Research Statement

website • github.com/josephmullins • mullinsj at umn dot edu

I am an economist who studies childhood, family dynamics, and labor markets. My research is motivated by one core question: how can policies, in the face of limited resources and market imperfections, best enable human potential and flourishing? In pursuit of this question, I innovate on frontier methodologies, bringing economic theory to data to deliver precisely articulated and quantitatively rich policy conclusions. These include prescriptions for the minimum wage (Flinn and Mullins 2015), the design of cash transfers for households with children (Mullins 2025; 2024), and legal frameworks intended to govern marital dissolution (Brown, Flinn, and Mullins 2024) and wage negotiation (Flinn and Mullins 2021).

Economic models should be quantified with the maximum possible *clarity* and *credibility*. In addition to addressing questions of first-order social importance, my research is a continuing effort to expand the frontier in these two dimensions. For example, in **Mullins (2025)**, I study the optimal design of income support for families with children. In doing so, I identify a small set of parameters as being particularly important for the ensuing policy conclusions, and show how to apply policy variation to a linear equation to estimate them. This step provides clarity in showing the exact source of empirical discipline for those parameters, and credibility by drawing on quasi-experimental evidence. In **Mullins (2024)**, I introduce a new conceptual approach to maximize the impact and utility of growing bodies of empirical evidence. This paper posits that economic models are particularly well-suited for evidence aggregation, and introduces – by way of an example applied to welfare reform experiments in the United States – the concept of a *structural meta-analysis*. Studies that aggregate evidence in this way can inherit the credibility of careful causal research while employing economic modeling to answer rich normative and positive policy questions.

With these themes in mind, my works to date are the first steps in a three-fold research agenda.

Research Agenda

- 1. To design effective policies that eliminate wasteful disadvantages and enrich society by investing in human development.
- 2. To improve the reach and impact of those policy analyses by expanding and combining the frontiers of knowledge in economic modeling and empirical research.
- 3. To pioneer tools that help conduct model-based research with the best possible clarity and credibility.

My research has been published or is under revision at the Journal of Political Economy, the Review of Economic Studies, the International Economic Review, and the Journal of Labor Economics.

Research on Skill Formation in Childhood

Designing Cash Transfers

In theory and in practice, the study of taxation and redistribution in Public Economics has focused on the role played by labor supply in shaping optimal policy. Governments must balance their desire to redistribute with their desire to minimize costly distortions of the labor and skills that individuals supply to the market. Analyses in this tradition forefront behavioral elasticities (such as labor supply) as the main objects of theoretical and empirical interest. In *Designing Cash Transfers in the Presence of Children's Human Capital Formation* (Mullins 2025) I show that accounting for the skill formation of children yields theoretically novel insights with quantitatively large implications for the design of optimal policies.

The paper develops a model in which single mothers work, make program participation choices, and invest time and money in each of their children. Children's behavioral and cognitive skills evolve according to a technology of skill formation that takes time investments, money investments, and childcare quality as inputs. The model has two tractable properties that allow for clear analysis of optimal policies and a transparent mapping between parameters and data: (1) an indirect utility function that incorporates dynamics and can be identified directly from data on mothers' work, participation, and earnings; and (2) a linear outcome equation that maps income, work hours, and childcare quality to skill outcomes. Estimation leverages these properties into a two-stage approach that first determines parameters of indirect utility using panel data on single mothers in the Panel Study of Income Dynamics (PSID), making use of historical policy variation that manipulates mothers' work and participation incentives. The second stage combines the outcome equation with this same policy variation to form instruments, which deliver credible estimates of the technology of skill formation using data from the PSID's Child Development Supplement (CDS). The result is a quantitative model in which both behavioral responses and long-run skill outcomes are articulated with respect to a wide class of policy reforms, disciplined by similar historical variation.

The model's tractability, beyond estimation, also permits clear theoretical analysis of a social planning problem in which the government ranks different policies in terms of their impacts on mothers' welfare and the net present value of total economic resources. The latter includes both the short-run costs of the program and the long-run returns through child human capital formation. A number of useful insights follow. First, the generosity of the optimal program (as measured by net income allocations) is not only increasing in the planner's overall taste for redistribution toward this group, but also in the factor share of money investments for child skill production. Second, optimal allocations are characterized by a wedge term that features a short-run component (the effect of mothers' labor supply responses on short-run costs) and a long-run component (the net effect of mothers' labor supply responses on long-run resources through child skill formation).

A series of quantitative policy experiments using model estimates mobilizes these theoretical insights and demonstrates that the novel mechanisms introduced by child skill formation quantitatively dominate the traditional fiscal tradeoffs in shaping optimal policies. For example, when the planner's taste for redistribution is calibrated to match the net income of single mothers without children in the benchmark, the model's estimates suggest that the optimal system of transfers is 50% more generous than the benchmark set of policies from the year 2000. It also features much higher marginal and participation tax rates. A set of decompositions that follow the paper's theoretical analysis show that the long-run effect on children's human capital is overwhelmingly the most important quantitative factor that determines the observed shape of the optimal schedule. A robustness exercise that significantly scales down the generosity of the optimal program verifies that this result remains true even for much smaller policy reforms.

Several features of the exercise emphasise the importance of particular modeling decisions in the paper. In particular, model estimates indicate a significant degree of latent heterogeneity among single mothers that turns out to be important for matching widely varying differences in the skill and welfare impact of each reform. Failing to capture the extent of this heterogeneity would result in researchers confusing redistribution for insurance. Likewise, the estimated model suggests that skill impacts on children of particular reforms are widely varying, with permanent differences in labor market productivity and childcare quality largely determining them.

Overall, the paper delivers important lessons about the first order importance of child skill formation when designing social insurance. Although questions remain regarding the suitability of these data for projecting the best policies for current cohorts of children, this paper lays the platform analytically and methodologically for future policy analyses. In particular, given rapidly accumulating evidence on the short and long-run effects of various historical reforms, the results of the paper ought to motivate a more rigorous approach to aggregating existing evidence in the spirit of **Mullins (2024)**.

A Structural Meta-analysis of Welfare-to-work

A key step in producing high quality model-based research is to provide a transparent and convincing mapping between the main question of interest and the features of the data that will ultimately determine the answer. More often than not, this requires a convincing and clear argument for identification of the model's key causal mechanisms. Finding appropriate empirical evidence for these mechanisms, in addition to the analytical and computational demands of writing and estimating a model, can be an overly burdensome prospect for individual researchers. In *A Structural Meta-analysis of Welfare Reform Experiments and their Impacts on Children* (Mullins 2024), I offer an alternative methodological approach by pointing to the rapid accumulation of evidence supplied by applied researchers, suggesting that economic models are a natural tool for aggregating this evidence for the purposes of policy and prediction. The paper defines a *structural meta-analysis* as any such exercise that holds a quantitative economic model accountable – for the purposes of aggregation, policy and prediction – to multiple studies or sources of evidence, and applies the approach to experimental evaluations of welfare-to-work programs conducted in the United States in the 1990s.

The paper combines individual data on work, participation, childcare use, and child outcomes from six experimental sites as well as a representative sample of single mothers from the Survey of Income and Program Participation (SIPP). Each site consists of individuals randomly assigned to either a control group or a treatment group. Each treatment group faced some version of a reformed welfare program, each differing in their benefit formulae, childcare subsidies, time limits, and work requirements. Control group members faced the prevailing welfare program rules given their state of residence.

Aggregation of evidence across these experiments without a theoretical framework is undermined by a number of challenges. Each treatment is a complicated bundle of program features, making comparison across sites difficult. Individuals entered the sample as either new or returning welfare applicants and are potentially quite different across sites and more different still from the general population of potential applicants. Several expansions through the 1990s in the Earned Income Tax Credit led to changes in the incentives to work for both control and treatment group members. Finally, the nature of selection and the magnitude of treatment effects are potentially influenced by local labor market conditions and other initial conditions.

In order to interpret treatment effects on the outcomes of interest, the paper introduces a dynamic model in which mothers make work, program participation, and childcare choices, in addition to making investments in their child. Each mother belongs to one of a finite number of types that are not directly observable but dictate their labor market productivity, preferences, child rearing capabilities, and the relative quality of the paid and unpaid care that they can access. The model also features latent and stochastically evolving job opportunities that, along with local unemployment rates, determine wage offers in each period. A maximum likelihood routine uncovers the parameters of the model that best fit individual decisions across all data sources, as well as the distribution of types and job opportunities in each location.

The model also delivers an outcome equation, similar to **Mullins (2025)**, in which each child's behavioral and cognitive skills are a linear function of total income, work decisions, and whether parents use paid or unpaid care while working. In principle, experimental assignment to treatment provides enough variation to identify these parameters with instrumental variables. In practice, those estimates do not deliver sufficient precision. A control function approach that uses all variation, conditional on type, produces more precise estimates that nonetheless still replicate treatment effects on skills.

The estimated model embodies an aggregation of the evidence across all sites, and provides several important lessons. First, estimates validate concerns about the experimental samples being highly selected: relative to the SIPP, the experimental sites oversample individuals with lower labor market productivity and with temporarily worse job opportunities. These differences are important in that they imply substantial differences between treatment effects on experimental populations relative to treatment effects of the same reforms on the general population. Second, estimates of the quality of paid and unpaid childcare options (relative to maternal care) differ by type: for some types, paid care is of lesser relative quality and unpaid care is of higher relative quality, reflecting important differences in mothers' skills as well as the quality of affordable paid care in the marketplace. Third, counterfactual results indicate that treatment effects on program participation are explained overwhelmingly by the introduction of time limits, while effects on labor supply are mostly explained by work requirements, with a smaller role played by changes in financial incentives and time limits. Finally, while the model replicates the fact that the experiments mostly did not have statistically significant impacts on skill outcomes, it also forecasts statistically significant negative impacts on behavioral skills when each site's experimental welfare program is applied to the general population (as determined by the distribution of types and job opportunities in the SIPP).

Overall, this paper makes important conceptual progress toward using economic theory for aggregating rapidly accumulating bodies of causal evidence. The paper's application to welfare reform demonstrates the utility of economic modeling for providing a unifying framework in which to understand results from different samples, locations, and institutional contexts. As an organizing concept, the paper identifies *well-articulated variation* – observable changes that can be modeled without additional parameters, such as changes to prices or endowments – as being particularly well-suited for structural meta-analysis. The paper opens many promising future avenues for future research in this style.

Accounting for parental and market-based investments in child skill production

The *technology of skill formation* – the causal mapping from parental choices and inputs to child skill outcomes – is the central object in all economic studies of childhood and child-related policies. A number of important and influential works have studied how parents spend their time and money at home and in the marketplace. These studies all rely on different combinations of restrictive assumptions on the substitutability of these inputs, or on the underyling econometric assumptions that allow estimation and identification.

In *Child skill production: Accounting for parental and market-based time and goods investments* **(Caucutt et al. 2025)**, we consider identification and estimation of a child skill production function that takes inputs of parental time and money at home, in addition to market-based childcare services. Estimation accounts for measurement error in inputs and outputs, and allows for differing elasticities of substitution of inputs inside and outside the home.

The paper focuses on two main issues: the role played by parental human capital in child development through wages, productivity differences, and preferences; and the impact of subsidies

on the composition of investments and child skill outcomes. In both cases, the substitutability of inputs plays an important theoretical and quantitative role.

Central to the paper's analysis is the formalization of a *two-stage budgeting* approach in order to separate the intratemporal problem (in which households decide on total investments and consumption today) and the intertemporal problem (in which households decide the optimal combination of inputs given prices and total allocated expenditures). The paper establishes formal conditions under which the household's decision problem can separate in this way. Separation of the intertemporal problem in particular leads to a number of theoretical and econometric results that are central to the paper. Theoretically, we analytically characterize the relationship between parental human capital and family investment inputs, as well as investment responses to input price changes. We find, for example, that wage subsidies (such as the Earned Income Tax Credit) can lead to increases in all inputs when there is sufficient complementarity in production. Empirically, we pursue a revealed preference approach to identification and estimation of the factor shares and elasticities of substitution for each input. Although this is a common approach in the estimation of production functions, we undertake estimation with a more rigorous treatment of measurement error, selection, and endogeneity of inputs.

We next show how to combine the intratemporal revealed preference approach with intertemporal dynamics of inputs and investment to estimate the full production technology. By relaxing the cross-equation restrictions between the demand system and the production function, this approach facilitates "direct" estimation of the production function from skill dynamics. In practice, we show with Monte Carlo experiments that direct estimation of intratemporal parameters works very poorly and requires unreasonable amounts of relative price variation compared to what is available in the data.

Using PSID-CDS data on children aged 5-12, we find robust evidence of moderately strong complementarity across all inputs, which produces strong comovements of inputs in response to price changes. Counterfactual policy simulations suggest that reductions in the price of home goods and services inputs or market-based childcare cause families to expand all types of investment, with much stronger own-price elasticities than cross-price elasticities. A decline in parental wages leads to reductions in all types of investment, including small reductions in parental time investments, due to diminished family resources. The co-movement of inputs in response to price changes contrasts with the commonly assumed Cobb-Douglas specification, which we show would lead one to overestimate the costs of providing subsidies for childcare or goods/services investments. These findings underscore the importance of allowing flexibile substitution of inputs in future work.

Family Law

A consistent theme throughout my research is that any government policy that affects household behavior or economic circumstances can have indirect effects on parenting and generate powerful spillovers on future generations. Although this logic applies more directly to programs that provide social insurance or subsidize child-specific inputs, it applies generally to institutional frameworks that seek to govern or regulate household interactions. In *Family Law Effects on Divorce, Fertility and Child Investment* (Brown, Flinn, and Mullins 2024), we consider one very salient example: laws that govern marital dissolution. The paper studies three parameters of such frameworks: (1) mututal consent vs no-fault divorce standards; (2) physical custody arrangements; and (3) child support requirements.

The welfare and child-skill impacts of policies in this space are difficult to forecast because there is a paucity of suitable variation with which to identify the complex, dynamic effects – on fertility, divorce decisions, parental investment, and child outcomes – directly from data. The paper develops

a non-cooperative model where married couples decide every year whether to have children and whether to dissolve the marriage. Wives and ex-wives in each pair also decide how much labor to supply to the market, while couples with children each decide how much time to invest in their child. A marriage quality variable that evolves probabilistically over time enters couples' utility as well as being an input in the production of child cognitive skills. A technology of skill formation takes the quality of the marriage (relative to divorce) and parental time investment as inputs and maps them to cognitive skill outcomes. The productivity of women in the labor market depends on a persistent stochastic component as well as their cumulative total work hours. Thus, both the child's and the mother's human capital are endogenous to the prevailing legal framework.

A multi-stage estimation procedure produces parameter estimates by matching observed statistics on wages and earnings, fertilty and divorce by year of marriage, time investment by child age and marital status, and child test scores. The paper uses counterfactual scenarios to study the long-run impact of changes to divorce standards, changes in custody allocation rules, and changes in rates of paternal child support.

An overarching lesson from the counterfactual studies is that the complicated patterns of dynamic selection and anticipation are crucial determinants of welfare and skill outcomes. Fertility, female labor supply, and marital dissolution all depend on the couple's current marriage quality, the endogenous and exogenous components of the wife's human capital, and the age of all household members. Importantly, these decisions also depend on couples' expectations of how any potential marital dissolution would be governed, with several empirical implications. First, marital law affects the outcomes of children in both married and unmarried households by shaping parental behavior dynamically. For example, married mothers may work more in order to insure themselves when divorce is more likely, or when child support is less generous. As a consequence, children receive less time investment, leading to lower cognitive skill outcomes. Second, marital law can affect average skill outcomes by changing *who* decides to have children and *when* they choose to divorce. Different legal frameworks can cause either advantageous or disadvantageous selection of fertility and divorce on marriage quality. Third, the long-run impacts of changes in divorce law is likely to differ from the short-run impact, since household positions are determined by many years of behavior under the prior regime. This insight underscores the potential limitations of using popular causal inference strategies in isolation, since at best they can typically only identify the effect of policy changes on the population of individuals who have lived under the prior regime up until that point.

The specific results of the counterfactuals bear out these lessons. For example, the estimated model indicates that unilateral divorce results in only a marginal (and statistically insignificant) increase in the rate of divorce because it also changes the rates at which married couples have children, which would otherwise make them more vulnerable to divorce. Similarly, increases in the generosity of child support lead to decreases in the divorce rate primarily through reductions in fertility. In both counterfactuals, changes in average marital quality (itself dictated by dynamic selection on fertility and divorce) are the primary determinants of child skill impacts. In the long-run, increases in child support lead to reductions in child skills through this selection channel.

This paper makes an important first step toward understanding the intricate relationship between fertility, divorce, child skill accumulation, and marital law. It lays out a path for future work to better combine available causal evidence on divorce and marital law with a robust, theoretically grounded modeling approach that enables clear policy analysis.

Research on Labor Markets

Firms' Choices of Wage-Setting Protocols

Labor market models featuring *search frictions* – where firms and workers can only form productive relationships through a random matching process – have offered significant advances in our understanding of individual trajectories and aggregate outcomes. These models are particularly compelling because they are imperfectly competitive. Positive rents are portioned unequally among otherwise identical workers and create a role for welfare-improving labor market interventions. Much however hangs on the particular assumptions that researchers make regarding how workers and firms divide these positive rents. In *Firms' Choices of Wage-Setting Protocols* (Flinn and Mullins 2021), we study a frictional labor market in which these wage-setting protocols are *endogenous* to features of each market through firm choices. We show that this flexibility is crucial both for understanding known facts about differences in wage-setting by individual and market characteristics (such as education, gender, and market tightness) but also for accurately forecasting the effects of labor market interventions on wages and employment.

In the model, firms choose between two wage-setting protocols that undergird a large fraction of studies in the literature. The first option is to make a non-negotiable wage offer that is fixed throughout the employment relationship. The alternative is to bargain and split the surplus with the worker, based on the worker's current outside option, and later renegotiate this wage when competing offers arrive. When choosing to make fixed wage offers over negotiating, firms claim a larger fraction of the surplus initially, but in failing to later renegotiate, cannot retain workers as effectively. The first result in the paper characterizes the overall rate of bargaining and renegotiation in a market in terms of one statistic that summarizes this trade-off: the average number of competing offers a worker will receive before the match is terminated for exogenous reasons. By matching data on bargaining and renegotiation by age, gender, and education to the Current Population Survey, we show that the ratio of employer-employer transitions to employment-unemployment transitions (a proxy for the theoretical statistic suggested by the model) does indeed robustly predict both bargaining and renegotiation in the cross-section.

Having validated one of the model's key predictions in data, the paper then develops a quantitative version of the model and uses the wage equation to discuss recently established empirical facts about wage dynamics. Allowing for flexibility in wage-setting – the paper's key innovation – is crucial for incorporating these new facts. An estimated version of the model then provides several important insights through counterfactual analyses. For example, a mandate that forces all firms to either make fixed wage offers or to bargain and renegotiate (intended to mimic increasingly popular policies that regulate wage-setting) reveals the role played by bargaining and renegotiation in shaping wage dispersion and gender gaps. Enforcing homogenous wage-setting practices can reduce the wage gap between men and women by up to 7%, the gap between college and non-college workers by 11%, and the overall wage dispersion within demographic groups by 1%. Worker welfare, measured in units of consumption, decreases when all firms are forced to post non-negotiable wages. Enforcing wage bargaining on the other hand leads to increases in worker welfare by 5 to 10% across groups.

The paper makes important progress for our understanding of what shapes heterogeneity in wagesetting practices, the role these differences play in determining individual wage trajectories and overall wage inequality, and the potential role for regulatory policies to improve welfare and market outcomes.

Labor Market Search and Schooling Investment

The overwhelming majority of analyses of education and labor market policies do not study the interaction between these two crucial phases of workers' life-cycles. In Labor Market Search and Schooling Investment (Flinn and Mullins 2015), we develop a model of imperfectly competitive labor markets where individuals with heterogeneous abilities make costly investments in their human capital prior to labor market entry. Upon labor market entry, workers search for vacant positions and, after meeting, negotiate their wages with firms. The presence of these search frictions in the labor market endows firms with market power that they can use to appropriate a fraction of the surplus generated by workers' educational investments. This *holdup problem*, in conjuction with more traditional imbalances in labor markets with search frictions (such as congestion externalities), creates a complex interplay of forces that lead to inefficiencies in the laissez-faire economy. We consider two policies that could potentially relieve these inefficiencies and result in welfare improvements: a minimum wage that targets inefficiencies in the labor market, and an education subsidy that targets inefficient under-investment in education. Minimum wages can potentially relieve these inefficiencies by transferring surplus from the firm to the worker, but will have ambiguous distortions on education decisions through its effect on wages and job creation in each market. Education subsidies on the other hand can help with inefficiently low college attendance, but cannot resolve remaining inefficiencies in the labor market.

We quantify these mechanisms by estimating the model on the Survey of Income and Program Participation (SIPP), which provides a panel of wage and employment outcomes for college and noncollege educated workers. We find that once endogenous job creation and schooling choices are taken into account, a minimum wage of \$17/hour (in 2004 USD) maximizes average worker welfare. Implementing this policy however would lead to large welfare losses for the least able workers in the market, who are effectively excluded from work by the imposition of this minimum wage. Since it cannot resolve any of the ex-post inefficiencies in each labor market, an optimal education subsidy delivers only a fraction of the welfare gains, but it also does not feature the same difficult welfare trade-offs.

This paper is also the first published work to quantitatively evaluate the minimum wage's ability to resolve inefficiencies in imperfectly competitive labor markets in the presence of heterogeneous workers.

Frictional Adjustment to Income Tax Incentives

Much of Public Economics is built on the neoclassical model of labor supply, in which workers freely and frictionlessly adjust their hours or earnings in response to changes in tax incentives. This model yields behavioral elasticities that govern optimal taxation formulae and calculations of deadweight loss. In *Frictional Adjustment to Income Tax Incentives: An Application to the Earned Income Tax Credit* (Mancino and Mullins 2023), we document that changes in tax incentives also lead to increases in employer-employer (EE) transitions, and changes in the wages of jobs accepted out of unemployment. To accommodate these two empirical findings, we estimate a labor market model with search frictions and hours constraints. The model implies that adjustments to tax incentives take time, particularly at the intensive margin, and short-run elasticities can differ substantially from long-run elasticities. This has important implications for positive and normative policy conclusions.

The paper borrows an empirical strategy based on a variable that proxies for county-level awareness of the Earned Income Tax Credit (EITC). The strategy infers a causal affect by comparing the difference between eligible (single mothers) and ineligible (single women without children) individuals in counties with different levels of awareness. This difference-in-differences approach replicates previous findings on the employment effects of the EITC while also documenting that it leads to an increase in EE transitions, and a shift in the distribution of wages accepted out of unemployment.

Next, the paper develops a model in which individuals find jobs through a random search technology. Jobs are characterized by a combination of wages and hours (full or part-time), both of which remain fixed throughout the match. Workers encounter job offers based on their employment status and rank them according to the net utility derived from the arrangement. Changes in tax incentives therefore distort the job ladder by changing workers' utility rankings over jobs and changing the kinds of jobs they are willing to accept while unemployed. In the case of the EITC, this leads to an increase in the EE rate by re-ordering worker's rankings over jobs (making them more likely to accept new jobs) and by decreasing their reservation wage (thereby lengthening the job ladder). We estimate the model by first estimating a finite mixture model at the county level with different and latent levels of awareness, and constructing labor market moments from the Current Population Survey for each awareness type. This incorporates the model.

The estimated model not only replicates the empirical evidence in the short-run, it also provides a projection of long-run impacts. Here, results indicate that the long-run employment effect of the EITC is about 8 percentage points higher than the short-run impact. Likewise, the welfare impacts in the short-run are slightly less than half compared to the long-run. The paper then demonstrates how these two findings are connected by deriving a sufficient statistic for the welfare effects of a marginal expansion in the EITC that depends only on elasticities of the earnings distribution. Echoing the previous findings, evaluating this formula with short-run elasticities results in welfare effects that are slightly less than half of long-run impact.

Ongoing and Future Work

In ongoing and future work I pursue my three-fold agenda by continuing to study the relationship between policies, institutions, and human capital, and by developing tools to help researchers with similar interests.

Ordinal Complementarity

Much of the literature on human capital development has organized around the notion of *dynamic complementarity*: if past inputs can increase the productivity of future inputs, this has important implications for the timing of interventions and the mechanisms of intergenerational inequality. This version of complementarity is defined by supermodularity of the production function of child skills. If one assumes – as most of the literature does – that skills are a latent variable with no inherent cardinality, researchers must combine the technology of skill formation with an *anchoring equation*: a function that maps latent skills to a cardinal outcome such as wages, earnings, or education. In *Ordinal Complementarity* (Almazura, Garcia-Vazquez, and Mullins, n.d.), we point out a serious issue with this approach. While all relationships between latent skills and outcomes are equally "good" in terms of what can be identified in data, supermodularity is not itself invariant to these monotonic transformations. In other words, it is always possible to write and estimate a production function with dynamic complementarity in latent skills that features no such complementarity in producing the outcome of interest.

To solve this problem we propose a stricter notion of supermodularity based on first order stochastic dominance ("ordinal complementarity") and show that it produces policy conclusions and inference that is nonparametrically robust to arbitrary transformations of the latent variables. We show a number of seemingly innocuous examples where ordinal complementarity does not hold, and hence conclusions regarding complementarity of production can be undone by alternative assumptions on

the anchoring relationship. We propose a test of ordinal complementarity based on moment inequalities, and demonstrate the performance of the test in a Monte Carlo study.

Loose Constraint Estimation

Estimating economic models can be computationally demanding, particularly when the estimation routine involves a nested solution of the model for each evaluation of the estimation criterion. A very widely-used set of tools has emerged that allows researchers to avoid nesting one computionally demanding task (model solution) inside another (function minimization in many dimensions). These tools combine two necessary ingredients: (1) reliable estimates of agents' decision rules; and (2) a set of constraints that relate these decision rules to the model parameters. One can combine these two ingredients to derive a set of computationally light, multi-step estimators, commonly referred to as the conditional choice probability (CCP) approach. Despite their appealing properties, the tools are limited to dynamic discrete choice models, and to cases when the data provide estimates of choice probabilities with reasonable statistical precision. In many cases, these estimators will inherit the variance and / or bias of the first stage estimates.

In *Estimating Structural Models with Loose Constraints* (Mullins and Yu, n.d.), we propose a *loose constraint* estimator that chooses model parameters and endogenous model objects (such as policies, prices, and distributions) to simultaneously minimize an estimation criterion and a sum of squared deviations from the set of constraints that define the model solution. The method applies very broadly while maintaining computational and statistical advantages over existing approaches. It is particularly useful when one has access to direct estimates of endogenous objects that are nonetheless heavily biased or statistically imprecise. This can arise due to high model dimensionality relative to sample size, or due to incomplete data.

Theoretically, we show that for any estimator that imposes that the model constraints hold exactly, one can compute a weighting scheme that allows the loose constraint estimator to achieve the same asymptotic distribution. In addition, we show that the loose constraint estimator can be embedded inside an expectation-maximization algorithm that is guaranteed to move up the likelihood in each iteration, giving it a significant advantage over iterative CCP methods. Finally, we illustrate how the method is particularly well-suited to estimation of high-dimensional models when combined with machine learning solution approaches.

We demonstrate the desirable properties of the estimator for three sets of models: a single agent dynamic discrete choice model, an industry equilibrium model, and a search model.

A Structural Meta-Analysis of the Child Safety Net

An array of social programs provide a safety net to at-risk children and their families in the United States. These include food stamps, cash welfare, Medicaid, tax credits, and public provision of – and subsidies for – childcare. A vast literature has sought to assess, on a case-by-case basis, the short and long-run effects of these programs on child skill development and life-cycle outcomes.

Despite the existence of a large body of evidence, many policy-relevant questions remain unanswered. What does the evidence collectively tell us about what inputs and resources matter for skill development? What would the effects be of marginally expanding these policies in different combinations, and what does the evidence suggest about how these policies could be optimally combined?

In *Aggregating Evidence on the Child Safety Net for Policy Conclusions* (Mullins, n.d.), I adopt the conceptual and methodological innovations of Mullins (2024) and conduct a structural metaanalysis of experimental and quasi-experimental evidence on the short and long-run effects of each safety net program on child outcomes. In order to do this, the paper develops a model of child development with inputs of parental time, medical care, nutrition, childcare, and other money inputs. The model allows for heterogeneity across individuals and cohorts in order to interpret evidence across shifting economic, demographic, and policy landscapes. By aggregating the evidence, the estimated model can provide answers to normative and positive policy questions not directly addressed by the evidence alone.

Bibliography

- Almazura, Martin, Martin Garcia-Vazquez, and Joseph Mullins. n.d. "Ordinal Complementarity". *Working Paper*.
- Brown, Meta, Christopher Flinn, and Joseph Mullins. 2024. "Family Law Effects on Divorce, Fertility and Child Investment". *Accepted at Journal of Labor Economics*.
- Caucutt, Elizabeth, Lance Lochner, Joseph Mullins, and Youngmin Park. 2025. "Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments". *Accepted at Journal of Political Economy*.
- Flinn, Christopher, and Joseph Mullins. 2015. "Labor Market Search and Schooling Investment". International Economic Review, no. 56.
- Flinn, Christopher, and Joseph Mullins. 2021. "Firms' Choices of Wage-Setting Protocols". *R&R at Review of Economic Studies*.
- Mancino, Antonella, and Joseph Mullins. 2023. "Frictional Adjustment to Income Tax Incentives: An Application to the Earned Income Tax Credit". *R&R at Journal of Political Economy: Microeconomics.*
- Mullins, Joseph. n.d. "Aggregating Evidence on the Child Safety Net for Policy Conclusions". *Working Paper.*
- Mullins, Joseph. 2024. "A Structural Meta-Analysis of Welfare Reform Experiments and Their Impacts on Children". *Accepted at Journal of Political Economy.*
- Mullins, Joseph. 2025. "Designing Cash Transfers in the Presence of Children's Human Capital Formation". *R&R at Journal of Political Economy*.
- Mullins, Joseph, and Dan Yu. n.d. "Estimating Structural Models with Loose Constraints". *Working Paper*.