Perfectly Still

Stability and Policy Response Capacity

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

Earlier interventions can change the response to a later policy. This paper studies that response as a state-dependent map from current instruments to future targets. Its evolution reflects entry, exit, and changes among units already reached. Adoption and effective-cap models provide mechanisms through which policy history changes the map. Sequential randomization and factorial assignment identify distinct history and configuration contrasts under their stated designs. In the myBPmyLife trial, the fitted current diet-notification response is 1.75 percentage points lower at the 90th than at the 10th percentile of fourteen-day assigned history, conditional on the recorded pre-state and time controls. A separate factorial experiment shows how an app-limit assignment changes the response to a temporary use-reduction bonus. Rank and finite-action calculations describe the directions available for the targets and horizons represented in each experiment.

JEL Codes: C23, C61, D12, I15, I18

Keywords: policy response capacity, randomized policy history, response matrix, finite policy actions, factorial experiments

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

1. Introduction

What does an intervention leave behind for the next policy decision? A subsidy can change experience with a product, repeated notifications can change engagement, and a usage limit can alter the response to a later incentive. The earlier policy changes the state through which the next one operates. Evaluating this legacy requires a target that measures the response to the later intervention.

The paper calls that map *policy response capacity*.[1] For continuous instruments it is a finite-horizon derivative from the current policy vector to target outcomes. For discrete interventions it is the corresponding finite contrast. Its columns describe the available target responses under stated instrument units.[2] The map can change even when the baseline outcome or participation rate returns to an earlier level.

The analysis starts with accounting and the experiment needed to measure a response change. A transition identity allocates aggregate capacity changes to entrants, exits, and continuing units. Sequential and factorial assignment then identify particular history or configuration contrasts. The stillness theorem shows that repeated reach can drive the measured response to zero while a target gap remains strictly positive. The result permits heterogeneous conversion rates and requires only finite initial reachable mass with positive rates. Stable baseline motion therefore leaves the later response map undetermined. Once the contrasts are defined, an effective-cap model and a two-date adoption model give mechanisms that can generate them. Instrument-menu and option-value extensions study the feasible responses available to a planner.

The empirical target varies two aspects of policy. In a sequential design, an earlier assigned history changes the state confronting a current randomized intervention. In a factorial design, one assigned configuration changes the contrast generated by another intervention. Following one treatment over several horizons describes a different object: the outcome path of that treatment.[3] The response-capacity matrices here are assembled from the stated two-policy contrasts, with their outcomes, horizons, normalizations, and support retained in rank inference.

The first application uses the public myBPmyLife micro-randomized trial (Golbus et al. 2026). Across 186,893 active-phase decision points, the fitted current diet-notification response is 1.75 percentage points lower at the 90th than at the 10th percentile of fourteen-day assigned history, with participant-clustered standard error 0.60.[4] This comparison

[1] Public-administration research uses policy capacity for the analytical, operational, and political competences of policy actors (Wu et al. 2015). Here the term denotes a target-response map under stated policy instruments.

[2] Rescaling an instrument rescales its column. Every rank and magnitude comparison therefore retains the instrument units, outcome normalization, and horizon.

[3] The sequential contrast conditions on the recorded pre-state. A regime contrast that changes the pre-state distribution uses an additional averaging step.

[4] The two history percentiles correspond to 10 and 18 successful assignments over the preceding 56 decision

holds the recorded pre-state and response-time controls fixed. It describes heterogeneity in the current randomized response across prior histories; a total effect of changing the earlier assignment regime also includes changes in the distribution of that pre-state. For four exploratory engagement targets and two notification types, the leading singular values of standardized seven- and fourteen-day response-change operators are 0.142 and 0.136. Simultaneous bands reject rank zero at these windows; a lower-rank-uniform test at fourteen days retains rank at most one. Endpoint classification and screening families follow the stated source design.

The second application uses the factorial experiment of Allcott et al. (2022). A temporary bonus reduces digital use by 62.9 minutes per day in the limit-control group and by 51.5 minutes under the limit assignment. Their interaction measures a change in the bonus response across randomized configurations.[5] A finite planner evaluates the assigned combinations at attainable reduction targets and implementation costs. This calculation keeps the experimental action set explicit.

Together the models and applications distinguish a system's baseline motion from the response directions available to policy. The evidence concerns the reported randomized designs and horizons. The accounting identities explain how responses can evolve, while the identification conditions determine which part of that evolution a particular experiment recovers.

This perspective connects dynamic treatment-response research with the targets-and-instruments tradition. Event-study and sequential-treatment methods organize how effects vary across histories and horizons; micro-randomized trials supply sequential experimental variation; factorial designs identify treatment interactions; and dynamic structural models emphasize the state variables through which current choices change future opportunities. Dynamic regression-discontinuity methods recover long-term direct effects from repeated eligibility events under stated mean-equivalence, selection, smoothness, and overlap conditions (Hsu and Shen 2024). The target here is the later-policy response map and its evolution across randomized histories and configurations. The contribution is to treat that map as an evolving state object and to ask which response directions are available to a later policy decision.

The targets-and-instruments tradition motivates the map from available interventions to objectives (Tinbergen 1952), while multi-instrument field experiments and factorial designs study the directions spanned by treatment combinations (Hahn et al. 2016; Muralidharan et al. 2025; Pless 2026). State-dependent multiplier research emphasizes that policy responses vary with slack, credit conditions, balance sheets, and histories (Auerbach and Gorodnichenko 2012; Ramey and Zubairy 2018; Eichenbaum et al. 2022). Work on

points.

[5] Limit assignment bundles access, enforcement, chosen tightness, and snoozing. The estimand concerns that randomized bundle.

take-up and experience goods supplies mechanisms through which earlier policy changes later participation and response (Bhargava and Manoli 2015; Finkelstein and Notowidigdo 2019; Deshpande and Li 2019; Bergemann and Välimäki 2006). Dynamic controllability provides the corresponding multi-date systems perspective (Buiter and Gersovitz 1981).

The paper develops response capacity, identification, endogenous changes, and the randomized-history and configuration applications in that order. Additional proofs, diagnostics, and benchmark calculations appear in the online appendix.

2. Policy History and Response Capacity

Let the full state be $s_t = (x_t, r_t)$, where x_t contains the conventional economic state and r_t contains the units and response parameters through which policy operates. For an admissible policy sequence $U = (u_0, \ldots, u_{H-1}) \in \mathcal{U}^H$, let $\Phi_H(s, U)$ denote the terminal state and let τ select a q-dimensional target. The endpoint map is

$$\Psi_H(s, U) = \tau(\Phi_H(s, U))\,. \tag{1}$$

For continuous instruments, the response matrix and response space are

$$M_H(s) = D_U \Psi_H(s, 0), \qquad \mathcal{V}_H(s) = \operatorname{Im} M_H(s). \tag{2}$$

The matrix has q rows and mH columns when there are m instruments per date. Its rank records the number of locally reached target directions. Under stated target and instrument units, its singular values record response strength along orthogonal directions.

For a discrete policy $b \in \{0, 1\}$, define the response-capacity contrast

$$K_H(s) = \mathbb{E}[\Psi_H(s, b = 1) - \Psi_H(s, b = 0) \mid s]. \tag{3}$$

Equation (3) supplies the experimental counterpart of a derivative. Continuous and discrete instruments can be stacked after their units are reported.

For a finite assignment set $\mathcal{A}$, let $k(a) \in \mathbb{R}^q$ be the target change generated by action a relative to a baseline action, and let $c(a)$ be its activation cost. The attainable response set is $\mathcal{K} = \{k(a) : a \in \mathcal{A}\}$. A planner facing target change z chooses

$$L_{\mathcal{A}}(z) = \min_{a \in \mathcal{A}} \left\{ \frac{1}{2}(z - k(a))' Q(z - k(a)) + c(a) \right\}. \tag{4}$$

PROPOSITION 1 (Finite-action policy choice). *Suppose Q is positive definite. Equation* (4) *compares the target losses of the attainable actions. If the planner can randomize assignment shares and target changes aggregate linearly, the attainable set becomes* $\operatorname{conv}(\mathcal{K})$. *With zero activation costs, replacing either set by* $\operatorname{span}(\mathcal{K})$ *gives a weakly lower loss. Equality holds exactly* *when a closest point in the span belongs to the relevant attainable set.*

The finite set applies to mutually exclusive experimental arms. Its convex hull applies to population assignment shares. The span applies to continuously scalable instruments whose admissible coefficients include the projection used by the targeting problem.

DEFINITION 1 (Perfect stillness). *A baseline fixed point $s^\star$ is perfectly still for a stated target, horizon, and policy class when it is locally attracting under the baseline law and every response derivative and discrete contrast in that class equals zero at $s^\star$.*

2.1. Units and exact transition accounting

Consider N units with weights $\omega_i > 0$. Let $I_{it} \in \{0, 1\}$ indicate attachment to the policy system and let $B_{it} \in \mathbb{R}^{q\times m}$ be unit i's target response conditional on attachment. The aggregate response matrix is

$$M_t = \sum_{i=1}^{N} \omega_i I_{it} B_{it}. \tag{5}$$

Attachment can mean enrollment, an active account, a claim with remaining entitlement, a refinancing incentive, or availability to an offer. The conditional matrix can include direct and equilibrium responses under a stated exposure mapping.

PROPOSITION 2 (Transition accounting). *The change in equation* (5) *satisfies*

$$M_{t+1} - M_t = E_t - X_t + S_t, \tag{6}$$

$$E_t = \sum_i \omega_i (1 - I_{it}) I_{i,t+1} B_{i,t+1}, \tag{7}$$

$$X_t = \sum_i \omega_i I_{it} (1 - I_{i,t+1}) B_{it}, \tag{8}$$

$$S_t = \sum_i \omega_i I_{it} I_{i,t+1} (B_{i,t+1} - B_{it}). \tag{9}$$

The three terms collect entrants, exits, and response changes among units attached at both dates.

The accounting identity uses the joint transition $(I_{it}, I_{i,t+1})$. When attachment is latent and only marginal probabilities $p_{it} = \Pr(I_{it} = 1)$ are available, a gross-flow decomposition additionally requires the transition probabilities $\pi_{it}^{ab} = \Pr(I_{it} = a, I_{i,t+1} = b)$ and response matrices indexed by transition type. The expected terms then replace the indicators in equation (6) with π^{01}, π^{10}, and π^{11}. Marginal probabilities determine $p_{i,t+1} - p_{it} = \pi_{it}^{01} - \pi_{it}^{10}$ and leave the two gross flows separate only after π^{01} or π^{10} is supplied.

Let an earlier policy a_t change next-period attachment and conditional response. For the probability representation

$$M_{t+1}(a_t) = \sum_i \omega_i p_{i,t+1}(a_t) B_{i,t+1}(a_t), \tag{10}$$

differentiation gives

$$D_{a_t} \operatorname{vec} M_{t+1} = \sum_i \omega_i \Big[\operatorname{vec}(B_{i,t+1}) D_{a_t} p_{i,t+1} + p_{i,t+1} D_{a_t} \operatorname{vec}(B_{i,t+1}) \Big]. \tag{11}$$

The first term operates through attachment. The second operates through response conditional on attachment. Panel records identify the transition terms directly; repeated cross sections require a model for gross flows and composition.

Additional instruments and their value. A dictionary-constrained response-addition problem uses a specified target loss, feasible instrument menu, and activation costs. The empirical applications keep the experimentally assigned action sets distinct from the linear span used in that calculation.

Baseline motion and policy responses. Stability of the untreated state and the rank of its policy-response map concern different parts of the transition law. A shared-baseline counterexample separates them, while additional conditions connect state convergence to declining response strength. Economic mechanisms can then move the response map. Sequential and factorial designs identify different parts of this change.

3. Identification and Rank Inference

The empirical object involves two interventions. An earlier policy A_t changes attachment, experience, or remaining demand. A later policy B_{t+1} acts through the resulting response state. Let H_t contain pretreatment histories, let L_{t+1} collect variables observed between interventions, and let $Y_{t+1+h}(a, b)$ be the potential target at horizon h.

3.1. Continuous sequential policies

For scalar policies, define the later response and the capacity effect by

$$m_h(a, h_t) = \frac{\partial}{\partial b} \mathbb{E}[Y_{t+1+h}(a, b) \mid H_t = h_t] \Big|_{b=0}, \tag{12}$$

$$\theta_h(a, h_t) = \frac{\partial}{\partial a} m_h(a, h_t). \tag{13}$$

The second derivative θ_h measures how the earlier policy changes the response to the later policy.

ASSUMPTION 1 (Sequential assignment and smooth support). *Conditional on H_t, A_t is independent of the potential intermediate variables and outcomes indexed by (a, b). Conditional on (H_t, A_t, L_{t+1}), B_{t+1} is independent of the potential outcomes indexed by b. The treatment values used in each comparison have positive conditional density. The outcome regression and the conditional distribution of L_{t+1} are differentiable in the relevant treatment values, and an integrable envelope permits differentiation under the integral.*

PROPOSITION 3 (Sequential identification). *Under Assumption 1 and consistency,*

$$m_h(a, h_t) = \int \frac{\partial}{\partial b}\mathbb{E}[Y_{t+1+h} \mid H_t = h_t, A_t = a, L_{t+1} = \ell, B_{t+1} = b]\Big|_{b=0} dF(\ell \mid h_t, a), \tag{14}$$

and differentiating the right side with respect to a identifies $\theta_h(a, h_t)$.

The formula requires variation in both policies over the reported support. A rollout chosen in response to backlog or shock severity places those assignment variables in H_t or calls for a design that fixes assignment. Panel records can then divide the total capacity effect using the transition terms in Proposition 2.

3.2. Sequential randomization and policy history

At decision point t, let D_{it} and A_{it} indicate mutually exclusive diet and activity notifications with known probabilities p_D and p_A. Let S_{it} be the target state recorded before assignment and define centered current assignments $\widetilde{D}_{it} = D_{it} - p_D$ and $\widetilde{A}_{it} = A_{it} - p_A$. For a window of w prior decision points, randomized policy history is

$$H_{it}^{(w)} = \sum_{\ell=1}^{w} \Big(D_{i,t-\ell} + A_{i,t-\ell} - p_D - p_A\Big). \tag{15}$$

The empirical projection is

$$Y_{it} = \alpha_i + \lambda_{d(t)} + \mu_{q(t)} + \beta_D^{(w)}\widetilde{D}_{it} + \beta_A^{(w)}\widetilde{A}_{it} + \gamma^{(w)}H_{it}^{(w)} + \theta_D^{(w)}\widetilde{D}_{it}H_{it}^{(w)} + \theta_A^{(w)}\widetilde{A}_{it}H_{it}^{(w)} \tag{16}$$

$$+ \delta^{(w)}S_{it} + \psi_D^{(w)}\widetilde{D}_{it}S_{it} + \psi_A^{(w)}\widetilde{A}_{it}S_{it} + e_{it}, \tag{17}$$

where $d(t)$ indexes relative enrollment day and $q(t)$ indexes time of day. The empirical specification also interacts both current assignments with centered linear and quadratic functions of relative enrollment day. Centering makes $\beta_D^{(w)}$ the current diet-notification response at mean randomized history, at the reference pre-state, and at the centered response-time terms. The coefficient $\theta_D^{(w)}$ gives the slope of the conditional current response with respect to assigned notification history, holding the recorded pre-state and response-time profile fixed.

PROPOSITION 4 (Randomized-history response). *Suppose current assignment is independent* *of current potential outcomes conditional on the observed decision history, the assignment probabilities equal* $(p_D, p_A, 1 - p_D - p_A)$ *on retained decision points, consistency holds, and the conditional current-treatment response is linear in* $(H_{it}^{(w)}, S_{it})$ *over the reported support. Then weighted centered least squares identifies* $(\beta_D^{(w)}, \theta_D^{(w)}, \psi_D^{(w)})$ *and the corresponding activity coefficients. With constant assignment probabilities, the weights are common and equation* (16) *supplies the estimator.*

Proof. Conditional random assignment makes the centered current-assignment vector orthogonal to every function of the observed history and gives it a positive-definite conditional second-moment matrix when all three assignment cells have positive probability. Multiplying the linear conditional response by the centered assignment vector and by its products with $(H_{it}^{(w)}, S_{it})$ yields the normal equations for the response coefficients. Consistency replaces the potential outcomes in those moments with the observed target. Constant probabilities make the weighting factor common across decision points. □

The proposition identifies moderation of the current randomized response by observed history. The pre-state S_{it} can itself be affected by earlier assignments. Consequently, conditioning on S_{it} changes the history comparison: a total response contrast between earlier assignment regimes averages over the regime-specific distribution of intermediate states, as in equation (14).[6] Such a contrast also requires a specified assignment regime and support for its histories. The empirical projection below reports conditional response moderation on the retained decision points. Its history slope has that interpretation even though the notifications entering the history were randomized.

Micro-randomized trials generate the repeated assignment variation used by equation (16) (Boruvka et al. 2018). Participant clustering preserves dependence created by overlapping histories and repeated outcomes. A cluster-multiplier maximum statistic gives simultaneous coverage across a declared collection of windows, outcomes, and current notification types. A lead-history placebo replaces equation (15) with future assignments; sequential randomization places the corresponding interaction at zero in expectation.

For q targets, let $\theta_{Dj}^{(w)}$ and $\theta_{Aj}^{(w)}$ denote the two current-assignment-by-history coefficients for target j, and let Δh_w be the difference between the 90th and 10th percentiles of realized randomized history. The window-specific response-change operator is

$$\Theta^{(w)} = \Delta h_w \begin{pmatrix} \theta_{D1}^{(w)} & \theta_{A1}^{(w)} \\ \vdots & \vdots \\ \theta_{Dq}^{(w)} & \theta_{Aq}^{(w)} \end{pmatrix}. \tag{18}$$

[6] Randomization of the earlier notifications identifies their regime-specific state distributions. Conditioning on a post-assignment state instead holds that intermediate state fixed and yields a different response contrast.

Each column records the fitted difference in the conditional current response to one notification type between the two history percentiles, across the target vector. The empirical singular values divide each row by its common-sample target standard deviation and carry the same transformation into the participant-cluster influence covariance. A Gaussian cluster multiplier supplies a joint operator-norm radius across the five window-specific matrices. Weyl's inequality then gives simultaneous lower endpoints $\max\{0, \sigma_j(\widehat{\Theta}^{(w)}) - c_{1-\alpha}\}$. The lower-rank-uniform procedure described below tests the rank of a selected window, while a maximum statistic over the five matrices protects the rank-zero test against window selection.

3.3. Factorial assignment of intervention and configuration

Let $D \in \{0, 1\}$ assign a policy configuration and let $B \in \{0, 1\}$ assign an intervention whose response depends on that configuration. For target horizon h, define

$$K_h(d) = \mathbb{E}[Y_h(d, 1) - Y_h(d, 0)], \qquad \Delta K_h = K_h(1) - K_h(0). \tag{19}$$

PROPOSITION 5 (Factorial capacity contrast). *If $(D, B) \perp \{Y_h(d, b) : d, b \in \{0, 1\}\} \mid G$, every assignment cell occurs in each included stratum, and consistency holds, then*

$$\Delta K_h = \mathbb{E}_G[\mathbb{E}(Y_h \mid D = 1, B = 1, G) - \mathbb{E}(Y_h \mid D = 1, B = 0, G)] \tag{20}$$
$$- \mathbb{E}_G[\mathbb{E}(Y_h \mid D = 0, B = 1, G) - \mathbb{E}(Y_h \mid D = 0, B = 0, G)]. \tag{21}$$

In a regression containing B, D, their interaction, and saturated stratum indicators, the interaction coefficient estimates equation (20). Stacking the coefficients across horizons supplies a capacity-shift vector. Replacing D with mutually exclusive randomized configurations yields one intervention response for each configuration.

Let $h^\star$ denote a horizon in which both instruments are active and let $\mathcal{C}$ collect comparison horizons. For weights $w_h \geq 0$ satisfying $\sum_{h \in \mathcal{C}} w_h = 1$, define the overlap capacity shift

$$\Gamma_{h^\star, \mathcal{C}} = \Delta K_{h^\star} - \sum_{h \in \mathcal{C}} w_h \Delta K_h. \tag{22}$$

PROPOSITION 6 (Factorial identification across policy timing). *Suppose the factorial assignment condition in Proposition 5 holds jointly for the outcome vector indexed by $\{h^\star\} \cup \mathcal{C}$. Then equation* (22) *is identified by the same linear contrast of the horizon-specific interaction coefficients. Its covariance is the corresponding quadratic form in their cross-horizon covariance matrix.*

Random assignment identifies each component of equation (22); linearity identifies the timing contrast. The effective-cap model assigns its mechanism when D changes the cap and B changes the unconstrained choice during $h^\star$. A weekly version can replace the period means and compare the payment weeks with the remaining weeks under the same assignment.

Let $\widehat{\theta}$ stack J horizon interactions, let $\mathcal{A} = \{a_1, \dots, a_L\}$ be a finite collection of timing contrasts, and let $\widehat{\Sigma}$ estimate the joint covariance of $\widehat{\theta}$. Define $s_j^2 = a_j'\widehat{\Sigma}a_j$ and let $c_{1-\alpha}$ be the $1-\alpha$ quantile of $\max_{j\le L} |a_j'Z|/s_j$ for $Z \sim N(0, \widehat{\Sigma})$.

PROPOSITION 7 (Finite timing family). *Under joint asymptotic normality and consistent covariance estimation, the intervals*

$$a_j'\widehat{\theta} \pm c_{1-\alpha}s_j, \qquad j = 1, \dots, L, \tag{23}$$

have asymptotic simultaneous coverage $1-\alpha$. *The interval indexed by any data-dependent selector taking values in* $\{1, \dots, L\}$ *inherits that coverage.*

The empirical timing analysis uses two finite families. The period family compares each of four periods with the equal-weight mean of the other three. The weekly family compares each of ten consecutive three-week windows with the other nine weeks. Gaussian max-t critical values use the joint HC2 covariance. A conditional randomization analysis applies the same four period contrasts to every re-randomization draw. Each draw re-estimates the interaction and its HC2 covariance from the observed outcome vector, then forms studentized statistics. This procedure targets weak null hypotheses for average factorial contrasts while allowing unit-level treatment-effect heterogeneity, following the studentization argument in Wu and Ding (2021). The payment-window comparison is reported as exploratory and receives simultaneous inference over both families. This treatment follows the inference-after-selection problem studied by Andrews et al. (2024).

3.4. Prospective external benchmarks

An external benchmark enters when its trial record predates assignment, its published estimand maps into a configuration or history contrast in this framework, and the analysis carries the source estimate with its reported uncertainty and multiplicity adjustment. The benchmark retains the source trial's outcome definition, horizon, and analysis family. It supplies evidence about the response mechanism across a separate population and implementation. Cross-trial differences in target units, intervention intensity, and follow-up remain part of the comparison.

3.5. A randomized early policy and a common later offer

The application has a discrete early assignment and a common later offer. Let $A \in \mathcal{A}$ be randomly assigned within block G, let $B = 1$ indicate receipt of the later offer, and let $Y_2(a, b)$ be later purchase. Suppose the offered product is unavailable through the comparison channel when $b = 0$, giving $Y_2(a, 0) = 0$. Define

$$K_2(a) = \mathbb{E}[Y_2(a,1) - Y_2(a,0)] = \mathbb{E}[Y_2(a,1)]. \tag{24}$$

PROPOSITION 8 (Randomized capacity effect). *If $A \perp \{Y_2(a,1) : a \in \mathcal{A}\} \mid G$ and each compared assignment occurs within every included block, then*

$$K_2(a) - K_2(a') = \mathbb{E}_G\big[\mathbb{E}(Y_2 \mid A = a, G) - \mathbb{E}(Y_2 \mid A = a', G)\big]. \tag{25}$$

The later offer can be common across households because the zero-offer outcome is fixed by product availability. Randomized variation in the earlier assignment identifies how policy history changes the response generated by that offer. A within-block permutation of early assignments supplies finite-sample inference under the corresponding sharp null.

With interference summarized by an exposure mapping $Q_i(A_{-i})$, random assignment also generates variation in neighbors' policy exposure. The reported estimand then conditions on the chosen mapping and includes own assignment and mapped exposure as separate instruments. Spatial covariance estimates account for overlapping neighborhoods in the outcome equations.

3.6. Response matrices and rank

Let $M \in \mathbb{R}^{q\times m}$ stack target responses to assigned exposures and let $k = \min\{q, m\}$. For $r < k$, the rank null is $H_0^{(r)} : \operatorname{rank}(M) \le r$. The empirical statistic is

$$T_n(r) = n \sum_{j=r+1}^{k} \sigma_j^2(\widehat{M}). \tag{26}$$

The singular values use the target and instrument scales stated for each application. The digital matrix divides each target row by its sample outcome standard deviation. Its covariance receives the same row transformation.

The lower-rank-uniform calculation follows the CFA two-step procedure in Chen and Fang (2019). Set $\kappa_n = n^{-1/4}$ and estimate a lower rank r_0 by counting singular values of $\widehat{M}$ above κ_n. A value $r_0 > r$ rejects in the first step. For $r_0 \le r$, write the singular-vector matrices as U and V, retain the columns associated with indices above r_0, and draw Z^* from the estimated law of $\sqrt{n}\,\mathrm{vec}(\widehat{M} - M)$. The bootstrap statistic applies equation (26) to the singular values of $U_0' Z^* V_0$ beginning at index $r - r_0 + 1$. The applications use 49,999 Gaussian draws from the joint coefficient covariance. The digital covariance is HC2 across participants; the bed-net covariance uses the spatial estimator in equation (14). Robin and Smith (2000) give characteristic-root and sequential rank procedures.

When instruments have different exposure schedules, let S contain the known schedule vectors in target units and let R_S project onto their orthogonal complement. The schedule-adjusted matrix is $R_S M$. Its rank records response variation remaining after the scheduled activation profiles are removed. The application reports raw and schedule-adjusted singular values and carries the transformed coefficient covariance into the second rank calculation.

For a two-by-two matrix, the regular rank-one diagnostic also uses $\det M = 0$. At a rank-one matrix, $\nabla \det(M) = \operatorname{vec}(\operatorname{adj}(M)')$ is nonzero, and

$$\operatorname{Var}(\det \widehat{M}) \simeq \nabla \det(M)' \operatorname{Var}(\operatorname{vec} \widehat{M}) \nabla \det(M). \tag{27}$$

The determinant calculation supplies a comparison at regular rank-one points. The two-step calculation covers the lower-rank points contained in $H_0^{(1)}$.

4. Instrument Overlap and Future Policy Response Capacity

The preceding identification results determine what a design can recover about response capacity. The models in this section give economic explanations for such changes. Interpreting an empirical contrast through one explanation additionally uses that model's restrictions on attachment, experience, or a binding cap.

Two mechanisms shift capacity. A configuration can change the contemporaneous response to an incentive when both instruments operate on the same choice. Earlier adoption can change the response to a later offer through experience and remaining demand. Repeated campaigns carry the resulting transition forward.

4.1. An incentive under an effective cap

Let $x(b) \in \mathbb{R}_+$ be a household's chosen use under incentive status $b \in \{0, 1\}$, with $x(1) \leq x(0)$. A configuration imposes an effective cap $\ell \in \mathbb{R}_+$, producing observed use

$$y(b, \ell) = \min\{x(b), \ell\}. \tag{28}$$

The incentive response under the configuration is $K(\ell) = y(1, \ell) - y(0, \ell)$.

PROPOSITION 9 (Capacity substitution under an effective cap). *For $\ell_1 \leq \ell_0$,*

$$K(\ell_1) - K(\ell_0) = \operatorname{len}\left((\ell_1, \ell_0] \cap (x(1), x(0)] \right) \geq 0, \qquad K(\ell) \leq 0. \tag{29}$$

For coupled configured and comparison caps $L^D \le L^C$ that remain fixed across incentive states, and random potential uses $X(1) \le X(0)$, the population capacity change is

$$\mathbb{E}[K(L^D) - K(L^C)] = \mathbb{E}\Big[\mathrm{len}\left((L^D, L^C] \cap (X(1), X(0)]\right)\Big]. \tag{30}$$

When the outcome is minutes of use, equation (30) gives the mean use margin covered by the change in effective caps. Compliance and snoozing enter through the distribution of effective caps. If the effective configured cap can depend on incentive status, write it as L_b. The configured incentive response relative to the response under an unrestricted cap obeys

$$\begin{aligned} &\min\{X(1), L_1\} - \min\{X(0), L_0\} - [X(1) - X(0)] \\ &\quad = \mathrm{len}\left((L_0, \infty) \cap (X(1), X(0)]\right) + [X(1) - L_0]_+ - [X(1) - L_1]_+. \end{aligned} \tag{31}$$

The first term is common-cap overlap. The second records the change in the effective cap across incentive states. Factorial assignment identifies their sum. The empirical design observes the limit throughout the study and pays the use-reduction bonus during one three-week interval. This timing yields an overlap contrast between the capacity shift during payment and the capacity shift in the surrounding horizons. User-selected tightness remains a post-assignment variable, so the application uses the timing result for the configured bundle and leaves the two terms in equation (31) combined.

4.2. A two-date adoption model

A household has baseline value $v \in \mathbb{R}$ for the product and a net experience effect $g \in \mathbb{R}$ after first adoption. The distribution of (v, g) has marginal density f_v for v. The net effect g combines learning about product quality with the decline in marginal value from owning a durable unit. In period 1 the posted price is p_1 and the subsidy is a. Purchase is

$$D_1(a) = \mathbf{1}\{v \ge c(a)\}, \qquad c(a) = p_1 - a. \tag{32}$$

In period 2 the product is offered at price p_2. The household's response to the offer is

$$D_2(a, p_2) = \mathbf{1}\{v + gD_1(a) \ge p_2\}. \tag{33}$$

The product is unavailable without the second offer, so the period-2 capacity contrast equals

$$K_2(a, p_2) = \mathbb{E}[D_2(a, p_2)]. \tag{34}$$

PROPOSITION 10 (Capacity formation by marginal adopters). *Suppose f_v is continuous at*

$c(a)$ and the conditional distribution of g given v is continuous in v at that point. At values where the displayed indicator is locally constant,

$$\frac{\partial K_2(a, p_2)}{\partial a} = f_v(c(a))\left[\Pr\{g \geq p_2 - c(a) \mid v = c(a)\} - \mathbf{1}\{c(a) \geq p_2\}\right]. \tag{35}$$

An increase in the early subsidy moves the adoption cutoff downward. The bracket in equation (35) compares the marginal household's response to the later offer after early adoption with its response after passing through period 1 without adoption. Experience can move a marginal household above the second purchase threshold. Existing ownership can move it below that threshold. The density $f_v(c(a))$ measures how much population the early subsidy moves into this comparison.

Equation (35) also distinguishes current reach from future response. The current adoption derivative is

$$\frac{\partial \mathbb{E}[D_1(a)]}{\partial a} = f_v(c(a)). \tag{36}$$

Two early subsidies with the same current adoption effect can generate different later capacity effects through the conditional distribution of g among marginal adopters.

4.3. Own and neighborhood exposure

Let a_i be household i's subsidy and let $q_i(a_{-i})$ be an exposure mapping from neighboring assignments. Baseline value and experience can depend on q_i through product information, disease transmission, or local availability. The two-target response to own and neighborhood exposure is

$$M^A = D_{(a,q)} \begin{bmatrix} \mathbb{E}[D_1] \\ \mathbb{E}[D_2] \end{bmatrix}. \tag{37}$$

The first row measures current adoption. The second row measures the response supplied to the later offer. Column proportionality gives rank one; differing ratios across dates produce a second response direction. The empirical application estimates the discrete-assignment version of equation (37).

4.4. Subsidy choice

Let w_1 be the value of first-period adoption, let w_2 be the value of purchase from the later offer, and let $C(a)$ be the fiscal and administrative cost of the early subsidy. A two-date planner chooses

$$\max_{a \in \mathcal{A}} \{ w_1 \mathbb{E}[D_1(a)] + \beta w_2 K_2(a, p_2) - C(a) \}. \tag{38}$$

PROPOSITION 11 (Capacity-adjusted subsidy). *Suppose the objective in equation* (38) *is differentiable and has an interior maximizer a^D. Then*

$$C'(a^D) = w_1 f_v(c(a^D)) + \beta w_2 \frac{\partial K_2(a^D, p_2)}{\partial a}. \tag{39}$$

The second term is given by equation (35).

The myopic subsidy condition contains the first term on the right side. The dynamic condition also values the later purchase response created or removed at the period-1 cutoff. Fiscal cost, learning, and remaining demand therefore enter the same marginal calculation.

4.5. Repeated reach and perfect stillness

Consider a population divided into three shares at date t: adopters x_t, reachable nonadopters r_t, and an excluded share $e \geq 0$, with $x_t + r_t + e = 1$. A campaign reaches and converts a fraction $\alpha \in (0, 1)$ of the reachable nonadopters. The law of motion is

$$x_{t+1} = x_t + \alpha r_t, \qquad r_{t+1} = (1 - \alpha) r_t. \tag{40}$$

The response to another campaign at date t is the discrete capacity contrast

$$K_t = \alpha r_t. \tag{41}$$

THEOREM 1 (Stillness with a persistent target gap). *Under equations* (40)–(41)*,*

$$r_t = (1 - \alpha)^t r_0, \qquad x_t = 1 - e - (1 - \alpha)^t r_0, \qquad K_t = \alpha (1 - \alpha)^t r_0. \tag{42}$$

The state converges to $(x^\star, r^\star) = (1 - e, 0)$ and the response to the campaign converges to zero. The fixed point is perfectly still for that campaign. When $e > 0$, the limiting target gap is e, and every addition drawn from the campaign's existing response span leaves that gap unchanged.

The same conclusion holds with heterogeneous conversion rates. Let $r_0(a)$ be a finite initial mass over $a \in (0, 1]$. Then

$$K_t = \int_0^1 a(1 - a)^t r_0(a)\, da \tag{43}$$

is weakly decreasing and converges to zero, while the reached share converges to $1 - e$. The decline *is strict at date t whenever responsive mass remains.*

The parameter e separates two fixed-point compositions. At $e = 0$, zero capacity accompanies full coverage. At $e > 0$, zero capacity accompanies a population remaining outside reach.[7] A delivery method that attaches part of the excluded group adds a response vector to the dictionary in equation (15). The cost of closing the remaining gap is then determined by the feasible-addition problem.

The heterogeneous clause permits early campaigns to select high-response units and later campaigns to face the surviving low-response types. Its conclusion uses finite initial mass and positive conversion rates while allowing campaign effects to vary across units.

5. Randomized Policy History, Configuration, and Assigned Actions

5.1. Repeated randomized notifications

5.1.1. Design and response estimand

The myBPmyLife intervention followed 287 participants with hypertension for six months and randomized diet and activity notifications at four daily decision points (Golbus et al. 2026). At each point, the probability of a diet notification was 0.125, the probability of an activity notification was 0.125, and the probability of receiving either notification was 0.25. The public core file contains 187,517 randomization points. The analysis uses 186,893 active-phase points with a recorded dice roll.[8] Diet notifications directed participants to the mobile application to find lower-sodium alternatives. The source article classifies 60-minute app engagement as an exploratory post hoc endpoint and reports that diet notifications raise engagement. The present analysis retains that outcome status and asks how the current randomized response varies with prior assigned history at a common recorded pre-state.

Equation (16) is estimated for histories spanning 4, 12, 28, 56, and 112 preceding decision points, corresponding to about 1, 3, 7, 14, and 28 days. The outcome is an indicator for app use during the 60 minutes after assignment, and the pre-state is app use during the preceding 30 minutes. Every specification includes centered current diet and activity assignments, centered randomized history, both current-assignment-by-history interactions, the pre-state, both current-assignment-by-pre-state interactions, centered linear and quadratic relative-day terms interacted with both current assignments, participant effects, relative-day effects, and time-of-day effects. Thus, the history coefficient compares

[7] Application costs, incomplete take-up, and screening can generate persistent exclusion in practice; see Bhargava and Manoli (2015); Finkelstein and Notowidigdo (2019); Deshpande and Li (2019). The parameter e records the resulting excluded mass while leaving its source mechanism unrestricted.

[8] Decision points outside the active phase or lacking the randomization-engine draw are excluded before the history variables are constructed.

current policy responses at the same recorded app-use state and response-time controls. Standard errors cluster by participant.

Inference reports two finite families. The displayed family contains the five diet-history coefficients for binary app use. A broader family contains 40 coefficients spanning five history windows, four engagement outcomes, and two current notification types. Gaussian participant-cluster multiplier statistics provide simultaneous coverage. A retrospective multivariate diagnostic arranges each window's eight coefficients as the four-target-by-two-notification operator in equation (18) and applies simultaneous inference across all five matrices.

5.1.2. State transition and response history

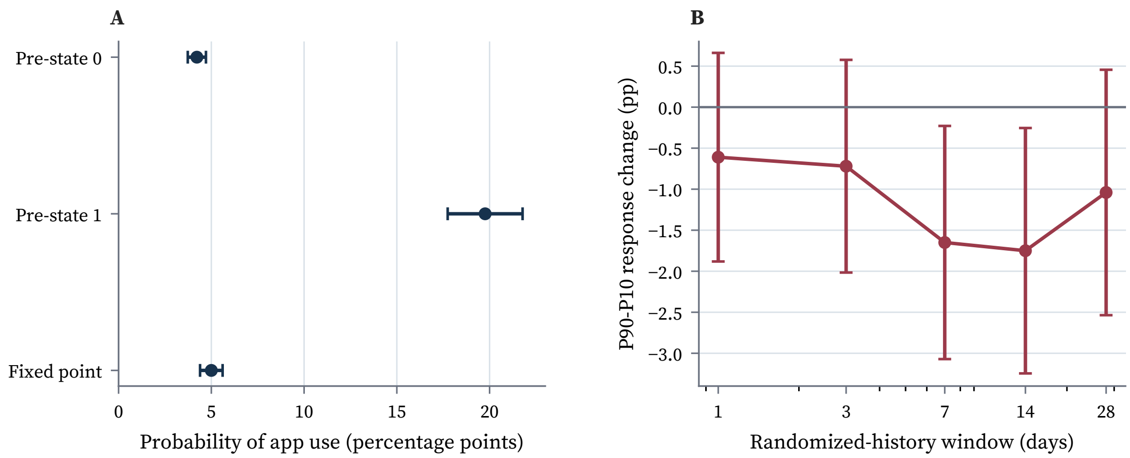

FIGURE 1. App-use transition and response under randomized notification history

Notes. Panel A uses 140,397 active-phase decision points assigned to the 0.75 control condition. Points report app use during the next 60 minutes after each preceding 30-minute app-use state and the fixed point implied by the estimated two-state transition; horizontal intervals are participant-clustered 95 percent intervals. Panel B reports the fitted difference in the conditional current diet-notification response between the 90th and 10th percentiles of assigned history, at common pre-state and time controls. Vertical intervals are simultaneous over the five displayed windows, use the 2.493 family critical value, and are based on 99,999 Gaussian participant-cluster multiplier draws.

At decision points assigned to the control draw, app use during the next 60 minutes equals 4.22 percent after a zero pre-state and 19.76 percent after a one pre-state. Their difference is 0.155 with a participant-clustered standard error of 0.010. Iteration of the corresponding two-state transition gives a fixed point of 5.00 percent. These quantities summarize the passive transition observed in the same randomized decision panel; Table 2 reports their uncertainty.

TABLE 1. Current app-use response by randomized notification history

History (days)	Current diet effect at mean history (pp)	Past assignments P10–P90	Response change P10–P90 (pp)	Primary adjusted p	40-coefficient family adjusted p
1	5.31	0–2	-0.61 (0.51)	0.6031	0.9896
3	5.21	1–5	-0.72 (0.52)	0.4681	0.9539
7	5.03	4–10	-1.65 (0.57)	0.0159	0.0878
14	4.84	10–18	-1.75 (0.60)	0.0136	0.0766
28	4.44	22–34	-1.04 (0.60)	0.2668	0.7856

Notes. The sample contains active-phase decision points for the 287-person core cohort. Current diet and activity assignments are centered at their design probabilities of 0.125. History is the number of successful notification assignments during the preceding window centered at 0.25 times the number of decision points. The current effect is evaluated at mean randomized history after a zero app-use pre-state and at the centered time terms. The conditional response difference multiplies the history-interaction coefficient by the difference between the 90th and 10th percentiles of realized assignment history; clustered standard errors appear in parentheses. Each row includes both current assignments, history, both current-assignment-by-history interactions, app use during the preceding 30 minutes, both current-assignment-by-pre-state interactions, both current-assignment-by-linear-and-quadratic-day interactions, participant effects, relative-day effects, and time-of-day effects. Primary adjustment covers the five displayed diet-history coefficients. The 40-coefficient adjustment covers five windows, four outcomes, and two current notification types. Gaussian cluster-multiplier inference uses the joint participant-score covariance.

At mean randomized history following a zero pre-state, a current diet notification raises 60-minute app use by 4.44 to 5.31 percentage points across the five specifications. Over the preceding seven days, the 10th and 90th percentiles of assigned notification history are 4 and 10. The fitted current diet-notification response is 1.65 percentage points lower at the upper endpoint of that range, with a clustered standard error of 0.57. The interaction coefficient has $p = 0.0041$, adjustment across the five displayed windows gives $p = 0.0159$, and adjustment across all 40 screened history coefficients gives $p = 0.0878$. Over the preceding 14 days, the corresponding history range is 10 to 18 assignments and the fitted conditional response difference is −1.75 percentage points, with a standard error of 0.60. Its pointwise, displayed-family, and screened-family values are 0.0035, 0.0136, and 0.0766.

TABLE 2. Policy-response operator across randomized notification history

Target	Diet notification	Activity notification
App use	-1.754 pp	-0.307 pp
App events	-0.211	-0.002
Searches	-0.027	0.001
Low-salt choices	-0.038	0.004
Standardized singular values	0.136, 0.014	
Simultaneous 95% lower endpoint for σ_1	0.017	
14-day rank-at-most-zero p	0.0059	
Five-window adjusted rank-zero p	0.0171	
14-day rank-at-most-one p	0.5774	

Notes. Entries are fitted differences in the conditional current notification response between the 90th percentile (18 assignments) and the 10th percentile (10 assignments) of preceding 14-day history, holding recorded pre-state and time controls fixed. The matrix uses the eight 14-day interaction coefficients already included in the 40-coefficient family. Rows are scaled by their common-sample standard deviations for the singular-value calculation. The simultaneous lower endpoint subtracts the 95 percent quantile of the largest operator norm across all five window-specific perturbation matrices, using 99,999 Gaussian participant-cluster multiplier draws. The adjusted rank-zero value also covers all five windows. Lower-rank tests use the Chen–Fang two-step procedure with 49,999 draws. Row scaling changes the metric and leaves algebraic rank unchanged. The source study classifies these engagement targets as exploratory post hoc outcomes.

The row-standardized 14-day operator has singular values 0.136 and 0.014. A simultaneous operator-norm band covering all five history windows gives a 0.017 lower endpoint for its leading singular value. The corresponding lower endpoints are 0.023 at seven days and zero at the other three windows. Rank-zero tests adjusted across the five matrices give $p = 0.012$ at seven days and $p = 0.017$ at 14 days. Replacing Gaussian multipliers with Rademacher participant-cluster multipliers gives adjusted values of 0.0084 and 0.0135 and lower endpoints of 0.024 and 0.019. At 14 days, the lower-rank-uniform tests give $p = 0.0059$ for rank at most zero and $p = 0.577$ for rank at most one. The evidence therefore supports at least one history-sensitive response direction across these engagement targets and leaves a second direction unresolved. The operator construction remains a retrospective diagnostic over outcomes classified as exploratory post hoc by the source study.

The history variable uses assignment success recorded by the randomization engine, so the coefficients are intention-to-treat response changes among retained active-phase decision points. The pre-state interactions allow the current diet-notification effect to vary with current app use, and centered linear and quadratic relative-day interactions allow both current notification effects to evolve over the study. The history interaction therefore records response variation remaining after state and study-time response heterogeneity enter the equation. Replacing the preceding 28 assignments with the following 28 assignments gives a diet interaction of 0.00119 with a standard error of 0.00094 and $p = 0.206$. Sequential assignment gives the future-history placebo a zero benchmark.

The distinction is visible across the two panels. Panel A estimates a passive transition for the target state. Panel B estimates history-related variation in the current policy-to-target response, conditional on the observed pre-state. Earlier assignments can also affect that pre-state, a channel held fixed in this projection. The result concerns digital engagement among participants retained in the active decision panel. The source outcome, population, and history windows delimit its transport to other policy systems.

5.2. A bonus under randomized app-use limits

5.2.1. Design and estimand

Allcott et al. (2022) recruited 1,933 Android users and recorded use of Facebook, Instagram, Twitter, Snapchat, browsers, and YouTube in five three-week periods. At the start of period 2, participants received random assignment to a temporary bonus that paid for lower use in period 3. They also received independent assignment to a limit-control group or to an app-limit configuration. The configuration assignments used snooze delays of 0, 2, 5, or 20 minutes and a block without snooze. The limit remained available across the study, while cash bonus payments occurred in period 3. The source article estimates habit formation and self-control and reports bonus-by-limit interactions by period in its online appendix.

Let B_i denote bonus assignment and D_i assignment to any limit configuration. The outcome Y_{ih} is average daily FITSBY use in period h. Equation (13) estimates the bonus response $K_h(0) = \beta_h$ under limit control, the response $K_h(1) = \beta_h + \theta_h$ under a limit assignment, and the capacity shift $\Delta K_h = \theta_h$. The period-level overlap estimand applies equation (22) with $h^\star = 3$, $\mathcal{C} = \{2, 4, 5\}$, and equal comparison weights. A weekly-window analysis compares weeks 7–9, when bonus payments occur, with weeks 4–6 and 10–15. Every regression includes baseline FITSBY use and the reported randomization-stratum indicators.

The smallest bonus-by-limit cell contains 194 participants. FITSBY use is observed for more than 99 percent of participants in each period.[9] A period-1 balance regression gives a bonus-by-limit coefficient of −2.68 minutes with a standard error of 6.50.

Period 3 is the primary horizon, and periods 2–5 form the horizon family. The reported results include the period-3 estimate, the joint family test, and the six enforcement-rule estimates. The finite-family comparison contrasts each period with the mean of the other three and each consecutive three-week window with the other nine weeks. The payment-overlap result is exploratory and receives simultaneous inference over these families.

[9] The 80-percent-power minimum detectable effect for the period-3 interaction is 0.187 outcome standard deviations.

Prospectively registered configuration comparisons illustrate the corresponding design requirements.

5.2.2. Response estimates

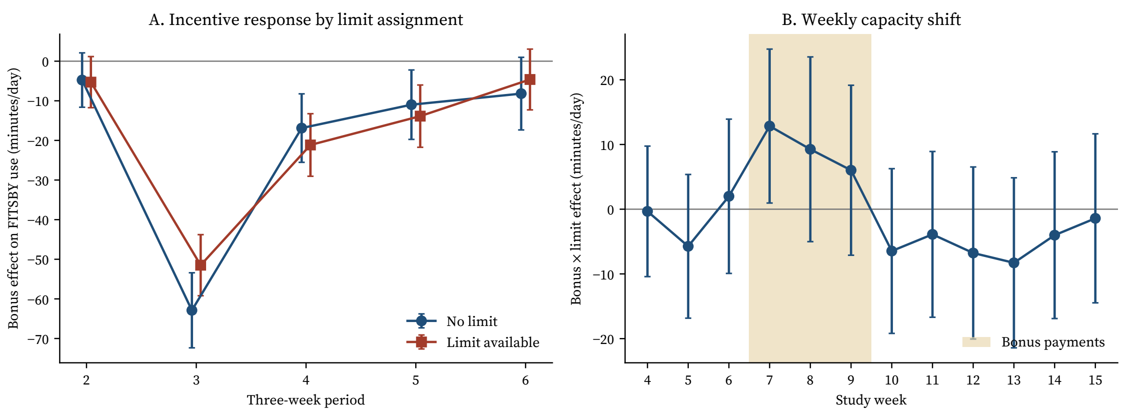

FIGURE 2. Policy overlap and the incentive response

Notes. Panel A reports the estimated bonus effect on average daily FITSBY use by three-week period and binary limit assignment. Panel B reports the weekly bonus-by-limit coefficient from weeks 4 through 15; shading marks the bonus-payment weeks 7 through 9. Bars give pointwise 95 percent HC2 intervals.

Figure 2 follows the incentive response across policy timing. In period 3, the bonus reduces use by 62.88 minutes per day in the limit-control group and by 51.49 minutes under a limit assignment. The period-3 capacity shift is 11.39 minutes with a standard error of 6.24 and a two-sided p-value of 0.068. Its sign agrees with the common-cap term in Proposition 9. The assignment bundles access, enforcement, chosen tightness, and snoozing, so equation (31) gives the interpretation when effective caps vary with bonus status.

The capacity-shift estimates for periods 2–5 are −0.40, 11.41, −4.64, and −3.03 minutes on the common 1,903-person sample. The payment-period shift minus the equal-weight mean of the surrounding shifts is 14.10 minutes, with a standard error of 5.08 and p = 0.0054. Across all four one-period-versus-rest contrasts, Gaussian max-t adjustment gives p = 0.020 and a simultaneous interval from 1.62 to 26.59 minutes. A 9,999-draw conditional re-randomization follows the recorded assignment sequence and re-estimates the HC2 standard error in every draw. The weak-null values are 0.0059 for the payment contrast and 0.0216 after adjustment over the four-period family. At weekly frequency, the mean shift during weeks 7–9 minus the mean shift in the other nine included weeks is 13.25 minutes, with a standard error of 4.96 and p = 0.0075. Across the ten consecutive three-week windows, its adjusted value is 0.047 and its simultaneous interval runs from 0.10 to 26.41 minutes. The adjacent weeks 6–8 window has the largest studentized estimate in that family and an adjusted value of 0.027. The four horizon coefficients have a joint Wald

TABLE 3. Randomized limit assignment and the response to a temporary use-reduction bonus

	Period 2	Period 3	Period 4	Period 5
Bonus effect without limit	-4.76	-62.88	-16.89	-10.98
	(3.50)	(4.83)	(4.41)	(4.47)
Bonus effect with limit	-5.29	-51.49	-21.15	-13.87
	(3.29)	(3.93)	(4.03)	(4.01)
Capacity shift: bonus × limit	-0.52	11.39	-4.26	-2.89
	(4.81)	(6.24)	(5.98)	(6.01)
Limit effect	-24.51	-25.62	-21.86	-20.58
	(2.29)	(2.89)	(2.93)	(2.89)
Observations	1930	1932	1931	1933

Notes. Each column reports HC2 estimates from a factorial regression of average daily FITSBY use on bonus assignment, limit assignment, their interaction, baseline FITSBY use, and randomization-stratum indicators. Period 3 is the bonus-payment period. Standard errors appear in parentheses.

statistic of 9.12 with four degrees of freedom and $p = 0.058$.

The limit coefficient ranges from −20.57 to −25.59 minutes per day across periods 2–5. After each horizon is scaled by its outcome standard deviation, the four-by-two matrix of bonus and limit coefficients has raw singular values 0.821 and 0.293. The bonus payment occurs in period 3 and limit access spans periods 2–5. Projecting the response columns away from these two exposure schedules gives singular values 0.094 and 0.024. The lower-rank-uniform schedule-adjusted test gives $p = 0.066$ for rank at most one. The raw rank records the assigned policies together with their timing; the schedule-adjusted calculation leaves the residual response rotation unresolved at the 5 percent level.

The six configuration estimates range from a 44.27-minute to a 71.47-minute period-3 bonus response. Figure 2 reports Gaussian max-t intervals across the six estimates; a companion family covers the five comparisons with limit control. Each comparison interval contains zero. The binary configuration contrast supplies the period and timing estimands.

The released author code specifies the same four equations as $Y_h \sim Y_1 + L + B + L \times B + S$. Its assignment program draws bonus status within eight baseline strata, draws limit status within stratum-by-bonus cells, and draws snooze delay within the corresponding delayed-snooze cells. Re-estimation with that equation changes every bonus, limit, and interaction coefficient by less than 2×10^{-10}. The conditional re-randomization preserves treatment counts within these cells and uses observed outcomes with an HC2-studentized statistic. The period-3 weak-null value is 0.069. Studentization permits heterogeneous treatment effects under the weak-null procedure in Wu and Ding (2021).

5.2.3. Choice over assigned actions

The experiment supplies four attainable actions: neither assignment, bonus alone, limit alone, and both assignments. Their target changes relative to the control cell are 0, $\widehat{\beta}$, $\widehat{\delta}$, and $\widehat{\beta} + \widehat{\delta} + \widehat{\theta}$. The planner declares a common daily use-reduction target for periods 2–5 and evaluates equation (4) over these four points. Each horizon is weighted by the inverse of its outcome variance in the 586-person control cell.

TABLE 4. Feasible factorial actions for declared use-reduction targets

Assigned action	15 minutes	30 minutes	50 minutes
None	0.049	0.195	0.543
Bonus	0.124	0.121	0.269
Limit	0.015	0.011	0.158
Bonus and limit	0.272	0.123	0.076
Estimated choice	Limit	Limit	Bonus And Limit
Draw probability of estimated choice	0.998	1.000	0.999

Notes. Each column declares the same daily reduction target for periods 2 through 5. The feasible set contains the four randomized assignment cells. The combined action uses the estimated bonus, limit, and interaction coefficients. Loss is one half of the squared target gap after each row is divided by the corresponding control-cell outcome standard deviation. The final row gives the frequency with which 49,999 Gaussian draws from the joint HC2 coefficient law select the reported action. Costs enter through the boundaries reported in the text.

With zero activation cost, the estimated action is no assignment for targets below 11.61 minutes, limit alone from 11.61 through 41.53 minutes, and both assignments above 41.53 minutes over the reported zero-to-80-minute range. Table 4 evaluates targets of 15, 30, and 50 minutes. The limit is selected in 99.8 percent and 100 percent of the 49,999 coefficient draws at 15 and 30 minutes. Both assignments are selected in 99.9 percent of draws at 50 minutes. Let λ denote dollars per standardized target-loss unit and let c_L be the resource cost of limit implementation. At a 30-minute target, limit assignment beats no assignment when $c_L < 0.184\lambda$ and beats bonus assignment when $c_L < 76.42 + 0.109\lambda$. At a 50-minute target, adding the bonus after the limit requires $\lambda > 929.9$ when the payment is counted as an agency outlay. Mean predicted bonus earnings are \$76.42. Mean bonus valuation is \$64.07, and mean limit valuation is \$4.18 among 1,149 observed responses. A social-welfare objective assigns the transfer incidence and implementation resources before applying the same four-action comparison.

A sequential comparison. The two-stage bed-net pricing experiment in Dupas (2014) provides a history design. Assignment to a Phase 1 price at or below 50 Kenyan shillings raises Phase 1 purchase by 37.2 percentage points and purchase from a common 150-shilling offer one year later by 7.2 percentage points. The within-area randomization values are 0.0001 and 0.0742. A response matrix adds randomized neighborhood exposure; its lower-rank-uniform test of rank at most one gives $p = 0.077$. The later estimate records the total effect of early assignment on response to the common offer.

6. Conclusion

A history of intervention can change the response to the next intervention. Policy response capacity records this change for a declared target, instrument set, and horizon. Joint unit transitions separate entry, exit, and changes among continuing units. The baseline law and the response map carry different information: systems with the same attracting baseline law can have different response ranks.

The randomized applications give this distinction empirical content at specific margins. In myBPmyLife, the fitted current app-use response to a diet notification differs by -1.75 percentage points between the 90th and 10th percentiles of 14-day assigned history, conditional on recorded pre-state and time controls. This is a conditional response contrast; the total effect of an earlier assignment regime also includes its effect on intermediate states. The displayed five-window and screened 40-coefficient families give different multiplicity references, and the exploratory status of the source outcomes remains relevant. In the digital-use experiment, the response to a temporary bonus varies with the randomized app-limit configuration. The analysis keeps its post-estimation period contrast separate from prospectively declared external comparisons.

Response rank also depends on the exposure schedules in the design. The schedule-adjusted digital-use test gives $p = 0.066$ for rank at most one. A second residual direction therefore requires additional evidence about independently varying treatment timing. Policy choice in the current exercise compares assigned action points and their population mixtures; the feasible-action condition specifies when a linear-span relaxation has the same targeting loss.

A stable target gap and a small current response must be read together with the intervention that is available. The scope applications distinguish population coverage, delivery-route exhaustion, attrition, and missing measurement. The framework makes those distinctions explicit, while the online appendix develops additional instruments, continuation value, and baseline-stability extensions.

References

Hunt Allcott, Matthew Gentzkow, and Lena Song. Digital addiction. *American Economic Review*, 112(7):2424–2463, 2022. doi: 10.1257/aer.20210867.

Isaiah Andrews, Toru Kitagawa, and Adam McCloskey. Inference on winners. *Quarterly Journal of Economics*, 139(1):305–358, 2024. doi: 10.1093/qje/qjad043.

Alan J. Auerbach and Yuriy Gorodnichenko. Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27, 2012.

Dirk Bergemann and Juuso Välimäki. Dynamic pricing of new experience goods. *Journal of Political Economy*, 114(4):713–743, 2006. doi: 10.1086/506923.

Saurabh Bhargava and Dayanand Manoli. Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11): 3489–3529, 2015.

Audrey Boruvka, Daniel Almirall, Katie Witkiewitz, and Susan A. Murphy. Assessing time-varying causal effect moderation in mobile health. *Journal of the American Statistical Association*, 113 (523):1112–1121, 2018. doi: 10.1080/01621459.2017.1305274.

Willem H. Buiter and Mark Gersovitz. Issues in controllability and the theory of economic policy. *Journal of Public Economics*, 15(1):33–43, 1981.

Qihui Chen and Zheng Fang. Improved inference on the rank of a matrix. *Quantitative Economics*, 10(4):1787–1824, 2019. doi: 10.3982/QE1139.

Manasi Deshpande and Yue Li. Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11(4):213–248, 2019.

Pascaline Dupas. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228, 2014. doi: 10.3982/ECTA9508.

Martin Eichenbaum, Sergio Rebelo, and Arlene Wong. State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, 112(3):721–761, 2022.

Amy Finkelstein and Matthew J. Notowidigdo. Take-up and targeting: Experimental evidence from SNAP. *Quarterly Journal of Economics*, 134(3):1505–1556, 2019.

Jessica R. Golbus, Michael P. Dorsch, Yuxuan Chen, Tanima Basu, Evan Luff, Predrag Klasnja, Mark W. Newman, Lesli E. Skolarus, Walter Dempsey, and Brahmajee K. Nallamothu. Impact of push notifications on physical activity and sodium intake among patients with hypertension: Microrandomized trial of a just-in-time adaptive intervention. *Journal of Medical Internet Research*, 28:e78218, 2026. doi: 10.2196/78218.

Robert Hahn, Robert D. Metcalfe, David Novgorodsky, and Michael K. Price. The behavioralist as policy designer: The need to test multiple treatments to meet multiple targets. Working Paper 22886, National Bureau of Economic Research, 2016.

Yu-Chin Hsu and Shu Shen. Dynamic regression discontinuity under treatment effect heterogeneity. *Quantitative Economics*, 15(4):1035–1064, 2024. doi: 10.3982/QE2150.

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.

Jacquelyn Pless. Are complementary policies substitutes? evidence from R&D subsidies in the United Kingdom. *American Economic Journal: Economic Policy*, 18(1):92–126, 2026. doi: 10.1257/pol.20220356.

Valerie A. Ramey and Sarah Zubairy. Government spending multipliers in good times and in bad: Evidence from U.S. historical data. *Journal of Political Economy*, 126(2):850–901, 2018.

Jean-Marc Robin and Richard J. Smith. Tests of rank. *Econometric Theory*, 16(2):151–175, 2000. doi: 10.1017/S0266466600162012.

Jan Tinbergen. *On the Theory of Economic Policy*. Contributions to Economic Analysis. North-Holland, Amsterdam, 1952.

Jason Wu and Peng Ding. Randomization tests for weak null hypotheses in randomized experiments. *Journal of the American Statistical Association*, 116(536):1898–1913, 2021. doi: 10.1080/01621459.2020.1750415.

Xun Wu, M. Ramesh, and Michael Howlett. Policy capacity: A conceptual framework for understanding policy competences and capabilities. *Policy and Society*, 34(3–4):165–171, 2015. doi: 10.1016/j.polsoc.2015.09.001.