WORKING PAPERS

N° 1786

OCTOBER 2026

"Hidden Eligibility and the Credibility of Adoption"

Georgy Lukyanov and Mikhail Zolotarev

Toulouse School of Economics

Hidden Eligibility and the Credibility of Adoption

Georgy Lukyanov* Mikhail Zolotarev†

Abstract

Public support for early adoption can affect what later users learn from observed choices. An adopter may have favourable private information, or may have received a rebate. We study a government that randomly assigns eligibility for a fixed rebate and chooses the program's coverage and reporting rule. We characterize when disclosing individual eligibility changes optimal coverage from zero to a positive rate under general convex program costs. This requires both a favourable marginal return to a disclosed program and a bound on the direct benefit available from a large opaque intervention. In the benchmark, disclosure allows the government to increase early adoption and welfare while leaving expected total adoption unchanged. We then allow imperfect compliance and private information for the late user. An eligibility label can be strictly informative without changing anyone's decision, in which case disclosure has no welfare value. The policy reversal survives when both extensions are present. The results identify when an eligibility record is a useful part of the design of an adoption program.

Keywords: social learning; adoption rebates; policy disclosure; early adoption; public program design.

JEL classification: D83; H23; D62.

Governments often support the early adoption of a new technology because its benefits extend beyond those received by the first users. A rebate can encourage households to install an energy-saving device, for example, and thereby produce an additional environmental benefit. Later households may also learn from these installations. An early adopter may have favourable information about the device, but she may instead have received a rebate large enough to make adoption worthwhile despite an unfavourable assessment. The intervention thus affects both the decision to adopt and the information conveyed by that decision.

This creates a question about the administration of the program. Suppose that later households know the rebate amount and the assignment procedure, but do not know which early users were eligible. Announcing the program does not tell them which adoptions occurred without support. An eligibility record identifies these choices and allows later users to interpret them separately. We ask when making this record public changes the government's choice from no intervention to a positive rebate program.

*Toulouse School of Economics, Toulouse, France.

†International College of Economics and Finance, HSE University, Moscow, Russia

We consider two users who decide sequentially whether to adopt. The early user receives a private signal about a common payoff state. The late user observes the early choice and decides whether to adopt himself. Both face the same adoption cost. Before the early user receives her signal, the government randomly assigns eligibility for a fixed rebate. In the benchmark, the payment induces an eligible user to adopt after either signal. The government chooses coverage and commits either to report eligibility or to leave it unobserved. Its objective includes the users' material payoffs and an additional social benefit from early adoption.

An adoption by an ineligible user reveals favourable information. An adoption by an eligible user does not, since the rebate makes her action independent of her signal. Disclosure therefore improves the late user's information at any given coverage rate. Our main result concerns the optimal policy once coverage can also change. We show that zero coverage is uniquely optimal under opacity and positive coverage is optimal under disclosure if and only if two conditions hold. The marginal return at zero must be positive under disclosure and non-positive under opacity. In addition, the best direct return from early adoption must be smaller than the late surplus lost when the program eliminates social learning. The second condition rules out a profitable large opaque program that a comparison of marginal returns would miss.

The welfare effect is not explained by an increase in total adoption. With full compliance, expected total adoption under disclosure is one at every coverage rate. The program creates more early adopters and fewer late adopters. These are different users, so this is a change in the allocation of adoption across cohorts.[1] The extra social benefit from early adoption must compensate for the private-surplus loss and the cost of the program. Some restriction on expansion is also needed: with linear costs and unrestricted coverage, a profitable disclosed program implies that full coverage is profitable under opacity as well. Our characterization allows general increasing, strictly convex costs.

We next consider imperfect compliance. If eligibility rarely changes a user's action, even an eligible adoption remains persuasive. In this case the late user follows every adoption with or without disclosure. The label is strictly informative but has no effect on decisions or welfare. Once compliance exceeds a threshold, an eligible adoption ceases to justify imitation and disclosure becomes valuable. We characterize this threshold and extend the optimal-coverage result. We also allow the late user to receive a private signal. An exact example shows that the policy comparison survives when both imperfect compliance and informative late private information are present.

The paper is related to the literature on learning from observed actions, beginning with Banerjee (1992) and Bikhchandani et al. (1992); see Bikhchandani et al. (2024) for a survey. The closest connection is to interventions that change the actions from which consumers learn. Kircher and Postlewaite (2008) study preferential treatment of informed consumers whose choices guide others, while Bose et al. (2006) study pricing under observational learning. Chen and Papanastasiou (2021) allow an informed seller to seed learning with a fake purchase. Their intervention is strategic and may convey the seller's information. Here the government is uninformed about the state, assignment is random, and the additional early action is real. Chen et al. (2025)

[1] Each user has one opportunity to adopt. The model does not include a decision to wait and adopt later.

study secret price reductions that attract favourable ratings. Their buyers learn from evaluations aggregated by a rating system; our late user observes the adoption choice itself.

Peres et al. (2020) provide a useful comparison. Their sequence contains occasional *revealers*, who disregard earlier actions and follow their own private signals. The revealing probabilities are known, but individual revealer identities are unobserved. These actions supply new information and can break wrong cascades; the paper characterizes the optimal asymptotic learning rate. A fully responsive rebate recipient in our model instead adopts independently of her signal. Disclosure identifies an action that contains no private information, but cannot recover the information that the action has ceased to convey. Our finite model concerns the welfare and scale of an intervention, rather than the long-run rate of learning.

Bénabou and Tirole (2006) and Ali and Bénabou (2020) offer a broader connection between incentives, visibility, and inference. Their agents care about image or societal preferences, whereas ours learn about a common payoff state. A recent working paper by Brandl et al. (2026) studies transfers and disclosure in an infinite sequence of decisions. We focus on a finite program with an additional benefit from early adoption and compare two reporting rules.

The rest of the paper is organized as follows. Section 1 describes the model. Section 2 characterizes adoption and optimal coverage. Sections 3 and 4 consider imperfect compliance and late private information. Section 5 discusses implementation. Proofs and extensions to heterogeneous costs and several early users are in the Appendix.

1 Model

1.1 Users and information

There are two users and two decision dates. The common component of the adoption payoff is a state $\theta \in \{0, 1\}$, with either state equally likely. Both users pay the same cost c if they adopt; their material adoption payoff is $\theta - c$. Non-adoption gives zero.

The early user observes a private signal $S \in \{h, \ell\}$. It is correct with probability q:

$$\mathbb{P}(S = h \mid \theta = 1) = \mathbb{P}(S = \ell \mid \theta = 0) = q.$$

Throughout the benchmark,

$$\frac{1}{2} < c < q < 1. \tag{1}$$

Her posterior is q after the favourable signal and $1 - q$ after the unfavourable one. Thus, without a rebate, she adopts only after h. The late user has no private signal in the benchmark. He observes the early action and the prescribed disclosure before deciding. Section 4 relaxes this information assumption.

The inequality $c > 1/2$ means that adoption is unattractive at the prior. The inequality $c < q$ means that favourable private information is enough to justify it. Equal costs and the absence of payoff spillovers keep the informational mechanism separate from conformity, network effects, or

differences in willingness to pay.

1.2 The rebate and the report

The planner chooses a coverage rate $r \in [0, 1]$. The early user is eligible with probability r, independently of the state and her signal. Let $R = 1$ denote eligibility and $R = 0$ ineligibility. The rebate amount s is fixed throughout the policy comparison and satisfies

$$s > c - (1 - q). \tag{2}$$

It is paid upon adoption. Consequently, an eligible user adopts after either signal. An ineligible user adopts only after h. Writing A for the early action, this rule is

$$A = \begin{cases} 1, & \text{if } S = h \text{ or the user is eligible,} \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

The fixed amount represents an existing rebate whose coverage and reporting arrangements can be changed. The model does not optimize a continuously varying subsidy schedule. Within the class of rebates that induce adoption after either signal, lowering the payment toward the bound in (2) changes behaviour only through any associated reduction in fiscal cost.[2]

The planner commits to one of two reports. Under *opacity*, the late user observes only the adoption decision A. Under *disclosure*, he observes both A and eligibility R. We use subscripts O and T, respectively; the latter refers to tagging eligibility. A tag records the offer, not whether it caused adoption. Some eligible users would have adopted without it.

The state is unknown to the planner. The planner knows the signal precision, adoption costs, benefit of early adoption, and program cost. It commits to coverage and reporting before assignment and before the early signal arrives. The early user then learns her eligibility and signal, chooses her action, and the late user receives the public report. Beliefs are Bayesian at every history with positive probability. The late user does not adopt when indifferent. Off-path beliefs at impossible histories have no effect on the results.

1.3 The planner's objective

An early adoption creates an additional social benefit $e \geq 0$, which is not received by the early user. It can represent an extra period of environmental benefit or another advantage of adoption during the initial rollout. We keep it constant across states. The government values both users'

[2] If the early user adopts when indifferent, the minimum payment in this class is $c - (1 - q)$. Without that tie convention, it is an infimum. A lower payment that never changes the early action cannot improve welfare when transfers are redistribution and implementation costs are non-negative. With heterogeneous responsiveness, jointly choosing the payment and coverage is a different policy problem.

material payoffs and this additional benefit:

$$W_D(r) = \mathbb{E}\big[A(\theta - c + e) + a_L(\theta - c)\big] - C(r), \qquad D \in \{\mathrm{O}, \mathrm{T}\}. \tag{4}$$

Here a_L is the late adoption decision. The rebate is a transfer and cancels in utilitarian welfare. The function C records real administration and implementation costs, together with any excess burden of raising public funds. For the fixed rebate, these costs are identical under the two reports. Disclosure itself is initially costless.

We assume $C(0) = 0$, C is increasing and continuously differentiable, and, for the exact policy characterization, strictly convex on $[0, 1]$. A quadratic cost is used only for an illustration. Convexity may reflect increasing outreach or delivery costs. It is not needed for the information or adoption calculations.

2 Eligibility disclosure and optimal coverage

2.1 What an adoption reveals

Under opacity, an observed adoption comes from either a favourable signal or an eligible user with an unfavourable signal. Bayes' rule gives

$$\mu(r) := \mathbb{P}(\theta = 1 \mid A = 1) = \frac{q + r(1 - q)}{1 + r}. \tag{5}$$

This posterior falls continuously from q at zero coverage to $1/2$ at full coverage. Non-adoption, whenever observed, reveals the unfavourable signal and gives posterior $1 - q$.

Under disclosure, an ineligible adopter reveals the favourable signal. An eligible adopter conveys only the prior: adoption was certain under either signal. The late user therefore follows an ineligible adopter and does not follow an eligible one. Under opacity, he follows an adopter only when $\mu(r) > c$.

Proposition 1. *Let*

$$\bar{r} = \frac{q - c}{q + c - 1} \in (0, 1). \tag{6}$$

Under opacity, the late user follows early adoption if and only if $r < \bar{r}$. Under disclosure, he adopts only after an ineligible adoption. Expected early adoption is $(1 + r)/2$, while expected total adoption is

$$U_{\mathrm{T}}(r) = 1, \qquad U_{\mathrm{O}}(r) = \begin{cases} 1 + r, & r < \bar{r}, \\ (1 + r)/2, & r \geq \bar{r}. \end{cases} \tag{7}$$

Below the threshold, opacity induces imitation after every adoption. Above it, adoption has become too weak a recommendation and imitation stops. The jump in total adoption comes from the late user's discrete choice, not from a discontinuity in beliefs. Welfare will remain continuous at the threshold because the late user is indifferent there.

Disclosure changes which actions are followed. It does not preserve all the information that would have existed without intervention. Let V denote the early action without a rebate: $V = 1$ after h and zero after ℓ. For $0 < r < 1$, the experiments satisfy

$$V \succ_B (A, R) \succ_B A,$$

where $X \succ_B Y$ means that Y can be obtained from X by adding state-independent noise, while the reverse transformation is impossible. One generates the tagged record from V by drawing eligibility independently and applying (3); dropping the tag gives opacity. The strictness argument is in Appendix A. This standard information comparison leaves open whether the additional information changes adoption or optimal coverage.

The adoption identity under disclosure is more informative about the policy's purpose. Increasing coverage by one unit raises expected early adoption by one half and lowers expected late adoption by one half. In the unfavourable-signal histories it creates an early adoption; in the favourable-signal histories it suppresses a late adoption because eligibility conceals the early signal. These changes occur in different histories. Thus, even a policy that leaves expected total adoption unchanged can alter its timing and welfare.

2.2 Welfare

To express welfare under the two reporting rules, define

$$b = \frac{1 - q - c + e}{2}, \qquad h = \frac{q - c}{2}, \qquad k = \frac{q + c - 1}{2}. \tag{8}$$

The coefficient b is the expected direct gain from a unit increase in coverage: only the half of users with an unfavourable signal change their early action. The quantity h is expected late surplus without intervention. Finally, k measures how quickly that surplus falls under opacity while the late user still follows adoption. Notice that $k > h > 0$ and $\bar{r} = h/k$.

At zero coverage, welfare is $W_0 = q - c + e/2$. Under disclosure, increasing coverage gives the direct gain br, but eliminates late surplus in the fraction r of histories with eligibility. Under opacity, late surplus falls until the late user stops adopting, after which it remains zero. Therefore

$$W_{\mathrm{T}}(r) - W_0 = (b - h)r - C(r), \tag{9}$$

$$W_{\mathrm{O}}(r) - W_0 = br - C(r) + [h - kr]_+ - h. \tag{10}$$

The notation $[x]_+$ means $\max\{x, 0\}$. It appears because the late user can always choose not to adopt. These formulas are derived from the users' material payoffs; total adoption is not itself the welfare criterion.

At any fixed interior coverage, disclosure raises welfare by

$$W_{\mathrm{T}}(r) - W_{\mathrm{O}}(r) = \begin{cases} r(c - 1/2), & r < \bar{r}, \\ (1 - r)h, & r \geq \bar{r}. \end{cases} \tag{11}$$

Below the threshold, it prevents the late user from following an eligible adoption whose average payoff is negative. Above the threshold, it preserves his response to an ineligible adoption. The gain follows from the value of information to the late user, whose payoff enters the government's objective. We next ask when it changes the optimal policy.

2.3 When disclosure changes the policy

Write

$$a_{\mathrm{T}} = b - h, \qquad a_{\mathrm{O}} = b - k$$

for the gross marginal returns to coverage at zero under disclosure and opacity. Their difference is $c - 1/2$. To compare global optima, define

$$\Pi(x) = \max_{0 \leq r \leq 1} \{xr - C(r)\}. \tag{12}$$

This is the largest net return from a program with a constant gross marginal benefit x. In particular, $\Pi(b)$ is the best direct return from early adoption when all downstream benefits are set aside.

Theorem 1. *Under the assumptions in Section 1,*

$$\max_r W_{\mathrm{T}}(r) - W_0 = \Pi(a_{\mathrm{T}}), \tag{13}$$

$$\max_r W_{\mathrm{O}}(r) - W_0 = \max\{\Pi(a_{\mathrm{O}}), \Pi(b) - h\}. \tag{14}$$

The opaque optimum is uniquely zero and the disclosed optimum is strictly positive if and only if

$$a_{\mathrm{O}} \leq C'(0) < a_{\mathrm{T}}, \qquad \Pi(b) < h. \tag{15}$$

The disclosed optimum is unique. It satisfies $C'(r_{\mathrm{T}}^*) = a_{\mathrm{T}}$ and lies strictly below one whenever (15) holds.

The first condition compares the marginal returns to starting the program under the two reports. The second concerns its scale. Even the most attractive program based only on the direct early benefit must fail to compensate for the loss of the baseline late surplus h. Otherwise the planner may prefer a large opaque program despite the damage to learning. Hence a negative marginal return at zero is insufficient to establish that zero coverage is optimal.

Theorem 1 also identifies what early adoption must be worth. With equal costs,

$$a_{\mathrm{T}} = \frac{e - (2q - 1)}{2}. \tag{16}$$

The private-surplus loss from replacing late adoption after a favourable signal with early adoption after an unfavourable signal is $(2q - 1)/2$ per unit of coverage. The additional early benefit must cover this loss and the marginal program cost. If $e \leq 2q - 1$, no positive rebate program can improve welfare under either reporting regime when C is non-negative.

Strict convexity ensures that the disclosed optimum is unique. The reason for assuming increasing marginal costs is that they can make further expansion unprofitable. With linear cost $C(r) = \kappa r$, disclosed welfare is linear. If some disclosed program improves on no intervention, then $a_{\mathrm{T}} > \kappa$, and full coverage improves welfare under both reports. Thus a zero-opaque, positive-disclosed comparison requires a force limiting expansion, such as increasing marginal cost, a capacity constraint, or diminishing early benefits. Our theorem treats the first case explicitly.

If the inequalities are strict also at the left of (15), they survive small changes in the parameters. For the quadratic family below, this means an ordinary neighbourhood of (q, c, e, κ, d) in which all maintained inequalities continue to hold. For a general cost function, the analogous statement uses small changes in both the function and its first derivative on $[0, 1]$.

2.4 An illustration and the cost of reporting

Take $q = 0.90$, $c = 0.75$, $e = 1.20$, and

$$C(r) = 0.05r + \frac{1}{2}r^2.$$

All values are expressed in normalized adoption-payoff units. Here $b = 0.275$, $h = 0.075$, and $k = 0.325$. The direct surplus of an induced low-signal early adoption is $1 - q - c + e = 0.55 > 0$. Nevertheless, the optimal opaque policy has zero coverage. Disclosure supports coverage $r^*_{\mathrm{T}} = 0.15$. Indeed, $a_{\mathrm{O}} = -0.05 < 0.05 < a_{\mathrm{T}} = 0.20$, while $\Pi(b) = 0.0253125 < h$.

Table 1: Adoption and welfare under alternative policies

Policy	Early	Late	Total	Welfare
No rebate	.500	.500	1.000	.75000
Opaque, $r = .15$	.575	.575	1.150	.72375
Opaque, $r = .50$	.750	.000	.750	.66250
Disclosed, $r = .15$	.575	.425	1.000	.76125

Table 1 illustrates why adoption counts alone do not determine the welfare ranking. A small opaque program raises total adoption and lowers welfare. A larger opaque program can lower both. The optimal disclosed program raises welfare while leaving expected total adoption unchanged. Figure 1 shows the full comparison, including the threshold at which opaque imitation stops.

Reporting may require a real fixed cost F, for example to create a verifiable eligibility record.

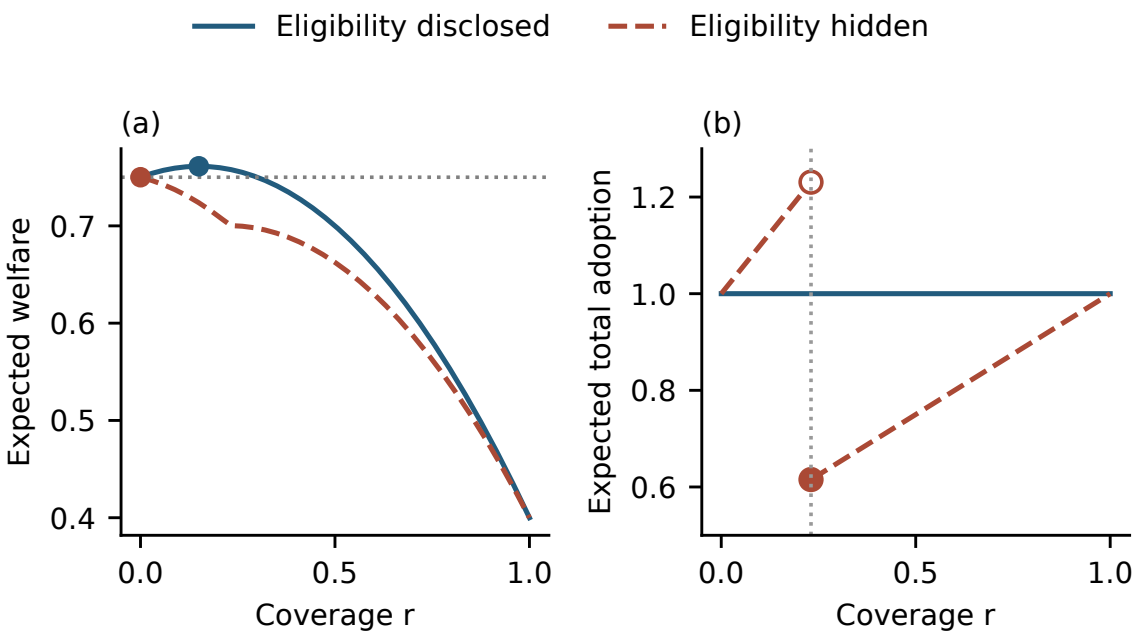

Figure 1: Coverage, welfare, and adoption in the illustrative economy. Solid lines show eligibility disclosure; dashed lines show opacity. The horizontal dotted line in panel (a) is welfare without intervention. The vertical dotted line in panel (b) marks $\bar{r} = 3/13$. At that point the late user chooses non-adoption; the upper point is open

When Theorem 1 applies, the largest cost worth paying is $\Pi(a_{\mathrm{T}})$: this is the improvement of the optimal disclosed program over the optimal opaque policy. In the example it equals 0.01125. A fixed reporting cost below this amount preserves the positive program; a cost above it makes no intervention preferable. This calculation presumes that reporting has no additional effect on behaviour, such as stigma or a privacy loss that differs across users.

3 Imperfect compliance

Suppose now that eligibility does not always induce an early user with an unfavourable signal to adopt. Let $\rho \in [0, 1]$ be that user's adoption probability conditional on eligibility. An eligible user with a favourable signal still adopts, and ineligible users continue to follow their signals. We assume the additional randomization is independent of the state conditional on the signal.

One implementation is random rebate activation. An assigned offer is usable with probability ρ, independently of the state and private signals. The early user observes whether it is usable. A usable offer pays the rebate in (2); otherwise she acts without a rebate. The public label records assignment, while activation is not reported. This produces partial compliance without asking a strictly optimizing user to mix between unequal payoffs. We hold ρ fixed and include the resulting expected real implementation costs in C. The model does not cover compliance that directly

depends on the state after conditioning on the private signal.

Under opacity, the posterior after adoption is $\mu(r\rho)$: the effective probability of inducing a low-signal adoption is $r\rho$. Under disclosure, an ineligible adoption still gives posterior q, but an eligible adoption now gives posterior $\mu(\rho)$. An observed non-adoption gives posterior $1-q$ under either report. Thus the informativeness of a treated action depends on responsiveness to the rebate.

Proposition 2. *For $0 < r < 1$ and $0 < \rho \leq 1$, the no-rebate action strictly Blackwell-dominates the tagged record, which strictly Blackwell-dominates opaque adoption. Nevertheless:*

(i) If $0 \leq \rho < \bar{r}$, the late user follows every adoption under both reports. Adoption and welfare coincide at every coverage rate. At $\rho = \bar{r}$, welfare still coincides, although the tie convention can make adoption differ.

(ii) If $\rho > \bar{r}$, the late user does not follow eligible adoptions under disclosure. The welfare gain from disclosure is strictly positive for $0 < r < 1$ and equals

$$W_{\mathrm{T}}(r) - W_{\mathrm{O}}(r) = \begin{cases} r(\rho k - h), & r\rho < \bar{r}, \\ (1-r)h, & r\rho \geq \bar{r}. \end{cases} \tag{17}$$

Part (i) distinguishes informativeness from decision value. When the rebate changes behaviour infrequently, an eligible adoption remains sufficiently good news. Learning eligibility changes the late user's posterior but not his decision. An informative label then has no value in this particular adoption problem. At $\rho = 0$, the label is itself uninformative about the state and there is no strict information ordering.

Part (ii) gives the condition under which disclosure changes decisions. Once eligible adoption is no longer persuasive, opacity either induces the late user to follow eligible adopters or prevents him from following ineligible ones. Thus the value of an accurate eligibility label depends on how strongly users respond to the rebate.

For $\rho > \bar{r}$, the exact policy comparison is obtained by replacing the direct coefficient b with ρb, the initial opaque slope with $\rho(b-k)$, and the disclosed slope with $\rho b - h$. In particular,

Corollary 1. *If $\rho > \bar{r}$, zero coverage is uniquely optimal under opacity and positive coverage is optimal under disclosure if and only if*

$$\rho(b-k) \leq C'(0) < \rho b - h, \qquad \Pi(\rho b) < h. \tag{18}$$

The marginal interval has width $\rho k - h$ and disappears exactly at the compliance boundary. Full compliance is therefore sufficient but unnecessary for the policy reversal. It is necessary for the benchmark's exact conservation of expected total adoption. When $\rho > \bar{r}$, disclosed total adoption is instead $1 - r(1-\rho)/2$, and hence falls with coverage if compliance is imperfect. The welfare improvement in this case is accompanied by a reduction in expected total adoption.

4 Private information for the late user

We now allow the late user to have private information of his own. Let the late user additionally observe a finite-valued signal Z, independent of the early signal and assignment conditional on θ. Its state-conditional probabilities are $f_1(z)$ and $f_0(z)$. Starting from a public belief m, his expected optimized payoff is

$$J(m) = \sum_z \left[m f_1(z)(1-c) - (1-m) f_0(z) c \right]_+ . \tag{19}$$

For each private-signal realization, the expression inside the brackets is the probability-weighted payoff from adoption. Taking its positive part implements the choice between adoption and non-adoption. In the absence of private information, $J(m) = [m - c]_+$.

Under full compliance, expected late surplus becomes

$$L_{\mathrm{T}}^{W}(r) = \frac{1-r}{2}\{J(q) + J(1-q)\} + rJ(1/2), \tag{20}$$

$$L_{\mathrm{O}}^{W}(r) = \frac{1+r}{2} J(\mu(r)) + \frac{1-r}{2} J(1-q). \tag{21}$$

Unlike the benchmark, the late user may adopt after an eligible action or after non-adoption, depending on his private information. We therefore do not retain the single imitation cutoff or the constant-total-adoption claim in this general extension.

Proposition 3. *At fixed coverage, disclosure weakly raises welfare with a late private signal as specified above. Its gain is*

$$\frac{1-r}{2} J(q) + rJ(1/2) - \frac{1+r}{2} J(\mu(r)) \geq 0. \tag{22}$$

For $0 < r < 1$, the inequality is strict if and only if J is not affine on $[1/2, q]$. If the late signal reveals the state perfectly, the gain is zero.

The intuition is as follows. Information about eligibility is valuable when it changes the action chosen after at least one private-signal realization. If the late user's own information makes the same decision optimal throughout the relevant interval of public beliefs, the label adds no decision value. This is the same distinction that appears in Proposition 2.

The comparison of optimal coverage can also survive a late private signal. Retain the illustrative parameters of Section 2, let compliance be $\rho = 0.9$, and let Z be a symmetric binary signal with precision $p = 0.6$. The global optima are

$$r_{\mathrm{O}}^{*} = 0, \qquad r_{\mathrm{T}}^{*} = \frac{49}{400} = 0.1225, \qquad W_{\mathrm{T}}^{*} - W_0 = \frac{2401}{320000} = 0.007503125. \tag{23}$$

Both optima are unique. Appendix C.1 proves the global comparison by maximizing all four quadratic branches of opaque welfare. The adoption inequalities and the positive gaps between branch maxima are strict, so the example persists under small simultaneous perturbations of its

parameters. A perfectly revealing late signal eliminates the informational reason for the policy difference.

5 Implementation and scope

Consider an illustrative municipal pilot for an energy-saving device. A publicly announced lottery assigns private rebate offers among pilot participants before they receive their private assessments. Later households see installations, but do not see the offer letters or any performance outcomes. The municipality could report installations alone, or record which installations occurred among households that had not been assigned a rebate. The model compares these two records. It presumes that assignment is genuinely random and that the municipality can commit to an accurate report. The example is hypothetical; it specifies the institutional conditions to which the model applies.

With several exchangeable early users, disclosure need not identify individuals. Under independent random assignment and full compliance, the total number assigned and the number of adopters among the unassigned contain the same state information as the complete tagged history. Under partial compliance, a sufficient anonymous report is the full table of adoption and non-adoption by assignment. Appendix E proves these statements. They concern information about the common state, not the causal effect of the rebate on each adopter. Correlated signals, differing signal quality, or informative assignment may require a richer report.

The random-assignment assumption limits the policy interpretation. If a government preferentially supports users with strong private information or unusually high expected gains, eligibility can itself be evidence about quality. If everyone receives the same subsidy, the policy changes an adoption threshold rather than randomly overwriting actions. In general the resulting experiments need not be Blackwell ordered: lowering a common cutoff can make adoption less favourable news while making non-adoption more unfavourable news. The present result applies most directly to randomized pilots, private coupons, and invitation-based rebates with exogenous assignment.

The analysis also separates learning from choices from learning from experience. We observe the decision to adopt, not a review, a realized payoff, or a verifiable performance measure. Where later users can inspect outcomes, subsidizing early adoption may create new information as well as conceal existing private information. This additional benefit is absent here. Similarly, we do not model a strategic government that signals the state through its chosen policy, privately informed assignment, endogenous participation in the pilot, or a seller's price response.

Heterogeneous late adoption costs smooth the sharp response in the benchmark. The planner then knows the distribution of costs, while each late user knows her own realization. Appendix D gives the adoption and welfare formulas and a sufficient condition preserving the policy comparison. The extension shows how dispersion in adoption costs smooths the late response.

6 Conclusion

We have studied a rebate program in which the choices of early users provide information to later decision makers. The government can improve the use of this information by reporting eligibility. Under the conditions we identify, the record changes optimal policy: a positive program becomes worthwhile even though no intervention is optimal when eligibility is hidden.

The result depends on both the benefit of early adoption and the cost of expanding the program. Disclosure does not recover the private information suppressed by the rebate, so the early benefit must compensate for this loss. Increasing marginal costs can then prevent the government from preferring a large opaque intervention. Finally, the record must affect a later decision. With sufficiently low compliance, an eligible adoption remains persuasive and reporting eligibility has no welfare value.

For a government administering a randomized pilot, these results suggest treating the eligibility record as part of the program's design. Its value depends on how users respond to support and what later users can learn independently. The amount of adoption and the quality of later decisions should therefore be considered together when assessing the program.

A Information and the benchmark calculations

Proof of Proposition 1 and the information order. Conditional adoption probabilities are $q + r(1-q)$ in state one and $1 - q + rq$ in state zero. Their sum is $1 + r$, giving (5). For $r < 1$, non-adoption reveals $S = \ell$. Under disclosure, $R = 0, A = 1$ reveals h; $R = 0, A = 0$ reveals ℓ; and $R = 1$ carries the prior because both signals lead to adoption. Solving $\mu(r) > c$ gives $r < \bar{r}$. Expected early adoption is $(1+r)/2$. Expected late adoption is $(1-r)/2$ under disclosure, $(1+r)/2$ under opacity below the threshold, and zero otherwise. Adding the two decisions proves (7).

For the information order, generate eligibility independently of the no-rebate action V and apply (3). This is a state-independent transformation from V to (A, R); dropping R is a transformation to A. Each is strict when $0 < r < 1$. Conditional on a treated adoption, the no-rebate posterior takes the distinct values q and $1-q$ with positive probability. Conditional on an opaque adoption, the tagged posterior takes the distinct values q and $1/2$ with positive probability. The posterior-mean identities for these transformations, followed by strict Jensen inequality for the square function, show a strict decrease in expected squared posteriors at each step. Neither experiment can therefore be recovered from its garbling. □

To verify (9)–(10), expected early surplus is

$$\frac{q - c + e}{2} + \frac{r}{2}(1 - q - c + e).$$

Under disclosure, expected late surplus is $(1-r)(q-c)/2 = (1-r)h$. Under opacity, its value is

the positive part of the probability-weighted adoption payoff:

$$\left[\frac{1+r}{2}\{\mu(r)-c\}\right]_+ = \left[\frac{q-c-r(q+c-1)}{2}\right]_+ = [h-kr]_+ .$$

Subtracting $W_0 = (q-c+e)/2 + h$ proves both formulas. Their difference gives (11); at $r = \bar{r}$, $h = k\bar{r}$, so the two expressions coincide.

B Optimal coverage

Proof of Theorem 1. Equation (10) can be written as the maximum of two expressions:

$$W_{\mathrm{O}}(r) - W_0 = \max\{(b-k)r - C(r),\ br - C(r) - h\}.$$

Maximizing this identity over the compact interval $[0,1]$ gives (14); (13) follows directly from (9).

Strict convexity and differentiability of C imply that $xr - C(r)$ has a unique maximizer. Its maximizer is zero if and only if $x \le C'(0)$, and it is positive if and only if $x > C'(0)$. Thus the first opaque expression is uniquely maximized at zero with value zero precisely when $a_{\mathrm{O}} \le C'(0)$. The second must have a strictly negative maximum for zero to be the unique opaque optimum. If its maximum were positive, zero would not be optimal. If its maximum were zero, its maximizer would be positive because its value at zero is $-h < 0$, creating a second opaque optimum. Hence unique optimality of zero is equivalent to $a_{\mathrm{O}} \le C'(0)$ and $\Pi(b) < h$.

The disclosed optimum is positive exactly when $a_{\mathrm{T}} > C'(0)$. Combining these conditions proves (15). It cannot be one: at full coverage the two information records coincide, so $W_{\mathrm{T}}(1) = W_{\mathrm{O}}(1) < W_0$, whereas the positive disclosed optimum exceeds W_0. The interior first-order condition follows. Finally, maxima of continuous functions on a fixed compact interval vary continuously under uniform changes in those functions. Strict inequalities therefore persist under the perturbations described in the text. □

For $C(r) = \kappa r + dr^2/2$, $d > 0$, the function in (12) is explicitly

$$\Pi(x) = \begin{cases} 0, & x \le \kappa, \\ (x-\kappa)^2/(2d), & \kappa < x < \kappa + d, \\ x - \kappa - d/2, & x \ge \kappa + d. \end{cases}$$

This expression verifies the numerical example without a search over coverage. It also shows why checking only the derivative near zero can miss a profitable large opaque program.

C Partial compliance and private information

Proof of Proposition 2 and Corollary 1. Generate an activation indicator independently, with success probability ρ. Conditional on a low signal, adoption occurs only if both assignment and activation succeed. Opacity therefore gives posterior $\mu(r\rho)$; assignment plus adoption gives $\mu(\rho)$. Non-adoption always reveals the low signal. The same state-independent transformations as in Appendix A give the weak information order. For $r, \rho > 0$, a treated adoption pools both signals, so the first comparison is strict. For $r < 1$, opaque adoption pools an ineligible posterior q with a treated posterior $\mu(\rho) < q$, so the second is strict.

The late user follows eligible adoption if and only if $\rho < \bar{r}$. His expected surplus under disclosure and opacity, respectively, is

$$(1-r)h + r\,[h - \rho k]_+\,, \qquad [h - r\rho k]_+\,.$$

When $\rho \le \bar{r}$, both expressions equal $h - r\rho k$. For the strict inequality $\rho < \bar{r}$, late actions also agree after every adoption; at equality, an eligible adoption is a zero-surplus event and the tie convention can change uptake. When $\rho > \bar{r}$, subtracting the two expressions proves (17) and its strict positivity for interior r.

Expected early surplus equals $(q - c + e)/2 + r\rho b$. In the latter regime, disclosed welfare relative to W_0 is $(\rho b - h)r - C(r)$, while opaque welfare is

$$\max\{\rho(b-k)r - C(r),\ \rho b r - C(r) - h\}.$$

The proof of Theorem 1 applies directly and yields (18). □

Proof of Proposition 3. The function J is convex because each summand in (19) is the positive part of an affine function. Under disclosure the public posterior has values q, $1-q$, and $1/2$, with probabilities $(1-r)/2$, $(1-r)/2$, and r. Under opacity it has values $\mu(r)$ and $1-q$, with probabilities $(1+r)/2$ and $(1-r)/2$. Conditional independence lets us apply the same continuation function J to each public posterior, giving (20)–(21).

The identity

$$\mu(r) = \frac{1-r}{1+r}q + \frac{2r}{1+r}\frac{1}{2}$$

and convexity prove (22). With both weights positive, equality in this chord inequality holds if and only if J is affine on the interval between the two endpoints. If Z reveals θ, then $J(m) = m(1-c)$, which is affine. □

C.1 Exact certificate with both extensions

Take $(q, c, e, \kappa, d, \rho, p) = (.9, .75, 1.2, .05, 1, .9, .6)$. A late favourable private signal after a public low signal gives posterior $1/7 < c$. After an ineligible adoption, even an unfavourable late signal gives posterior $6/7 > c$. After an eligible adoption, even a favourable late signal gives posterior

$297/479 < c$. Thus tagged actions are the same as in the high-compliance case of Section 3:

$$W_{\mathrm{T}}(r) - W_0 = \frac{49}{400}r - \frac{1}{2}r^2.$$

Its unique optimum and value are exactly (23).

Under opacity, an observed non-adoption never induces the late user to adopt, even with a favourable signal. Following adoption, the probability-weighted payoffs for the favourable and unfavourable late signals are, respectively,

$$\frac{21}{400} - \frac{459}{4000}r, \qquad \frac{9}{400} - \frac{711}{4000}r.$$

Consequently,

$$W_{\mathrm{O}}(r) - W_0 = \frac{79}{400}r - \frac{1}{2}r^2 - \frac{3}{40} + \left[\frac{21}{400} - \frac{459}{4000}r\right]_+ + \left[\frac{9}{400} - \frac{711}{4000}r\right]_+. \tag{24}$$

The sum of two positive parts is the maximum over the four subsets of included terms. Maximizing each resulting quadratic on $[0, 1]$ gives the gains

$$-\frac{17759}{320000}, \quad -\frac{610439}{32000000}, \quad -\frac{1673759}{32000000}, \quad 0,$$

where the last case includes both terms and is uniquely maximized at zero. All other branch maxima are strictly negative, proving the unique global opaque optimum. The strict posterior inequalities, negative derivative of the last branch at zero, and strict gaps for the remaining branch maxima persist under small parameter changes. This supplies an open neighbourhood with both imperfect compliance and informative late private information.

D Heterogeneous late costs

Keep the early cost c, but let the late user's cost T be drawn independently of the state, early signal, and assignment from a distribution G with continuous density. Its support is $[\underline{t}, \overline{t}] \subset (1/2, q)$. The planner and the early user know the distribution; only the late user observes her realization before choosing. The planner commits before that draw. Equivalently, the formulas give per-capita outcomes for a unit mass of late users who see the same early action.

Define the expected optimized payoff at public belief m by $J_G(m) = \int [m - t]_+ \, dG(t)$. Under full compliance, expected late surplus is

$$L_{\mathrm{T}}^W(r) = \frac{1-r}{2}(q - \mathbb{E}[T]), \qquad L_{\mathrm{O}}^W(r) = \frac{1}{2}\int [q - t - r(q + t - 1)]_+ \, dG(t).$$

All types decline after non-adoption or an eligible adoption under disclosure. Expected late uptake under opacity is $(1 + r)G(\mu(r))/2$, which varies continuously in coverage. Expected total

uptake is

$$U_{\mathrm{O}}(r) = \frac{1+r}{2}\{1 + G(\mu(r))\}.$$

If G has density g, differentiation at an interior posterior gives

$$U'_{\mathrm{O}}(r) = \frac{1 + G(\mu(r))}{2} - \frac{2q-1}{2(1+r)} g(\mu(r)).$$

Total uptake therefore falls only when enough late users lie near the marginal threshold. Heterogeneity removes the discrete jump; it does not imply that every distribution produces a decline.

For a simple policy condition, write $H = (q - \mathbb{E}[T])/2$, $K = (q + \mathbb{E}[T] - 1)/2$, and $r_{\min} = (q - \bar{t})/(q + \bar{t} - 1)$. Under a quadratic cost, suppose

$$b - K < \kappa < b - H, \qquad \kappa + d r_{\min} > b.$$

Then opaque welfare is strictly decreasing. Before $r_{\min}$, every late type follows adoption and its derivative is $b - K - \kappa - dr < 0$. After $r_{\min}$, the derivative of late surplus is non-positive and $b - \kappa - dr < 0$. Disclosed welfare has a unique positive optimum $(b - H - \kappa)/d < r_{\min}$. In particular, distributions sufficiently concentrated around $c = .75$ preserve this comparison in the illustrative economy: at the degenerate limit, $.05 + 3/13 > .275$, and the remaining inequalities are strict. This establishes persistence for continuous heterogeneity without imposing the quadratic form on the main theorem.

E Several early users and anonymous reporting

Let there be n simultaneous early users with signals that are independent conditional on the common state, each with precision q. Assignment is independent across users and independent of all signals and the state. With full compliance, let M be the number assigned and J the number of adopters among the $n - M$ unassigned. Conditional on $M = m$, the count J is binomial with success probability q in state one and $1 - q$ in state zero. Thus

$$\mathbb{P}(\theta = 1 \mid M = m, J = j) = \frac{q^j (1-q)^{n-m-j}}{q^j (1-q)^{n-m-j} + (1-q)^j q^{n-m-j}}.$$

The likelihood of the full labelled record depends on the state only through these two counts. They therefore contain the same information about θ as the complete record. When total adoption is already public, publishing M is enough to recover J, since every assigned user adopts.

Under partial compliance, let K_1 and K_0 be adopters among the assigned and unassigned. The sufficient anonymous record is (M, K_1, K_0), equivalently the complete assignment-by-adoption table. Conditional adoption probabilities among the assigned are $q + \rho(1-q)$ and $1 - q + \rho q$; among the unassigned they are q and $1 - q$. Conditional independence makes the full likelihood a product of four powers, one for each cell of this table. The individual identities add no state

information under these exchangeability assumptions. Reporting only total assignment and total adoption is no longer generally sufficient.

For completeness, under full compliance the opaque adoption count is a garbling of the no-rebate count: conditional on j no-rebate adopters, it equals $j + B$, where B is binomial with $n - j$ trials and success probability r. This finite-sample information comparison does not establish permanent learning failure. For fixed $r < 1$, the state-conditional adoption rates differ, and an arbitrarily large independent early cohort reveals the state through its empirical adoption rate.

References

Ali SN, Bénabou R (2020) Image versus information: Changing societal norms and optimal privacy. American Economic Journal: Microeconomics 12(3):116–164.

Banerjee AV (1992) A simple model of herd behavior. The Quarterly Journal of Economics 107(3):797–817.

Bénabou R, Tirole J (2006) Incentives and prosocial behavior. American Economic Review 96(5):1652–1678.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100(5):992–1026.

Bikhchandani S, Hirshleifer D, Tamuz O, et al (2024) Information cascades and social learning. Journal of Economic Literature 62(3):1040–1093.

Bose S, Orosel GO, Ottaviani M, et al (2006) Dynamic monopoly pricing and herding. RAND Journal of Economics 37(4):910–928.

Brandl F, Huang W, Jain A (2026) On the Inefficiency of Social Learning. arXiv:2602.08812

Chen CH, Chen KP, Ishida J (2025) Social learning and strategic pricing with rating systems. American Economic Journal: Microeconomics 17(4):147–180.

Chen L, Papanastasiou Y (2021) Seeding the herd: Pricing and welfare effects of social learning manipulation. Management Science 67(11):6734–6750.

Kircher P, Postlewaite A (2008) Strategic firms and endogenous consumer emulation. The Quarterly Journal of Economics 123(2):621–661.

Peres Y, Rácz MZ, Sly A, et al (2020) How fragile are information cascades? The Annals of Applied Probability 30(6):2796–2814.