. In 200 control villages, they followed the usual government “automatic screening” procedure with program administrators visiting potential beneficiaries at their homes to determine eligibility.

The RCT revealed that the different enrollment schemes result in very different patterns of program participation. Per capita household consumption is lower for participating households in the treated villages than in the control villages. In fact, the very poorest households, as measured by per capita consumption, were twice as likely to receive benefits in the self-targeting scheme. However, only about 60% of eligible households apply under self-targeting, so the coverage rate is also lower.

To better understand the mechanisms generating the enrollment patterns, the authors develop and estimate a discrete choice model of the household’s program application deci-

sion under uncertainty about whether they will they will pass the asset test. In the model, households weigh the expected benefits of applying against the costs, inclusive of any distance travel costs. The model incorporates two types of households - sophisticated and unsophisticated - with sophisticated households being better informed about the income components that comprise the asset-based eligibility test. As seen in the table, the discrete choice model is estimated only using treatment group data and model fit is assessed with within-sample fit tests.

Simulations from the estimated model show that a key factor driving the selection of poorer households into the program under self-targeting is that rich households forecast that they have a small likelihood of receiving benefits and therefore do not apply when there is an application cost. The estimates show that a small distance cost is effective in targeting the program to the poorest households and that further increasing the distance cost has no additional targeting benefit. The authors also use the estimated model to examine how application decisions change when the fraction of sophisticated households is increased and when households change their expectation of receiving benefits. Lastly, they compare how the two types of enrollment schemes influence the poverty gap. They find that it is possible to achieve a 29 to 41 percent greater reduction in the poverty gap under self-targeting than under automatic screening with an identical budget.

Evaluating effects of an earnings supplement in Canada

Lise, Seitz and Smith (2003) use data from an RCT, the Canadian Self-Sufficiency Project (SSP), to evaluate the performance of a job search/matching model in forecasting the effect of a wage subsidy given to welfare recipients on employment. The SSP provided an earnings supplement for up to 3 years for individuals receiving Income Assistance (IA), the Canadian welfare program, to get full-time employment within a 12-month time period. The data contain information on 5,685 recipients: 2,827 control group members and 2,858 treatment group members. The study focuses on a subsample of 3,346 single women who were regularly included in follow-up surveys.

First, the authors calibrate a standard job search model, in the style of Pissarides (2000), only using data from the control group. The key model parameters are the discount factor, search friction parameters, and exogenous job separation rates. Second, they use the model

to simulate the behavior of the treatment group and they compare the predictions with data obtained from the RCT. They simulate the behavior in partial equilibrium, because the experiment only affected a small subset of the economy and is therefore not expected to have equilibrium impacts. Lastly, the authors recalibrate their model combining the data from the control and treatment groups and examine changes in the parameters and model fit. The analysis is done separately for the provinces of New Brunswick and British Columbia, because the labor markets and unemployment benefits programs differ across provinces.

The authors examine outcomes related to job search intensity, job destruction and earnings. In British Columbia, they find that the impacts of SSP on the IA-to-work transition rates predicted by the model match very well the transition rates observed under the experiment. However, the predictions are less accurate for New Brunswick, for which the model predicts a higher transition rate than observed in the data. The study finds that the search effort cost must be higher in New Brunswick than in British Columbia to match the data. With regard to job destruction, the authors find support for the assumption of a constant job destruction rate, because there is no statistically significant difference in the employment survival rates for the treatment and control groups and also no change observed when the treatment group stops receiving supplemental payments. In terms of earnings, the hourly earnings rate did not differ for the treatment/control groups but the treatment group worked longer hours. Lastly, including the treatment group in calibrating the model changes the parameter estimates for New Brunswick but not for British Columbia and leads to a better within-sample model fit.

Evaluating effects of welfare policy changes in Minnesota and Vermont

Choi (2018) uses data from two state welfare reform experiments conducted by MDRC during the mid 1990s - the Minnesota Family Investment Project (MFIP) and the Vermont Welfare Restructuring Project (WRP)) - to assess a structural model's performance in forecasting the effects of welfare rule changes. The paper develops and estimates static discrete choice models of labor supply and welfare participation that incorporate heterogeneity in preferences, fixed costs of working and disutility associated with welfare take-up. The welfare policy impacts estimated under the RCT are used as a benchmark for the structural model predictions.

The model is a static labor supply/welfare participation model. The model assumes that individuals face a finite and discrete choice set of choices.³² The utility function is quadratic in hours and consumption and includes an interaction term (to allow consumption and leisure to be complements or substitutes). Consumption depends on earned income, taxes, EITC, and welfare benefits.

In the two state experiments, individuals assigned to the control group received the standard AFDC (Aid to Families with Dependent Children) program, which has a 100% welfare benefit reduction rate for every dollar earned. Individuals assigned to the treatment groups faced lower benefit reduction rates - 62 percent in MN and 75 percent in VT. The lower benefit reduction rate generates an income effect and a wage effect and will increase work if the wage effect dominates.

The MFIP and WRP samples include 14170 and 7691 individuals in the three program groups. Baseline survey data were collected prior to random assignment and at two follow-up surveys, 36 months and 42 months after random assignment. The model is estimated using data from the control group in MN. Model parameters are identified from cross-section variation across individuals in hours of work and welfare participation. The stigma effect is identified, because some of the eligible controls choose not to participate in AFDC.

After estimating six different model specifications, Choi (2018) uses the estimated models to predict welfare policy impacts both in MN (within-state) and in VT (cross-state). The RCT in MN ensures that the distribution of unobservables for the control and treatment groups in MN are comparable. However, performing the cross-state prediction requires an additional assumption that any unobservable factors governing labor supply and welfare participation decisions are similar in MN and VT.

Choi (2018) finds that some of the specifications she estimates provide a very good within-sample fit to labor supply and welfare participation patterns, particularly the specifications that incorporate a fixed cost of working. However, the model's out-of-sample predictions

³²The discrete choice assumption avoids the analytical difficulties created by nonlinear budget constraints with convex and nonconvex kinks. Similar models have been estimated by Fraker and Moffitt (1988), van Soest (1995), Hoynes (1996), Keane and Moffitt (1998), Gong and van Soest (2002), Creedy and Kalb (2005), Brewer et al. (2006), and Blundell and Shephard (2012).

of the policy treatment effects are not good, either within-state or cross-state. Specifically, the RCT in MN found that the decrease in the welfare benefit reduction rate induced a substantial decrease in hours of work, while the estimated models predict either a small decrease or an increase. In VT, the RCT showed no change in welfare participation patterns, whereas the model predicts increases. The study concludes that a good within-sample fit is not necessarily indicative of good out-of-sample predictions.

5.2.2 Quasi-experimental studies

Evaluating effects of a welfare policy change in Canada

Hansen and Liu (2015) estimate a model of labor supply and welfare participation to perform an *ex ante* evaluation of a 1989 Canadian welfare reform. Prior to the reform, welfare benefits were much less generous for people less than 30 years of age than for similar people 30 or older. The reform eliminated age discrimination in benefit levels and increased the average monthly benefit for younger individuals from \$185 to \$507. The authors estimate a static discrete choice model where individuals choose among 7 different hours of work options and whether to participate in welfare. The model also includes a stigma effect of participating in welfare. It accounts for the detailed budget sets for each welfare work combination as well as income taxes. Model parameters are estimated by maximum likelihood using a sample of single men from Quebec collected prior to the reform (from the 1986 Canadian Census).

The authors perform an out-of-sample fit test of the model by comparing the model's predictions of the reform impacts to those obtained using a regression discontinuity (RD) estimator applied to post-reform data, exploiting the age discontinuity. They find that the estimated model predicts the employment reduction and the increase in welfare participation associated with the reform. The largest policy effects are for lower income individuals for whom there is a 4.5% decrease in employment and a 4.9% increase in welfare participation.

The authors also use the model to study how employment, welfare use and hours of work would change as social assistance benefits are further increased. They find the responses to be highly nonlinear with respect to benefit increases. In addition, they use the model to explore how labor supply and welfare participation changes in response to changes in the income tax system.

Evaluating effects of welfare policy among states in the U.S.

There is a large literature on structurally estimating models to assess the impact of welfare programs in the U.S. on economic and social outcomes. Of particular interest in the context of this survey is the approach adopted for model validation by Keane and Wolpin (2007). In the other papers discussed above, the use of a holdout sample was based on the design of the program, either one of the RCT groups or either the group before or after the implementation of a program. In both of these cases, the treatment and control groups are thought of as comparable in terms of the sample distributions of unobservables (in the second, at least conditional on observables). Keane and Wolpin (2007) instead explicitly choose a non-random sample as the comparison group, specifically a sample with a considerably different level of treatment. In this case, the treatment level varies across states by the benefit generosity of the welfare program (AFDC) in the state. The holdout sample is a state (Texas) that, relative to the set of states in the "treatment" sample, provides considerably less generous benefits. The notion is that forecasting well the effect of a program far outside the range of the program parameters of the estimation sample should be a more demanding out-of-sample validation criterion. The authors conclude that their DCDP behavioral model produced plausible forecasts, more plausible than a purely statistical model. Their follow-up paper, Keane and Wolpin (2011) provided an analysis of the impact of the AFDC program and counterfactual policies on program take-up, labor supply, wages, fertility and marriage.

5.3 Early childhood programs

5.3.1 RCT studies

A home visitation/parenting program in Colombia

Orazio Attanasio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina (2020) study the effects of a randomized early childhood intervention in Columbia that was offered to households participating the Columbian CCT program *Familias en Accion*. The intervention was targeted at children age 12-24 months and consisted of weekly home visits (one hour per week) aimed at improving parenting skills and providing micronutrient supplementation.³³ The data were gathered by a household survey, tests that were

³³This type of intervention was shown to be effective in the Jamaica Study (Grantham-

administered to the children, and by interviewer observations. 1429 children in total were randomized into four groups: (i) one group received only the psychosocial stimulation program, (ii) one group received only the micronutrient intervention, (iii) one group received both (i) and (iii), and (iv) a control group. Attanasio et. al. (2018) reports the RCT impact estimates that showed significant effects of the psychosocial intervention on child outcomes but no effects of micronutrient supplementation. Therefore, in Attanasio et. al. (2020), groups (i) and (iii) and groups (ii) and (iv) are combined.

The goal of the study is to elucidate the mechanisms underlying the observed treatment impacts. To this end, the authors develop a model of the cognitive and socioemotional skill production technology along with parental investment decision-rules. The inputs in the production function model are baseline child skills, maternal skills, and material and quality time investments in the child. The production function also includes the presence of other siblings in the family who might reduce attention available for the focal child. The model incorporates a latent factor structure to combine multiple outcome measures and input measures and also allows for measurement error in inputs and outputs.³⁴ Some of the estimated specifications allow for material and time investments to be endogenous, using as instruments prices of toys and food and maternal exposure to violence.

The empirical analysis has two primary goals. The first is to understand the nature of the production function in this high poverty context. The second is to ascertain whether the positive treatment effects occurred because of changes in the production function, changes in parental investment decisions, or changes the mother's characteristics (e.g. rates of depression or socio-emotional skills).³⁵ The paper also decomposes production function changes into changes in TFP, changes in other parameters and a direct effect of the treatment, possibly operating through the one-hour home visits.

With regard to the skill production technology, the study finds that the current stock of

McGregor, Powell, Walker, and Himes, J., 1991) and in the Perry Preschool Program (Heckman, Moon, Pinto, Savelyev, and Yavitz, A., 2010). A difference in the Colombian program, however, was that the home visits were conducted by local women who received training but did not otherwise have expertise in child development.

³⁴The approach is similar to that of Cunha, Heckman, and Schennach (2010).

³⁵Because the treatment is allowed to affect model parameters, it is not possible to estimate the model using only the control group data and ex ante evaluation is not possible.

cognitive (socio- emotional) skills strongly affects the development of future cognitive (socio- emotional) skills. This is called self-productivity of skills using the terminology of Cunha, Heckman, and Schennach (2010). Second, the estimates show that the current stock of cognitive skills fosters the development of future socio-emotional skills, but not the reverse.

The treatment intervention increased the cognitive development of the children by 0.115 log points and socio-emotional development by 0.087 log points. The authors' preferred production function estimates imply that the parental investment increases (both in material and time) induced by the program account for around 91% of the intervention impact on cognition and at least 66% of its impact on socio-emotional skills. The parental investment increases were greater for children with higher initial baseline skills and for more highly skilled mothers. There is no evidence of a direct effect of the program and also no evidence that the program led to significant changes in the mothers' characteristics. The study concludes that the involvement of the parents and induced increases in parental investments were the key to the program's success.

An income and child care subsidy program in Wisconsin

Welfare programs with work requirements often require parents to make greater use of external child-care, raising concerns about how children are affected by such programs. Some of the best evidence on this issue comes from an RCT implemented by MDRC used to evaluate the New Hope program in Milwaukee, Wisconsin. The sample consisted of 1,357 individuals; 678 were randomly selected to the treatment group and 679 to the control group. Data were collected from the families at baseline and up to eight years after. The New Hope program randomly allocated an income subsidy, similar to the EITC, and a child care subsidy to disadvantaged families with a requirement to engage in full-time work. To be eligible, individuals had to be at least 18 years old and have a household income equal to or less than 150% of the federal poverty line and they received the subsidies for three years. The RCT showed significant positive program impacts on labor supply, family income, and child care use. Interestingly, the RCT also revealed significant positive impacts on children's cognitive achievement. The treatment consisted of a bundle of conditional and unconditional subsidies and it is not possible to know which of the components of the treatment were most important in generating the positive impacts.

A recent study by Rodriguez (2018) analyzes data from the New Hope RCT with the following goals: to understand the mechanisms that underlie the observed treatment impacts, to disentangle which of the program components was most important in generating the observed impacts, and to analyze impacts of modifying the program’s design. The paper estimates a dynamic discrete choice model of the household labor supply and child human capital formation. In the model, a unitary household with a single child chooses hours of work and child care types (informal home care or formal, center-based child care). Household choices and the current stock of child human capital are inputs in the child human capital production function. The specification of the household’s budget set accounts for different means-tested programs available to the household including AFDC, EITC, and New Hope. The model is estimated using a method of moments approach and only using non-experimental moments. The model’s predictions are compared to the experimental impact estimates.

The paper finds that New Hope’s effects on child human capital are entirely explained by the child care subsidy component, which led parents to take their children to center-based child care. Model simulations show that giving an average family an amount of money equal to the cost of child care increases child human capital by 0.8% of a standard deviation, but giving the same amount only for use in purchasing child care services increases child human capital by 52% of a standard deviation. The greater productivity of external child care in fostering human capital development accounts for the treatment effects that were observed on cognitive achievement. Rodriguez (2018) also uses the estimated model to estimate the effects of varying the program design to not include the full-time work requirement, which he finds would lead to an even greater increase in children’s human capital (by 0.04 standard deviations).

5.3.2 Quasi-experimental studies

Child care subsidy in Norway

Chan and Liu (2018), using data from a large-scale welfare reform in Norway, study effects of alternative child care policies on women’s life-cycle decisions and on long-term child cognitive outcomes. They develop and estimate a DCDP model of women’s decisions

with regard to labor supply, child care and fertility. The model allows children's cognitive development to be affected by childcare arrangements. The model is estimated using Norwegian administrative data that includes child test score data measured beyond age 10. The cognitive outcomes include scores on reading, math and English tests.

In estimating the model, the authors exploit a large-scale child care reform called "cash for care," which provided cash to families with young children who did not use formal child care options. They argue that this reform provides exogenous variation in the relative price of different child care options that is useful to identify model parameters. The empirical results show that "cash for care" reform had a significant impact in reducing the employment rates of lower education mothers. The authors find that the use of nonmaternal early child care leads to lower reading scores than formal care. Also, maternal care leads to lower reading scores than formal care for low education mothers but to similar scores as formal care for high education mothers. The estimated DCDP model is also used to evaluate the effects of counterfactual policies, such as tax policies and maternal leave policies.

5.4 Relocation/migration subsidies

5.4.1 RCT studies

Housing subsidy in Boston

Galiani, Murphy and Pantano (2015) study the effects of a housing rent subsidy on residential neighborhood choices. They use data from the Moving to Opportunity (MTO) housing subsidy experiment to estimate a model of household neighborhood choice and to analyze the effects of changing the design of the program subsidy. In the MTO experiment, low income households in six cities (Baltimore, Boston, Chicago, LA and NYC) were placed in three groups. One group received housing vouchers that could be used only in low-poverty areas (<10% poverty) for the first year in addition to counseling to help them find housing. After a year, they could use their vouchers anywhere. One group received vouchers that could be used anywhere but no counseling. A third control group did not receive vouchers but were eligible for any other government assistance for which they qualified. Prior studies examined the effects of the MTO intervention on labor market, educational and health

outcomes.³⁶ The focus of the Galiani et. al. (2015) study is instead on evaluating a range of counterfactual policies, such as changes in the neighborhood poverty threshold that is a condition for receiving the voucher. Their analysis sample includes 541 households in Boston, of which 165 are in the control group, 172 in the section 8 voucher group, and 204 in the conditional treatment experimental group.

The paper develops and estimates a model in which households choose a neighborhood according to their preferences for neighborhood characteristics and according to their own characteristics. They consider the choice over 585 tracts that represent different neighborhood. The model also incorporates a moving cost that depends on distance, which varies depending on the household's initial location of residence.

As noted in the paper, a challenge in estimating these kinds of location choice models is the potential endogeneity of rent prices, because neighborhoods may have unobserved amenities that are correlated with rent levels. The usual approach to addressing this endogeneity problem is to use instruments that come from imposing exclusion restrictions on the model.³⁷ Galiani et. al. (2015) show that the RCT provides another way of addressing this endogeneity problem, because it generates exogenous variation in rental prices across treatment and control groups and also within groups over time (before and after the intervention), which can be used to identify the model parameters without instruments.

Galiani et. al. (2015) use location, demographic, and rent data from the control group and from the experimental group that was subject to the low poverty restriction to estimate the location choice model.³⁸ For purposes of model validation, they hold out the treatment group that received the unrestricted voucher. They find that the estimated model is successful in replicating the mobility and neighborhood choice patterns of the held-out group. They also use the model to calculate households' willingness to pay for specific neighborhood attributes (such as the percentage of residents who are poor).

Lastly, they use the model to analyze the reasons the different take-up rates in the two

³⁶See e.g. Kling, Liebman and Katz (2007)

³⁷See, for example, Berry, Levensohn and Pakes (1985) and Bayer, Ferriera and McMillian (2007)

³⁸In estimation, they also use census tract data and require that the location shares predicted by the model match the location shares in the census data.

treatment groups, to consider counterfactual programs and to explore the question of optimal program design.³⁹ When the estimated model is used to simulate residential choices under a range of alternative poverty thresholds, ranging from 2.5% poor to 20% poor, the authors find that the program take-up rate is very sensitive to threshold level. Adopting a less stringent poverty cut-off threshold of 20% generates higher take-up and leads to an overall lower exposure of households to poor neighborhoods, arguably improving on the program design that was implemented.

Migration subsidies in Bangladesh

There have been multiple field experiments in developing countries that showed that small travel subsidies generate substantial migration along with increases in income and consumption over multiple years. Lagakos, Mobarak, and Waugh (2018) argue, however, that the experimental evidence is not enough to understand whether there is a spatial mismatch of workers, namely that workers are not living in the area where they would be most productive. They also note the impact estimates are not informative about welfare effects of such a program if individuals experience disutility from rural-urban migration.

For this reason, Lagakos et. al. (2018) develop a dynamic model of rural-urban migration in Bangladesh and use data from a migration subsidy field experiment by Bryan, Chowdhury, and Mobarak (2014) to obtain model parameter estimates. In their model, households are heterogeneous in their degree of permanent productivity advantage in the urban area, and they choose to locate in either an urban region or a rural region. The model incorporates seasonal income fluctuations and stochastic income shocks. It assumes that markets are incomplete and that agents insure themselves through a buffer stock of savings.⁴⁰ Individuals face both a monetary cost of migration and a non-monetary disutility from migration that depends on past migration experience. They can migrate permanently or temporarily.

Both treatment and control groups are used to obtain model parameters estimates, by fitting model moments to data moments obtained from the RCT. The main moments targeted are: (i) the increase in the seasonal migration rate resulting from the subsidy, which was 22

³⁹The program take-up rate was 63% for the treatment group that received the unrestricted vouchers in comparison to 55% for the group that was subject to the low poverty restriction.

⁴⁰As in Bewley (1977), Aiyagari (1994) and Huggett (1996)

percent; (ii) the consumption increase for those induced to migrate, which was 30 percent; and (iii) the increase in seasonal migration one year later, after the subsidies were removed, which was nine percent.

The authors find that the consumption gains from migration observed under the RCT are not due to permanent productivity gaps between urban and rural residents as the labor mismatch hypothesis might suggest. Rather, individuals from rural regions tend to migrate to urban areas at times when they face bad shocks as a form of insurance. The migrants are negatively selected on productivity and assets. The model estimates also reveal a high non-monetary disutility from migration, particularly for first-time migrants. The inference from the model presents a much more nuanced view about the welfare benefits of the migration subsidy policy.

5.5 Other programs

5.5.1 RCT studies

Firm-provided wage subsidies in British Columbia

Bellefleur and Shearer (2011) analyze how gift-giving, increases in compensation explained to workers as acts of kindness, affects the productivity of workers at a tree planting firm in British Columbia, Canada. The workers' output is observable and workers are typically compensated piece-rate (per tree planted) taking into account labor market conditions and the terrain in which the planting takes place. The firm implemented a field experiment in which a random sample of workers received one of two treatments—one that provided an increase of 20-28% in the piece-rate wage and one that provided a base wage payment of \$80 on top of the piece-rate (0.20 cents per tree planted). The increase in the base wage amounted to about 40% increase in the daily wage.

The authors analyze the RCT impact estimates for the two incentive designs that were implemented. In addition, they develop and structurally estimate a model of a worker's effort decisions given a particular gift-giving scheme. In the model, a worker's effort decision depends on two key parameters: one measuring the curvature of the effort cost function and another that measures the worker's response to monetary gifts from the firm, which

they called a kindness parameter.⁴¹ After using both the control and treatment groups to identify and estimate the model parameters, the authors use the model to calculate optimal gift-giving/piece-rate contracts.⁴²

The experimental results show that the base wage gift was not profitable. On the other hand, the increase in the piece-rate gift was profitable, but only when the labor market conditions led to low piece rates in the absence of the gift. The study also finds substantial heterogeneity among workers in how they respond to the firm's kindness with about half of the workers reciprocating by supplying greater effort and the other half not. The estimates indicate that reciprocity is associated with a longer tenure within the firm but this effect diminishes with age. The paper finds that the piece-rate gift is most profitable for workers with strongly reciprocal preferences - profit per worker increases by as much as 14% for certain types of workers.

Lastly, the authors use the estimated model to study the question of optimal contract design. In particular, they analyze the effects of composite gifts that combine a base wage increase and a piece-rate increase, even though the RCT did not include such a composite gift. They conclude, however, that workers respond much more strongly to piece rate gifts than to composite gifts. By analyzing the effect of differing magnitude increases in the piece-rate wage, they conclude that the firm could increase profits per worker by as much as 10% on average, and by up to 17% for workers exhibiting strongly reciprocal preferences.

Another study by Paarsch and Schearer (2009) analyzes data from the same experiment but only from the treatment arm where the piece rate was varied. The goal of the paper is to assess whether observed contracts are optimal and what types of contract changes if any could increase firm profits.

The paper develops a model where firms are choosing a contract to satisfy workers' participation constraints, without assuming that the firm is maximizing profits. The piece rate is chosen to satisfy the participation constraint of the marginal worker. Workers are assumed to supply effort and to maximize their income subject to an effort cost. Model parameters are estimated by maximum likelihood using both the control group and treatment

⁴¹The modeling approach was in part inspired by Rabin's (1993) theoretical work on fairness and reciprocity.

⁴²The model is estimated by a two-step nonlinear least squares procedure.

group data. The paper demonstrates that the randomized variation in the piece rate under the experiment permits identification of the elasticity of effort choice (as the piece rate is varied) under weaker conditions, without having to make assumptions on workers' alternative utility.

After estimating the model parameters, the authors derive the firm's optimal linear contract, consisting of a base rate and a piece rate, and compares profits under the optimal contract and under the observed piece rate contract (where the base rate was zero). The results show that the difference in profits is negligible, implying that the realized contract is close to optimal. Lastly, the paper considers the possibility of tailoring contracts to specific workers by offering different base wages to workers after their productivity types are revealed. It finds that firms could potentially increase their profits by 14% with a tailored wage scheme.

Active labor market programs (ALMP) in Denmark

In many European countries, participation in so-called active labor market programs (ALMP) is a requirement for receiving unemployment insurance (UI). ALMP takes various forms, but often it includes meetings, job search assistance and workfare/activation programs. If individuals view these arrangements as costly (e.g. a tax on their leisure), then measuring the effect of ALMP programs on the duration of unemployment can overstate the benefits of such programs.⁴³ A large literature estimates the impacts of ALMP programs on employment and earnings outcomes, but very few studies explore the mechanisms through which the treatment effects occur and the utility costs of such programs.

Maibom (2017) develops and estimates a dynamic discrete choice model of job search behavior using data from a Danish RCT in order to more fully understand the costs and benefits of such programs. In the model, individuals search for jobs and they choose a level of search intensity.⁴⁴ If they get a job offer, then they choose whether to accept the offer. They stochastically accumulate skills while employed. Job offer rates depend on the search intensity and on the unemployment duration. Individuals also receive UI benefits that may

⁴³Heckman, Lalonde and Smith (1999) note that it is problematic that program impact evaluation studies value labor supply at the market wage but value time spent in the non-market sector at a zero wage rather than a reservation wage.

⁴⁴The model is inspired by a model of Ferrell (2012).

require participation in ALMPs. Participation in ALMPs can affect utility but it can also affect job offer arrival rates. The model is estimated using both the control and treatment groups to identify differences between the groups in technology and preferences. Parameter estimates are obtained by simulated method of moments.

The RCT data analyzed include 3099 individuals (age 22-58) living in two regions. There was a control group and a treatment group in each region. The control group was required to attend caseworker meetings every 3rd month and to participate in a labor market activation program after 9 months of unemployment (6 months for persons under age 30) and thereafter every 26 weeks. Treatment in one region consisted of an intensified meeting schedule (every other week) and treatment in the other region consisted of earlier participation in activation. The RCT impacts showed that the employment rate was significantly higher in the treated regions but there was generally no effect on wages.

The estimated model parameters indicate substantial costs associated with ALMP participation. Model estimates are used to calculate the monetary compensation which would make individuals indifferent between participating in ALMP or not and the estimates show that individuals would give up about 50% of the UI benefit to avoid participation. This calculation allows assessment of whether the program is a worthwhile social investment by comparing the employment gains to costs, inclusive of those borne by participants. The model estimates are also used to analyze the heterogeneity in the compensating variation in relation to future prospects and the timing of treatment. The results show that traditional cost-benefit calculations which do not take the individual utility costs into account largely overstate the gains from these types of ALMP programs.

5.5.2 Quasi-experimental studies

Microfinance program in Thailand

Micro-finance programs are viewed as an important mechanism for stimulating investment in developing countries. However, there are few estimates of the economic returns from such programs. Kaboski and Townsend (2011) develop and estimate a model of credit constrained households and they use the model to compare microfinance programs to direct transfer schemes. In particular, they estimate the model using data collected prior to the

introduction of a large scale government microfinance program, the Thai Million Baht Village Fund Program, and then validate the model using post-program data.

The Thai Million Baht program, begun in 2001, involved the transfer of one million baht (about \$25,000) to each of almost 80,000 villages in Thailand to start village banks that lend to households. Kaboski & Townsend (KT) view the program as an unanticipated exogenous quasi-experimental increase in credit. The data analysis samples come from the Townsend Thai project, which gathered panel data on rural and semi-urban households and businesses from sixty-four villages in four Thai provinces from 1997 to the present.

The model is based on the standard buffer stock model of savings behavior under income uncertainty, along the lines of Aiyagiri (1994) and Deaton (1991). In the model, households start the first period with some level of permanent income and liquid wealth, and a potential investment project. Each period, the household makes a decision about whether to undertake an investment project of a given size. The household maximizes the expected discounted value of utility over an infinite horizon. The model is estimated by GMM using the first five years of "pre-experiment" data.

The validity of the estimated model is assessed by comparing the model's predictions of the effects of the Thai Million Baht program on consumption, investment and the probability of investing to the actual effects observed post program. The program is introduced into the model as a reduction in borrowing constraints by an amount that would increase the amount of total expected credit (as calculated from the model) in the village by one million baht. Impact estimates obtained using the model's simulated data are very close and, in fact, not statistically different from impact estimates obtained from regressions based on actual post-experiment data. One of the notable model predictions that is also borne out in the data is that the impact on consumption exceeds one million baht.

After finding support for the model's accuracy in predicting program impacts, the authors use the estimated model to compare the costs of the microfinance program to the costs of a direct transfer program that would provide the same utility benefit. They find that the cost of the microfinance program is 33 percent less, attributable to the fact that the microfinance program relaxes borrowing constraints which the transfer program does not do.⁴⁵ The results

⁴⁵Even households that do not use credit can be affected by the relaxation in borrowing

also indicate that the largest program impact is on consumption rather than investment.⁴⁶ In summary, KT demonstrate that microfinance programs are an effective means of increasing liquidity of credit constrained households, that they would positively impact both investment and consumption, and that they are more effective than a simple transfer program.

6 Evaluating effects of programs with spillover or general equilibrium effects

6.1 RCT studies

Inference from RCTs can be complicated when the treatment generates spillover effects on untreated persons or when there are general equilibrium effects.⁴⁷ For example, a vaccination program could have positive spillovers for people who do not receive the vaccination. Sometimes, the issue of spillover effects is addressed by using a place-based randomization design, where randomization is performed over larger units that do not interact with each other to avoid spillovers (e.g. schools rather than students within a school). Alternatively, some studies develop models that explicitly account for the spillover effects in assessing the treatment impacts. The issue of general equilibrium effects is addressed through the explicit modeling of all market participants (for example, workers and firms).

Spring protection in Kenya

Kremer, Leino, Miguel and Zwane (2011) implement an RCT to evaluate the effects of a water intervention in Kenya on outcomes related to water quality and child health. The spring protection intervention seals off the source of a naturally occurring spring and encases it in concrete so that water flows from a pipe instead of seeping from the ground, which helps to avoid contaminants from other individuals accessing the water source. In constraints, as it lowers their need for a buffer stock of liquidity and allows them to invest and increase consumption. Households who increase their borrowing are those who have the highest marginal valuation of liquidity, which makes the village fund program more cost effective than a simple transfer program.

⁴⁶Additionally, KT perform a counterfactual that limits the use of credit to investment rather than consumption. The restricted policy is found to be slightly more cost effective.

⁴⁷This violates the single unit treatment value (SUTVA) assumption commonly invoked in impact evaluations.

Kenya, water rights are communal and owners with a spring on their property are obliged to allow neighbors to use it without charge. This arrangement provides few incentives for investing in improvements.

The RCT randomized 184 viable unprotected springs into treatment and control groups.⁴⁸ A random selection of households that regularly used each spring was interviewed at baseline and also at follow-up rounds. Analysis of the experimental impacts showed that the intervention significantly improved water quality (as measured by E Coli contamination at the source and at the household) and also improved child health, reducing the incidence of child diarrhea by 25%.

As the study notes, many households access water from multiple sources and spring protection can generate spillover benefits on households in the comparison group. These households could decide to travel to a more distant protected water source rather than use a closer unprotected source. At baseline, 15.4% of comparison households get at least some of their drinking water from protected springs, but the percentage rises to 24.5% in follow-up rounds. To address the issue of comparison group households also obtaining water from protected springs, the authors perform a LATE analysis, using treatment assignment as an instrument from the fraction of trips taken to obtain water from a protected source. They find that more frequent access to protected water sources significantly improves household water quality.

In addition to obtaining the experimental estimates, the authors develop and estimate a mixed logit random utility model of households decisions about where to obtain water. Based on household reports on the trade-offs they face between money and walking time to collect water, the authors calculated an estimated mean annual valuation for spring protection equal to US\$2.96 per household. They use the estimate to derive an implied value of avoiding a statistical child death of \$769, which is substantially lower than the amounts typically used by policy makers. They interpret the estimates as evidence of a low willingness to pay for preventative health in this context. Lastly, the discrete choice model is used to simulate the welfare effects of counterfactual policies, such as giving the land owner private property rights over the spring. They find that welfare is greater with communal rights than with

⁴⁸The treatment was administered in multiple rounds.

private property rights.

Better informed school choice in Chile

Policy-makers are often concerned that low SES families are not investing enough in their children's human capital despite high returns to investment. One argument for why underinvestment occurs is that the parents are not well informed about their options or about the returns, raising the possibility that providing better information could lead to more efficient levels of investment. Allende, Gallego and Nielson (2019) examine the effects of an information provision RCT that targeted families of pre-K children in Chile who were soon to be entering elementary schools in Chile. The intervention consisted of a video and a personalized report card that compared different local schools. The intervention also included a general message about the importance of selecting a high quality school for children and the importance of schooling for labor market outcomes.

The RCT took place in 2010 in 133 preschools. There were 1612 parents who answered the baseline and follow-up surveys. The RCT impact estimates showed that the treatment intervention shifted parents' choices towards schools with higher average test scores, higher value-added test scores, higher prices, and longer distances from home. The children are followed five years later using administrative test score data, which shows that the positive treatment effects on academic achievement are sustained.

As the authors note, it is prohibitively costly to carry out the RCT on a large scale and it would be interesting to know the policy impacts were the intervention adopted on a large scale. The goal of the Allende et. al. (2019) study is to understand the implications of scaling up the information provision intervention, which they term *ex ante* aggregate policy evaluation. To this end, the authors develop and estimate an equilibrium model of school choice and competition among schools. The demand side of the model captures how parents make trade-offs between different relevant factors, such as quality of the schools, distance, and price.⁴⁹ The model assumes that families observe noisy signals of school characteristics and that providing them with the information treatment can shift the weight families put on price, distance and quality. The supply side model is a model of school competition in which schools choose price and quality and, over the longer term, capacity. Schools are assumed

⁴⁹The model builds on an earlier framework developed in Nielson (2013).

to maximize profits and a quality weighted average subject to technological constraints.⁵⁰ The authors use instruments to deal with the endogeneity of school characteristics, which are derived from variation of costs across markets and changes to the school voucher policy over time.

Using the estimated model, the paper evaluates the policy effects of an at-scale evaluation (extending the intervention to all families in the market) when schools do not react, students sort, and capacity constraints bind. It also evaluates the equilibrium effects under different assumptions on how public and private schools react, and how costs change. The effects on average school quality attended by low socioeconomic families is between $0.06\sigma - 0.22\sigma$. The general equilibrium effects of the policy are somewhat larger than the partial equilibrium effects. Also, the analysis shows that binding capacity constraints can greatly limit the effects of the policy.

6.2 Quasi-experimental studies

Active labor market program in Denmark

Gautier et. al. (2018) evaluate the effects of a Danish active labor market program (ALMP) on labor market outcomes (earnings, employment), allowing for the possibility that the program may have negative spillover effects on nonparticipating individuals. The program was implemented as an RCT in two Danish counties and provided job search assistance to randomly chosen newly unemployed workers.⁵¹ There were 1814 individuals in the treatment group and 1937 in the control group. The estimated impacts derived from the RCT show that the program participants found jobs more quickly than nonparticipants. (See, e.g., Graversen and van Ours (2008) and Rosholm (2008)).

If there are negative spillover effects of the program onto untreated individuals, then the impact estimates derived from the RCT have limited policy relevance. They do not give the

⁵⁰Chile has a nationwide school voucher system and more than half of children attend private schools, which can be for-profit schools.

⁵¹All individuals who started collecting unemployment benefits between November 2005 and February 2006 participated in the experiment. Individuals born on the first to the fifteenth of the month participated in the activation program, while individuals born on the sixteenth to the thirty-first did not receive this treatment.

average effect of the program on the treated, but rather combine positive impacts on the treated with negative impacts on the untreated. The presence of spillover effects violates the usual SUTVA (single unit treatment value) assumption that is commonly invoked in program evaluation settings. In this type of context, the RCT estimates cannot be used to examine the effects of a change in treatment intensity, such as the effects of a large-scale roll-out of the program.⁵² To get an idea of whether the program generated negative spillover effects, the authors perform a difference-in-difference analysis comparing the control group living in treatment counties to individuals living in similar counties where the program was not available, finding that the controls living in treatment counties have worse labor market outcomes.

To be able to address the question of how the treatment and treatment intensity affects both participants and nonparticipants, Gautier et. al. (2018) estimate the parameters of an equilibrium search model using the method of indirect inference. Their dataset combines information from the counties where the experiment took place with individuals from other comparison group counties. They argue that using data from the RCT in combination with nonexperimental data provides auxiliary moments to estimate congestion effects in the matching process and to analyze how the supply of job vacancies responds to an increase in the search intensity of program participants. The model exploits the fact that the program induces an exogenous increase in search intensity. The authors use the estimated model to understand the effects of counterfactual programs, such as one in which all newly unemployed workers receive the treatment.

7 Conclusion

Structural estimation was often seen as a rival approach to reduced-form analyses. This view was especially prominent in the context of program/policy evaluation, with the strongest contrast being between the experimental RCT approach and the structural modeling approach. However, as our examples illustrate and as the papers described in this survey show, the two approaches can usefully complement each other. When done well, a field experiment iden-

⁵²Blundell, Costa Dias, and Meghir (2003) and Ferracci, Jolivet, and van den Berg (2014) found evidence for spillover effects in the context of ALMP programs.

tifies as cleanly as possible and under minimal assumptions the average impact of a policy on outcomes of interest for the treated population. A limitation, though, is that RCTs tend to be implemented on a small scale because of cost and it is not usually feasible to extensively vary treatments within an experiment. Policy makers often need more information to guide their decision-making at the different stages of designing, implementing and evaluating programs. For example, prior to implementation, there is the question of how to optimally design the program to achieve particular targeting and outcome objectives and to meet cost criteria. After implementation, there is interest in understanding the mechanisms generating treatment effects, in drawing inferences about how treatment effects would vary if the program were modified in some ways and/or extended to other individuals, and in predicting treatment effects over longer terms of exposure. Lastly, there are situations where programs can generate spillover effects on control group individuals or general equilibrium effects that make it difficult to draw inferences about impacts from group mean differences.

In this paper, we surveyed over a dozen papers that combine experimental and structural modeling approaches to program/policy evaluation. These papers span a number of fields, including labor, development and urban economics, and a number of social/economic programs, including conditional cash transfer programs, welfare programs, relocation/moving subsidy programs, active labor market programs and early childhood development programs. As these studies illustrate, there are a variety of ways to fruitfully combine structural modeling with data from RCTs. A critical requirement, though, is that the experiment include data beyond simple measurement of the treatment and the outcomes. The structural approach, because it models agent behavior, specifies all the structural relationships in the agent's decision problem. Empirical implementation of the behavioral model requires that the variables that enter the structural components be measured.

Through this research approach, and by observing which models produce accurate forecasts or not, we are slowly gaining a broader understanding of what types of programs can be analyzed and with what types of models. Even the failure of models to accurately reproduce experimental benchmarks is valuable information that guides future developments. The recent recognition of the value of combining structural modeling with field experiments will likely spur further applications.

Table 1: Studies of CCT programs in education

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Todd & Wolpin (2006)</i> model of school-going, child labor and fertility used to evaluate effect of CCTs in Mexico	control	yes, using treated sample	yes, different subsidy designs compulsory schooling enforce child labor laws
<i>Attanasio, Meghir, & Santiago (2012)</i> model of school-going and child labor in Mexico used to evaluate effect of CCTs in Mexico	control & treatment	no	yes, different subsidy designs
<i>Leite, Narayan, & Skoufias (2015)</i> model of school-going and child work used to evaluate effect of CCTs in Mexico and Ecuador	control	yes, using treated sample	yes, different subsidy designs
<i>Dufló, Hanna, & Ryan (2008)</i> model of teacher attendance in India used to evaluate effect of teacher subsidies	treatment	yes, using control sample	yes, different subsidy designs
<i>Angelucci & Attansio (2013)</i> Engel curve model of food demand	control and pre-reform treatment	yes, compare to quasi-experimental diff-in-diff and matching estimates	no

Table 2: Studies of welfare programs

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Atalas et. al. (2016)</i> model of application decision to a subsidy program in India used to evaluate different targeting mechanisms	treatment	no	yes, change program application costs (time, distance), change expected prob of receiving benefits, and change fraction well-informed about eligibility rules
<i>Lise, Seitz, and Smith (2015)</i> job search and matching model in Canada used to evaluate effect of income supplement	control	yes using treated sample	no
<i>Choi (2018)</i> model of labor supply and welfare participation in US used to evaluate effect of changing benefit reduction rates	control	yes using two treated samples	no
<i>Keane & Wolpin (2007)</i> model of labor supply, fertility, welfare participation in US used to evaluate effect of changes in welfare rules	control	yes using held-out state	yes, changes in welfare benefits
<i>Hansen & Liu (2015)</i> model of labor supply and welfare participation in Canada used to evaluate effect of changes in welfare rules	pre-reform treatment and control	yes using post-reform treated	yes, changes in welfare benefits and income tax schedule

Table 3: Studies of early childhood programs

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Attanasio et. al. (2020)</i> model of early child skill formation used to evaluate effect of home visitation program in Colombia	control & treatment	no	yes
<i>Chan & Liu (2018)</i> model of female labor supply, child care choices, fertility used to evaluate effect of home care subsidies in Norway	control & treatment	no	yes, tax policies maternal leave
<i>Rodriguez (2018)</i> model of labor supply and child care choices used to evaluate effect of income and child care subsidies in the US	control & treated	yes, to experimental moments not used in estimation	yes, changes to subsidy design and conditionality requirements

Table 4: Studies of relocation/migration subsidy programs

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Lagakos, Mobarak, and Waugh (2018)</i> model of urban-rural migration	control & treatment	no	yes, different amenities upon migration
<i>Galiani, Murphy & Pantano, (2015)</i> discrete choice model of resid. location in Boston used to evaluate effect of conditional rent subsidy vouchers	control and one treatment arm	yes, using one treatment arm	yes, alternative poverty threshold conditionality requirements

Table 5: Other programs

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Kaboski & Townsend (2002)</i> model of consumption, investment and savings	pre-program treatment	yes	yes, alternative transfer programs
<i>Bellmare & Shearer (2018)</i> model of worker effort with firm gift-giving	treatment & control	no	yes, alternative payment schemes schemes
<i>Paarsch & Shearer (2009)</i> model of worker effort	treatment & control	no	yes, to study effect of alt. payment schemes on firm profits
<i>Maibom (2017)</i> model of how ALMP affects job search and labor market outcomes	treatment & control	no	yes, alternative timing of meetings/activation program interventions

Table 6: Studies of programs with spillover/GE effects

Study	Group used for estimation	Out-of-sample model validation?	Evaluate counterfactual programs?
<i>Kremer, Leino, Miguel, Zwane (2011)</i> discrete choice model of water source in Kenya	control & treatment	no	yes, private vs. communal property rights
<i>Allende, Gallego & Nielson (2011)</i> equil. model of school choice and program in GE and analyze effect school competition in Chile of binding capacity constraints	control & treatment	no	yes, extend to universal
<i>Gautier et. al. (2018)</i> job search model in Denmark	control & treatment	no	yes, universal program

References

- [1] Aiyagari, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," in *Quarterly Journal of Economics*, 109(3), 659-84.
- [2] Alatas, Vivi, Banerjee, Abhijit, Hanna, Rema, Olken, Benjamin A., Purnamasari, Ririn (2016): "Self-Targeting: Evidence from a Field Experiment in Indonesia" in *Journal of Political Economy*, vol. 124, no. 2, 371-427.
- [3] Allende, Claudia, Francisco Gallego, Christopher Neilson (2019): "Approximating the Equilibrium Effects of Informed School Choice," manuscript, Princeton University.
- [4] Andrews, Donald W. K. (1988). "Chi-Square Diagnostic Tests for Econometric Models: Theory." in *Econometrica*, 56(6): 1419-53.
- [5] Andrew, A., Attanasio, O., Fitzsimons, E., Grantham-McGregor, S., Meghir, C., and Rubio-Codina, M. (2018): "Impacts 2 years after a scalable early childhood development intervention to increase psychosocial stimulation in the home: A follow-up of a cluster randomised controlled trial in Colombia," *PLoS Med* 15(4): e1002556. <https://doi.org/10.1371/journal.pmed.1002556>.
- [6] Angelucci, Manuela and Orazio Attanasio (2013): "The Demand for Food of Poor Urban Mexican Households: Understanding Policy Impacts Using Structural Models," in *American Economic Journal: Economic Policy*, Vol. 5, No. 1, pp. 146-178.
- [7] Angrist, Joshua D., and Jfffdrn-Steffen Pischke. 2010. "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives*, 24 (2): 3-30.
- [8] Attanasio Orazio P., Costas Meghir and Ana Santiago (2012). "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, vol. 79(1), pages 37-66.
- [9] Attanasio, Orazio, Cattan, Sarah, Fitzsimons, Emla, Meghir, Costas and Rubio-Codina,

- Marta (2020). "Estimating the Production Function for Human Capital: Results from a Randomized Control Trial in Colombia," in *American Economic Review*, 110(1): 48-85.
- [10] Bayer, Patrick, Fernando Ferreira, and Robert McMillan (2007): "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." in *Journal of Political Economy*, 115 (4), 588-638.
- [11] Behrman, Jere R., Piyali Sengupta, Petra E. Todd (2005): "University of Pennsylvania Progressing through PROGRESA: An Impact Assessment of a School Subsidy Experiment in Rural Mexico," in *Educational Development and Cultural Change*, 54:1, 237-275.
- [12] Bellemare, Charles and Bruce Shearer (2011): "On the relevance and composition of gifts within the firm: evidence from field experiments," in *International Economic Review*, Vol. 52, No. 3, 855-882.
- [13] Berry, Steven, James Levinsohn, and Ariel Pakes (1995): "Automobile Prices in Market Equilibrium," in *Econometrica* 63 (4): 841-90.
- [14] Bewley, T. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," in *Journal of Economic Theory*, 16(2), 252-292.
- [15] Blaug, Mark (1980): *The Methodology of Economics*, Cambridge: Cambridge University Press.
- [16] Blundell, Richard and Andrew Shephard (2012): "Employment, hours of work and the optimal taxation of low income families," in *Review of Economic Studies*, Volume 79, Issue 2, April, Pages 481-510.
- [17] Blundell, Richard, Monica Costa Dias, and Costas Meghir (2003): "The impact of wage subsidies: A general equilibrium approach," manuscript, UCL.
- [18] Bourguignon, F., F. Ferreira, and P. Leite (2003). "Conditional cash transfers, schooling, and child labor: Micro-simulating Brazil's Bolsa Escola program." in *World Bank Economic Review*, 17(2), 229-254.

- [19] Brewer, Mike, Alan Duncan, Andrew Shephard, and Maria Jose Suarez, "Did working families tax credit work? The impact of in-work support on labour supply in Great Britain," in *Labour Economics*, 13 (6), 699-720.
- [20] Bryan, G. S., S. Chowdhury and M. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," in *Econometrica*, 82(5), 1671-1748.
- [21] Carneiro, Pedro, Karsten T. Hansen, and James J. Heckman (2003): "Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice," in *International Economic Review*, 44:2,361-422.
- [22] Chan, Mark K. and Kai Liu (2018): "Life-cycle and intergenerational effects of child care reforms," in *Quantitative Economics* 9, 659-706.
- [23] Choi, Eleanor Jawon (2018): "Evaluating a Structural Model of Labor Supply and Welfare Participation: Evidence from State Welfare Reform Experiments," manuscript, Hanyang University.
- [24] Creedy, John and Guyonne Kalb (2005): "Discrete Hours Labour Supply Modelling: Specification, Estimation and Simulation," in *Journal of Economic Surveys*, December, 19 (5), 697-734.
- [25] Cunha, F., Heckman, J., and Schennach, S. (2010). "Estimating the technology of cognitive and non-cognitive skill formation," in *Econometrica* 78(3), 883-931.
- [26] Cunha, Flavio, Heckman, James J. and Salvador Navarro (2007): "The Identification and Economic Content of Ordered Choice Models with Stochastic Thresholds," in *International Economic Review*, 48(4): 1273-1309.
- [27] Deaton A. (2009): "Instruments of Development: Randomization in the Tropics and the Search for the Elusive Keys to Economic Development." NBER Work. Pap. #14690.

- [28] Deaton A. (2010): "Instruments, Randomization, and Learning about Development," in *Journal of Economic Literature* 48: 424-455.
- [29] Deaton, Angus (1991) "Savings and Liquidity Constraints," in *Econometrica*, 59, 1221-1248.
- [30] Dehejia, Rajeev H. and Sadek Wahba (1999): "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," in *Journal of the American Statistical Association*, December 1999, 94(448), pp. 1053-62.
- [31] Dehejia, Rajeev H. and Sadek Wahba (2002): "Propensity Score-Matching Methods for Nonexperimental Causal Studies," in *Review of Economics and Statistics*, Volume 84, Issue 1, p.151-161.
- [32] Delavigna, Stefano D. (2018): "Structural behavioral economics," NBER working paper #24797.
- [33] Duflo Esther, Rema Hanna and Stephen P. Ryan (2012). "Incentives Work: Getting Teachers to Come to School," *American Economic Review*, vol. 102(4), pages 1241-78.
- [34] Ferracci, Marc, Grry Jolivet, and Gerard J. van den Berg (2014): "Evidence of treatment spillovers within markets," in *Review of Economics and Statistics*, 96, no. 5:812-23.
- [35] Ferrall, Christopher (2012): "Explaining and Forecasting Results of the Self-sufficiency Project," in *Review of Economic Studies* 79 (4), 1495-1526.
- [36] Fraker, Thomas and Robert Moffitt, "The effect of food stamps on labor supply : A bivariate selection model," in *Journal of Public Economics*, 1988, 35 (1), 25-56.
- [37] Galiani Sebastian, Alwyn Murphy and Juan Pantano (2015). "Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment," in *American Economic Review*, 105 (11): 3385-3415.
- [38] Gautier, Pieter A., Muller, Paul, van der Klaauw, Bas, Rosholm, Michael and Svarer, Michael (2018). "Estimating Equilibrium Effects of Job Search Assistance," in *Journal of Labor Economics*, Vol 36: vol. 4, 1073-1125.

- [39] Glewwe, Paul, Michael Kremer, Sylvie Moulin, and Eric Zitzewitz (2004): "Retrospective vs. prospective analyses of school inputs: the case of flip charts in Kenya," in *Journal of Development Economics*, 74, 251-268.
- [40] Gong, Xiaodong and Arthur van Soest (2002): "Family Structure and Female Labor Supply in Mexico City," in *The Journal of Human Resources*, 37 (1), 163-191.
- [41] Grantham-McGregor, S., Powell, C., Walker, S., and Himes, J. (1991): "Nutritional supplementation, psychosocial stimulation, and mental development of stunted children: the Jamaican study," in *Lancet*, 338(758):1-5.
- [42] Graversen, Brian K., and Jan C. van Ours (2008): How to help unemployed find jobs quickly: Experimental evidence from a mandatory activation program," in *Journal of Public Economics*, 92.
- [43] Griffen, Andrew S. and Petra E. Todd (2017): "Assessing the Performance of Nonexperimental Estimators for Evaluating Head Start," in *Journal of Labor Economics*, 35, no. S1 (July 2017): S7-S63.
- [44] Hansen, Jorgen and Xingfei Liu (2015): "Estimating labour supply responses and welfare participation: Using a natural experiment to validate a structural labour supply model," in *Canadian Journal of Economics*, Volume 48, Issue 5, 1831-1854.
- [45] Heckman James J. (1984). "The χ^2 Goodness of Fit Statistic for Models with Parameters Estimated from Microdata." *Econometrica* 52(6): 1543-47.
- [46] Heckman, James J. (1990): "Varieties of selection bias," in *American Economic Review: Papers and Proceedings*, 80:2, 313-318.
- [47] Heckman, James J. (2000): "Microdata, heterogeneity and the evaluation of public policy: Nobel lecture," in *Journal of Political Economy*, 109:4, 673-748.
- [48] Heckman, James J. (2010): "Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy," in *Journal of Economic Literature*, vol. 48, no. 2, June, pp. 356-98.

- [49] Heckman, James J. and Hotz, Joseph (1989): "Alternative methods for evaluating the impact of training programs," in *Journal of the American Statistical Association*, 84(408), 862-880.
- [50] Heckman, James J., Hidehiko Ichimura and Petra E. Todd (1997). "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme," *Review of Economic Studies*, 64(4), 605-654.
- [51] Heckman, James J., Hidehiko Ichimura and Petra E. Todd (1998). "Matching as an econometric evaluation estimator," *Review of Economic Studies*, 65(2), 261-294.
- [52] Heckman, James J., Hidehiko Ichimura, Jeffrey Smith, and Petra E. Todd (1998): "Non-parametric Characterization of Selection Bias Using Experimental Data", in *Econometrica*, Vol. 66, 1017-1098.
- [53] Heckman, J. J., Lalonde, R., Smith, J. (1999): "The Economics and Econometrics of ALMP," in Vol. 3 of *Handbook of Labor Economics*. North-Holland, Amsterdam.
- [54] Heckman James J. and Vytlacil, Edward J. (2007). "Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation," *Handbook of Econometrics*, J.J. Heckman and E.E. Leamer (ed.), Handbook of Econometrics, edition 1, volume 6, chapter 70 Elsevier.
- [55] Heckman, James and Sergio Urzua (2010): "Comparing IV with structural models: What simple IV can and cannot identify," in *Journal of Econometrics*, 106(1): 27-37.
- [56] Heckman, J., Moon, S., Pinto, R., Savellyev, P., and Yavitz, A. (2010): "Analyzing social experiments as implemented: A reexamination of the evidence from the HighScope Perry Preschool Program," in *Journal of Quantitative Economics* 1:146.
- [57] Hoynes, Hilary Williamson (1996): "Welfare Transfers in Two-Parent Families: Labor Supply and Welfare Participation Under AFDC-UP," *Econometrica*, 64 (2), 295-332.
- [58] Huggett, M. (1996): "Wealth Distribution in Life-Cycle Economies," in *Journal of Monetary Economics*, 38(3), 469-94.

- [59] Imbens, Guido W., and Joshua D. Angrist (1994): "Identification and Estimation of Local Average Treatment Effects," in *Econometrica*, 62(2): 467-75.
- [60] Imbens, G. (2010): "Better Late than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)," in *Journal of Economic Literature*, 48: 399-423.
- [61] Kaboski JP, Townsend RM. (2011). "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative" in *Econometrica*, 79: 5, 1357-1406.
- [62] Keane, Michael and Robert Moffitt (1998): "A Structural Model of Multiple Welfare Program Participation and Labor Supply," in *International Economic Review*, 39 (3), 553-589.
- [63] Keane Michael P. and Kenneth I. Wolpin (2007). "Exploring the Usefulness of a Non-Random Holdout Sample for Model Validation: Welfare Effects on Female Behavior", *International Economic Review*. 48: 1351-78.
- [64] Keane, Michael P. and Kenneth I. Wolpin (2010): "Structural vs. Atheoretic Approaches to Econometrics". *Journal of Econometrics* 156 (1): 3-20.
- [65] Keane, Michael P., Petra E. Todd and Kenneth I. Wolpin (2011): "The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications," in *Handbook of Labor Economics*, Volume 4, Part A, 331-461.
- [66] Keane Michael P. and Kenneth I. Wolpin (2010): "The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System on the Life Cycle Decisions of Black, Hispanic and White Women," in *International Economic Review*, 51(3): 851-892.
- [67] Keane, Michael and Robert Moffitt (1998): "A Structural Model of Multiple Welfare Program Participation and Labor Supply," in *International Economic Review*, 39 (3), 553-589.
- [68] Koopmans, T. C. (1947): "Measurement without theory" in *The Review of Economics and Statistics*, 29(3): 161-172.

- [69] Kling, Jeffrey R., Liebman, Jeffrey B. and Lawrence F. Katz (2007). "Experimental Analysis of Neighborhood Effects," in *Econometrica*, 75 (1): 83-119.
- [70] Kremer, Michael, Jessica Leino, Edward Miguel and Alix Peterson Zwane (2011): "Spring cleaning: rural water impacts, valuation and property rights institutions," in *Quarterly Journal of Economics*, 126, 145-205.
- [71] Lagakos, David, Ahmed Mushfiq Mobarak and Michael E. Waugh (2018): "The welfare effects of encouraging rural-urban migration," NBER Working Paper 24193.
- [72] Lalonde, R. (1986). "Evaluating the econometric evaluations of training programs with experimental data," in *The American Economic Review*, 76(4), 604-620.
- [73] Leamer, Edward (1978): *Specification Searches: Ad Hoc Inference with Nonexperimental Data*, New York: Wiley.
- [74] Leamer Edward (1983): "Let's take the Con out of Econometrics," in *American Economic Review*, 73(1): 31-43.
- [75] Leite, Phillippe, Amber Narayan, Emmanuel Skoufias (2015): "How do Ex Ante Simulations Compare with Ex Post Evaluations? Evidence from the Impact of Conditional Cash Transfer Programs" in *Journal of Poverty Alleviation and International Development*, 6(1),1-43.
- [76] Lise, Jeremy, Shannon Seitz and Jeffrey Smith (2015): "Evaluating search and matching models using experimental data," in *IZA Journal of Labour Economics*, 4:16, 1-35.
- [77] Low, Hamish, and Costas Meghir. 2017. "The Use of Structural Models in Econometrics," in *Journal of Economic Perspectives* 31 (2): 33-58.
- [78] Lumsdaine RL, Stock JH, Wise DA. (1992): "Three Models of Retirement Computational Complexity versus Predictive Validity," in *Topics in the Economics of Aging*, ed. David Wise, University of Chicago Press.

- [79] Lumsdaine RL, Stock JH, Wise DA. (1994): "Pension plan provisions and retirement: men and women, medicare, and models," in *Studies in the Economics of Aging*, ed. DA Wise. 2004. University of Chicago Press, Chicago.
- [80] Maibom, Jonas (2017): "Assessing Welfare Effects of ALMPs: Combining a Structural Model and Experimental Data," manuscript, Aarhus University.
- [81] Marschak, J. (1953). "Economic measurements for policy and prediction" in Hood,W., Koopmans, T. (Eds.), *Studies in Econometric Method*. Wiley, New York, pp. 1-26.
- [82] McFadden, Daniel, Talvitie, A.P., associates (1977): Validation of disaggregate travel demand models: Some tests. Urban demand forecasting project, final report. Volume V, Institute of Transportation Studies, University of California, Berkeley.
- [83] Moffitt, Robert (1979): "The Labor Supply Response in the Gary Experiment," in *The Journal of Human Resources*, Autumn, Vol. 14, No. 4, pp. 477-487.
- [84] Neilson, Christopher A. (2013): "Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students," Job Market Paper, Yale University.
- [85] Paarsch, Harry J. and Bruce S. Shearer (2009): "The response to incentives and contractual efficiency: Evidence from a field experiment Author links open overlay panel," in *European Economic Review*, Volume 53, Issue 5, July 2009, Pages 481-494.
- [86] Pissarides CA (2000) *Equilibrium Unemployment Theory*, 2nd ed.,The MIT Press.
- [87] Rodriguez, Jorge (2018): "Understanding the Effects of a Work-Based Welfare Policy on Child Human Capital," manuscript.
- [88] Rosenbaum, Paul and Rubin, Donald (1983): "The Central Role of the Propensity Score in Observational Studies for Causal Effects." in *Biometrika*, April, 70(1), pp. 41-55.
- [89] Rosholm, Michael (2008): "Experimental evidence on the nature of the Danish employment miracle," IZA Discussion Paper no. 3620, Institute of Labor Economics, Bonn.

- [90] Schady, N. and Araujo, M. (2006). "Cash transfers, conditions, school enrollment, and child work: Evidence from a randomized experiment in Ecuador," Policy Research Working Paper No. 3930, Impact Evaluation Series No. 3. Washington, DC: World Bank.
- [91] Schorfheide, Frank and Kenneth I. Wolpin (2012): "On the Use of Holdout Samples for Model Selection," in *Papers and Proceedings. American Economic Review*, 102(5): 477-481..
- [92] Schorfheide, Frank and Kenneth I. Wolpin (2016): "To Hold Out or Not to Hold Out?" in *Research in Economics*, 70(2): 332-345
- [93] Schultz, T. Paul (2004): "School Subsidies for the Poor: Evaluating a Mexican Strategy for Reducing Poverty." *Journal of Development Economics* 74, no. 1:199-250.
- [94] Smith, Jeffrey and Petra E. Todd (2005). "Does matching overcome LaLonde's critique of nonexperimental estimators?" in *Journal of Econometrics*, 125, 305-353.
- [95] Stone, Richard (1954): "Linear expenditure systems and demand analysis: An application to the pattern of British demand," in *The Economic Journal*, 64 (255): 511-527.
- [96] Todd Petra E. and Kenneth I. Wolpin (2006). "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*, vol. 96(5), 1384-1417.
- [97] Todd Petra E. and Kenneth I. Wolpin (2008): "Ex ante Evaluation of Social Programs" in *Annales D'Economie Et De Statistique*, no. 91/92, 263-291.
- [98] Todd Petra E. and Kenneth I. Wolpin (2010): "Structural Estimation and Policy Evaluation in Developing Countries," in *Annual Review of Economics*, vol. 2(1), 21-50.
- [99] van Soest, Arthur (1995): "Structural Models of Family Labor Supply: A Discrete Choice Approach," in *The Journal of Human Resources*, 30 (1), 63-88.
- [100] Vytlacil, Edward J. (2002): "Independence, Monotonicity and Latent Index Models: An Equivalence Result," in *Econometrica*, 70:1, 331-341.

- [101] Wise, DA. (1985): "Behavioral Model verses Experimentaion: The Effects of Housing Subsidies on Rent" in *Methods of Operations Research* 50, ed. P Brucker and R Pauly, Konigstein: Verlag Anton Hain, 441-89.
- [102] Wolpin, Kenneth I. (2013): *The Limits of Inference without Theory*. MIT press: Cambridge, MA