Subject to Availability

Option Production under Complementary Inputs

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

Policies create options by supplying complementary inputs. Each option can be described by the minimal input paths that make it actionable. A package surface has a path representation exactly when it is monotone; nonnegative Möbius coefficients characterize the case with one exact repair set per option and recover the mass attached to every set. Incomplete package support yields sharp linear-program bounds. A minimum-enablement dual prices the supplied inputs and assigns expenditure across providers. Childcare and work provide the application. In the 2024 ACS, model-assisted job-compatibility mass rises from 0.83 percent under an 8 a.m.–noon window to 23.74 percent under a 7:30 a.m.–5:30 p.m. package, and 93.1 percent of repaired mass requires at least two timing blocks. A randomized childcare–cash factorial bounds short-term response mass attributable to either supply route between 3.17 and 5.15 percentage points. A cross-fitted transport design links the trial to the 2019–2020 Uganda National Panel Survey. The childcare-factor effect on full-day use is 0.442 in the trial and 0.429 in the supported nine-district target at short term; the corresponding long-term effects are 0.313 and 0.204. Direct set measurements, randomized outcomes, population transport, and provider accounts distinguish option production, use, and resource provision.

JEL Codes: C14, C61, D11, D61, J13, J22

Keywords: actionable sets, policy packages, partial identification, childcare, labor supply, option value

---

Yonsei University. Email: eunixon@yonsei.ac.kr. This draft is circulated for discussion; comments and suggestions are welcome.

1. Introduction

A parent can accept a job when care covers departure through return, the workplace can be reached, the parent meets the job's requirements, and an employer offers the position. A public program chooses coverage hours, a firm chooses a schedule, a household member supplies care, and a transit system determines travel time. The economic object behind a package effect is the collection of input paths that turns a particular job into an action the parent can take.

Option production and choice among produced options are separate empirical objects. An active-set measure records which job cells become feasible under each package. Package outcomes record whether households then move into those cells. A policy design can therefore measure the inputs that create options, the providers that supply them, and the behavioral response that follows.

The application considers childcare coverage and job compatibility. Household-specific job cells combine occupation, work interval, commute, work location, and credentials.[1] Public data record these components and support an estimate of model-assisted job-compatibility mass.

In the 2024 ACS, model-assisted compatibility mass rises from 0.83 percent under an 8 a.m.–noon window to 23.74 percent under a 7:30 a.m.–5:30 p.m. package.[2] Sets containing at least two timing blocks account for 93.1 percent of repaired mass.

A randomized childcare–cash factorial supplies two routes into full-day care.[3] In the short-term use surface, response mass attributable to either route lies between 3.17 and 5.15 percentage points. Linking the trial to the 2019–2020 Uganda National Panel Survey separates the randomized-population effect from a supported nine-district target.[4] The childcare-factor effect on full-day use is 0.442 in the trial and 0.429 in that target at short term; the long-term estimates are 0.313 and 0.204.

The two applications use different evidence to identify related objects. The ACS construction measures job-compatibility mass and its repair composition. Adding employer offers, vacancies, care slots, and prices would extend that object to actionable-job mass. The Uganda experiment identifies assigned supply paths, delivered care, use, and the subsequent response in its randomized population. Its transport analysis declares a separate target law, retains common-support households, and carries trial and survey uncertainty

[1]A cell is counted as compatible when its measured timing, commute, location, and credential constraints are jointly satisfied. Employer offers, vacancies, care slots, and prices enter the larger actionable-job object developed below.

[2]The denominator is household-specific local reference-job mass. It differs from the employment rate and from the share of employed mothers whose chosen job fits the window.

[3]The childcare arm pays a registered center for one year of care, breakfast, and lunch. The cash arm transfers the same resource value to the mother.

[4]The transported target retains households on the declared common support and carries uncertainty from both the trial and survey samples.

into simultaneous intervals. The package representation connects the designs while keeping their populations, outcomes, and assignment mechanisms distinct.

For agent i and option j, let $\mathcal{D}_{ij}$ be the antichain of its minimal repair paths.[5] Package $S \subseteq \{1, \ldots, K\}$ activates the option when it contains some $D \in \mathcal{D}_{ij}$. With fixed nonnegative option weights w_{ij}, active-set mass is

$$M(S) = \mathbb{E}\left[\sum_j w_{ij} \mathbb{1}\{D \subseteq S \text{ for some } D \in \mathcal{D}_{ij}\}\right].$$

The minimal-path representation theorem shows that M has a nonnegative mixture representation over path antichains if and only if it rises with package inclusion. This is an exhaustive representation of monotone package surfaces, though its path distribution can remain set-valued. In the single-path case, D_{ij} is one exact repair set and

$$M(S) = \sum_{T \subseteq S} \theta(T), \qquad \theta(T) = \sum_{S \subseteq T} (-1)^{|T|-|S|} M(S).$$

Nonnegative Möbius coefficients characterize this case and identify the exact repair masses uniquely. An order-r saturated package contrast then measures option mass requiring exactly the corresponding r inputs. A negative coefficient inside a monotone package surface places the surface in the alternative-path class; with two inputs, its magnitude gives the lower bound on mass attributable to either route.

Observed policies may cover a subset of the package lattice. Linear programs over compatible path distributions deliver sharp bounds for repair mass and package counterfactuals on that support, and the row-space condition characterizes which linear functionals are point identified for every feasible surface. Direct set measurement and excluded input variation connect package outcomes to option production. Randomized package assignment then identifies movement into the option cells associated with a frozen repair map.

The resource problem asks what each repair path costs and who supplies it. A minimum-enablement program selects the inputs that relax an option's binding constraints. Its dual multipliers price those constraints, assign expenditure across providers, and provide the marginal values used in pre-choice welfare accounting.

The analysis draws on work on opportunity-set welfare, latent choice sets, institutional constraints, and factorial policy design (Kreps 1979; Sen 1993; Pattanaik and Xu 1990; Masatlioglu et al. 2012; Cattaneo et al. 2020; Abaluck and Adams-Prassl 2021; Barseghyan et al. 2021; Agarwal and Somaini 2026; Mbiti et al. 2019; Muralidharan et al. 2025). Random-set cores, submodular minimization, and transport formulations characterize sharp iden-

[5] An antichain removes redundant supersets: no retained repair path contains another retained path for the same option.

tified sets in multiple-equilibrium models (Galichon and Henry 2011). Möbius inversion and complete monotonicity characterize exact repair sets (Rota 1964; Chateauneuf and Jaffray 1989), while minimal component sets connect the representation to sufficient-cause models (VanderWeele and Robins 2008). The unit of analysis is an option together with the input technology that produces it.

Actionable sets and constraint ownership lead to the prices of minimum enablement. Repair paths then connect package identification to randomized transitions and trial-to-target transport. The childcare applications carry these objects into welfare, provider accounting, the coverage frontier, and the repair distribution. Proofs, construction formulas, sensitivity results, and data construction details appear in the online appendix.

2. Actionable Sets and Constraint Ownership

2.1. Timing

There are beneficiaries $i \in \{1, \ldots, I\}$ and input providers $k \in \{1, \ldots, K\}$. Availability state q_i is observed. Providers choose $e_{ki} \in \mathbb{R}_+^{m_{ki}}$ and pay $c_{ki}(e_{ki})$. State ω is then realized, the active set forms, and beneficiary i selects an action. This timing allows inputs to carry value across states that leave current realized choice unchanged.

Agent i has a universe of active actions $\mathcal{X}_i^+$ and an outside action 0. Write $e_i = (e_{1i}, \ldots, e_{Ki})$. Given (q_i, ω, e_i), the active set is

$$\mathcal{F}_i(q_i, \omega; e_i) = \left\{a \in \mathcal{X}_i^+ : g_{i\ell}(a, q_i, \omega; e_i) \le 0,\ \ell = 1, \ldots, L_i\right\}.$$

Observed choice satisfies

$$a_i^*(q_i, \omega; e_i) \in \underset{a \in \{0\} \cup \mathcal{F}_i(q_i, \omega; e_i)}{\operatorname{argmax}} u_i(a, \omega).$$

Active-set collapse is the event $\mathcal{F}_i(q_i, \omega; e_i) = \varnothing$. The observed action remains 0 on this event. The outside action isolates collapse from the existence of the choice maximizer.

2.2. Constraint ownership

Each input enters a recorded subset of constraints through

$$g_i(a, q_i, \omega; e_i) = \bar{g}_i(a, q_i, \omega) - \sum_{k=1}^{K} H_{ki}(a, q_i, \omega) e_{ki}.$$

The map H_{ki} records the constraints affected by provider k. A household member can cover care hours, a firm can alter a shift, a transit provider can reduce travel time, and a public program can supply a slot. The cost c_{ki} belongs to the provider's objective; H_{ki} enters the beneficiary's active set.

The one-period indirect value is

$$v_i(q_i, \omega; e_i) = \max_{a \in \{0\} \cup \mathcal{F}_i(q_i, \omega; e_i)} u_i(a, \omega), \qquad W_i(q_i; e_i) = \mathbb{E}_\omega[v_i(q_i, \omega; e_i)].$$

Provider k's payoff takes the form

$$\Pi_{ki}(e_i) = r_{ki}(e_i) + \alpha_{ki} W_i(q_i; e_i) - c_{ki}(e_{ki}),$$

where r_{ki} collects own operating returns and α_{ki} records the share of beneficiary value captured by the provider. The planner includes W_i, provider returns, and resource costs under declared welfare weights.

2.3. Actionability and consideration

The baseline sets the considered set equal to the active set. An extension introduces $\mathcal{C}_i \subseteq \mathcal{F}_i$ through attention or search.[6] This ordering separates actionability from the decision process that selects which active alternatives receive consideration. For childcare and work, a feasible combination must first cover the required times and satisfy the other recorded constraints. Whether the parent prefers that combination, notices it, or substitutes it for another arrangement determines subsequent choice. Joint use of care and work can therefore reflect both the input technology and demand; the measured repair map supplies the production-side object.

3. The Cost of Creating an Option

3.1. Primal and dual programs

Fix a beneficiary and a realized state. Fold the fixed restrictions on active actions into the constraint system. Consider the linear technology $Ba \le b + He$, with free action vector $a \in \mathbb{R}^d$, enabling vector $e \in \mathbb{R}^m_+$, and resource prices $p \in \mathbb{R}^m_{++}$. The minimum expenditure that creates an active action is

$$C^* = \min_{a \in \mathbb{R}^d,\, e \ge 0} p'e \quad \text{s.t.} \quad Ba - He \le b.$$

[6] For revealed-attention and stochastic consideration-set models, see Masatlioglu et al. (2012); Cattaneo et al. (2020); Abaluck and Adams-Prassl (2021); Barseghyan et al. (2021); Caplin et al. (2019). Those models organize selection from available alternatives; the active-set object here records which alternatives are usable before consideration.

PROPOSITION 1 (Minimum enablement duality). *Suppose the primal program is feasible. Its value is finite, an optimum exists, and*

$$C^* = \max_{\lambda \geq 0} -\lambda' b \quad s.t. \quad B'\lambda = 0, \quad H'\lambda \leq p.$$

For any primal-dual optimum (a^*, e^*, λ^*)*,*

$$C^* = p' e^* = \lambda^{*\prime} H e^* = -\lambda^{*\prime} b.$$

Every input r with $e_r^* > 0$ *satisfies* $H_r' \lambda^* = p_r$.

The primal selects an action and a repair vector. The dual combines the constraints whose incompatibility creates the repair cost and assigns a weight to each of them. If $H = (H_1, \ldots, H_K)$ groups inputs by provider, the assignment

$$S_k^* = \lambda^{*\prime} H_k e_k^*$$

satisfies $\sum_k S_k^* = C^*$. This equality records resource expenditure by constraint owner in the units of p.

3.2. Collapse and repair certificates

At $e = 0$, active-set collapse in the linear system is $\{a : Ba \leq b\} = \varnothing$. Farkas' lemma supplies $\lambda \geq 0$ with $B'\lambda = 0$ and $-\lambda' b > 0$. This vector records a weighted contradiction among the current constraints.[7] The minimum-enablement dual adds $H'\lambda \leq p$, which prices the technologies available to resolve that contradiction.

For a finite family of convex constraint sets in $\mathbb{R}^d$, Helly's theorem gives a subfamily with at most $d + 1$ members whose intersection is empty. The certificate dimension therefore follows the action-space dimension. Mixed-integer action cells and nonconvex schedule restrictions require separate certificates; the empirical algorithm reports the formulation used for each job cell (Basu et al. 2021).

3.3. Comparative statics

Let $C(b, p, H)$ denote the value function. At points with a unique dual solution, $\partial C / \partial b_\ell = -\lambda_\ell^*$. For an input column H_r, a marginal increase in its effectiveness changes cost at rate $-\lambda^{*\prime}(dH_r)e_r^*$. For a price change, $\partial C / \partial p_r = e_r^*$ at a unique primal optimum. These derivatives supply the constraint, technology, and resource-price margins used in the application.

[7] A Farkas certificate is scale-free and may be nonunique. The price inequalities in the dual program normalize the economically relevant certificate through the available repair technologies.

4. Repair Sets and Policy Packages

4.1. Minimal repair paths

Let $\mathcal{T} = \{1, \ldots, K\}$ index inputs that a policy package can supply. For every agent–option pair (i, j), let $\mathcal{D}_{ij}$ be an antichain of minimal repair paths: each $D \in \mathcal{D}_{ij}$ activates option j, and no member contains another. Package $S \subseteq \mathcal{T}$ activates the option according to

$$A_{ij}(S) = \mathbb{1}\{D \subseteq S \text{ for some } D \in \mathcal{D}_{ij}\}.$$

The antichain records alternative ways to produce the same option while preserving its identity and weight. A household member and a care center can provide alternative care paths; a bus route and remote work can provide alternative commute paths. Fix nonnegative weights w_{ij} before package assignment and define

$$M(S) = \mathbb{E}\left[\sum_j w_{ij} A_{ij}(S)\right].$$

Let $\mathfrak{A}_K$ denote the nonempty antichains of $2^{\mathcal{T}}$, let $B_{S,\mathcal{D}} = \mathbb{1}\{D \subseteq S \text{ for some } D \in \mathcal{D}\}$, and let $\psi(\mathcal{D})$ be weighted option mass with path antichain $\mathcal{D}$.

THEOREM 1 (Minimal-path representation theorem). *Fix candidate option mass $\bar{M}$. A package surface $M : 2^{\mathcal{T}} \to \mathbb{R}$ has a minimal-path representation*

$$M(S) = \sum_{\mathcal{D} \in \mathfrak{A}_K} B_{S,\mathcal{D}} \psi(\mathcal{D}), \qquad \psi \geq 0, \qquad \mathbf{1}'\psi \leq \bar{M},$$

if and only if $0 \leq M(S) \leq M(S') \leq \bar{M}$ whenever $S \subseteq S'$. The representation can contain several path distributions with the same surface. It has a representation supported on singleton antichains $\mathcal{D} = \{T\}$ if and only if every Möbius coefficient of M is nonnegative. In that case the singleton-path masses are unique.

The two representation classes impose different empirical restrictions. Monotonicity is sufficient for option production through arbitrary alternative paths. Requiring nonnegative Möbius coefficients narrows the model to one exact missing-input set per option. With two inputs, an option activated by either singleton has surface $(0, 1, 1, 1)$ and a second-order Möbius coefficient of -1. The negative coefficient records the alternative paths shared by the singleton packages.

4.2. Exact repair sets

Suppose each path antichain is a singleton, $\mathcal{D}_{ij} = \{D_{ij}\}$, and define

$$\theta(T) = \mathbb{E}\left[\sum_j w_{ij}\mathbb{1}\{D_{ij} = T\}\right].$$

PROPOSITION 2 (Repair-set inversion). *For every $S \subseteq \mathcal{T}$,*

$$M(S) = \sum_{T \subseteq S} \theta(T).$$

Consequently, every exact repair mass is recovered by

$$\theta(T) = \sum_{S \subseteq T} (-1)^{|T|-|S|} M(S) \geq 0.$$

Let $\Theta_r = \sum_{|T|=r} \theta(T)$. The sequence $(\Theta_0, \ldots, \Theta_K)$ gives active-option mass by the number of missing inputs. For any T, the saturated package contrast over the arms contained in T equals $\theta(T)$.

The inversion coefficients provide two diagnostics. A negative population coefficient violates the exact-repair model; in a sample, a negative estimate may also reflect sampling error, displacement across option weights, package-induced changes in the candidate universe, or a gate with alternative repair paths. The lowest order with positive Θ_r gives the package order at which activation begins. Singleton effects measure $\theta(\{k\})$ and exclude options whose repair sets contain both k and additional inputs.

For disjoint packages $R, S \subseteq \mathcal{T}$, define the mixed package increment

$$\Delta_R M(S) = \sum_{Q \subseteq R} (-1)^{|R|-|Q|} M(S \cup Q).$$

THEOREM 2 (Package-surface characterization). *Let $M : 2^{\mathcal{T}} \to \mathbb{R}$. The following statements are equivalent: (i) there is a unique vector $\theta \geq 0$ satisfying $M(S) = \sum_{T \subseteq S} \theta(T)$ for every package S; (ii) every Möbius coefficient $\sum_{S \subseteq T} (-1)^{|T|-|S|} M(S)$ is nonnegative; and (iii) $M(\varnothing) \geq 0$ and $\Delta_R M(S) \geq 0$ for every disjoint pair (R, S). Under these conditions,*

$$\Delta_R M(S) = \sum_{V \subseteq S} \theta(R \cup V).$$

Thus the marginal active-set mass produced by package R rises weakly as the background package expands. Every order of a saturated package contrast has a repair-mass interpretation, and any negative population contrast rejects the conjunction technology, fixed option weights, and stable candidate universe.

The equivalence is the complete-monotonicity characterization of a set function with a nonnegative Möbius transform (Rota 1964; Chateauneuf and Jaffray 1989). Its use here turns the package response surface into a specification test that separates a single missing-input bundle from alternative repair paths. A direct active-set measure produced by one exact path lies in this cone. A direct active-set measure with several paths lies in the larger monotone cone from Theorem 1. Employment, income, and other realized outcomes can leave both cones because choice, substitution, and equilibrium responses occur after activation.

The active-set measure remains separate from outcomes generated after activation. Let $R(S)$ be the probability of choosing an active option and $V(S)$ be indirect utility. Preferences, substitution, and equilibrium responses map $M(S)$ into $R(S)$ and $V(S)$. Empirical tables report package-surface diagnostics, repair mass, movement into repaired cells, and welfare calculations as distinct objects.

4.3. Limited package support

Let $\mathcal{O} \subseteq 2^{\mathcal{T}}$ be the package arms observed in an experiment or policy record. Index columns by $T \subseteq \mathcal{T}$, rows by $S \in \mathcal{O}$, and define the incidence matrix

$$(A_{\mathcal{O}})_{S,T} = \mathbb{1}\{T \subseteq S\}.$$

Let $\bar{M} = \mathbb{E}[\sum_j w_{ij}]$ denote the candidate option mass. Mass outside the declared input universe appears as slack in the total-mass restriction.

THEOREM 3 (Sharp repair bounds). *Suppose the exact-repair representation holds and $M(S) = m_S$ is observed for $S \in \mathcal{O}$. The identified set for repair masses is*

$$\mathcal{H}(m_{\mathcal{O}}, \mathcal{O}) = \left\{\theta \in \mathbb{R}_+^{2^K} : A_{\mathcal{O}}\theta = m_{\mathcal{O}},\ \mathbf{1}'\theta \le \bar{M}\right\}.$$

For any declared functional $q'\theta$, its sharp lower and upper bounds solve

$$\underline{q} = \min_{\theta \in \mathcal{H}(m_{\mathcal{O}},\mathcal{O})} q'\theta, \qquad \overline{q} = \max_{\theta \in \mathcal{H}(m_{\mathcal{O}},\mathcal{O})} q'\theta.$$

The functional is point identified for every feasible $m_{\mathcal{O}}$ when q belongs to the row space of $A_{\mathcal{O}}$. Full repair-vector identification requires full column rank. If active-set measurements satisfy $L_S \le M(S) \le U_S$, replacing $A_{\mathcal{O}}\theta = m_{\mathcal{O}}$ by $L_{\mathcal{O}} \le A_{\mathcal{O}}\theta \le U_{\mathcal{O}}$ gives the sharp set under those intervals.

The bounds use package incidence, nonnegative option mass, and a fixed candidate universe. A downward-closed arm family identifies $\theta(T)$ whenever it contains every subset of T. If an arm is missing, the linear programs show which repair bundles remain pooled.[8] Appending the incidence row of a proposed arm and recomputing the target interval shows how much that arm would add to identification.

For the path model, use antichain columns. The sharp path-distribution set is

$$\mathcal{H}^P(m_{\mathcal{O}}, \mathcal{O}) = \left\{\psi \in \mathbb{R}_+^{|\mathfrak{A}_K|} : B_{\mathcal{O}}\psi = m_{\mathcal{O}},\ \mathbf{1}'\psi \le \bar{M}\right\}.$$

Every functional of path types receives LP bounds over $\mathcal{H}^P$. Comparing feasibility of $\mathcal{H}$ and $\mathcal{H}^P$ distinguishes a single exact path, alternative paths, and violations of package monotonicity. Full package support identifies the surface and the unique exact-set vector when $\mathcal{H}$ is feasible; it generally leaves the antichain distribution set-valued.

4.4. Package choice under the identified set

Let a_S be the incidence row for a candidate package and let $c(S)$ be its resource cost in the same units as the option weights. Net active-set value is $a_S'\theta - c(S)$. Under limited support, a decision maker can rank S against S' for every repair distribution consistent with the data when

$$\min_{\theta\in\mathcal{H}(m_{\mathcal{O}},\mathcal{O})} \left[(a_S - a_{S'})'\theta - c(S) + c(S')\right] \ge 0.$$

The criterion is an LP. A robust package solves

$$S^R \in \operatorname*{argmax}_{S\in\mathcal{P}} \min_{\theta\in\mathcal{H}(m_{\mathcal{O}},\mathcal{O})} \{a_S'\theta - c(S)\},$$

for a declared feasible package family $\mathcal{P}$. Costs, capacity, and provider incidence enter $\mathcal{P}$ and $c(S)$; the repair distribution supplies the benefit side. Singleton accounting selects packages using only $\theta(T)$ with $|T| \le 1$. When positive mass lies at higher orders, that accounting assigns zero activation benefit to the corresponding options.

For a target $q'\theta$, let θ^L and θ^U attain its current lower and upper bounds. A candidate arm S has extremizer-separation score

$$G_q(S) = \left|a_S'(\theta^U - \theta^L)\right|.$$

Every arm with $G_q(S) > 0$ distinguishes the two displayed endpoint distributions. To evaluate every observationally equivalent distribution, define the fiber $\mathcal{H}_S(\theta^0) = \{\theta \in \mathcal{H} : a_S'\theta = a_S'\theta^0\}$ and the worst-case post-arm width

$$W_q(S) = \sup_{\theta^0\in\mathcal{H}} \left\{\max_{\theta\in\mathcal{H}_S(\theta^0)} q'\theta - \min_{\theta\in\mathcal{H}_S(\theta^0)} q'\theta\right\}.$$

[8] Point identification requires only that the target lie in the observed incidence row space.

Equivalently, $W_q(S)$ is the value of the linear program

$$\max_{\theta^1,\theta^2\in\mathcal{H}} q'(\theta^1-\theta^2) \quad \text{s.t.} \quad a_S'(\theta^1-\theta^2)=0.$$

The design selects a feasible arm minimizing $W_q(S)$, uses $G_q(S)$ to screen tied or costly candidates, and recomputes the post-arm row space. The width equals zero precisely when the target is constant on every fiber generated by the added arm. Membership of q in the row space of the augmented incidence matrix is a sufficient condition. The same criterion applies to $\mathcal{H}^P$ with antichain incidence rows.

4.5. Threshold gates

The repair-set representation can be generated by input thresholds. Input k satisfies its requirement for option j when $x_{jk}(e,\omega)=\mathbb{1}\{e_k\geq\tau_{jk}(\omega)\}$. A conjunction gate activates the option when each required status equals one. For finite input increments, its order-r cross-difference equals the weighted probability that all r thresholds cross inside their assigned increments while the remaining requirements are satisfied. The online appendix proves the formula and gives the sign change for alternative activation paths. The threshold model links continuous coverage hours and travel times to the binary repair sets used in a package design.

5. Measurement and Identification

5.1. Choice equivalence and the measurement requirement

Choice probabilities can remain fixed while the active set changes. The following result states the resulting identification barrier before imposing the repair technology.

PROPOSITION 3 (Choice-equivalent availability systems). *Fix a finite option universe $\mathcal{J}$, an outside action 0, observed state x, and any conditional choice distribution $p(a\mid x)$ on $\mathcal{J}\cup\{0\}$. The distribution is generated by each of two systems: (i) every option is active and utility types select action a with probability $p(a\mid x)$; (ii) all agents share utilities satisfying $u(j)>u(0)$ for $j\in\mathcal{J}$, while the active set equals $\{j\}$ with probability $p(j\mid x)$ and is empty with probability $p(0\mid x)$. The systems have identical choice probabilities and different active-set distributions. Conditional choice probabilities, including randomized-arm probabilities obtained by including assignment in x, therefore leave active-set mass and repair mass unidentified without an excluded restriction or an auxiliary active-set measure.*

This identification barrier motivates the empirical sequence used below. A reference sample fixes option cells and repair types. Package assignment then changes the declared inputs, direct follow-up measures which frozen cells are actionable, and observed transitions record selection from the measured set. Arm outcomes remain behavioral objects throughout the analysis.

5.2. Package effects and repair masses

The repair estimand requires an active-set measure for each observed package. In a factorial experiment, random assignment identifies $M(S)$ from arm means when the option weights and candidate universe are fixed before assignment. A saturated arm regression returns the same cell means and their Möbius contrasts. The analysis reports every prespecified interaction order because dropping interaction terms changes the estimand (Muralidharan et al. 2025).

Policy records often supply a nested or irregular family $\mathcal{O}$. Estimated arm means, their joint covariance matrix, and measurement intervals enter the bound program in Theorem 3. Joint resampling of arm means and generated option weights produces a confidence region that covers the union of feasible repair polytopes across a confidence set for $M_{\mathcal{O}}$. Arm selection, target functionals, and the order of package contrasts are fixed before outcome inspection.

COROLLARY 1 (Projection inference). *Let $\mathcal{C}_M(1-\alpha)$ be a joint $(1-\alpha)$-confidence set for the supported active-set masses $m_{\mathcal{O}}$. Define*

$$\mathcal{C}_\theta(1-\alpha) = \bigcup_{m \in \mathcal{C}_M(1-\alpha)} \left\{\theta \geq 0 : A_{\mathcal{O}}\theta = m,\ \mathbf{1}'\theta \leq \bar{M}\right\}.$$

Then $\mathcal{C}_\theta(1-\alpha)$ covers the true repair vector with probability at least $1-\alpha$. Minimizing and maximizing any prespecified collection of linear functionals over this single set gives simultaneous coverage for the collection. Replacing $A_{\mathcal{O}}$ and θ by antichain incidence $B_{\mathcal{O}}$ and path mass ψ gives the corresponding path-distribution confidence set.

Randomization inference supplies $\mathcal{C}_M$ in experiments, and survey replicate weights or joint resampling of the reference and target samples supplies it in constructed-set applications. The reported confidence object states which source of sampling variation enters $\mathcal{C}_M$.

5.3. Generated repair sets

Let $\widehat{D}_{ij}^R$ be the missing-input set predicted for target household i and job cell j from a reference sample R. Reference-worker schedules, local occupation weights, commute mappings, credential rules, care windows, and price bounds determine this map. The active indicator under package S is

$$\widehat{A}_{ij}^R(S) = \mathbb{1}\{\widehat{D}_{ij}^R \subseteq S\}.$$

The application excludes target mothers from R, freezes the resulting state schedule cells and PUMA occupation weights, and evaluates the map on target households. Person-weighted variation across target households supplies the reported sampling intervals. Pooling rules, schedule assignments, and time-bin bounds supply construction sensitivity.

Every input component has a lower and upper classification. A lower repair set contains requirements recorded as missing under every admissible schedule, commute, price, offer, and slot value. An upper classification contains requirements missing under at least one admissible value. These two constructions generate interval restrictions on $M(S)$, which enter Theorem 3. The held-out sample compares predicted compatible cells with reported work intervals and care disruptions.

5.4. Targeted transitions

For a frozen repair map, let $\mathcal{G}^R_{iT} = \{j : \widehat{D}^R_{ij} = T\}$ collect agent i's option cells assigned to exact repair type T. Let $C_i(S)$ denote the option chosen after assignment to package S, and define the targeted-response probability

$$R_T(S) = \mathbb{P}\Big(C_i(S) \in \mathcal{G}^R_{iT}\Big).$$

The active-set estimand $\theta(T)$ records option mass, while $R_T(S)$ records movement into that mass. Positive option value and current movement remain separate objects when agents substitute among active options.

PROPOSITION 4 (Repair-targeting restrictions). *Suppose the frozen repair map equals the realized input technology over the declared candidate universe, package assignment supplies exactly the inputs in S, and choice belongs to the active set together with an outside action. Then*

$$R_T(S) = 0 \qquad \textit{whenever } T \nsubseteq S.$$

Let $H_T = \mathbb{P}(\mathcal{G}^R_{iT} \neq \varnothing)$. The minimal-package response satisfies

$$0 \leq R_T(T) \leq H_T.$$

Under randomized package assignment, arm means identify every supported $R_T(S)$. Positive movement into $\mathcal{G}^R_{iT}$ under an arm with $T \nsubseteq S$ measures a violation of the frozen repair map, package exclusion, or candidate-universe stability.

The corresponding design has three stages. A reference sample fixes option cells and their repair types. Random assignment or a policy event supplies package variation. Follow-up data record direct actionability for the frozen cells and realized movement into each repair class. Möbius inversion applies to the direct actionability measure, while the targeted-response matrix reports which repaired cells reach behavior. Repair types containing an input omitted by the assigned arm provide constraint-specific placebo cells.

5.5. Coverage frontiers

For opening time o and closing time c, define household i's measured compatibility frontier by

$$F_i(o,c) = \sum_{j \in \mathcal{J}_m} N^0_{jm} A^{\text{cred}}_{ij} \mathbb{1}\{t^-_j \geq o,\ t^+_j \leq c\}.$$

For $c_2 > c_1$, its closing-time increment is

$$F_i(o,c_2) - F_i(o,c_1) = \sum_{j \in \mathcal{J}_m} N^0_{jm} A^{\text{cred}}_{ij} \mathbb{1}\{t^-_j \geq o,\ c_1 < t^+_j \leq c_2\}.$$

The increment measures job-cell mass whose return time lies in the added coverage interval. For $o_2 < o_1$, the opening-time increment equals the mass with $o_2 \leq t^-_j < o_1$ and $t^+_j \leq c$. A grid of openings and closings therefore gives the marginal schedule mass attached to each added time block.

Partition the interval between the baseline and full package into blocks e, d, l. For every continuous job-travel interval inside the full window, its exact repair set contains the blocks intersected by the interval outside baseline coverage. Summing cell weights by this classification estimates $\theta(T)$ directly. Comparing these direct repair masses with the bounds generated by a restricted package family separates information supplied by policy arms from information supplied by measured schedules.

5.6. Validation

The construction is checked against three implications. The frontier rises weakly as the closing time advances and as the opening moves earlier. Exact repair masses are nonnegative and sum to full-package compatibility mass; a continuous interval that intersects both the early and late blocks also intersects the midday block. Finally, an auxiliary survey with reported start and end bins should place the chosen-job compatibility rate inside its interval estimate. The code records these identities, pooling sensitivity, source hashes, and held-out SIPP statistics.

5.7. Randomized alternative supply paths

The Uganda factorial supplies two routes into target-child full-day care. The childcare arm pays a registered center for one year of care, breakfast, and lunch. The cash arm transfers the same resource value to the mother. Let c and k index the two assignments and let $M(S)$ denote regression-adjusted target-child use or entry in assignment cell S. This surface is a behavioral response measure. Its path decomposition describes observationally equivalent response types under the monotone surface restriction and remains separate from the directly measured compatibility surface in the ACS application. For two inputs, the path mixture contains masses $\psi_c, \psi_k, \psi_{c\vee k}, \psi_{c\wedge k}$, and the package equations imply

$$\psi_{c\wedge k} - \psi_{c\vee k} = M(\{c,k\}) - M(\{c\}) - M(\{k\}) + M(\varnothing).$$

When the four cell means are monotone, alternative-path mass satisfies

$$\max\{0, -\Delta_{ck}M(\varnothing)\} \le \psi_{c\vee k} \le \min\{M(\{c\}) - M(\varnothing), M(\{k\}) - M(\varnothing)\}.$$

TABLE 1. Randomized supply paths into full-day childcare

Outcome	Ctl	Care	Cash	Both	Either-route bounds	Exact joint (p-value)
Full-day childcare use, short term	0.371	0.828	0.423	0.848	[3.17, 5.15]	-3.17 (0.482)
Entry into full-day childcare, short term	0.364	0.824	0.413	0.843	[2.99, 4.86]	-2.99 (0.514)
Full-day childcare use, long term	0.491	0.840	0.579	0.869	[5.94, 8.86]	-5.94 (0.192)

Notes: Cell means standardize the saturated treatment regression over the estimation sample's covariates. Alternative-path bounds are percentage points and apply when the displayed four-cell surface is monotone. Exact joint is the Möbius coefficient in percentage points; its two-sided robust p-value appears in parentheses. The three displayed surfaces satisfy package monotonicity at their point estimates.

Short-term full-day use gives an alternative-path interval of [3.17, 5.15] percentage points. Short-term entry gives [2.99, 4.86], and long-term use gives [5.94, 8.86]. Their exact joint coefficients have two-sided p-values of 0.482, 0.514, and 0.192. Any-care use and entry produce exact joint coefficients of −8.53 and −13.26 points with p-values 0.003 and 0.002. Their combined-arm means lie 1.08 and 1.10 points below the childcare-arm means, with lower-tail p-values 0.232 and 0.291. These rows separate sampling evidence about package monotonicity from evidence about the single-path restriction.

The experiment also supplies a resource account for each route. The annual childcare payment averaged 411,752 Ugandan shillings and went from the program to the selected registered center. The cash path delivered the same amount to the mother by mobile money.

TABLE 2. Resource and provider accounts for two supply paths

Supplied path	Constraint row	Resource recipient	Annual program resource
Full-day childcare	Target-child care time	Registered childcare center	411,752 UGX
Cash grant	Household liquidity	Mother by mobile money	411,752 UGX

Notes: The childcare resource covers tuition, breakfast, and lunch for one year. Household transport remains outside the observed program account. The cash grant has the same face value and enters the household-liquidity constraint. Amounts follow the experimental protocol reported by Bjorvatn et al. (2025).

5.8. A randomized incomplete-repair design

Bjorvatn et al. (2025) assigned mothers of children ages three through five in Uganda to a control arm, a one-year offer of free full-day childcare for one target child, an equivalent cash grant, or both. The childcare package covered tuition, breakfast, and lunch at a nearby registered center; the study team assisted enrollment and paid the center directly. Assignment therefore supplies a recorded childcare-access input for the target child. A baseline indicator records whether the target child has a younger sibling. In those households, care for the target child leaves an additional childcare requirement for work actions that require care for every young child.

Let C_i and K_i denote randomized childcare and cash assignment and let Z_i indicate a younger sibling. A saturated regression includes the four assignment cells, Z_i, their interactions, randomization strata, and the outcome's baseline value when collected. Define the average marginal childcare effect across cash states by

$$\tau_C(z) = \frac{1}{2}\{\mu(1,0,z) - \mu(0,0,z) + \mu(1,1,z) - \mu(0,1,z)\}.$$

The repair-targeting contrast is $\delta_C = \tau_C(1) - \tau_C(0)$. The target-child input is offered in both classes; the class $Z_i = 1$ retains a second-child care requirement. Target-child use records delivery in each class, while δ_C measures the difference in economic response associated with that residual requirement. The original study reports the younger-sibling split in Online Appendix Table A.9; Table 4 displays all nine maternal earnings, employment, and hours outcomes in that table.

TABLE 3. Delivery of the randomized childcare input

Outcome	Target youngest	Younger sibling	Interaction	p	N
Full-day childcare use, short term	0.440 (0.027)	0.444 (0.043)	0.004 (0.051)	0.943	1,430
Entry into full-day childcare, short term	0.444 (0.027)	0.447 (0.043)	0.003 (0.051)	0.954	1,404
Full-day childcare use, long term	0.326 (0.027)	0.303 (0.044)	-0.023 (0.051)	0.657	1,358

Notes: Each effect is the average marginal effect of randomized childcare assignment across the two cash-assignment states. The columns split households by the target child's baseline status as the youngest child or as having a younger sibling. Regressions include the saturated assignment cells, the younger-sibling interaction, randomization strata, the baseline outcome and its missing indicator where applicable, and HC1 standard errors. Interaction is the second column's effect minus the first column's effect.

Table 3 shows that the childcare factor raises short-term full-day use by 0.440 when the target child is youngest and by 0.444 when a younger sibling is present. The interaction is 0.004 with standard error 0.051, p-value 0.943, and 95 percent confidence interval $[-0.095, 0.103]$. Full-day entry gives 0.444 and 0.447, with interaction p-value 0.954. At long-term follow-up, the effects are 0.326 and 0.303, with interaction p-value 0.657. These estimates record the delivery point estimate and its uncertainty for each baseline repair class.

TABLE 4. Economic response when an additional childcare requirement remains

Outcome	Target youngest	Younger sibling	Interaction	p	N
	11.09	-4.41	-15.50		
Self-employment profits	(4.28)	(5.90)	(7.29)	0.034	1,414
	-2.34	-5.17	-2.84		
Wage earnings	(2.71)	(4.16)	(4.94)	0.566	1,414
	8.97	-8.24	-17.21		
Total labor income	(5.28)	(7.51)	(9.16)	0.060	1,414
	-0.01	0.00	0.01		
Self-employed	(0.03)	(0.05)	(0.06)	0.889	1,414
	-2.14	4.43	6.58		
Self-employment hours	(9.48)	(13.96)	(16.84)	0.696	1,414
	-0.02	-0.01	0.02		
Wage-employed	(0.02)	(0.03)	(0.04)	0.649	1,414
	-10.34	4.28	14.62		
Wage-employment hours	(4.55)	(6.37)	(7.83)	0.062	1,414
	-0.02	0.01	0.03		
Employed	(0.03)	(0.05)	(0.06)	0.589	1,414
	-12.85	6.89	19.74		
Total work hours	(9.84)	(14.69)	(17.66)	0.264	1,414

Notes: Each effect is the average marginal effect of randomized childcare assignment across the two cash-assignment states. The table reports all nine maternal economic outcomes in the original younger-sibling heterogeneity table. Monetary outcomes are thousands of Ugandan shillings in the prior month, winsorized at the 99th percentile following the original code; hours refer to the prior month. Regressions use the specification in Table 3. Interaction is the effect for households with a younger sibling minus the effect for households whose target child is youngest.

The childcare factor raises monthly self-employment profits by 11.09 among households whose target child is youngest and changes profits by −4.41 among households with a younger sibling. Their difference is −15.50, with p-value 0.034. The corresponding total-labor-income effects are 8.97 and −8.24; their difference is −17.21, with p-value 0.060. The other seven interactions have p-values from 0.062 to 0.889; wage-employment hours produce the 0.062 value. A covariance-based joint test of the nine interactions yields $\chi^2(9) = 20.14$ and p-value 0.017. Recomputing that Wald statistic over 4,999 assignments within the 156 randomization strata gives p-value 0.017. The full outcome vector places a 0.004 target-child-use interaction beside the allocations of earnings and hours across the two baseline repair classes.

The sensitivity analysis varies upper-tail rules, controls, districts, inference, outcome family, and treatment component. Across raw outcomes and caps from the 99.5th through 95th percentiles, the profit contrast ranges from −17.92 to −9.28, and the total-income contrast ranges from −19.21 to −11.01. The profit contrast is −11.88 with assignment-cell controls, −13.39 after adding randomization strata, and −15.50 after adding the baseline outcome; the respective p-values are 0.122, 0.069, and 0.034. Leave-one-district estimates retain the negative sign in every case. Stratified randomization-inference p-values are

0.035 for profits and 0.053 for total income. Max-t adjustment gives p-values 0.193 and 0.311, and Holm adjustment gives 0.303 and 0.484. The childcare-minus-cash heterogeneity contrasts are −18.76 for profits and −21.11 for total income, with p-values 0.083 and 0.114. A sum-of-squared-studentized-effects randomization omnibus gives p-value 0.187. The covariance-weighted randomization Wald gives p-value 0.017. The family tests answer separate aggregation questions, and the adjusted component tests retain uncertainty about the outcome allocation.

This design observes an assigned enabling input and its realized use. It also fixes a baseline class in which an additional care component remains. Because job-level actionability, work schedules, and care arrangements for the younger sibling are unmeasured, the exercise identifies an incomplete-package response diagnostic. Direct job-cell repair mass still requires the auxiliary mapping in Proposition 3. Arm-by-arm enrollment transitions, the original factorial sign diagnostic, outcome construction, and replication checks accompany the sensitivity record.

5.9. Trial and population targets

Let $Q \in \{E, T\}$ index the experimental population and a survey target, let $A = (C, K) \in \{0, 1\}^2$ denote childcare and cash assignment, and let R denote a roster-mapped repair class. For outcome Y, define the childcare-factor effect in population Q as

$$\tau_r^Q = \frac{1}{2}\mathbb{E}_Q[\{Y(1,0) - Y(0,0)\} + \{Y(1,1) - Y(0,1)\} \mid R = r].$$

The experimental target averages over the randomized sample. The population target contains children ages four through six in the nine experimental districts in the 2019–2020 Uganda National Panel Survey (Uganda Bureau of Statistics 2019). Following the trial archive's population comparison, survey age is shifted back one year to align with the trial's baseline ages three through five. A household-size support rule retains 161 children in 137 households, representing 488,386 children under survey weights. It removes 3.66 percent of the initial nine-district target weight. The class $R = 1$ records another roster member younger than the target child and is constructed identically in both files; this measure agrees with the registered trial stratum for 92.2 percent of trial children. The $R = 0$, age-five cell has no trial observations, so its class-specific target covers aligned ages three and four and removes 17.9 percent of the initial $R = 0$ target weight.

The common covariate basis contains child-age indicators, district indicators, child sex, current school attendance, the roster-mapped repair class, household size, and household size squared. Five folds keep each of 389 trial villages and each of 137 target households together. Separate fold training estimates trial participation, arm-specific outcome regressions, and outcome-observation probabilities. Entropy calibration matches the retained survey-weighted covariate moments. For arm a, the transported mean uses the cross-fitted score

$$\widehat{\mu}_{a,r}^{T} = \mathbb{P}_{T,r}^{w}\widehat{m}_{a}(X) + \mathbb{P}_{E,r}\left[\widehat{\omega}_{r}(X)\frac{\mathbb{1}\{A=a\}O_{Y}}{\widehat{e}_{a}\widehat{\pi}_{a}(X)}\{Y-\widehat{m}_{a}(X)\}\right],$$

where O_Y records outcome observation, $\widehat{e}_a$ is the randomized-arm probability, and $\widehat{\omega}_r$ is the calibrated trial-to-target density ratio. The basis and regularization constants are fixed before evaluation. Trial-village and target-household influence sums provide the covariance matrix. A Gaussian max-t critical value of 3.348 covers the trial, transported, and target-minus-trial versions of six package contrasts across two outcomes and three repair classes.

TABLE 5. Transport support and covariate balance

Repair class	Trial N	Target N	Target HH	Weight ESS	Max share	Max |SMD|
All target children	1,481	161	137	314.6	0.023	0.000
No younger household member (supported ages 3–4)	967	46	42	174.9	0.023	0.000
Younger household member	514	102	88	89.6	0.044	0.000

Notes: Target N counts UNPS children after the declared class-specific support rule. Target HH counts households and supplies the target clustering unit. Weight ESS is the effective trial sample size under calibrated trial-to-target weights. Max share is the largest normalized trial weight. Max |SMD| ranges over the declared age, district, sex, school-attendance, repair-class, and household-size basis after calibration. UNPS public-use files supply survey weights and omit the PSU and stratum identifiers required to reproduce the full complex-survey variance; target inference clusters co-resident children by household.

The aggregate transport weights have effective sample size 314.6, maximum normalized share 0.023, and maximum absolute standardized mean difference below 6×10^{-6}. The two supported repair-class targets have effective sample sizes 174.9 and 89.6 and maximum normalized shares 0.023 and 0.044. These diagnostics define the population range used in the estimates.

TABLE 6. Childcare-factor effects in the trial and supported target

Outcome	Repair class	Trial	Nine-district target	Difference	Target share
Full-day childcare use, short term	All target children	0.442 (0.024)	0.429 (0.052)	-0.014 (0.046)	1.000
Full-day childcare use, short term	No younger household member (supported ages 3–4)	0.457 (0.030)	0.542 (0.065)	0.084 (0.059)	0.371
Full-day childcare use, short term	Younger household member	0.415 (0.041)	0.220 (0.095)	-0.195 (0.088)	0.549
Full-day childcare use, long term	All target children	0.313 (0.024)	0.204 (0.052)	-0.109 (0.046)	1.000
Full-day childcare use, long term	No younger household member (supported ages 3–4)	0.341 (0.029)	0.219 (0.078)	-0.122 (0.072)	0.371
Full-day childcare use, long term	Younger household member	0.262 (0.042)	0.147 (0.074)	-0.115 (0.067)	0.549

Notes: Each entry is a childcare-factor effect, averaging the childcare contrast across randomized cash states; cluster standard errors appear in parentheses. Trial effects average over the common-variable trial sample. Nine-district target effects use cross-fitted augmented inverse-odds transport and UNPS survey weights. Difference is target minus trial. Target share divides the class-specific supported survey weight by the supported aggregate target weight. The two reported shares sum below one because the first target omits its unsupported aligned-age-five cell.

The short-term childcare-factor effect is 0.442 with standard error 0.024 in the trial and 0.429 with standard error 0.052 in the supported nine-district target. The 95 percent simultaneous target interval is [0.254, 0.603]. At long term, the trial effect is 0.313 with standard error 0.024 and the target effect is 0.204 with standard error 0.052; the simultaneous target interval is [0.031, 0.378]. The target-minus-trial differences are −0.014 at short term and −0.109 at long term, with simultaneous intervals of $[-0.168, 0.141]$ and $[-0.264, 0.046]$.

Within the supported $R = 0$ target, the short- and long-term effects are 0.542 and 0.219. The $R = 1$ target gives 0.220 and 0.147. The simultaneous interval excludes zero for the short-term $R = 0$ effect and contains zero for the other three class-specific target effects. Outcome-regression, selection-weighted Hájek, and augmented estimates agree within 0.025 for each aggregate target. Class-specific dispersion, outcome response rates, worst-case attrition bounds, and the balance vector accompany these estimates. These transported outcomes measure target-child full-day use. Maternal labor outcomes remain tied to the experimental population because the UNPS mother–child link is incomplete.

5.10. Outside choice and latent constraints

Observed use of the outside action has two sources in the model. The active set may collapse, and the agent may select the outside action from a menu containing active alternatives. Their probability decomposition is

$$\mathbb{P}(a_i = 0 \mid x) = \mathbb{P}(\mathcal{F}_i = \varnothing \mid x) + \mathbb{P}\left(\mathcal{F}_i \neq \varnothing,\ u_i(0) \geq \max_{a \in \mathcal{F}_i} u_i(a) \mid x\right).$$

Choice data combine these events. Preference shifters, availability shifters, and auxiliary constraint measurements supply separation conditions in revealed-consideration and latent-choice-constraint models (Masatlioglu et al. 2012; Cattaneo et al. 2020; Abaluck and Adams-Prassl 2021; Barseghyan et al. 2021; Agarwal and Somaini 2026). The job application adds measured care, schedule, commute, qualification, price, slot, and offer requirements. The resulting repair-set distribution records which input packages create each job cell.

6. Pre-Choice Welfare Accounting

The resource account and the behavioral surfaces supply different inputs to a welfare calculation. Payments identify who finances an input and who receives the transfer. Beneficiary value additionally depends on preferences over the resulting options, while the planner's comparison subtracts resource costs and includes effects on other beneficiaries. A larger compatibility set or higher care use alone leaves those monetary values undetermined. The following conditions state the additional structure under which constraint relaxation can be valued.

6.1. Envelope representation

Fix (q_i, ω) and consider an active optimum that solves

$$v_i(e_i, \omega) = \max_{a \in \mathcal{X}_i^+} u_i(a, \omega) \quad \text{s.t.} \quad g_i(a, q_i, \omega; e_i) \leq 0.$$

Let $\lambda_i(e_i, \omega) \geq 0$ be the multiplier vector under the Lagrangian $u_i - \lambda_i' g_i$.

PROPOSITION 5 (Pre-choice envelope). *Suppose the optimizer and multiplier are unique almost surely, the constraint qualification holds, and differentiation can pass through the expectation. If enabling input e_{ki} has zero direct utility effect and enters the constraint system, then*

$$\frac{\partial W_i(q_i; e_i)}{\partial e_{ki}} = -\mathbb{E}_\omega \left[\lambda_i(e_i, \omega)' \frac{\partial g_i(a_i^*, q_i, \omega; e_i)}{\partial e_{ki}} \right].$$

When the input relaxes every affected constraint, $\partial g_i / \partial e_{ki} \leq 0$, the derivative is weakly positive.

The expression can be positive across future states while the realized base-state multiplier equals zero. Current choice, earnings, and utility can remain fixed at the base state, while the expected value changes through states in which an affected constraint binds.

6.2. Availability dynamics

Let next period's availability state satisfy $q_i' = G_i(q_i, a_i, e_i, \xi_i')$, and let $J_i(q_i')$ be continuation value. The current problem includes $\beta\mathbb{E}[J_i(q_i')]$. Under the same envelope conditions, provider k's input has derivative

$$\frac{\partial W_i}{\partial e_{ki}} = -\mathbb{E}\left[\lambda_i' \frac{\partial g_i}{\partial e_{ki}}\right] + \beta\mathbb{E}\left[J_{iq}(q_i')\frac{\partial G_i}{\partial e_{ki}}\right].$$

The first term prices current constraint relaxation across states. The second term values persistence in care arrangements, credentials, transport access, work history, and other availability states.

6.3. Provider and planner conditions

Let provider k's captured beneficiary share be α_{ki}, and let r_{ki} denote the provider's operating return. At an interior choice, the provider condition is

$$c_{ki}'(e_{ki}) = \frac{\partial r_{ki}}{\partial e_{ki}} + \alpha_{ki}\frac{\partial W_i}{\partial e_{ki}}.$$

With unit welfare weight, the planner condition is

$$c_{ki}'(e_{ki}) = \frac{\partial r_{ki}}{\partial e_{ki}} + \frac{\partial W_i}{\partial e_{ki}} + \sum_{h \neq i} \frac{\partial W_h}{\partial e_{ki}}.$$

These equations locate the wedge in beneficiary value and cross-beneficiary effects. They also permit negative values for inputs that tighten constraints or remove active alternatives.

PROPOSITION 6 (Implementation by shadow prices). *Suppose input e_{ki} is observable and the provider and planner objectives are locally concave. At a planner solution e^*, define the marginal input payment*

$$s_{ki}^* = (1 - \alpha_{ki})\frac{\partial W_i(e^*)}{\partial e_{ki}} + \sum_{h \neq i} \frac{\partial W_h(e^*)}{\partial e_{ki}}.$$

Adding a payment schedule with derivative s_{ki}^* at e_{ki}^* to the provider payoff implements the *planner's first-order condition. In the one-period constraint model,*

$$s_{ki}^{*} = -(1-\alpha_{ki})\mathbb{E}\left[\lambda_i' \frac{\partial g_i}{\partial e_{ki}}\right] - \sum_{h\neq i}\mathbb{E}\left[\lambda_h' \frac{\partial g_h}{\partial e_{ki}}\right].$$

The dynamic formula adds the corresponding continuation terms.

The formula assigns the input payment to observed constraint channels. A public care-hour extension, a firm schedule adjustment, a household care input, and a transit improvement can carry different rates because they enter different rows of g, affect different beneficiaries, and permit different captured shares.

6.4. Option and transfer spanning

Let menu A expand by option b, and define its state-contingent gain

$$\Delta_b(\omega) = \left[u_i(b,\omega) - \max_{a\in A} u_i(a,\omega)\right]_+.$$

Let $Z = (z_1, \ldots, z_K)$ collect the state payoff profiles generated by a declared transfer system.

PROPOSITION 7 (Transfer-spanning condition). *A transfer portfolio $\theta \in \mathbb{R}^K$ reproduces the option's state-by-state gain exactly if and only if $\Delta_b \in \operatorname{col}(Z)$. In a finite state space, exact replication exists if and only if* $\operatorname{rank}(Z) = \operatorname{rank}([Z, \Delta_b])$.

The condition fixes the transfer technology and state space. Money-metric welfare continues to value changes in choice sets (Baqaee and Burstein 2023). The empirical comparison uses cash instruments observed in the program or survey and reports the spanning residual within that class.

7. Childcare Coverage and Job Compatibility

7.1. Job cells and household constraints

Define a job cell $j = (o, h, s, r, w)$ by occupation o, weekly hours h, start-time band s, commute band r, and work location or remote mode w. Let N_{jm}^0 be the cell's local employment weight frozen at a base-period year. For household i, an offered cell j is actionable under coverage regime d when

$$A_{ijmt}(d) = A_{jmt}^{\text{offer}} A_{ij}^{\text{cred}} A_{ijmt}^{\text{time}}(d) A_{ijmt}^{\text{slot}} A_{ijmt}^{\text{price}} = 1.$$

The offer indicator records a vacancy and employer acceptance. The credential indicator compares the household member's education and licenses with occupation requirements. The time indicator places outward travel, work, and return travel inside the coverage window. The slot indicator records a care place covering that interval. The price indicator places the required care expenditure inside the household resource bound.

Let $\mathcal{T}$ index coverage, commute, credential, price, offer, and slot inputs. The repair set is $D_{ijmt} = \{k \in \mathcal{T} : A^k_{ijmt} = 0\}$. Package S activates cell j when it supplies every component in D_{ijmt}. The empirical distribution weights these household–cell pairs by N^0_{jm}. Coverage-duration increments are separated into midday and shoulder components when policy arms permit that distinction.

The household's actionable-job mass and activation gain are

$$M_{imt}(d) = \sum_{j \in \mathcal{J}_m} N^0_{jm} A_{ijmt}(d), \qquad \Delta M_{im} = M_{im}(1) - M_{im}(0).$$

The ACS construction observes employment composition, credentials, departure, hours, commute, and remote work. Vacancies, employer acceptance, care slots, and prices lie outside the constructed indicator. Define $\widetilde{A}_{ijmt}(d) = A^{\text{cred}}_{ij} A^{\text{time}}_{ijmt}(d)$ and $\widetilde{M}_{imt}(d) = \sum_j N^0_{jm} \widetilde{A}_{ijmt}(d)$. The reported estimand is model-assisted job-compatibility mass $\widetilde{M}$. The full actionable-set indicator satisfies

$$A_{ijmt}(d) = \widetilde{A}_{ijmt}(d) A^{\text{offer}}_{ijmt} A^{\text{slot}}_{ijmt} A^{\text{price}}_{ijmt}.$$

Minimum repair solves

$$C_{imt} = \min_{j \in \mathcal{J}_m,\, e \geq 0} p'_{imt} e \quad \text{s.t.} \quad A_{ijmt}(e) = 1.$$

For linearized cell constraints, the minimum-enablement program supplies the dual certificate and provider assignment. For discrete credentials and remote-work status, the calculation enumerates admissible cells and solves the continuous repair inside each cell.

7.2. Public data

TABLE 7. Measured objects and public sources

Object	Source	Construction
Target households	2024 ACS PUMS	Children, education, person weights, employment, chosen occupation, hours, departure time, commute, work location, and remote work
Reference job cells	2024 ACS PUMS	Occupation weights, work intervals, commute, remote status, and schooling thresholds among workers outside the target group
Held-out validation	2024 SIPP	Reported workdays, start and end bins, work-at-home days, childcare interference, and person weights for calendar year 2023
Randomized incomplete repair	Uganda childcare experiment	Childcare and cash assignment, target-child enrollment and program duration, younger-sibling stratum, and labor and business outcomes for 1,496 households
Population target	Uganda National Panel Survey 2019–2020	Survey-weighted children in the nine experimental districts, household rosters, child age and sex, school attendance, and household size
Policy-delivery gate	2005 ACS and 2001–2022 CPS	Prepolicy occupation–hours cells, full-day kindergarten enrollment, maternal labor outcomes, event time, and interstate migration

ACS PUMS supplies target-household characteristics and the reference-worker distribution used to construct local job cells. The reference sample excludes target mothers. Public SIPP job records supply reported workdays, start and end bins, work-at-home days, childcare arrangements, and care interference in a survey held outside the ACS construction. The Uganda replication archive supplies randomized package assignment, the baseline younger-sibling repair class, target-child use, and economic outcomes. The Uganda National Panel Survey supplies the distinct population target used for transport.[9] Source

[9]The survey requires user registration, so replication code records the expected source hash and reads a user-acquired archive. Row-level survey records are excluded from the release package.

archives, variable dictionaries, sample rules, and file hashes accompany the inputs.

7.3. National construction benchmark

The coverage comparison holds the reference job-cell distribution fixed and changes the hours supplied by the care package. An employer's response in vacancies or shifts, a provider's response in slots or prices, and a household's substitution across care arrangements would alter other parts of the option-production problem. The fixed-reference comparison isolates the schedule requirements that a proposed extension would cover.

The implementation uses the 2024 ACS one-year national PUMS person file. The target sample contains 112,533 women ages 25–54 with own children under age six, representing 12,599,726 people under person weights. Reference workers outside this target sample form state cells by two-digit occupation, weekly-hours band, departure-time band, commute band, and remote status. A cell's credential requirement is the weighted tenth percentile of schooling among its workers. PUMA occupation weights distribute each state's occupation-specific schedule and commute cells across local markets. PUMAs with at least 150 reference workers and 10 occupation groups supply local composition; the state distribution supplies the remaining markets. This rule assigns PUMA composition to 86.7 percent of the weighted target sample.[10]

Figure 1 reports three care windows. The 8 a.m.–noon window admits 0.83 percent of model-assisted local employment mass, and the 8 a.m.–3 p.m. window admits 2.36 percent. Extending coverage to 7:30 a.m.–5:30 p.m. raises compatibility mass to 23.74 percent. The full-day expansion adds 1.53 percentage points, while the subsequent early-and-late coverage adds 21.38 percentage points. Each window uses the same occupation, schedule, commute, credential, and geographic composition cells.

[10] Usual weekly hours divided by five gives daily duration, and remote cells receive an 8 a.m. start. The schedule-sensitivity bounds vary these assignments.

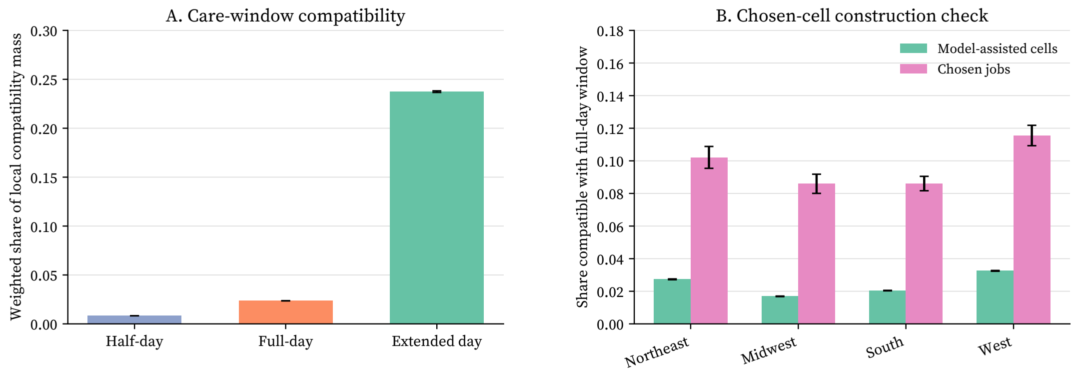

FIGURE 1. Care windows and job-compatibility mass

Notes: Panel A reports person-weighted job-compatibility mass for women ages 25–54 with own children under age six in the 2024 ACS national benchmark. Panel B compares model-assisted local mass with the share of employed mothers' chosen jobs whose departure-to-return interval fits the 8 a.m.–3 p.m. window, by census region and nationally. Reference workers determine state schedule cells, PUMA occupation weights, and credential thresholds. Error bars extend 1.96 reported design-based standard errors above and below each estimate. Sensitivity checks vary the PUMA pooling rule.

Figure 2 evaluates 26 windows. With an 8 a.m. opening, compatible mass reaches 12.05 percent at 4 p.m., 18.34 percent at 5 p.m., 19.55 percent at 5:30 p.m., and 21.65 percent at 6 p.m. Moving the opening to 7:30 a.m. raises the corresponding masses to 12.88, 21.33, 23.74, and 26.02 percent. The 3:30–4 p.m. increment contributes 9.65 percentage points under the 7:30 opening, and the 4:30–5 p.m. increment contributes 6.93 points. These increments locate the return-time mass attached to each added half hour under the schedule-cell construction.

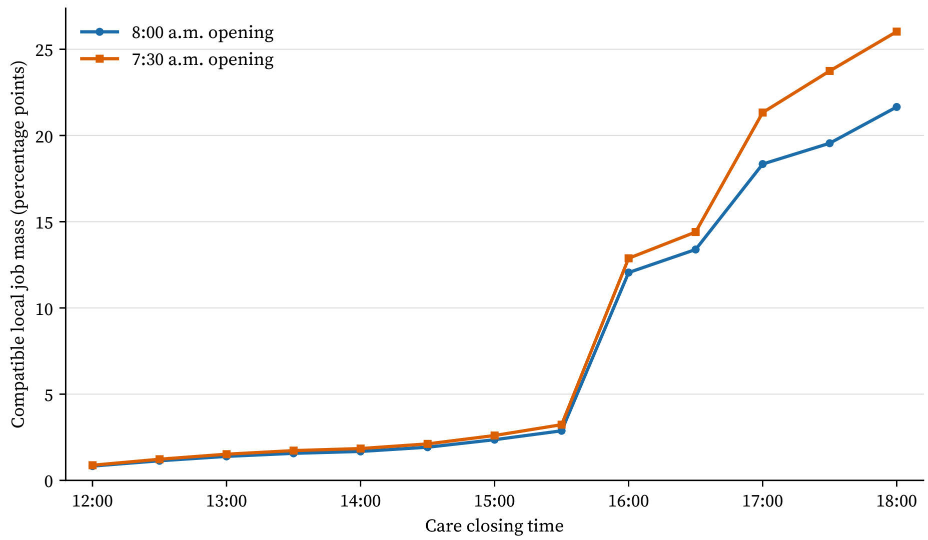

FIGURE 2. Coverage frontier over opening and closing times

Notes: Each curve reports household-weighted model-assisted local job-compatibility mass over closing times from noon through 6 p.m. The two curves hold the opening at 8 a.m. and 7:30 a.m. Reference-worker cells, credential thresholds, PUMA occupation weights, and the target sample remain fixed across windows. Shaded bands show 95 percent intervals from target-household variation conditional on the cell construction.

Table 8 applies Theorem 3 to the nested windows. Let d supply noon-to-3 p.m. coverage and let w supply the early and late shoulders. The observed arms are $\varnothing$, $\{d\}$, and $\{d, w\}$. They identify baseline mass and the midday repair mass. Shoulder-only and joint repair masses sum to 21.38 percent, and each component spans the corresponding nonnegative interval. The schedule measure classifies cells by departure and return: 0.05 percentage points require a shoulder alone and 21.33 require both components.

TABLE 8. Repair mass under nested coverage packages

Exact repair set	Package lower	Package upper	Schedule measure
Already compatible, $\varnothing$	0.83	0.83	0.83
Midday extension, $\{d\}$	1.53	1.53	1.53
Shoulder extension, $\{w\}$	0.00	21.38	0.05
Midday and shoulder, $\{d, w\}$	0.00	21.38	21.33
Shoulder or joint repair	21.38	21.38	21.38

Notes: Entries are percentage points of model-assisted local job mass. The observed package family is $\{\varnothing, \{d\}, \{d, w\}\}$, where d extends noon coverage to 3 p.m. and w adds 7:30–8 a.m. and 3–5:30 p.m. shoulders. Package bounds use arm means alone. The schedule measure assigns state-cell departure and return times to exact repair sets. The two component rows obey $\theta_w + \theta_{dw} = 21.38$. Residual mass outside the measured duration package is 76.26 percent.

The schedule cells also reveal the exact time blocks required by each job interval. Let e denote 7:30–8 a.m., d denote noon–3 p.m., and l denote 3–5:30 p.m. Table 9 reports the full package surface, its Möbius coefficients, and simultaneous intervals. Baseline-compatible cells account for 0.83 percentage points. Among the 22.91 points repaired by added coverage, d alone accounts for 1.53, e alone for 0.05, $\{e, d\}$ for 0.19, $\{d, l\}$ for 17.19, and $\{e, d, l\}$ for 3.95. Repair sets containing at least two blocks account for 93.1 percent of newly compatible mass. All eight coefficients and all 27 disjoint mixed increments satisfy the package-surface sign restrictions. The 95 percent simultaneous interval is [23.55, 23.93] for full-package mass and [3.87, 4.03] for the all-three-block repair mass.

TABLE 9. Joint inference for the timing-block package surface

Package	Active mass (pp)	Exact repair mass (pp)
$\varnothing$	0.83 [0.79, 0.87]	0.83 [0.79, 0.87]
e	0.88 [0.84, 0.92]	0.05 [0.04, 0.06]
d	2.36 [2.29, 2.43]	1.53 [1.48, 1.59]
l	0.83 [0.79, 0.87]	0.00 [0.00, 0.00]
ed	2.60 [2.53, 2.68]	0.19 [0.18, 0.21]
el	0.88 [0.84, 0.92]	0.00 [0.00, 0.00]
dl	19.55 [19.38, 19.72]	17.19 [17.03, 17.35]
edl	23.74 [23.55, 23.93]	3.95 [3.87, 4.03]

Notes: Brackets give 95 percent simultaneous intervals using a Gaussian max-t critical value of 2.632. The 80 official ACS replicate weights reweight reference job cells, PUMA occupation composition, and target households together; cell definitions are frozen before reweighting. Job mass is measured in percentage points of model-assisted local employment mass. Continuous work intervals imply zero mass for $\{e, l\}$ without d.

Table 10 gives the four Census regions and the national aggregate. Full-day compatibility mass ranges from 1.7 percent in the Midwest to 3.3 percent in the West. Chosen jobs fit that window at rates from 8.6 percent in the Midwest and South to 11.6 percent in the West. Nationally, the chosen-cell rate exceeds household-specific local compatibility mass by 7.1 percentage points. Employment selection contributes to this gap, so the comparison serves as a construction check. The held-out SIPP benchmark supplies the validation step for work intervals and childcare friction.

The held-out 2024 SIPP file covers calendar year 2023 and contains 1,280 target reference mothers, representing 12.9 million people.[11] Among 780 mothers with reported job schedules, the lower and upper work-interval shares for the 8 a.m.–3 p.m. window are 9.3

[11]The held-out survey is used as a measurement diagnostic for work intervals and childcare interference; it is separate from the ACS cell construction.

and 13.4 percent. The ACS chosen-job rate of 9.5 percent lies in this interval. Weekly hours divided by reported days differ from the midpoint duration between SIPP start and end bins by 0.99 hours on average under person weights. Childcare-limited work is reported by 6.1 percent of the target sample, with rates of 10.7 percent in the lower-compatible group and 6.4 percent in the group whose reported bins lie outside the window. Employment and care selection enter these group means, so they remain measurement diagnostics.

TABLE 10. National job-compatibility construction benchmark

	Half-day mass	Full-day mass	Expansion gain	Extended mass	Chosen compatible	Chosen–local gap
Northeast	0.009	0.027	0.018	0.251	0.102	0.074
Midwest	0.005	0.017	0.011	0.221	0.086	0.068
South	0.008	0.020	0.013	0.229	0.086	0.065
West	0.012	0.033	0.021	0.257	0.116	0.082
United States	0.008	0.024	0.015	0.237	0.095	0.071

Notes: Mass columns divide weighted compatible cells by total model-assisted local employment weight. Activation gain is the change from 8 a.m.–noon to 8 a.m.–3 p.m. Chosen compatible is the weighted share of employed target mothers whose observed job fits the full-day window. Chosen–local subtracts each mother's local compatibility mass from the chosen-job indicator.

7.4. Interpretation and scope

The frontier measures the schedule and credential portion of job compatibility. Its denominator is local reference-job mass, so each point on the curve gives the share of that mass fitting a declared care window for a household with a given schooling level. The chosen-job comparison uses a different denominator and reflects employment selection. The held-out SIPP interval evaluates the work-interval construction among employed mothers with reported schedules. Every table and figure keeps these three sample definitions separate.

The repair distribution assigns coverage blocks to job intervals. It records the package composition required to fit a job-travel interval inside a care window. Employer acceptance, current vacancies, care-slot availability, and price eligibility enter the full actionable-set measure as additional components. The 93.1 percent multiple-block share therefore describes timing complementarity inside the measured coverage technology.

Public-school duration reforms and childcare experiments provide settings in which the repair map can be joined to delivered coverage and realized job transitions (Mbiti et al. 2019; Berthelon et al. 2024; Gibbs et al. 2026b,a; Simintzi et al. 2026). The implementation covers 12 policy cohorts adopted from 2007 through 2017 and 14 states without a policy through 2022. A 2005 ACS map fixes state, two-digit occupation, and weekly-hours cells before every included adoption. Full-day-friendly policies raise the CPS full-day enrollment share by 0.226, with a 95 percent simultaneous interval of [0.107, 0.344]. The corresponding intervals are [-0.014, 0.020] for maternal employment and [-0.022, 0.023] for employment in the frozen predicted gain set. The registered event-time lead bands fall outside their equivalence criteria. The extension is therefore reported as a design assessment, while the paper's causal transition claims rest on the randomized and cross-sectional designs.

8. Conclusion

Economic choice begins from a set of actions produced by complementary inputs. Minimal input paths describe how each agent–option pair becomes actionable. Package surfaces can then reveal three objects: the mass of options created by each repair path, the resource cost and provider value of supplying that path, and movement into the option after it becomes available. Monotone package surfaces admit mixtures of minimal-path antichains. Nonnegative Möbius coefficients recover the exact repair distribution in the single-path case, and observed package arms yield sharp bounds when support is incomplete. Minimum-enablement duality attaches input prices and provider assignments to the same repair paths.

Childcare hours illustrate the distinction between a service and the options it creates. In the 2024 ACS, model-assisted job-compatibility mass rises from 0.83 percent under an 8 a.m.–noon window to 23.74 percent under a 7:30 a.m.–5:30 p.m. package. Repair sets containing at least two timing blocks account for 93.1 percent of the additional compatible mass. A randomized childcare–cash factorial supplies two routes into full-day care and places short-term response mass attributable to either route between 3.17 and 5.15 percentage points. The childcare-factor effect on full-day use is 0.442 in the trial and 0.429 in a supported nine-district target at short term; the corresponding long-term effects are 0.313 and 0.204. The ACS evidence measures schedule compatibility and its repair composition. The randomized factorial measures assigned supply paths, delivered care, and subsequent use, while the transport design states the population and support over which those use effects extend.

These objects organize policy evaluation around the production of choice sets. Direct active-set measurement records which options a package creates. Randomized package outcomes record which produced options enter behavior. Provider accounts value the inputs used in production. Employer offers, vacancies, care slots, and prices can extend schedule compatibility to actionable employment, while targeted transition designs can connect repaired job cells to realized choice. The resulting accounting follows a policy package from supplied inputs to produced options and from produced options to behavior.

Data and code availability. Replication files rebuild the ACS coverage frontier, timing-block repair masses, joint reference–target intervals, package bounds, held-out SIPP benchmark, Uganda randomized path surfaces, childcare-use first stages, repair-targeting estimates, provider accounts, and the public-school duration analysis, with source hashes, processed outputs, and analysis summaries. The Uganda National Panel Survey requires registration and its source terms prohibit redistribution; the transport code reads a user-acquired archive whose hash is recorded in the source manifest and excludes the archive and row-level survey data from release packages.

References

Jason Abaluck and Abi Adams-Prassl. What do consumers consider before they choose? identification from asymmetric demand responses. *Quarterly Journal of Economics*, 136(3): 1611–1663, 2021.

Nikhil Agarwal and Paulo Somaini. Demand analysis under latent choice constraints. *Review of Economic Studies*, 93(4):2215–2249, 2026. doi: 10.1093/restud/rdaf093.

David R. Baqaee and Ariel Burstein. Welfare and output with income effects and taste shocks. *Quarterly Journal of Economics*, 138(2):769–834, 2023.

Levon Barseghyan, Francesca Molinari, and Matthew Thirkettle. Discrete choice under risk with limited consideration. *American Economic Review*, 111(6):1972–2006, 2021.

Amitabh Basu, Tongtong Chen, Michele Conforti, and Hongyi Jiang. Helly systems and certificates in optimization, 2021.

Matias Berthelon, Diana Kruger, Catalina Lauer, Luca Tiberti, and Carlos Zamora. Longer school schedules, childcare and the quality of mothers' employment. *Economic Policy*, 39(120): 817–846, 2024.

Kjetil Bjorvatn, Denise Ferris, Selim Gulesci, Arne Nasgowitz, Vincent Somville, and Lore Vandewalle. Childcare, labor supply, and business development: Experimental evidence from uganda. *American Economic Journal: Applied Economics*, 17(2):75–101, 2025.

Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, 86(3):1061–1094, 2019.

Matias D. Cattaneo, Xinwei Ma, Yusufcan Masatlioglu, and Elchin Suleymanov. A random attention model. *Journal of Political Economy*, 128(7):2796–2836, 2020.

Alain Chateauneuf and Jean-Yves Jaffray. Some characterizations of lower probabilities and other monotone capacities through the use of möbius inversion. *Mathematical Social Sciences*, 17(3): 263–283, 1989. doi: 10.1016/0165-4896(89)90056-5.

Alfred Galichon and Marc Henry. Set identification in models with multiple equilibria. *Review of Economic Studies*, 78(4):1264–1298, 2011. doi: 10.1093/restud/rdr008.

Chloe R. Gibbs, Esra Kose, and Maria Rosales-Rueda. More hours, more work: Head start expansions boost maternal employment. Working Paper 34831, National Bureau of Economic Research, 2026a.

Chloe R. Gibbs, Jocelyn Wikle, and Riley Wilson. A matter of time? measuring effects of public schooling expansions on families. *American Economic Journal: Economic Policy*, 2026b. Forthcoming.

David M. Kreps. A representation theorem for preference for flexibility. *Econometrica*, 47(3): 565–577, 1979.

Yusufcan Masatlioglu, Daisuke Nakajima, and Erkut Y. Ozbay. Revealed attention. *American Economic Review*, 102(5):2183–2205, 2012.

Isaac Mbiti, Karthik Muralidharan, Mauricio Romero, Youdi Schipper, Constantine Manda, and Rakesh Rajani. Inputs, incentives, and complementarities in education: Experimental evidence from tanzania. *Quarterly Journal of Economics*, 134(3):1627–1673, 2019. doi: 10.1093/qje/qjz010.

Karthik Muralidharan, Mauricio Romero, and Kaspar Wüthrich. Factorial designs, model selection, and (incorrect) inference in randomized experiments. *Review of Economics and Statistics*, 107(3):589–604, 2025. doi: 10.1162/rest_a_01317.

Prasanta K. Pattanaik and Yongsheng Xu. On ranking opportunity sets in terms of freedom of choice. *Recherches Economiques de Louvain*, 56(3–4):383–390, 1990.

Gian-Carlo Rota. On the foundations of combinatorial theory i: Theory of möbius functions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2:340–368, 1964. doi: 10.1007/BF00531932.

Amartya Sen. Markets and freedoms: Achievements and limitations of the market mechanism in promoting individual freedoms. *Oxford Economic Papers*, 45(4):519–541, 1993.

Elena Simintzi, Sheng-Jun Xu, and Ting Xu. Childcare access and the allocation of female human capital across firms. Working Paper 33835, National Bureau of Economic Research, 2026.

Uganda Bureau of Statistics. Uganda national panel survey 2019–2020, wave 7. World Bank Microdata Library, 2019. URL `https://microdata.worldbank.org/catalog/3902`.

Tyler J. VanderWeele and James M. Robins. Empirical and counterfactual conditions for sufficient cause interactions. *Biometrika*, 95(1):49–61, 2008. doi: 10.1093/biomet/asm090.