Almost a Policy

Transmission through Administrative Handoffs
Online Appendix

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

This online appendix contains descriptive timing evidence, implementation alternatives, general rank and frontier inference, and additional Oregon checks. Core proofs, the local-summary derivation, and the interpretation of the diagonal rank intervals accompany the main paper.

JEL Codes: C14, C31, C44, H51, I13, I18

Keywords: administrative handoffs, partial identification, information retention, record design, Oregon Health Insurance Experiment, policy transmission

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

Contents

Appendix A. Delivery timing and descriptive reception

The arrival analysis then conditions on lottery selection and links 13,069 responses with observed survey-return and notification dates. Figure A1 treats the survey answer as a current-status measure of arrival. Reported receipt is 15.0 percent during days 0–28 after notification, 49.4 percent during days 29–42, 84.0 percent during days 43–55, and 89.9 percent during days 56–84. The 12.9 percent report before the focal notification date records household timing and reporting error contained in this measure. Delivery time remains interval-censored between dated observations.

Among survey responses after day 55 with complete administrative outcomes, 7,383 people report receipt and 867 report nonreceipt. Later application rates are 77.7 and 43.7 percent, and later approval rates are 41.4 and 14.4 percent. Survey response timing and reported receipt are selected outcomes, so the panel describes continuation through the administrative chain. Among 2,569 selected respondents who had yet to submit and answered the reason question, 40.4 percent report unfinished work, 8.8 percent report paperwork hassle, 5.7 percent report missing citizenship documents, and 6.5 percent report missing other documents; the survey permits several reasons.

FIGURE A1. Reported arrival and subsequent administrative outcomes

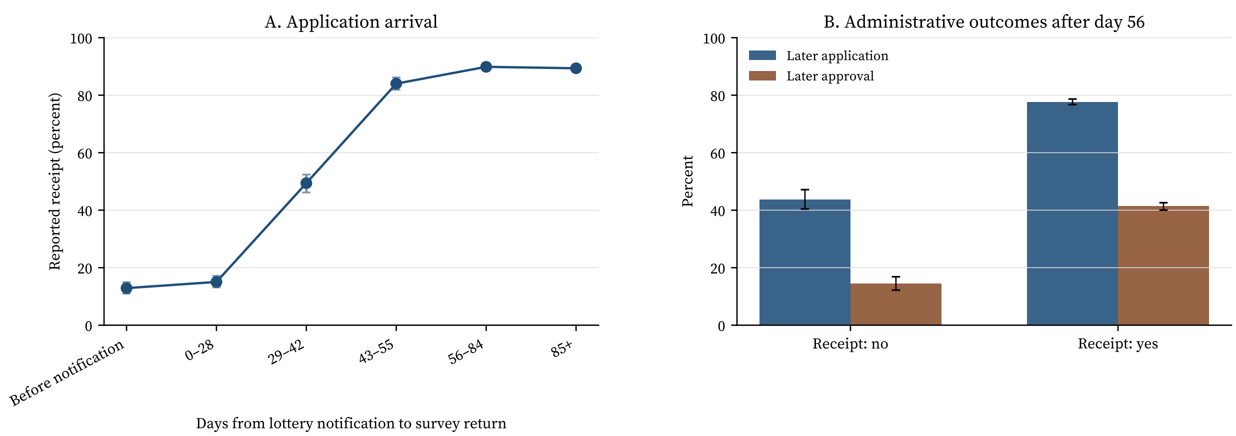

Notes. Panel A reports application receipt as of the initial mail survey among 13,069 lottery-selected respondents with observed survey-return and notification dates. Panel B reports later administrative application and approval among respondents observed after day 55, separated by reported receipt. Bars give household-bootstrap 95% percentile intervals from resampling households within lottery draw. The comparisons describe selected survey respondents.

The administrative decision clock is observed for 18,100 selected applicants with a decision date on or after notification. Its twenty-fifth, fiftieth, seventy-fifth, and ninetieth percentiles are 55, 76, 98, and 120 days. An event study aligned to lottery notification combines arrival, submission, review, and post-decision response. Proposition A1 separates the arrival distribution from the response after effective receipt when linked timing data identify the former.

FIGURE A2. Application and conditional approval after lottery selection

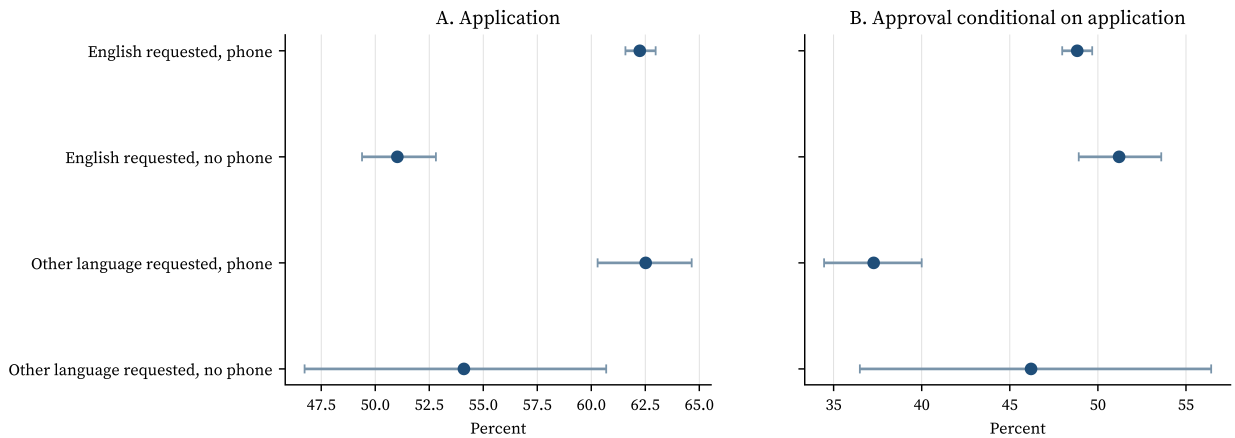

Notes. Points are group rates among 29,799 selected people with complete application and approval records. Bars are household-bootstrap 95% percentile intervals from resampling households within lottery draw. Language request and phone availability come from the reservation-list record.

A.1. Arrival before propagation

Let a nominal intervention occur at date zero, let D_i be the delay before it becomes effective for receiver i, and let g_k be the causal response k periods after effective arrival. If $f_\ell = \mathbb{P}(D_i = \ell)$, the response aligned to the nominal event date is

$$m_h = \sum_{\ell=0}^{h} f_\ell g_{h-\ell}.$$

PROPOSITION A1 (Arrival convolution). *The nominally aligned response is the convolution $m = f * g$. Over unrestricted real response sequences, a given sequence m admits multiple pairs (f, g). Knowledge of f with $f_0 > 0$ yields*

$$g_0 = \frac{m_0}{f_0}, \qquad g_h = \frac{m_h - \sum_{\ell=1}^{h} f_\ell g_{h-\ell}}{f_0}.$$

The two sequences describe distinct clocks. The arrival distribution follows delivery, application, notification, account-credit, or processing records. The post-arrival response follows behavior after effective receipt. When linked timing observations identify the arrival law, the stated common-response convolution permits recursive recovery. Current-status survey reports at selected dates require their own observation assumptions for that step.

A.2. Proof of the arrival convolution

PROOF OF PROPOSITION A1. Conditioning on $D_i = \ell$ and aligning the response by the nominal event date gives $m_h = \sum_{\ell=0}^{h} f_\ell g_{h-\ell}$. For any sequence f with $f_0 > 0$, the equation at $h = 0$ gives $g_0 = m_0/f_0$. At horizon $h \geq 1$, separating the term with $\ell = 0$ and dividing by f_0 gives the displayed recursion. The point mass $f_0 = 1$ yields $g = m$. The distinct probability law $f_0 = f_1 = 1/2$ and the same recursion yield another response sequence with the same m, establishing multiplicity over the unrestricted real response-sequence class. $\square$

Appendix B. Implementation and alternative releases

PROPOSITION A2 (Pilot implementation). *Suppose the upstream score and the repair loading $R(\theta)$ are continuous in quadratic mean near θ_0. Assume the rank of $R(\theta)$ is locally constant at θ_0. Let an independent pilot produce $\widetilde{\theta} \to_p \theta_0$, and apply the channel with loading $R(\widetilde{\theta})$ to the production records. Its conditional information converges in probability to $Q_{R(\theta_0)}$. When the relevant information and target-image ranks are also locally constant, with their positive spectra separated from zero, the estimated channel has the same score rank with probability approaching one, and its target ambiguity and covariance converge to the corresponding repair benchmark.*

The pilot and production samples together define a parameter-independent rule. Local rank stability keeps the estimated projection continuous. The link-sampling construction in the main text instead implements its menu directly through recorded sampling probabilities.

B.1. Proof of pilot implementation

Conditional on the pilot, the production channel is a Markov kernel with loading $R(\widetilde{\theta})$. Quadratic-mean continuity of the score, continuity of the loading, and local constancy of its rank give $Q_{R(\widetilde{\theta})} \to_p Q_{R(\theta_0)}$. Eigenvalues and singular subspaces separated from zero preserve the selected score rank, while the associated target ambiguity subspace and covariance converge to their limiting values. Integrating over the pilot yields the stated unconditional result.

The Gaussian summary benchmark uses adjudication precision that scales with the application rate. A finite-sample construction fixes the number of adjudication outcomes and places the random record-reading cost upstream. The stopping rule is inverse binomial sampling (Feller 1968). Let $\tau_r = \inf\{s : \sum_{i=1}^{s} A_i = r\}$ be the location of the rth application in an iid record stream.

PROPOSITION A3 (Exact separation and record cost). *Fix $m \geq 1$ and release $T_m = \sum_{r=1}^{m} Y_{\tau_r}$ while retaining $(\tau_1, \ldots, \tau_m)$ upstream. Then*

$$T_m \sim \text{Binomial}(m, b),$$

so the release has Fisher information $\operatorname{diag}\{0, m/[b(1-b)]\}$ *for (a, b) and $\mathbb{E}[\tau_m] = m/a$. Any release whose distribution is invariant to a over an interval has b information invariant to a over that interval. The full adjudication information from n invitations is $na/[b(1-b)]$, so coincidence with full adjudication precision occurs for at most one value of a within that interval.*

The proposition describes two production choices. A fixed-grid rate attains full adjudication precision in the local experiment. A fixed-decision release separates adjudication from application in finite samples and prices production by the number of records read. Within their stated models, both carry a positive b score. Their precision and production costs differ.

B.2. Proof of exact separation and record cost

At every application time τ_r, conditional independence gives $Y_{\tau_r} \sim \text{Bernoulli}(b)$. The strong Markov property of the iid sequence makes these outcomes independent across r, so their sum is binomial with parameters (m, b). Its score for b is $(T_m - mb)/\{b(1-b)\}$ and has variance $m/[b(1-b)]$. Its law is constant in a, which gives zero a score. The waiting increments $\tau_r - \tau_{r-1}$ are geometric with mean $1/a$, hence $\mathbb{E}[\tau_m] = m/a$.

For the last claim, write the released law as Q_b. Its b score and Fisher information are functions of (Q_b, b) and therefore have the same value at every a in the interval. The linked n-invitation experiment has b information $na/[b(1-b)]$, which is strictly increasing in a. Equality between these two information quantities therefore holds at most one value of a.

B.3. The fixed-decision illustration in Oregon

The supplied exercise orders invitation records externally and reads until $m_g = \lfloor n_g/4 \rfloor$ applications have appeared in each group. It reports the approval count and keeps the stopping location upstream.

TABLE A1. Exact separation through a fixed-decision release

Group	Decisions	Records read	Approval rate
English requested, phone	5,804	39.7%	47.9%
English requested, phone absent	915	48.7%	51.5%
Other language requested, phone	663	39.4%	35.9%
Other language requested, phone absent	67	45.1%	50.7%

Notes. The release places invitation records in an external random order, reads records until one quarter of each group count has applied, and reports the approval count. Records read are realized percentages of group invitation counts. Expected reading shares range from 40.0 to 49.0 percent and equal the shares of linked adjudication information retained.

The theorem describes an iid stream. Table A1 records one finite-archive ordering, while its expected-cost comparison plugs group rates into the stream benchmark. A finite-archive production rule also needs a protocol for exhaustion before the required number of applications. The fixed-grid benchmark and this fixed-decision illustration answer related release questions under different observation schemes.

B.4. Scalar specialization

For $q = 1$, write the target gradient as g. Suppose $I > 0$, and omit the cap $Q \preceq L$ temporarily. Define $s = I^{-1/2}g$ and $H = I^{1/2}WI^{1/2}$. The cap-free solution at $b > 0$ is rank one. With $z_b = (H + b\mathbf{I}_p)^{-1}s$ and $x_b = \{b/(z_b'Hz_b)\}^{1/2}z_b$, it is

$$Q_b = I^{1/2}x_b x_b' I^{1/2}, \qquad \Phi(b) = g'I^{-1}g - b\,s'(H + b\mathbf{I}_p)^{-1}s.$$

When $Q_b \preceq L$, this pair also solves the capped problem. To see the formula, transform $Q = I^{1/2}RI^{1/2}$ and use

$$s'(\mathbf{I}_p + R)^{-1}s = \max_y\{2s'y - y'y - y'Ry\}.$$

For fixed y, maximizing $y'Ry$ over $R \succeq 0$ with $\operatorname{tr}(HR) \le b$ produces $bH^{-1}yy'H^{-1}/(y'H^{-1}y)$. Substitution and maximization over y give the displayed expressions.

Appendix C. Rank and frontier inference

C.1. Rank estimation under separated spectra

Let P_{j-1} and P_j^- denote the orthogonal projections onto N_{j-1} and N_j^-. For a matrix B, write $\operatorname{rank}_c(B)$ for the number of singular values above c. With estimators $\widehat{I}_{j-1}$, $\widehat{I}_j^-$, and $\widehat{C}$, construct the corresponding small-eigenvalue projections $\widehat{P}_{j-1}(c)$ and $\widehat{P}_j^-(c)$ and set

$$\widehat{\ell}_j(c) = \operatorname{rank}_c\{\widehat{C}\widehat{P}_j^-(c)\} - \operatorname{rank}_c\{\widehat{C}\widehat{P}_{j-1}(c)\}.$$

PROPOSITION A4 (Separated-rank estimation). *Allow the parameter and target dimensions to vary with n. Suppose $\|\widehat{I}_{j-1} - I_{j-1}\| + \|\widehat{I}_j^- - I_j^-\| + \|\widehat{C} - C\| = O_p(a_n)$, $\|C\|$ is bounded, and the positive eigenvalues that define the two null spaces and the positive singular values of CP_{j-1} and CP_j^- exceed η_n. If*

$$\frac{a_n}{\eta_n} = o(c_n), \qquad c_n = o(\eta_n),$$

then $\widehat{\ell}_j(c_n) \to_p \ell_j$.

The proposition describes point estimation at a parameter with separated ranks. Its rate conditions cover a growing target list when operator-norm estimation error remains below the spectral separation scale. Projection of the estimated information pair onto $\widehat{I}_{j-1} \succeq \widehat{I}_j^- \succeq 0$ enforces the population ordering before the calculation. The Oregon target list contains four pre-specified reception groups, so its reported rank set uses the finite-dimensional confidence-region construction below.

C.2. Confidence sets near rank changes

Let $\beta_j = (I_{j-1}, I_j^-, C)$ and let $\mathfrak{P}_j(r)$ collect primitive triples that satisfy the information ordering and have loss rank r. Suppose $\mathfrak{C}_{n,j}(1-\alpha)$ is a joint confidence region for β_j with coverage uniform over a class $\mathfrak{P}_j$ that permits eigenvalues and target singular values of order $n^{-1/2}$. Define

$$\mathfrak{R}_{n,j}(1-\alpha) = \{r \in \{0, \ldots, q\} : \mathfrak{C}_{n,j}(1-\alpha) \cap \mathfrak{P}_j(r) \neq \varnothing\}.$$

PROPOSITION A5 (Weak-rank confidence set). *If $\mathfrak{C}_{n,j}(1-\alpha)$ has uniform asymptotic coverage, then*

$$\liminf_{n\to\infty} \inf_{\beta_j \in \mathfrak{P}_j} \mathbb{P}_{\beta_j}\{\ell_j \in \mathfrak{R}_{n,j}(1-\alpha)\} \geq 1-\alpha.$$

The construction propagates coverage from the primitive region to the rank set. A singleton conclusion additionally requires that this region separate the admissible rank classes. Spectral separation for a point estimator and such separation of a confidence region are distinct requirements. A Wald or bootstrap region can supply $\mathfrak{C}_{n,j}$ when its required joint coverage has been justified for the maintained class. In moderate dimensions, computation enumerates r and minimizes distance to $\mathfrak{P}_j(r)$. When the null-space bases are separated, the induced quotient map can be tested directly with the sum of its trailing squared singular values; bootstrap rank tests developed by Chen and Fang (2019) provide critical values for that calculation. The Oregon application has a diagonal quotient map, and inversion of simultaneous confidence intervals gives the rank set in closed form.

C.3. Inference for the frontier

Let $\beta = (I, L, W, C, \Omega)$ and $\phi_b(\beta) = \Phi(b; \beta)$. Suppose $\sqrt{n}(\widehat{\beta} - \beta) \rightsquigarrow Z$ and a cluster or block bootstrap consistently estimates the law of Z. Suppose also that, on the compact budget set used for inference, target identification holds with a uniform margin and the parametric semidefinite value map is continuous and Hadamard directionally differentiable uniformly over the reported finite budget grid. Then $\sqrt{n}\{\phi_b(\widehat{\beta}) - \phi_b(\beta)\}$ converges to the directional derivative $\phi'_{b,\beta}(Z)$. Active caps can give this derivative a nonlinear form.

For a bootstrap draw $\widehat{\beta}^*$, let $Z_n^* = \sqrt{n}(\widehat{\beta}^* - \widehat{\beta})$ and choose $\varepsilon_n \downarrow 0$ with $\sqrt{n}\varepsilon_n \to \infty$. The numerical delta draw is

$$D_{n,b}^* = \frac{\phi_b(\widehat{\beta} + \varepsilon_n Z_n^*) - \phi_b(\widehat{\beta})}{\varepsilon_n}.$$

PROPOSITION A6 (Frontier bands). *Under the preceding conditions, the conditional law of $D_{n,b}^*$ consistently estimates the law of $\phi'_{b,\beta}(Z)$ at each budget. On a finite budget grid, quantiles of the maximum centered numerical-delta process give simultaneous bands. If the weak-rank confidence region includes primitives that violate target identification at a budget, the corresponding upper endpoint is $+\infty$.*

The numerical delta construction follows the resampling logic for directionally differentiable maps in Fang and Santos (2019) and Hong and Li (2018). Value-function inference supplies related uniform procedures (Firpo et al. 2023). In the Oregon specialization, positive group coefficients and a scalar capacity constraint yield a unique allocation. The analysis uses the ordinary household bootstrap and a simultaneous max-t band over the displayed capacity grid. Supplemental Appendix D.4 reports ordinary-bootstrap and numerical-delta coverage in a design calibrated to the four reception groups.

C.4. Proofs for rank and frontier inference

For Proposition A4, Weyl's inequality places estimated eigenvalues within $O_p(a_n)$ of their population values. The threshold conditions recover both null-space dimensions, and Davis–Kahan perturbation bounds give $\|\widehat{P}-P\| = O_p(a_n/\eta_n)$. Boundedness of $\|C\|$ then gives $\|\widehat{CP}-CP\| = O_p(a_n/\eta_n)$. Singular-value perturbation, $a_n/\eta_n = o(c_n)$, and $c_n = o(\eta_n)$ recover each image rank, and their difference yields ℓ_j.

For Proposition A5, coverage of $\mathfrak{C}_{n,j}(1-\alpha)$ implies that the true primitive belongs to the region with asymptotic probability at least $1-\alpha$, uniformly over the stated class. On that event, the region intersects $\mathfrak{P}_j(\ell_j)$ at the true primitive, so $\ell_j \in \mathfrak{R}_{n,j}(1-\alpha)$. Coverage alone gives the inclusion statement. In the Oregon diagonal model, the simultaneous coefficient box excludes zero in every coordinate and thus separates rank four from the lower ranks.

For Proposition A6, generalized Schur complements place the value problem in a finite-dimensional convex program with compact Q set. The uniform identification margin bounds feasible covariance matrices over the budget set. Hadamard directional differentiability is an explicit assumption of the result. Under that assumption, the stated value-function sensitivity is available. The directional delta method maps Z through $\phi'_{b,\beta}$. The numerical derivative with $\varepsilon_n \downarrow 0$ and $\sqrt{n}\varepsilon_n \to \infty$ consistently estimates that map on bootstrap draws (Fang and Santos 2019; Hong and Li 2018). Applying the same construction jointly over a finite grid yields the maximum-process band.

Appendix D. Oregon Data and Additional Results

D.1. Source and construction

The source is the Oregon Health Insurance Experiment Public Use Data archive distributed by the National Bureau of Economic Research. The analysis links the descriptive, initial mail-survey, and state-program files. The arrival sample contains 13,069 selected people with reported receipt, survey-return date, and lottery-notification date; application and approval fields are complete for 29,799 of 29,834 selected people.

TABLE A2. Lottery selection and reported application receipt

Outcome	Control	Selected	Adjusted difference
Receipt item response	44.6	44.2	-0.4 [-1.3, 0.5]
Documented receipt	6.8	31.3	24.7 [24.1, 25.4]
Reported receipt	15.1	70.9	56.2 [55.2, 57.2]
Lee trimming interval	trim 1.04% of controls		55.6–56.6

Notes. The initial survey sample contains lottery-selected and control reservation-list members. Means are percentages. Adjusted differences are coefficients on lottery selection in linear probability models with lottery-draw and household-size fixed effects and household-clustered 95% intervals. Documented receipt equals one when a sampled member answers the item and reports receipt, so its assignment contrast uses the full survey frame. Reported receipt conditions on item response. The trimming interval imposes aggregate monotone response selection and matches the lower item-response rate by trimming the assignment group with the higher rate.

TABLE A3. Finite-file bounds for incomplete administrative outcomes

Reception group	Missing	Approval after invitation	Approval conditional on application
English, phone	30	30.35–30.48	48.73–48.93
English, no phone	3	26.10–26.19	51.12–51.28
Other, phone	2	23.29–23.36	37.23–37.35
Other, no phone	0	25.00–25.00	46.21–46.21

Notes. Bounds assign each incomplete selected record any logically admissible application and approval outcome, with approval requiring application. They are sharp finite-file bounds conditional on the observed group labels. Sampling uncertainty is treated separately.

Arrival lag is the number of days from lottery notification to survey return, grouped as pre-notification, 0–28, 29–42, 43–55, 56–84, and 85 days onward. The subsequent-outcome panel begins at day 56. Reason shares use 2,569 selected respondents with an outstanding submission and a multi-response reason item. Percentile intervals resample households within lottery draw. The survey records current status at response; the administrative file supplies later outcomes.

For person i in group g, the linked log likelihood contribution is

$$A_i \log a_g + (1 - A_i)\log(1 - a_g) + A_i\{Y_i \log b_g + (1 - Y_i)\log(1 - b_g)\}.$$

The two score components are $(A_i - a_g)/\{a_g(1-a_g)\}$ and $A_i(Y_i - b_g)/\{b_g(1-b_g)\}$. Their cross moment is zero by iterated expectations. Summation yields $I_{a,g}$ and $I_{b,g}$ in Section IV. The conditional-rate release maps the group records to $\widehat{b}_g = \sum_i A_i Y_i / \sum_i A_i$ and rounds this rate to the public grid $\delta_g = n_g^{-3/4}$. Proposition 3 gives the statistic-level Gaussian limit used for the local summary benchmark. The fixed-decision release draws an external ordering and stops at application $m_g = \lfloor n_g/4 \rfloor$. Its approval count has information $m_g/[b_g(1-b_g)]$, zero a_g information, and expected scan length m_g/a_g. On the resulting application-score null space, the derivative of $d_g = a_g b_g$ is b_g, which gives the diagonal quotient map.

For a retained fraction q_g, application information is $q_g I_{a,g}$. The independent-person benchmark delta-method variance of $\widehat{d}_g$ is

$$\frac{b_g^2 a_g(1-a_g)}{n_g q_g} + \frac{a_g b_g(1-b_g)}{n_g}.$$

Multiplication by N and summation over groups gives the empirical frontier. The first-order condition on every uncapped group gives $q_g \propto b_g\{a_g(1-a_g)\}^{1/2}/p_g$; capped groups are set to one and the remaining budget is reallocated among the remaining groups.

TABLE A4. Reception-stage contrasts in the lottery records

Lottery-list contrast	Application	Approval | applied
Phone absent minus recorded, English requested	-11.2 [-13.0, -9.3]	2.4 [0.0, 4.9]
Other language minus English, phone	0.3 [-2.0, 2.5]	-11.6 [-14.4, -8.7]

Notes. Entries are percentage-point differences with household-bootstrap 95% percentile intervals. The records describe contact fields supplied at lottery signup. The contrasts measure stage-specific associations within selected households.

TABLE A5. Approval rates under weighting and draw deletion

Group	Selected cohort	Lottery-list weighted	Leave-one-draw range
English requested, phone	30.4	31.3	[29.3, 31.6]
English requested, phone absent	26.1	26.4	[25.2, 27.3]
Other language requested, phone	23.3	23.5	[22.1, 24.0]
Other language requested, phone absent	25.0	23.5	[22.0, 27.1]

Notes. Entries are percentages. Lottery-list weights invert the empirical selection rate within the number-of-household-members stratum. The final column gives the range from repeating the selected-cohort calculation after removing each of the eight lottery draws.

D.2. Information matrices and stage-category sensitivity

TABLE A6. Information at the invitation–application reporting handoff

Group	Application	Decision	b_g	95% interval
English requested, phone	98,790	57,836	0.488	[0.478, 0.499]
English requested, phone absent	14,658	7,480	0.512	[0.483, 0.541]
Other language requested, phone	11,317	7,091	0.373	[0.338, 0.407]
Other language requested, phone absent	1,079	583	0.462	[0.337, 0.587]
Target-loss rank	4; simultaneous 95% confidence set $\{4\}$			

Notes. Application information is $n_g/[a_g(1-a_g)]$. Decision information is $n_g a_g/[b_g(1-b_g)]$. The linked source contains both scores. The release reports a rounded conditional rate on a public grid with width $n_g^{-3/4}$ and is evaluated using the displayed decision information in its Gaussian summary limit. Raw and released conditional rates differ by at most 0.6 percentage point. The approval target is $d_g = a_g b_g$; the missing-score derivative is b_g. The simultaneous intervals use the household bootstrap and exclude zero in each coefficient of the maintained diagonal quotient map.

Let ε_g bound the share of records whose linked stage category can be replaced within group g. The three categories have counts $(n_g - K_g, K_g - R_g, R_g)$ and replacements respect approval requiring application. For $m_g = \lfloor \varepsilon_g n_g \rfloor$, the lower conditional-approval bound replaces up to m_g approvals by denials. The upper bound replaces denials by approvals first and then records without an application by approvals. The same category set gives the application and approval bounds. The group labels stay fixed.

TABLE A7. Sensitivity to one-percent stage-record contamination

Reception group	Replaced	Application	Conditional approval	Approval
English requested, phone	232	61.2–63.2	47.2–50.4	29.4–31.4
English requested, no phone	36	50.0–52.0	49.3–53.1	25.1–27.1
Other language requested, phone	26	61.5–63.5	35.7–38.8	22.3–24.3
Other language requested, no phone	2	53.4–54.9	44.8–47.6	24.3–25.7

Notes. For each reception group, up to one percent of linked records may be reassigned across application and approval categories, subject to approval implying application. Displayed intervals are sharp under this groupwise gross-error model. The analysis file also reports 0.1 and 0.5 percent radii and combines the contamination decrement with the simultaneous sampling lower bounds for conditional approval.

The supplied sampling-plus-contamination calculation subtracts the stated decrement from each simultaneous sampling lower bound. It assesses this fixed-group category model. A different model for group misclassification or linkage would change the uncertainty set.

D.3. Allocation stability across household partitions

TABLE A8. Outcome-separated validation of the retention allocation

Validation quantity	Result
Prespecified household partitions	20
Rotated train–evaluation exercises	40
Median evaluation capacity	25.0%
Median allocated/full risk	1.037
Worst allocated/full risk	1.044
Median uniform/full risk	1.830
Allocated beats uniform	40/40

Notes. Twenty seeds (41021–41040) partition households within lottery draw before outcomes enter the allocation. Each fold estimates the group-audit allocation at 25% capacity and the other fold evaluates its covariance bound; roles are then reversed. Uniform retention is evaluated at the realized held-out capacity. The exercise is a selection-validation check; the main frontier's simultaneous band supplies sampling inference.

The forty rotated comparisons share source households. They assess the stability of constructing the allocation from one subset and evaluating its model criterion on the other. Inference for the frontier continues to use the stated household-resampling procedure.

D.4. Monte Carlo

TABLE A9. Monte Carlo coverage near a frontier kink and a weak rank

Procedure and design	Nominal	Coverage	Mean interval/set size
Numerical delta, smooth budget	0.95	0.932	3.101
Ordinary bootstrap, smooth budget	0.95	0.938	3.112
Numerical delta, first cap	0.95	0.941	3.298
Ordinary bootstrap, first cap	0.95	0.942	3.289
Rank set, separated	0.95	1.000	1.000
Rank set, local fourth direction	0.95	1.000	1.837
Rank set, rank-three boundary	0.95	0.984	1.984

Notes. The design uses 100,000 observations, four group shares and transition probabilities calibrated to the Oregon sample, 1000 Monte Carlo replications, and 499 Gaussian score draws per replication. The budget is the expected share of person links retained. The first-cap budget is the population value at which one group reaches full retention. Frontier size is average interval length; rank size is average confidence-set cardinality.

References

Chen, Qihui and Zheng Fang, "Improved Inference on the Rank of a Matrix," *Quantitative Economics*, 2019, *10* (4), 1787–1824.

Fang, Zheng and Andres Santos, "Inference on Directionally Differentiable Functions," *The Review of Economic Studies*, 2019, *86* (1), 377–412.

Feller, William, *An Introduction to Probability Theory and Its Applications*, 3 ed., Vol. 1, New York: Wiley, 1968.

Firpo, Sergio, Antonio F. Galvao, and Thomas Parker, "Uniform Inference for Value Functions," *Journal of Econometrics*, 2023, *235* (2), 1680–1699.

Hong, Han and Jessie Li, "The Numerical Delta Method," *Journal of Econometrics*, 2018, *206* (2), 379–394.