When Tax Enforcement Changes: Social Learning and Compliance*

Jingnan (Cecilia) Chen† Yixin Chen‡ Zhixin Dai§
Tianqi Wei¶ Su Yang§

This version: 31 July 2026
Latest version

Abstract

Taxpayers rarely observe audit probabilities and must infer changes in enforcement from personal and social experience. We study this process in a laboratory tax-reporting experiment with 568 participants. Each participant faces hidden audit probabilities of 5 percent and 25 percent in randomized order, and we vary peer information across sessions from none to one or two preceding audit outcomes in a sparse network and three in a dense network. Before any peer outcome is transmitted, assignment to either network raises compliance by about 20% relative to no peer information. Once outcomes circulate, compliance in the dense network is approximately twice as responsive to enforcement as without peer information and about 60% more responsive than in the sparse network. This amplification is directional: relative to no peer information, the dense network raises compliance by about 12% after enforcement strengthens but lowers it by about 25% after enforcement weakens. Elicited belief distributions show that broader information reach improves learning about the changed enforcement environment. Higher perceived audit probabilities predict greater subsequent compliance, and the dense network's advantage comes from accumulating more peer signals rather than weighting each signal more heavily. Counterfactual policy exercises show that broader diffusion can reinforce deterrence under strong enforcement but erode it under weak enforcement; a model-based exercise suggests that full disclosure of the audit probability can reduce compliance under both weak and strong enforcement.

Keywords: tax compliance; tax enforcement; social learning; information networks; subjective beliefs.
JEL codes: H26; D83; C91.

†Corresponding author. Department of Economics, University of Exeter Business School, Exeter EX4 4PU, United Kingdom. Email: j.chen2@exeter.ac.uk.
‡Central University of Finance and Economics.
§Renmin University of China.
¶China University of Political Science and Law.
*The experiment was preregistered at the AEA RCT Registry (AEARCTR-0017009).

1 Introduction

Tax enforcement deters evasion only to the extent that taxpayers perceive a meaningful risk of detection. Yet the canonical model of tax evasion treats the audit probability as known (Allingham and Sandmo, 1972), while taxpayers in practice rarely observe the strength of the audit environment directly. Tax administrations may publish aggregate examination statistics and general descriptions of selection procedures, but these do not translate into a precise audit probability for current reporting decisions. In the United Kingdom, HM Revenue & Customs (HMRC) describes a risk-assessment process that reviews tax returns and other information to identify specific compliance risks, but it does not quantify the audit probability implied by that process (HM Revenue & Customs, 2026). The US Internal Revenue Service (IRS) similarly describes computer screening, random selection, and links to related examinations without quantifying the resulting audit probability (Internal Revenue Service, 2026). Taxpayers must consequently infer enforcement risk from incomplete evidence. Such opacity can produce substantial misperceptions of audit risk.[1]

This informational problem becomes especially consequential when enforcement changes. Audit activity varies as budgets and staffing expand or contract, tax administrations redirect resources toward emerging risks, and new data and technology alter audit targeting. Between 2010 and 2018, real appropriations for the IRS fell by about 20%, enforcement staffing by about 30%, and examinations by about 40% (Congressional Budget Office, 2020). In the opposite direction, the United Kingdom's 2025 Spending Review funded data modernization and 5,500 additional compliance staff to expand enforcement capacity (HM Treasury, 2025). Fiscal pressure can also strengthen enforcement incentives: after China's 2005 abolition of agricultural taxes reduced county revenues, county governments largely offset the loss through tougher tax enforcement (Chen, 2017). Such changes can alter effective enforcement even when statutory tax rules remain unchanged. Their behavioral consequences depend on whether taxpayers recognize that the enforcement environment has changed and how quickly their beliefs adjust. The fiscal consequences therefore depend not only on revenue recovered through audits but also on how voluntary reporting responds. Social learning can magnify the revenue gains from stronger enforcement or accelerate the losses when enforcement capacity falls.

Direct experience is sparse for any one taxpayer, making the audits of neighbours, clients, colleagues, or other social contacts a potentially important source of evidence. Field studies document precisely such spillovers. Enforcement communication to one household affects nearby households (Drago et al., 2020), while an audit of one client can alter the reporting of other clients who share the same tax preparer (Boning et al.,

[1]Bérgolo et al. (2018) find that surveyed Uruguayan firms reported an average three-year audit probability of about 40%, compared with an administrative rate of approximately 8%.

2020). These studies establish that enforcement can propagate beyond the directly treated taxpayer. They cannot, however, isolate how information reach changes learning because the relevant networks and beliefs are inferred.

We provide the first causal evidence that the reach of peer audit information shapes perceived audit probabilities and tax compliance. Our belief elicitations offer direct evidence that learning about enforcement underlies the compliance response. In a laboratory tax-reporting experiment, 568 participants make repeated declaration decisions in two audit probability blocks with undisclosed rates: 5 percent and 25 percent. Within each block, the same hidden audit probability applies to all participants, allowing us to isolate learning about a common enforcement environment. Every participant experiences both environments, with their order randomized at the beginning of the session. We assign sessions to one of three information environments and randomly match participants into groups of four. Participants in the condition without peer information receive no peer audit outcomes. The sparse network is a chain, and participants observe whether one or two directly connected neighbours were audited in the previous round. The dense network is complete, and participants receive the same information for all three other group members. We reassign groups and network positions every round, so the treatments vary the reach of audit information without creating persistent social relationships or repeated-group incentives. In every round, we also elicit an incentivized subjective distribution over the audit probability.

Our design separates two effects. First, the initial round of each block contains no peer signal. Treatment comparisons in these rounds identify a *network presence effect*: the effect of knowing that one is embedded in an information network before it transmits audit experience. Second, from round 2 onward, the treatments vary the amount of social evidence available for learning. Our primary measure of information transmission is *enforcement responsiveness*: the difference in compliance between the 25 percent and 5 percent audit environments. Comparing this compliance difference across information conditions shows whether broader social information makes declarations track the enforcement environment more closely. We then examine enforcement increases and decreases separately to show whether greater responsiveness reflects higher compliance under strong enforcement, lower compliance under weak enforcement, or both.

We report four main findings. First, before any peer outcome is transmitted, assignment to either network raises compliance by about 11 percentage points, approximately 20% relative to the no peer information benchmark, and raises the perceived audit probability by about 3 points. The similar effects in the sparse and dense networks show that this initial response reflects network embeddedness rather than the amount of information subsequently transmitted.

Second, once audit experiences circulate, enforcement responsiveness increases with network density. The dense network raises responsiveness by 8.1 percentage points relative to no peer information—approximately doubling the response—and by 6.0 points relative to the sparse network. The sparse-network estimate lies between the two information conditions, and an ordered specification yields a statistically significant increase of 2.7 percentage points in responsiveness per expected peer signal. The behavioral response therefore strengthens with information reach and is strongest in the dense network.

Third, examining enforcement decreases and increases separately shows what greater responsiveness means for compliance. After enforcement decreases from 25 percent to 5 percent, the dense network reduces compliance by 10.7 percentage points relative to no peer information and by 13.2 points relative to the sparse network. The latter difference represents a 29% reduction and appears in the early rounds and persists through the post-change block. After enforcement increases from 5 percent to 25 percent, the dense network raises compliance by 5.2 points relative to no peer information. Social diffusion therefore does not mechanically raise compliance: it can reinforce deterrence when enforcement strengthens and accelerate its erosion when enforcement weakens.

Fourth, the elicited belief distributions reveal the learning mechanism. After enforcement changes, peer information shifts probability mass toward the true audit rate and produces more concentrated beliefs, with the strongest accuracy gains in the dense network after enforcement weakens. A 10-percentage-point increase in the prior perceived audit probability predicts 4.3 percentage points higher compliance in the next declaration, about 9% of mean compliance. Participants respond only about half as strongly to any single peer outcome as to their own audit experience, while the response to each peer outcome is similar in the sparse and dense networks. The dense network's advantage therefore comes from accumulating more individually discounted signals rather than making each signal more influential. The overall pattern supports graded information accumulation, with the strongest behavioral effects in the dense network.

These findings make three contributions. First, we connect the public finance literature on uncertain enforcement and post-audit behavior (see, e.g., Alm et al., 1992, 2009; Snow and Warren, 2005; Choo et al., 2016; Mittone et al., 2017; Kasper and Rablen, 2023) to field evidence on tax-compliance spillovers. The experiment reveals the learning process that field evidence cannot directly observe: audit experiences change perceived audit risk, and the reach of those experiences governs how compliance adjusts when enforcement changes. We contribute to experimental tax compliance by keeping the audit probability hidden and repeatedly measuring each participant's subjective distribution over possible audit probabilities. Second, we contribute to the network and belief-updating literatures. Existing tax experiments show that peers' reporting choices and social context can change compliance (see, e.g., Fortin et al., 2007; Alm et al., 2017; Lefebvre et al., 2015); we transmit only audit outcomes and randomize how many of those outcomes reach a participant. This isolates social learning about a common hidden state and shows how several individually discounted signals can accumulate into a substantial behavioral response. Third, we show that the policy value of information diffusion depends on the enforcement environment. Networks do not simply raise compliance. They make behavior more responsive to enforcement, reinforcing deterrence when enforcement is strong but potentially weakening it when enforcement is low.

We further examine these policy implications through two counterfactual exercises. The first combines experimentally observed information and enforcement cells while keeping the audit probability hidden. It shows that broader diffusion is most useful under strong enforcement and can reduce compliance under weak enforcement. The ranking favors adapting information reach to the enforcement environment, but does not show that maximum reach is always required. The second models the effect of exact disclosure, which was not experimentally varied, using the estimated relationship between perceived audit probabilities and declarations. Participants substantially overestimate both audit probabilities. Replacing those beliefs with the true rate is predicted to reduce compliance by 11.1 percentage points (or 27% of observed compliance) in the 5 percent environment and by 5.5 points in the 25 percent environment. Together, the counterfactuals distinguish publicizing genuine enforcement activity from revealing its exact probability.

The remainder of the paper proceeds as follows. Section 2 relates the paper to the existing literature. Section 3 describes the experiment. Section 4 presents the conceptual framework and predictions. Section 5 reports the results. Section 6 presents the policy counterfactuals, and Section 7 concludes.

2 Related Literature

Our study connects three literatures that have largely developed separately: laboratory research on tax enforcement, evidence on social interactions and enforcement spillovers, and research on belief updating from private and social signals. We use the experiment to connect these literatures around a specific public-finance question: how taxpayers learn about a change in enforcement when the audit probability is hidden.

2.1 Experimental tax compliance and uncertain enforcement

Following the deterrence framework of Allingham and Sandmo (1972), a large experimental literature examines how declarations respond to the audit probability, fine rate, tax rate, public-good return, and institutional environment. Reviews and quantitative syntheses show that audit risk and penalties generally increase compliance, but that effect sizes vary across behavioral margins and experimental designs and that expected monetary payoffs alone do not explain observed reporting (see, e.g., Andreoni et al., 1998; Blackwell, 2010; Alm, 2019; Slemrod, 2019). This literature establishes enforcement as a central policy variable and experiments as a useful method for identifying behavioral responses to it.

The standard tax evasion game nevertheless makes enforcement more transparent than it is for many taxpayers: participants are typically told the audit probability and choose declarations under a known lottery (for a meta-analysis, see Alm and Malézieux, 2021; for a broader review, see Alm and Kasper, 2021). A smaller literature relaxes this assumption. Alm et al. (1992) experimentally vary institutional uncertainty, while Snow and Warren (2005) show theoretically that ambiguity about the audit probability can increase or reduce compliance depending on taxpayers' ambiguity preferences. Choo et al. (2016) directly compare compliance under disclosed and undisclosed audit probabilities within a tax evasion experiment. They find little average response to whether the audit probability is disclosed and do not elicit participants' beliefs. This evidence leaves open how taxpayers learn about a hidden audit probability and how such learning affects declarations.

A related dynamic literature shows that audit sequences and recent audit experience affect later declarations through misperceptions of audit risk (Kastlunger et al., 2009; Mittone et al., 2017) and through learning when audit selection responds to reported income (Kasper and Rablen, 2023). Most closely related to our information treatment, Alm et al. (2009) experimentally compare official reports of realized audit activity with informal communication among taxpayers. They show that audit information changes compliance and that its effect depends on whether taxpayers already know the audit probability. An audit is therefore not only a sanction; it is also information from which taxpayers may learn about future enforcement.

We build directly on this work by asking how the social diffusion of audit experiences changes learning about enforcement. Earlier experiments show that communicated audit information affects compliance, but leave both the learning process and the role of information reach within it unresolved. We keep the audit probability hidden, vary the number of observable peer audit outcomes, and elicit a complete subjective distribution in every round. We thereby identify how social information determines the speed and extent of compliance adjustment when an opaque enforcement environment changes.

2.2 Social interactions, networks, and enforcement spillovers

Tax compliance is socially embedded. Laboratory evidence shows that group tax conditions can generate fairness responses (Fortin et al., 2007), examples of low compliance can encourage further evasion (Lefebvre et al., 2015), and information about neighbours' filing, reporting, and audit outcomes can change compliance (Alm et al., 2017). Field evidence on public tax-return disclosure further shows that searches cluster within household and workplace networks, although being searched does not measurably increase reported income (Reck et al., 2022). In these settings, taxpayers typically observe what others report or pay. The resulting responses may therefore combine information, imitation, social comparison, and norms rather than isolate learning about the enforcement environment.

Our information treatment isolates enforcement learning in two ways. First, participants never observe peer declarations, income, earnings, or identities; they observe only whether peers were audited. Second, we randomize the reach of that information. Models show that network topology can shape tax beliefs and compliance dynamics (see, e.g., Hashimzade et al., 2014; Korobow et al., 2007; Di Gioacchino and Fichera, 2020), but their simulated outcomes jointly reflect assumptions about both network structure and behavior. Our design holds the tax decision fixed and changes only how many peer audit outcomes reach each participant.

The closest policy motivation comes from field spillovers. Drago et al. (2020) shows that an enforcement intervention can affect untreated households connected through neighbourhood exposure. Boning et al. (2020) finds that audits propagate across firms that share a tax preparer. These studies provide credible evidence that enforcement effects travel beyond directly treated taxpayers. The network is necessarily inferred from geography or a common intermediary, however, and perceived audit probabilities are not observed repeatedly. Consequently, the field evidence cannot determine whether spillovers reflect learning about enforcement, direct communication, preparer behavior, norms, or another channel. By randomizing information reach and observing both subjective distributions and declarations, we isolate the learning mechanism behind these spillovers.

2.3 Belief updating and social learning

Our mechanism relates to research on how individuals process noisy information. Bayes' rule provides the natural benchmark: a person with an unknown audit probability should combine a prior with the likelihood of observed audit outcomes (Snow and Warren, 2007). Grether (1980) documents base-rate neglect, with subjects placing too little weight on prior probabilities. Reviewing a broader experimental literature, Benjamin (2019) concludes that people often underweight both prior information and new signals. Bordalo et al. (2026) show that selective attention to salient features can generate both under- and overreaction and make judgments unstable across equivalent problem descriptions. These departures matter here because audits and non-audits are noisy signals: a taxpayer can easily experience no audit under either enforcement state, and informative outcomes need not receive their full Bayesian weight. When information travels socially, individual updating rules interact with network structure and can slow or distort aggregate learning (Möbius and Rosenblat, 2014; Chandrasekhar et al., 2020; Khandelwal, 2024).

The social learning literature typically studies how people infer a common state from others' actions or beliefs. Observational learning experiments find that participants often place too much weight on private information relative to social evidence (see, for example, Anderson and Holt, 1997; Weizsäcker, 2010; De Filippis et al., 2022). Because another person's choice reflects both the information they received and how they interpreted it, these designs cannot always distinguish imperfect inference about what others know from differential weighting of the information itself. Our experiment transmits peers' underlying audit outcomes directly and therefore separates these two channels.

The most closely related evidence comes from Conlon et al. (2022), who compare sensitivity to equally relevant information discovered by oneself and by others. They find weaker responses to others' information even when it is perfectly communicated or directly observed. Conlon et al. (2021) show that the size of this asymmetry depends on the relationship between the information sources. Together, these studies demonstrate that weaker responses to others' information cannot generally be reduced to communication failure or objective signal quality. This comparison motivates our analysis of how participants respond to their own and their peers' audit outcomes.

Our design adds an objectively defined but hidden enforcement state, exogenous variation in the number of social signals, and repeated elicitation of full belief distributions. These features allow us to distinguish greater responsiveness to each peer outcome from the cumulative effect of receiving more outcomes. We thereby connect the social learning literature to tax enforcement and test how information reach shapes belief and compliance adjustment when an opaque enforcement environment changes.

3 Experimental Design and Procedure

3.1 Information environments

We implement the no peer information, sparse network, and dense network conditions at the session level. In every round, the software randomly rematches participants into groups of four and assigns new network positions. We remind participants that both their group members and positions change. A participant's position determines only the number of preceding peer audit outcomes that can be observed.

In the no peer information condition, participants receive no information about other group members' audit outcomes. The sparse network is a chain: four positions lie on a line, so endpoint positions observe one neighbour and interior positions observe two. The dense network is complete: each participant observes all three other group members.

Beginning in round 2 of each block, the interface presents the participant's own preceding-round audit outcome and the preceding-round audit outcomes of their newly assigned neighbours simultaneously. A linked neighbour's node appears red if audited and green if not audited. A grey node indicates that no audit information is available because the position is not linked or no preceding-round outcome exists. Thus, all peer nodes are grey in round 1. From round 2 onward, the sparse network displays one or two red or green peer nodes, while the dense network displays three. The own and peer signals concern the same hidden audit probability, have the same audited-or-not-audited form, and neither source is observed before the other. Figure 1 illustrates the three environments from the perspective of a focal participant. Participants never observe peer declarations, income, payoffs, or identities. We therefore vary the density of the information network and the number of signals about the common enforcement environment while holding the declaration task fixed.

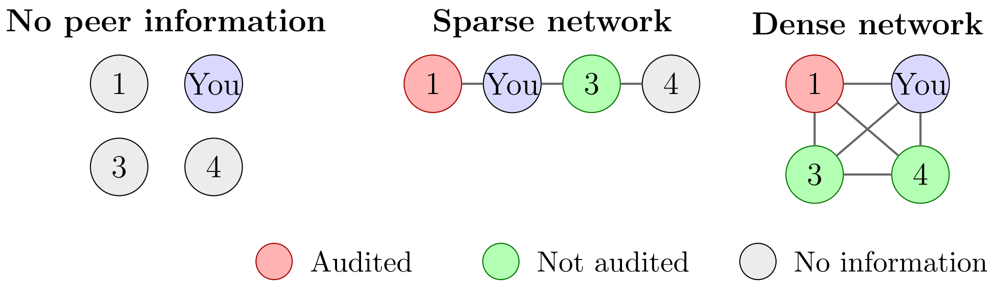

Figure 1: Experimental information environments

Notes: The diagrams show one possible position of the focal participant and an illustrative realization of preceding-round audit outcomes. Red indicates that an observed neighbour was audited, green that an observed neighbour was not audited, and grey that no audit information is available. Groups and positions are randomly reassigned every round. The no-peer-information condition provides no peer audit outcomes; the sparse network provides one or two; the dense network provides three. The sparse network is a minimally connected chain, whereas the dense network is complete.

3.2 Tax-reporting task and enforcement change

The experiment contains two consecutive ten-round blocks: an *initial block* (Part A) and a *post-change block* (Part B), which begins when the hidden audit probability changes. In each round, a participant receives an income drawn between 2,000 and 3,500 experimental currency units (ECU) and chooses declared income $D \in [0, Y]$. Declared income is taxed at rate $t = 0.25$. If the participant is not audited, earnings are

$$\pi^N = Y - tD. \tag{1}$$

If audited, undeclared income is detected and the participant pays the evaded tax plus a penalty equal to the evaded tax:

$$\pi^A = Y - tD - 2t(Y - D). \tag{2}$$

Our principal outcome is the compliance rate $C = D/Y$.

Every participant makes ten decisions under a 5 percent audit probability and ten under a 25 percent audit probability. The audit probability is constant throughout the initial block and then changes to the other probability at the start of the post-change block. At the beginning of each session, the software randomly determines the order, so some sessions move from 5 percent to 25 percent and others from 25 percent to 5 percent. Participants are told at the start of each block that its audit probability will remain fixed for ten rounds, but they are never told the probability itself. Before the post-change block begins, they are explicitly informed that the audit probability has changed and will again remain fixed for the next ten rounds.

The transition also resets the information history. Audit outcomes from the initial block are not displayed in the post-change block, so the first post-change round occurs after the enforcement change but before any outcome generated under the new environment can be transmitted. The first round of the initial block is likewise played before any peer outcome exists. We therefore refer to round 1 of each block as a *signal-free first round*. From round 2 onward, participants can observe the preceding-round outcomes generated within that block according to their assigned information environment.

3.3 Perceived audit probabilities and norm elicitation

In every round, we elicit participants' subjective distribution of the audit probability using the method proposed by Harrison et al. (2017). Participants allocate 100 tokens across ten mutually exclusive intervals: $[0, 10), [10, 20), \ldots, [90, 100]$. Their belief-elicitation payoff depends on the tokens assigned to the interval containing the true audit probability and on the allocation across all ten intervals.[2] Appendix Figure A1 reproduces the participant-facing elicitation interface in English translation. The interface makes the task intuitive: as participants move tokens across intervals, the number above each interval updates to show the payoff they would receive if the true audit probability fell in that interval. Participants can therefore see the payoff consequences of any allocation without calculating the scoring rule themselves.

The distributional elicitation is particularly useful for our design. Let π_{itk} denote the share of tokens that participant i assigns to interval k in round t, and let m_k denote that interval's midpoint. We define the participant's *perceived audit probability* as $\widehat{p}_{it} \equiv \sum_{k=1}^{10} \pi_{itk} m_k$. Two participants can report the same $\widehat{p}_{it}$ while differing substantially in confidence, and social information may change the location of the subjective distribution, its concentration around the true rate, or both. A point report cannot distinguish these responses. The full distribution allows us to measure each margin and also mitigates a familiar incentive problem in point-probability elicitation: under subjective expected utility, Harrison et al. (2017) show that empirically plausible risk aversion produces only small distortions in important features of the reported distribution, without requiring separate calibration for risk attitudes. We therefore summarize the elicited dis-

[2] We implement the quadratic scoring rule with payoff parameters $\alpha = \beta = 500$ ECU.

tributions using $\widehat{p}_{it}$, probability mass on the interval containing the true rate, and belief concentration measured by normalized entropy.

At the end of the initial block and again at the end of the post-change block, we conduct an incentivized elicitation of perceived social norms following Krupka and Weber (2013). Participants assess the social appropriateness of four alternative declaration choices in the enforcement environment they have just experienced. For each block, one declaration choice is randomly selected for payment, and a participant earns 500 ECU if their rating matches the modal rating among participants in the session. These data allow us to examine whether the information conditions change perceived norms, as distinct from beliefs about enforcement.

3.4 Experimental procedure

We conducted the experiment in the Behavioral and Experimental Economics Laboratory at Renmin University of China using oTree (Chen et al., 2016). We recruited 568 undergraduate and graduate students from a range of academic disciplines and ran 26 sessions.[3] Participants were 20.1 years old on average, and 66.9% were female.

Before data collection, we used balanced randomization to assign sessions to the three information conditions. Table 1 reports the realized allocation of sessions and participants. At the beginning of each session, oTree randomly determined whether participants encountered the high- or low-audit environment first. The sample is broadly balanced across the three information environments (see Appendix Table A1).

Table 1: Treatment assignment and sample composition

Information condition	Sessions	Participants	Peer audit outcomes
No peer information	9	196	0
Sparse network	8	184	1 or 2
Dense network	9	188	3
Total	26	568	

Notes: Information conditions were assigned at the session level. Peer audit outcomes refer to outcomes from the previous round that are observable from round 2 of each block. Participants at the endpoints of the sparse network observed one peer outcome, while those at interior positions observed two; participants in the dense network observed all three peer outcomes. Groups and network positions were randomly reassigned in every round.

Upon arrival, participants received a CNY 10 show-up fee and were randomly seated at individual computer terminals. They first read the instructions for the initial block, which explained the tax-reporting task, their assigned information environment, and the

[3] Data collection took place in two waves: 20 sessions with 420 participants in December 2025 and six sessions with 148 participants in March 2026.

belief-elicitation procedure. Participants then completed comprehension questions and an explicitly identified practice round using artificial data. The practice round familiarized them with the interface and payoff calculations and did not enter the analysis.

Figure 2 summarizes the session timeline and the sequence within each experimental round. Participants completed ten rounds and the norm elicitation in each block. Before the post-change block, they read new instructions explaining that the audit probability had changed but would remain fixed throughout the block. The task and interface otherwise remained unchanged, and no audit outcome carried over from the initial block. After both blocks, participants completed a questionnaire covering risk and social preferences, tax morale, a three-item cognitive reflection task, two probability-reasoning questions, and demographic characteristics.

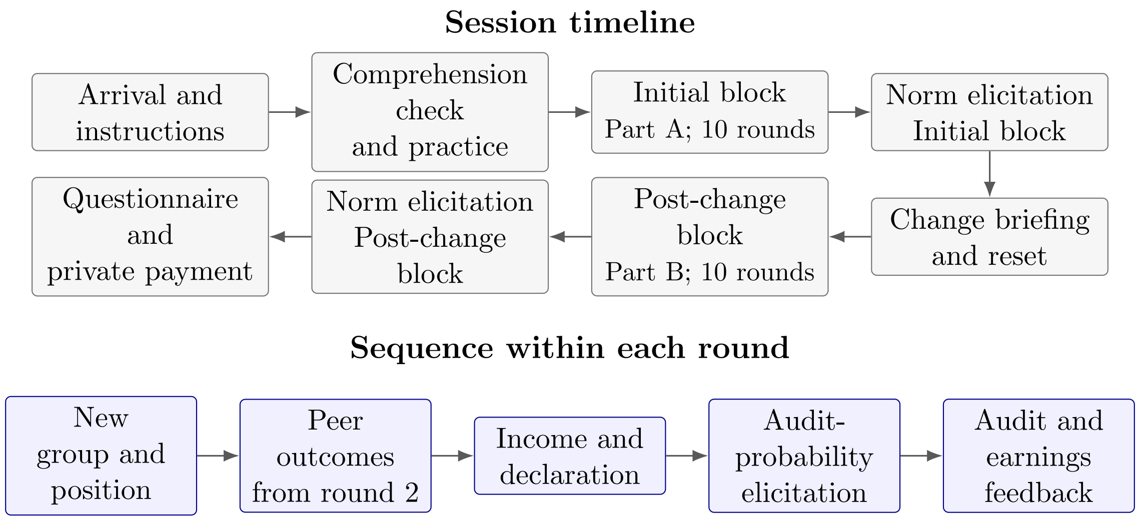

Figure 2: Experimental session and within-round sequence

Notes: Each block contains ten experimental rounds. Round 1 of each block is signal-free; peer audit outcomes are displayed from round 2 onward according to the assigned information condition. The post-change block begins with a refreshed information history.

In addition to the show-up fee, two rounds from each block were randomly selected for payment. Final earnings combined after-tax income and belief-elicitation payments from these four rounds, any norm-elicitation bonuses earned in the initial and post-change blocks, and 200 ECU for each correct answer in the three-item cognitive reflection task. We converted experimental currency units at a rate of 200 ECU to CNY 1 and paid participants privately in cash. Average earnings, including the show-up fee, were CNY 74.02. A session lasted approximately 70 minutes, including instructions, the experimental tasks, the questionnaire, and payment.

4 Conceptual Framework

We use a simple conceptual framework to guide the empirical analysis and clarify how the reach of audit experience shapes perceived audit probabilities and declarations. The framework distinguishes an initial response to visible network presence from the state-dependent response that emerges once audit outcomes circulate. These two channels yield testable implications for the signal-free network-presence effect, enforcement responsiveness as information reach expands, and the link from perceived audit probabilities to subsequent compliance.

4.1 Compliance and learning under hidden enforcement

Let the true audit probability be $p \in \{p_L, p_H\}$, where $p_L = .05$ and $p_H = .25$. Participant i begins round t with prior belief $\mu_{it} = \Pr_i(p = p_H)$. In this two-state representation, the corresponding perceived audit probability is $\widehat{p}_{it} = \mu_{it} p_H + (1 - \mu_{it}) p_L$. Given the tax rate, fine, income, and preferences, optimal compliance is increasing in $\widehat{p}_{it}$:

$$C^*_{it} = c(\widehat{p}_{it}), \qquad c'(\cdot) > 0, \tag{3}$$

The monotonicity is the standard deterrence implication of the tax-evasion model (Allingham and Sandmo, 1972). It yields a direct empirical link: the perceived audit probability formed from signals in round t should predict compliance in round $t + 1$.

Each audit outcome is a Bernoulli signal generated by the common state. Let a_{it-1} denote the participant's own preceding audit outcome, let $\mathcal{N}_{it}$ denote the set of observed peer outcomes, and define posterior log odds as $\ell_{it} = \log[\mu_{it}/(1 - \mu_{it})]$. Allowing own and peer outcomes to receive different weights, beliefs evolve according to

$$\ell_{it} = \ell_{i,t-1} + \omega_o \log LR(a_{it-1}) + \omega_s \sum_{j \in \mathcal{N}_{it}} \log LR(a_{jt-1}), \tag{4}$$

where $LR(a)$ is the likelihood ratio of outcome a under high rather than low enforcement, $\omega_o \geq 0$ is the weight on own experience, and $\omega_s \geq 0$ the weight on each peer outcome. Standard Bayesian updating is the nested benchmark $\omega_o = \omega_s = 1$. Experimental evidence commonly finds conservative and heterogeneous updating (Benjamin, 2019), while social-learning experiments show that individuals often place different weights on their own and others' information (Weizsäcker, 2010; De Filippis et al., 2022; Conlon et al., 2022). We therefore do not impose equality of the two weights. The information-reach prediction in Section 4.3 requires only that peer outcomes receive positive weight.

4.2 Network presence before signals

Equation (4) also clarifies the signal-free first round. Participants have not yet observed an audit outcome generated under the current enforcement environment, so updating cannot produce a density difference. Assignment to an audit-information network is nevertheless visible. Network embeddedness may make audits and the future transmission of audit outcomes more salient, raising perceived enforcement risk before any outcome is observed. Audit-related cues and experiences can increase subjective audit risk and compliance (Spicer and Hero, 1985; Bérgolo et al., 2018, 2023). If embeddedness, rather than the number of outcomes later transmitted, drives this salience response, the sparse and dense networks should have similar first-round effects. Equation (3) then maps higher perceived risk into higher initial compliance.

Prediction 1 (Network presence) *Before any peer audit outcome is transmitted, assignment to either audit-information network should raise perceived audit risk and compliance relative to no peer information. Because no peer outcome has yet been transmitted, this initial response should not increase with network density.*

4.3 Information reach and enforcement responsiveness

We designed the experiment to study this learning problem around a change in enforcement. As described in Section 3.2, each session experiences both hidden audit probabilities in randomized order. The post-change block begins when the probability switches to the other level and the displayed audit history is reset, so participants initially have no outcomes generated under the new environment. From round 2 onward, they can learn whether enforcement has strengthened or weakened from their own and their peers' audit outcomes. The dense network supplies three peer likelihood terms in equation (4), the sparse network one or two, and the no-peer-information condition none.

For a conditionally independent audit outcome, the expected log-likelihood ratio satisfies

$$\mathbb{E}[\log LR(a) \mid p_H] > 0, \qquad \mathbb{E}[\log LR(a) \mid p_L] < 0. \tag{5}$$

Thus, when $\omega_s > 0$, each additional peer outcome raises expected posterior log odds under high enforcement and lowers them under low enforcement. More peer signals therefore increase the expected separation of perceived audit probabilities across the two environments. Because compliance is increasing in $\widehat{p}_{it}$, this separation translates into greater enforcement responsiveness—a larger compliance difference across the two environments. Underweighting reduces the contribution of each peer outcome but does not change its expected direction, so several weakly weighted signals can still have a meaningful combined effect.

We measure enforcement responsiveness as compliance in the 25 percent audit environment minus compliance in the 5 percent audit environment. This comparison captures whether information reach sharpens behavioral differentiation between stronger and weaker enforcement without treating the experiment's 25 percent rate as a universal threshold for "high" enforcement. We also report the state-specific treatment effects to establish the direction of adjustment.

Prediction 2 (Network density and enforcement responsiveness) *As the number of observable peer audit outcomes increases, perceived audit probabilities and compliance should become more responsive to the true enforcement environment. Enforcement responsiveness should therefore increase from no peer information to the sparse network and then the dense network.*

To unpack this responsiveness prediction, we consider separately how peer information shapes adjustment after enforcement increases and after it decreases. When enforcement increases from 5 percent to 25 percent, subsequent audit outcomes are drawn from the stronger enforcement environment. Participants connected to peers observe additional draws from that environment and, on average, receive more evidence that audit risk has risen (equations (4) and (5)). Peer information therefore raises perceived audit risk and compliance (equation (3)), with a larger upward informational response as the number of observable peer outcomes increases. When enforcement decreases from 25 percent to 5 percent, the logic reverses. Peer outcomes are drawn from the weaker environment and provide evidence that audit risk has fallen, producing a downward informational response that also becomes larger with information reach.

The total network effect additionally includes the compliance response to network presence established in Prediction 1. Let $P_n \geq 0$ denote any part of this response that persists once signals begin, and let $L_n^{L\to H} > 0$ and $L_n^{H\to L} > 0$ denote the magnitudes of learning following a change from the low (L) to the high (H) enforcement environment and vice versa in network n. The total effects relative to no peer information are

$$\tau_n^{L\to H} = P_n + L_n^{L\to H}, \qquad \tau_n^{H\to L} = P_n - L_n^{H\to L}. \tag{6}$$

For these comparative predictions, we treat any persistent presence response as common across the sparse and dense networks, following the signal-free logic in Prediction 1. If the initial response fully dissipates, then $P_n = 0$ and equation (6) reduces to the pure learning case. Equation (4) implies that both learning responses increase with the number of peer outcomes. Network presence and learning therefore reinforce each other following an enforcement increase but oppose each other following a decrease.

Corollary 1a (Enforcement increase) *Following an increase from 5 percent to 25 percent, the total network effect on compliance should be positive and become more positive as information reach expands.*

Corollary 1b (Enforcement decrease) *Following a decrease from 25 percent to 5 percent, the total network effect should become less positive, or more negative, as information reach expands. Its sign depends on the relative magnitudes of the positive network-presence response and the negative learning response: the total effect remains positive when presence dominates, is close to zero when the two forces offset, and becomes negative when learning dominates.*

4.4 Perceived audit probability and subsequent compliance

The updating process in equation (4) yields a testable behavioral implication: the perceived audit probability formed from recent audit outcomes should carry forward into subsequent reporting decisions. A participant with a higher $\widehat{p}_{it}$ faces stronger expected deterrence and should therefore declare a larger share of income in the next round, as summarized by equation (3). This relationship should remain visible among participants facing the same enforcement environment and comparable recent audit information.

Prediction 3 (Perceived audit probability and subsequent compliance) *Holding the realized enforcement environment and recent audit signals fixed, a higher prior $\widehat{p}_{it}$ should predict greater compliance in the next round.*

5 Results

We organize the results around the predictions in Section 4. We first isolate the effect of network presence before any peer audit outcome is transmitted (Section 5.1). We then test whether broader information reach increases compliance responsiveness to the enforcement environment (Section 5.2) and use the post-change block to show how this response differs when enforcement strengthens and when it weakens (Section 5.3). Finally, we examine the mechanism by asking whether peer information improves perceptions of the audit probability and whether those perceptions predict subsequent compliance (Section 5.4). We also use the norm elicitation as a brief check on an alternative channel. Where tables report specifications with and without controls, we discuss the estimates with controls and report the treatment-only estimates as benchmarks.

5.1 Network presence before information transmission

To test Prediction 1, we pool the signal-free first round of the initial and post-change blocks. For compliance and the perceived audit probability, we estimate

$$Y_{ist} = \alpha + \beta_S Sparse_s + \beta_D Dense_s + \mathbf{X}_{ist}'\boldsymbol{\gamma} + \varepsilon_{ist}, \tag{7}$$

where Y_{ist} is either compliance, defined as declared income divided by income, or $\widehat{p}_{ist}$, the midpoint-weighted mean of the elicited subjective distribution. The indices i, s, and t denote participants, sessions, and experimental blocks. The coefficients β_S and β_D measure the sparse- and dense-network presence effects relative to no peer information; their difference tests whether the signal-free response increases with network density. The vector $\mathbf{X}_{ist}$ contains experiment-specific controls[4] and predetermined participant characteristics.[5] The subjective distribution is elicited after the declaration but before that round's audit outcome, so the two outcomes are measured under the same signal-free information set.

Table 2 shows that the sparse and dense networks raise compliance by 10.4 and 11.0 percentage points, respectively, relative to no peer information (both $p < .001$). At this signal-free stage, the effects are economically substantial: relative to mean compliance of 56.5% under no peer information, they represent an 18–19% increase in the share of income declared. The two network effects are statistically indistinguishable ($p = .785$).

The perceived audit probability shows a parallel pattern. The first-round $\widehat{p}_{ist}$ is 2.8 percentage points higher in both network conditions (both $p < .05$), and the two network effects are again statistically indistinguishable ($p = .978$). We find the same compliance pattern when we examine the blocks separately. Both networks raise compliance in each block (all $p < .01$), while the increase in perceived audit probability is concentrated in the initial block and weaker after participants have gained experience with the task (see Appendix Table A2).

The similarity of the two network effects matters for interpretation. Participants know their information environment in these rounds but have not received a peer signal through it, so the dense network has no informational advantage over the sparse network. The parallel increases in perceived audit probability and compliance are consistent with the salience channel in Prediction 1: an environment in which audit experiences will become socially connected makes enforcement more prominent before the network conveys evidence about its strength.

Result 1 *Before any peer outcome is transmitted, visible network presence raises com-*

[4] These are indicators for the audit environment, the post-change block, enforcement order, and experimental wave.

[5] We include only characteristics determined before treatment, as shown in Panel A of Appendix Table A1: age, gender, bachelor's-student status, Communist Party membership, monthly expenditure category, and time spent with friends.

Table 2: Network presence before peer information is transmitted

	Compliance		Perceived audit probability	
	(1)	(2)	(3)	(4)
Sparse network	0.106***	0.104***	0.027**	0.028**
	(0.026)	(0.021)	(0.013)	(0.010)
Dense network	0.118***	0.110***	0.029**	0.028***
	(0.021)	(0.016)	(0.012)	(0.009)
Constant	0.565***	0.722**	0.342***	0.302***
	(0.015)	(0.261)	(0.010)	(0.093)
Controls	No	Yes	No	Yes
Observations	1,136	1,136	1,136	1,136
Sessions (clusters)	26	26	26	26
R-squared	0.023	0.043	0.010	0.053
p-value: dense network = sparse network	0.651	0.785	0.846	0.978

Notes: Participant-round OLS using the signal-free first round of each experimental block. No peer information is omitted. Controls include indicators for the audit environment, experimental block, enforcement order, and wave, together with age, gender, education level, Communist Party membership, expenditure category, and time spent with friends. The final row reports p-values from tests of equality between the dense- and sparse-network coefficients. Perceived audit probability, $\widehat{p}_{it}$, is the midpoint-weighted mean of the elicited subjective distribution. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

pliance by 11 percentage points—approximately 20% relative to no peer information—and the perceived audit probability by 3 points. Neither response varies with network density.

The first-round results establish a level effect of visible network presence before any peer information arrives. We next ask whether broader information reach makes compliance more responsive to the enforcement environment after audit outcomes begin to circulate. Because every session moves from one hidden audit probability to the other, we can compare this response across information conditions.

5.2 Information-network density and enforcement responsiveness

To test Prediction 2, we match each participant's decisions at the same round number across the initial and post-change blocks. For signal rounds $r = 2, \ldots, 10$, define enforcement responsiveness as

$$\Delta C_{irs} = C_{irs}(25\%) - C_{irs}(5\%), \tag{8}$$

where i indexes participants and s sessions. A larger value of ΔC_{irs} indicates a wider compliance gap between the high- and low-enforcement environments.

Figure 3 presents the raw evidence on how information reach shapes enforcement responsiveness. Without peer information, mean compliance is 7.1 percentage points higher in the 25 percent than in the 5 percent audit environment. This gap increases to 9.9 points in the sparse network and 16.9 points in the dense network. Thus, when audit experiences can diffuse throughout the group, compliance responds more than twice as strongly to the enforcement environment as it does without peer information. The observations for individual sessions show the underlying variation, while the regression estimates below provide our preferred tests of these differences.

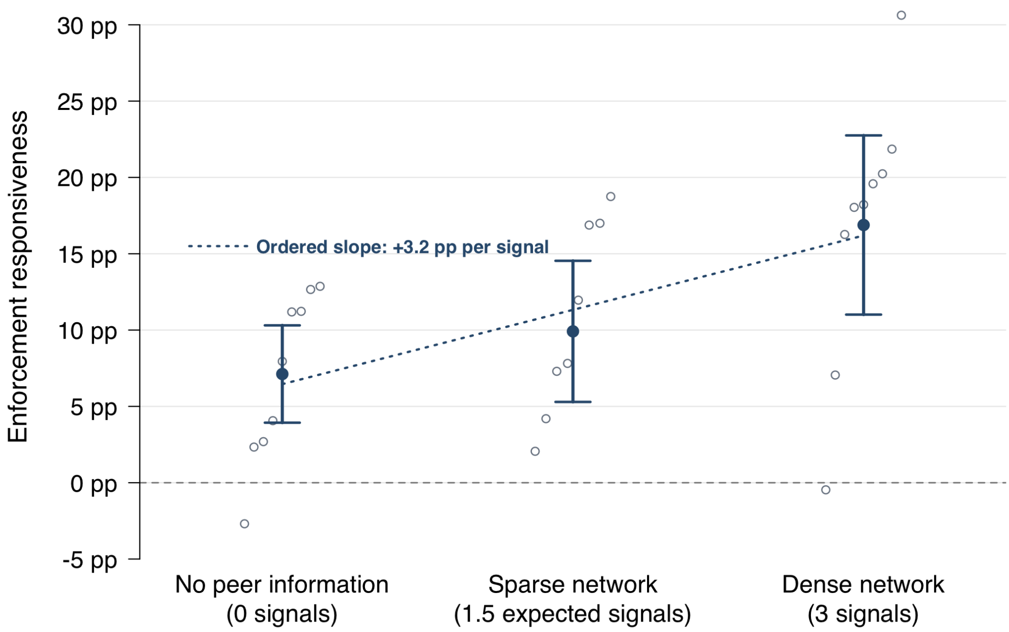

Figure 3: Enforcement responsiveness by information condition

Notes: Open circles show raw mean enforcement responsiveness for each session over signal rounds 2–10. Solid dots are raw means across the matched observations for each participant and round; vertical bars are 95% confidence intervals based on standard errors clustered by session. The dashed line shows fitted values from the ordered specification without controls in column (3) of Table 3, using expected peer signals of 0, 1.5, and 3.

To test these differences formally, we regress responsiveness on indicators for the two network conditions:

$$\Delta C_{irs} = \alpha + \delta_S \text{Sparse}_s + \delta_D \text{Dense}_s + \mathbf{X}'_{irs}\boldsymbol{\gamma} + u_{irs}. \tag{9}$$

The coefficients δ_S and δ_D measure how the sparse and dense networks change responsiveness relative to no peer information; $\delta_D - \delta_S$ compares the two networks directly. Columns (1) and (2) of Table 3 report estimates of equation (9), while columns (3) and (4) replace the treatment indicators with expected peer signals (0, 1.5, and 3). The treatment factor model is our preferred specification because it allows the two network effects to differ freely. The ordered model provides a complementary test of whether responsiveness increases systematically with information reach. The main results are robust to alternative clustering choices (see Appendix Table A3).

Table 3: Information-network density and enforcement responsiveness

	Dependent variable: enforcement responsiveness			
	Treatment factor model		Ordered model	
	(1)	(2)	(3)	(4)
Sparse network	0.028	0.021		
	(0.027)	(0.018)		
Dense network	0.098^{***}	0.081^{***}		
	(0.032)	(0.023)		
Expected peer signals			0.032^{***}	0.027^{***}
			(0.011)	(0.008)
Constant	0.071^{***}	0.070	0.065^{***}	0.066
	(0.015)	(0.140)	(0.016)	(0.140)
Controls	No	Yes	No	Yes
Observations	5,112	5,112	5,112	5,112
Sessions (clusters)	26	26	26	26
R-squared	0.010	0.039	0.009	0.039
p-value: dense network = sparse network	0.066	0.025		

Notes: Each observation pairs the same participant and round number across the two experimental blocks. Enforcement responsiveness is compliance in the 25 percent audit environment minus compliance in the 5 percent audit environment. No peer information is omitted in columns (1) and (2). Controls include indicators for enforcement order, experimental wave, and matched round, together with age, gender, education level, Communist Party membership, expenditure category, and time spent with friends. Expected peer signals equal 0, 1.5, and 3 in the no-peer-information, sparse-network, and dense-network conditions. The final row reports p-values from tests of equality between the dense- and sparse-network coefficients. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3 confirms the graphical pattern. The sparse network increases responsiveness by 2.1 percentage points relative to no peer information, but the difference is not statistically significant ($p = .260$). The dense network increases responsiveness by 8.1 points relative to no peer information ($p = .002$) and by 6.0 points relative to the sparse network ($p = .025$). The ordered model provides a complementary summary: responsiveness increases by 2.7 percentage points per additional expected peer signal ($p = .002$). The treatment-only estimates are similar but less precise. Together, the ordered estimate and the two dense-network comparisons show that responsiveness increases with information reach and is strongest in the dense network. These results support Prediction 2.

Result 2 *Enforcement responsiveness increases with the reach of peer-audit information. The dense network produces significantly greater responsiveness than either no peer information or the sparse network. The sparse-network estimate lies between the two but is not statistically different from no peer information.*

5.3 Adjustment after an enforcement change

To unpack the responsiveness result, we examine the two directions separately and test Corollaries 1a and 1b: how networks affect compliance after enforcement increases from 5 percent to 25 percent and after it decreases from 25 percent to 5 percent. We use post-change rounds 2–10. Let H_s equal one when enforcement increases to 25 percent and zero when it decreases to 5 percent. We estimate

$$\begin{aligned} C_{irs} &= \alpha + \beta_S \text{Sparse}_s + \beta_D \text{Dense}_s + \lambda H_s \\ &\quad + \theta_S(\text{Sparse}_s \times H_s) + \theta_D(\text{Dense}_s \times H_s) + \mathbf{X}'_{irs}\boldsymbol{\gamma} + \varepsilon_{irs}, \end{aligned} \tag{10}$$

where C_{irs} is compliance for participant i in post-change round r and session s. Table 4 reports the estimated network effects separately for the two directions of change. Following an enforcement decrease, the sparse and dense network effects are β_S and β_D. Following an increase, they are $\beta_S + \theta_S$ and $\beta_D + \theta_D$.

We begin with the total effect of each information network relative to no peer information. These contrasts capture the combined behavioral response to being embedded in a network and to receiving audit outcomes through that network. Figure 4 presents the estimates, which are reported in column (2) of Table 4.

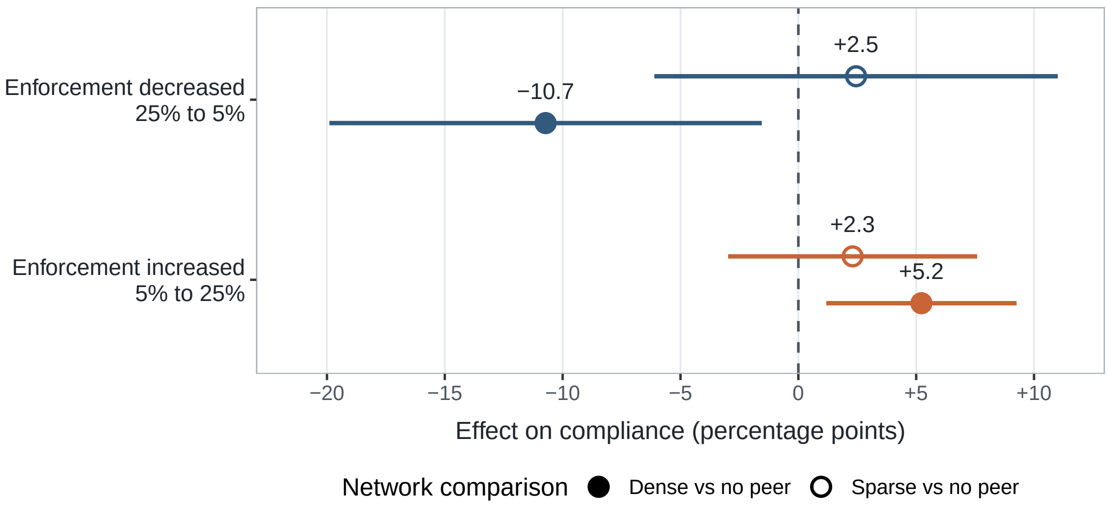

Figure 4: Total information-network effects following a change in enforcement

Notes: Points compare the sparse and dense networks with no peer information using the controlled participant-round specification in column (2) of Table 4; horizontal lines are 95% confidence intervals based on standard errors clustered by session. Negative values indicate lower compliance, and positive values indicate higher compliance. The sample contains post-change rounds 2–10.

Consider first the decrease from 25 percent to 5 percent. Corollary 1b predicts that the total network effect becomes less positive, or more negative, as information reach expands. Relative to no peer information, compliance is 2.5 points higher in the sparse network (p = .561) and 10.7 points lower in the dense network (p = .024; about a 25%

reduction relative to no peer information). This pattern supports Corollary 1b. The negative dense-network effect is especially notable because network presence initially raises compliance, implying that the downward response to information is large enough to dominate any remaining positive presence response.

Following the increase from 5 percent to 25 percent, Corollary 1a predicts positive total network effects that become larger as information reach expands. Consistent with this prediction, the estimate rises from a statistically insignificant 2.3 points in the sparse network to a statistically significant 5.2 points in the dense network ($p = .013$; approximately a 12% increase relative to the observed no-peer-information compliance rate). These total contrasts do not, however, isolate learning: relative to no peer information, they can combine responses to transmitted audit outcomes with any network-presence response that remains after the signal-free first round.

We next compare the dense network directly with the sparse network. Both conditions make the network visible and transmit peer audit outcomes, but the dense network supplies more signals. Their signal-free presence effects are also nearly identical and statistically indistinguishable (Section 5.1). If any lingering presence response is common across the two networks, the difference between them nets out that component and isolates the effect of broader information reach. Figure 5 presents this comparison.

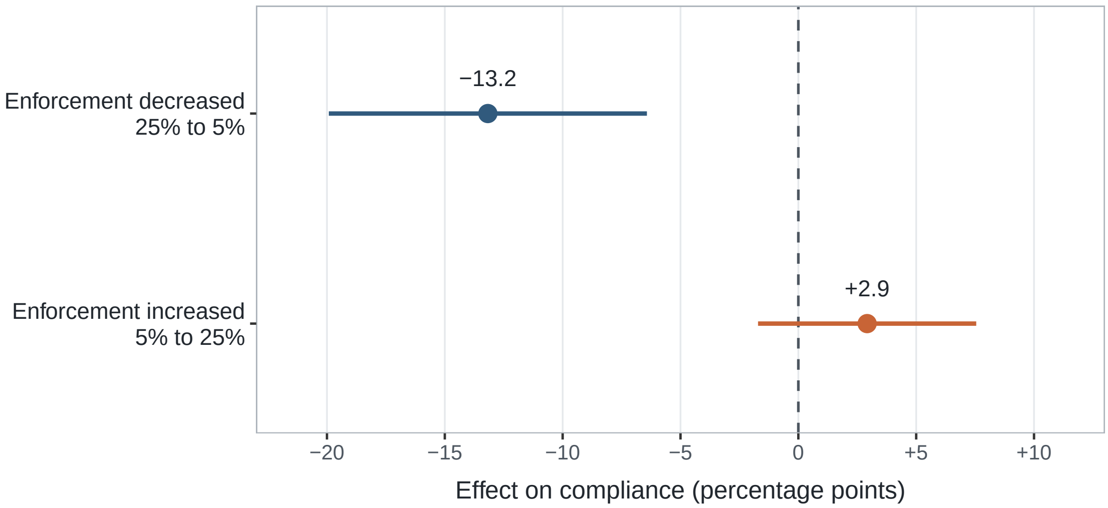

Figure 5: Additional effect of dense rather than sparse information reach

Notes: Points compare the dense with the sparse network using the controlled participant-round specification in column (2) of Table 4; horizontal lines are 95% confidence intervals based on standard errors clustered by session. Negative values indicate lower compliance, and positive values indicate higher compliance. The sample contains post-change rounds 2–10.

Following the enforcement decrease, the dense network lowers compliance by 13.2 percentage points relative to the sparse network ($p < .001$; a 29% reduction relative to the observed sparse-network compliance rate). The estimate is nearly unchanged without controls. This 13.2-point difference strongly supports Corollary 1b: broader information reach produces greater downward adjustment when enforcement weakens. Following the increase, compliance is 2.9 points higher in the dense network than in the sparse network under the controlled specification, but the difference is not statistically significant ($p =$ .206). Without controls, the difference is 4.7 points ($p = .041$). The positive sign supports Corollary 1a, although the attenuation after covariate adjustment makes the evidence that broader reach amplifies the upward response sensitive to specification.

To examine whether these responses are short-lived or persist, we split the post-change block into early rounds 2–5 and later rounds 6–10. Appendix Table A4 reports the network effects for each direction and window from a pooled model that extends equation (10) with interactions for the early rounds. After the enforcement decrease, the sparse network remains statistically indistinguishable from no peer information in both windows ($p = .928$ and $p = .428$, respectively). The dense network, by contrast, lowers compliance relative to the sparse network by 13.6 points in the early rounds ($p < .001$) and by 12.8 points in the later rounds ($p = .007$), reductions of approximately 27% and 31%, respectively, relative to observed sparse-network compliance in those windows. The two differences are statistically indistinguishable ($p = .865$). Broader information reach therefore produces an immediate, economically substantial downward adjustment that continues beyond the initial response to weaker enforcement. Following the enforcement increase, the dense network's total effect relative to no peer information attenuates from 8.4 points in the early rounds ($p = .003$) to 2.7 points in the later rounds ($p = .308$), although the early–later difference is only marginal ($p = .100$).

Table 4: Estimated information-network effects by direction of enforcement change

	(1)	(2)
Panel A: After enforcement decreases (25% to 5%)		
Effects relative to no peer information		
Sparse network	0.027	0.025
	(0.044)	(0.042)
Dense network	-0.111^{**}	-0.107^{**}
	(0.045)	(0.045)
Pairwise comparison		
Difference: dense − sparse	-0.137^{***}	-0.132^{***}
	(0.031)	(0.033)
Panel B: After enforcement increases (5% to 25%)		
Effects relative to no peer information		
Sparse network	0.015	0.023
	(0.024)	(0.026)
Dense network	0.063^{***}	0.052^{**}
	(0.020)	(0.020)
Pairwise comparison		
Difference: dense − sparse	0.047^{**}	0.029
	(0.022)	(0.022)
Controls	No	Yes
Observations	5,112	5,112
Sessions (clusters)	26	26
R-squared	0.024	0.058

Notes: The sample contains post-change rounds 2–10. Each column reports one pooled regression based on equation (10). The panels report network effects calculated separately for each direction of change rather than the underlying regression coefficients. Effects following an enforcement increase add the relevant network and interaction coefficients. Within each panel, the first two rows compare the sparse and dense networks with no peer information. The pairwise comparison reports the estimated difference between the two network effects. Controls include round and experimental-wave indicators, age, gender, education, Communist Party membership, expenditure category, and time spent with friends. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The directional evidence strongly supports Corollary 1b and provides weaker support for Corollary 1a. Greater enforcement responsiveness reflects adjustment in both directions rather than a uniform increase in compliance. The strongest evidence comes after enforcement decreases: the gap between the dense and sparse networks appears immediately and remains virtually unchanged in the later rounds.

Result 3 *Social diffusion does not mechanically increase compliance. Relative to no peer information, the dense network raises compliance by about 12% after enforcement strengthens and lowers it by about 25% after it weakens. The sparse-network estimates provide only partial evidence that this amplification grows with information reach.*

5.4 Learning about the audit probability as the mechanism

5.4.1 Information reach, belief learning, and compliance

The directional compliance results point to learning about the audit probability as the mechanism. This interpretation yields a direct empirical implication: perceived audit probabilities should move toward the new rate as audit outcomes accumulate, and the adjustment should be stronger when participants observe more peer outcomes. Figure 6 shows these belief dynamics. We plot the improvement in perceived probability accuracy relative to the signal-free first round of the post-change block. After enforcement falls from 25 percent to 5 percent, beliefs become steadily more accurate in all three information conditions, but the improvement is consistently largest in the dense network. After enforcement rises from 5 percent to 25 percent, the gains are smaller and the three conditions remain closer together. The visual pattern mirrors the directional compliance evidence: the strongest separation by information reach occurs after enforcement weakens.

Table 5 tests this learning pattern using three features of the elicited distributions: absolute error in the perceived audit probability,[6] probability mass assigned to the true rate interval,[7] and belief concentration, measured by normalized entropy.[8] As in Section 5.3, we estimate pooled interaction models and report the effects separately for enforcement increases and decreases. Here, the effects compare the change from the signal-free first round to subsequent signal rounds across information conditions.

---

[6] Recall that $\widehat{p}_{it} = \sum_{k=1}^{10} \pi_{itk} m_k$ is the midpoint-weighted mean of participant i's elicited probability distribution in round t. We calculate absolute error as $|\widehat{p}_{it} - p_t|$, where p_t is the true audit probability. Lower values indicate greater accuracy. The table reports the reduction in this error relative to the first post-change round, so positive values indicate improved accuracy.

[7] This is the share of tokens assigned to the interval containing the true audit probability: 0–10% when the true rate is 5 percent, and 20–30% when it is 25 percent.

[8] Normalized entropy captures how dispersed the subjective distribution is across audit-rate intervals and therefore how uncertain the participant is. Following Shannon (1948), normalized entropy is $H_{it} = -\sum_{k=1}^{10} \pi_{itk} \log(\pi_{itk}) / \log(10)$. It ranges from zero when all probability mass is placed in one interval to one when mass is distributed uniformly across all ten intervals. The table reports reductions in entropy, so positive values indicate more concentrated beliefs.

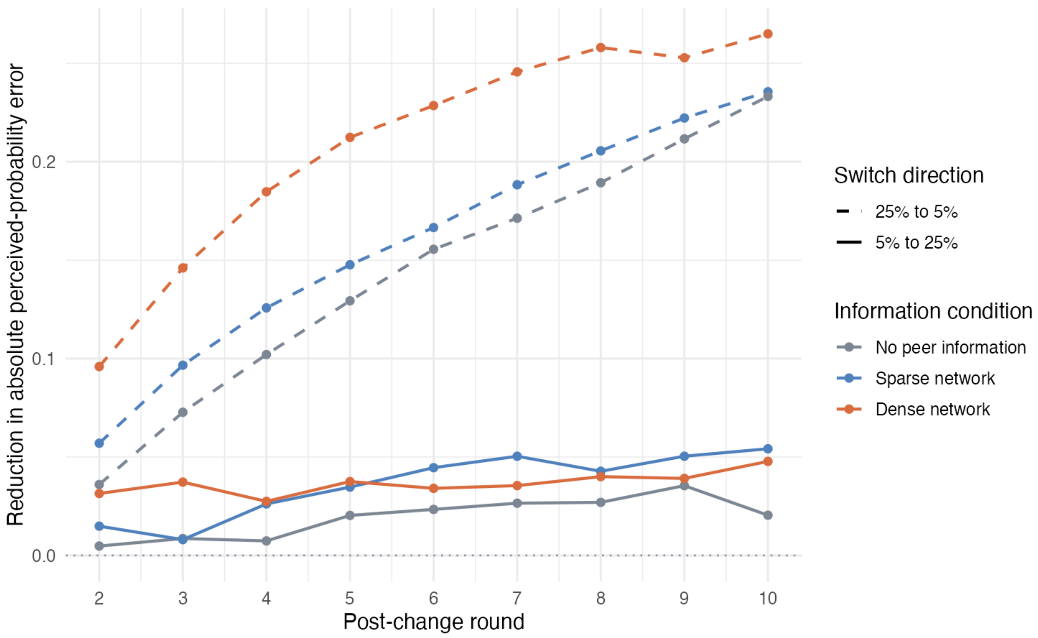

Figure 6: Accuracy of perceived audit probabilities after the enforcement switch
Notes: The lines are constructed directly from the elicited belief data. For each participant and round, the accuracy gain is the absolute error of $\widehat{p}_{it}$ in the first post-change round minus its absolute error in the indicated round. Positive values indicate improved accuracy. Solid and dashed lines distinguish switch directions; colors distinguish information environments.

Panel A shows the strongest learning gains in the dense network after enforcement decreases. Relative to no peer information, the dense network reduces absolute error by 6.4 percentage points, increases mass on the true rate interval by 15.6 points, and reduces belief entropy by 11.3 points. Each effect is statistically significant at $p < .001$. The sparse network significantly increases true rate mass and reduces belief entropy, although its improvement in absolute accuracy is not statistically significant. The ordered specification confirms that all three measures improve as the expected number of peer signals increases ($p \leq .002$).

Panel B shows a different pattern after enforcement increases. Both networks shift probability mass toward the true rate interval and produce more concentrated beliefs, but neither significantly improves the absolute accuracy of the perceived audit probability. Panel C confirms that the dense-network increase in true rate mass and reduction in belief entropy are significantly larger after the enforcement decrease ($p = .048$ and $p = .012$, respectively), while the cross-direction difference in absolute accuracy is not statistically significant ($p = .109$). The evidence therefore supports Prediction 2 in both directions but reveals a meaningful asymmetry: broader information reach improves distributional learning after either change, while its effect on the perceived audit probability is concentrated after enforcement weakens.

One possible explanation for the stronger raw accuracy gain after enforcement decreases is that participants have more scope to improve. In the signal-free first post-change round, mean absolute error is 34.2 percentage points after the decrease but 14.6 points after the increase. We therefore estimate baseline-adjusted specifications that condition on each participant's round 1 belief measure. Adjusting for these different starting points does not eliminate the asymmetry. Relative to no peer information, the dense-network accuracy gain is 5.2 percentage points after enforcement decreases and 1.7 points after it increases; the difference is statistically significant ($p = .029$). The ordered specification yields the same conclusion ($p = .026$), and the baseline-adjusted differences in true rate mass and belief concentration are also statistically significant. We report these results in Appendix Table A5.

Table 5: Estimated effects of information networks on learning after an enforcement change

	(1) Change in perceived-probability accuracy	(2) Change in true rate mass	(3) Reduction in belief entropy
Panel A: Enforcement decreases from 25% to 5%			
Estimated effects from the treatment factor model			
Sparse network	0.016	0.053**	0.038**
	(0.010)	(0.024)	(0.015)
Dense network	0.064***	0.156***	0.113***
	(0.017)	(0.030)	(0.014)
Estimated effect per expected peer signal			
Expected peer signals	0.022***	0.053***	0.038***
	(0.006)	(0.011)	(0.005)
Panel B: Enforcement increases from 5% to 25%			
Estimated effects from the treatment factor model			
Sparse network	0.016	0.058***	0.060***
	(0.017)	(0.011)	(0.020)
Dense network	0.019	0.077***	0.051***
	(0.020)	(0.020)	(0.018)
Estimated effect per expected peer signal			
Expected peer signals	0.007	0.027***	0.019***
	(0.007)	(0.007)	(0.006)
Panel C: Tests of equal effects across directions (p-values)			
Sparse network	0.991	0.859	0.392
Dense network	0.109	0.048	0.012
Expected peer signals	0.118	0.055	0.020
Controls	Yes	Yes	Yes
Participant fixed effects	Yes	Yes	Yes
Round fixed effects	Yes	Yes	Yes
Observations	5,680	5,680	5,680
Participants	568	568	568
Sessions (clusters)	26	26	26

Notes: Participant-round regressions use the post-change block. For each outcome, the treatment factor model allows the network effects to differ by the direction of enforcement change; the ordered model replaces the information-condition indicators with expected peer signals (0, 1.5, and 3). For each direction, Panels A and B report the sparse and dense network effects relative to no peer information and the effect per expected peer signal. Panel C reports p-values from tests that the corresponding effects are equal across directions. Positive values indicate greater perceived probability accuracy, more probability mass on the true rate interval, and a larger reduction in belief entropy (greater concentration). All specifications include participant and round fixed effects. Standard errors clustered by session are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Improved beliefs matter for compliance if they carry forward into reporting decisions. Table 6 tests Prediction 3 by regressing each declaration on the perceived audit probability elicited in the preceding round. An increase of 10 percentage points in the preceding perceived audit probability is associated with 4.3 percentage points higher compliance in the next declaration ($p < .001$), equivalent to about 9% of mean compliance in the estimation sample. Belief entropy has little additional predictive power once the perceived audit probability is included. The relationship between perceived audit probability and subsequent compliance is positive in every information condition and does not differ significantly across them (see Appendix Table A6). Information reach therefore changes the beliefs entering the compliance decision, while the relationship between a given belief and subsequent compliance remains similar across conditions.

Table 6: Perceived audit probability and subsequent compliance

	Next declaration (1)
Prior perceived audit probability, per 10 percentage points	0.043***
	(0.004)
Prior belief entropy, per 10 percentage points	0.005
	(0.004)
Controls	Yes
Participant fixed effects	Yes
Session-round fixed effects	Yes
Observations	10,224
Participants	568
Sessions (clusters)	26

Notes: The table pools the three information conditions and regresses each declaration on the perceived audit probability and belief entropy elicited in the preceding round. Controls include the participant's preceding audit outcome, the number of observed peer audits, and the number of peer signals. Standard errors clustered by session are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4.2 Why the dense network produces the strongest response

The analysis above establishes that learning improves with information reach, but it does not explain why the dense network produces the strongest response. Equation (4) provides two explanations: participants may place greater weight, ω_s, on each peer signal when information reach expands, or the per-signal weight may remain unchanged and the dense network may produce stronger learning simply because it supplies more peer signals. We distinguish these explanations by estimating how the round-to-round perceived audit probability responds to participants' own audit outcomes and to each observed peer outcome. Appendix Section B develops the empirical specification and reports the estimates in Appendix Table A9.

Participants respond roughly twice as strongly to their own audit outcome as to any single peer outcome. In the sparse network, the own outcome coefficient is .062 and the per-peer coefficient is .028, so the response to each peer outcome is 55% smaller. In the dense network, the corresponding coefficients are .059 and .031, making the response to each peer outcome 47% smaller.[9] The within-network differences between own and peer responsiveness are statistically significant ($p < .001$ and $p = .001$, respectively). The response to each peer signal is statistically indistinguishable across the sparse and dense networks ($p = .475$). The evidence therefore points to the accumulation of additional signals, rather than a larger response to each signal, as the source of the dense network's learning advantage.

The belief distributions and responses to individual signals trace a coherent mechanism from information reach to compliance. After enforcement weakens, broader reach improves all three belief measures, with the strongest gains in the dense network. After enforcement strengthens, broader reach shifts probability mass toward the true rate interval and produces more concentrated beliefs, but it does not significantly improve the accuracy of the perceived audit probability. Higher perceived audit probabilities predict higher compliance in the next round. The signal-response estimates explain why the dense network produces the strongest response: participants respond less to each peer outcome than to their own audit experience, but the dense network supplies more peer signals.

Result 4 *Social information improves belief concentration and the probability mass assigned to the true audit rate after enforcement changes in either direction. Gains in the accuracy of the perceived audit probability are concentrated after enforcement weakens. The dense network's advantage comes from receiving more signals, not from greater responsiveness to each peer outcome, and higher perceived audit probabilities predict greater subsequent compliance.*

The norm elicitation provides an additional check on the mechanism. Perceived appropriateness responds to the enforcement environment, but this response does not increase significantly with information reach; neither network treatment nor the ordered specification produces a statistically significant effect. Appendix Table A7 reports the results. Norms therefore respond to enforcement, but they do not explain the network pattern in compliance as closely as perceived audit probabilities do.

[9] This qualitative ranking is consistent with evidence that people often respond less to social information than to equally relevant private or self-generated information (Weizsäcker, 2010; De Filippis et al., 2022; Conlon et al., 2022).

6 Policy Counterfactuals

The compliance and belief results distinguish two margins of enforcement communication. An authority can influence how widely evidence of audit activity circulates without revealing the audit probability, or it can disclose the probability itself. We use two counterfactual exercises to examine these margins separately. The first uses the randomized information conditions observed under each enforcement environment to change information reach while keeping the audit probability hidden. The second holds the no-peer-information environment as the benchmark and uses the estimated belief–compliance relationship to predict the effect of exact disclosure. The first exercise therefore remains anchored in experimentally observed outcomes; the second is explicitly model based.

6.1 State-contingent information reach under opacity

The first exercise asks whether broader circulation of audit experiences should be encouraged regardless of the enforcement environment. We treat no peer information, the sparse network, and the dense network as progressively broader levels of information reach. Each counterfactual policy is a rule that maps the true enforcement environment, which the authority knows, to one of these experimentally observed conditions. The main state-contingent rule uses the dense network when enforcement is high and no peer information when enforcement is low. We compare it with rules that always use the same information condition and with the reverse rule. The audit probability remains hidden under every policy.

For comparability, each policy places equal weight on the 5 percent and 25 percent enforcement environments and averages predictions over the observed participant characteristics. Table 7 displays the information condition used in each enforcement environment. Its final two columns report predicted compliance and each alternative's difference from the state-contingent dense-network benchmark; a negative difference indicates lower compliance than under the benchmark rule.

Table 7 shows that more information reach is not uniformly compliance enhancing. The state-contingent dense-network rule has the highest point estimate, with predicted compliance of 46.2%. It exceeds always providing no peer information by 2.6 percentage points ($p = .013$) and always using the dense network by 5.4 points ($p = .024$). The reverse rule—no peer information under high enforcement and dense diffusion under low enforcement—performs worst, with predicted compliance 8.0 points below the benchmark ($p = .004$). These comparisons reflect the directional result in Section 5.3: broad diffusion reinforces deterrence when enforcement is high but erodes it when enforcement is low.

The remaining comparisons favor adapting information reach to the enforcement environment but do not establish that maximum density is necessary. Always using the sparse network produces predicted compliance of 46.0%, statistically indistinguishable from the

Table 7: Counterfactual information-reach policies under opacity

Policy rule	Information condition		Predicted compliance	Difference from benchmark
	High enforcement	Low enforcement		
State-contingent dense	Dense	None	0.462 (0.020)	—
Always sparse	Sparse	Sparse	0.460 (0.014)	−0.002 (0.025)
State-contingent sparse	Sparse	None	0.448 (0.020)	−0.015 (0.011)
Always none	None	None	0.436 (0.019)	-0.026^{**} (0.010)
Always dense	Dense	Dense	0.409 (0.013)	-0.054^{**} (0.022)
Reverse	None	Dense	0.382 (0.015)	-0.080^{***} (0.025)

Notes: "None" denotes no peer information. Predictions combine estimates for each randomized information condition and enforcement environment over post-change rounds 2–10, place equal weight on the high- and low-enforcement environments, and average fitted values over the observed covariate distribution. The benchmark is the state-contingent dense policy. These are policy contrasts constructed from experimentally observed information conditions within each enforcement environment, not separately randomized interventions. Controls include round and experimental-wave indicators, age, gender, education, Communist Party membership, expenditure category, and time spent with friends. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

state-contingent dense-network rule ($p = .925$). The analogous state-contingent sparse-network rule is also not significantly different from it ($p = .206$). Thus, indiscriminate broad diffusion can lower compliance, while making genuine audit activity more visible is most useful when enforcement is strong. Public communication and media coverage offer practical channels for increasing the visibility of genuine enforcement activity.

6.2 Exact disclosure of the audit probability

The first exercise changes the reach of audit information while preserving opacity about the true probability. The second asks what would happen if the authority instead disclosed the exact rate. Because exact disclosure was not an experimental treatment, we construct a model-based benchmark. We use the belief–compliance model in Table 6, replace each prior subjective distribution with a point mass at the true rate, set belief entropy to zero, and hold the other determinants of compliance fixed. Table 8 starts from the environment without peer information and reports the true rate, the average belief under hidden enforcement, observed compliance, predicted compliance under exact disclosure, and the difference between the two.

Table 8: Model-based exact disclosure counterfactual under no peer information

Regime	True rate	Average hidden belief	Observed compliance	Predicted compliance under disclosure	Disclosure effect
Low audit	0.05	0.248	0.415	0.304	-0.111^{***}
					(0.017)
High audit	0.25	0.314	0.486	0.431	-0.055^{**}
					(0.021)

Notes: The calculation replaces each participant's elicited prior subjective distribution with a degenerate distribution at the true audit probability and predicts compliance using the estimated relationship between prior perceived audit probabilities and declarations. The disclosure effect is predicted compliance under disclosure minus observed compliance under hidden enforcement. Exact disclosure was not randomized, so these are model-based counterfactuals. Standard errors for the disclosure effects are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$ in the primary perceived-probability-and-entropy specification.

Table 8 shows why exact disclosure can weaken deterrence in this setting. When the true audit probability is 5 percent, the average hidden belief is 24.8%, almost five times the true rate. Replacing that belief with exact disclosure lowers predicted compliance from 41.5% to 30.4%, a reduction of 11.1 percentage points ($p < .001$), or approximately 27% of observed compliance. When the true rate is 25 percent, the average hidden belief is 31.4%; exact disclosure lowers predicted compliance from 48.6% to 43.1%, a reduction of 5.5 points ($p = .016$), or approximately 11%. The low-enforcement result is stable across alternative prediction models, whereas the high-enforcement estimate is smaller and more specification sensitive.[10]

[10] Across the alternative prediction models, exact disclosure reduces predicted compliance by 7.1–11.1

This exercise separates belief accuracy from deterrence.[11] Exact disclosure corrects upwardly biased perceptions of audit risk and therefore improves accuracy, but the same correction reduces predicted compliance.

The two exercises reveal an important distinction between information reach and precision. The consequences of information diffusion depend on the enforcement environment: broader reach can reinforce a high-enforcement regime but undermine a low-enforcement regime. Revealing the exact rate can also reduce compliance when taxpayers otherwise overestimate audit risk. These findings provide a deterrence-based explanation for why tax agencies may publicize enforcement priorities and visible audit activity while withholding exact taxpayer-specific probabilities and selection thresholds.

7 Discussion and Conclusion

Tax enforcement operates through perceived audit probabilities as well as audits. When those probabilities are hidden, taxpayers must learn from sparse personal experience and from the experiences that reach them through social connections. Our experiment makes this process observable by varying the reach of peer audit outcomes, changing the underlying enforcement environment, and eliciting complete subjective distributions in every round. This design reveals the learning mechanism behind enforcement spillovers that is difficult to observe in the field.

We identify two distinct effects of an information network. Before any peer outcome is transmitted, knowing that one is embedded in a network raises compliance by approximately 11 percentage points, with no corresponding increase as the network becomes denser. Once audit outcomes begin to circulate, information reach matters. Compliance responds about twice as strongly to the enforcement environment in the dense network as it does without peer information. This responsiveness is directional rather than a uniform increase in compliance: the dense network reinforces deterrence after enforcement strengthens and reduces compliance after enforcement weakens. The latter response appears immediately and persists throughout the post-change block.

The belief evidence explains this behavioral pattern. Social information shifts probability mass toward the true audit-rate interval and produces more concentrated beliefs after enforcement changes in either direction; improvements in the accuracy of the perceived audit probability are strongest after enforcement weakens. Higher perceived audit

percentage points in the low-enforcement environment and by 2.4–5.5 points in the high-enforcement environment (see Appendix Table A8).

[11] Official audit communication may also affect compliance through salience. In a large field experiment, Bérgolo et al. (2023) find that audit-related letters increased tax payments even though they lowered perceived audit probabilities, with little sensitivity to the numerical audit probabilities communicated. Our exercise isolates the predicted consequence of replacing subjective beliefs with the true rate rather than the broader effect of receiving an official message.

probabilities, in turn, predict higher compliance in the next round. Participants respond roughly twice as strongly to their own audit outcome as to any single peer outcome, but the response to each peer signal does not differ significantly between the sparse and dense networks. The dense network's advantage therefore comes from receiving more individually discounted signals rather than placing greater weight on each one. The norm elicitation provides a useful boundary: perceived appropriateness responds to enforcement, but its response does not vary with information reach.

Taken as a whole, the evidence favors graded information accumulation over a literal threshold mechanism. Enforcement responsiveness increases in the ordered specification, the sparse network improves probability mass on the true rate interval and belief concentration even when its compliance effect is small, and responsiveness to each peer signal is similar across network densities. These findings indicate that information accumulates as additional peer outcomes arrive. The behavioral response is nevertheless nonlinear. Sparse-network compliance estimates often lie between the no-peer and dense-network estimates but are statistically indistinguishable from no peer information, whereas the dense network produces clear behavioral separation. The coexistence of a positive network-presence response and modest learning can leave the sparse network's net compliance effect near zero, making adjustment most visible once several signals accumulate. With three levels of information reach, we cannot distinguish a steep continuous response from an exact threshold. We therefore interpret the pattern as graded learning with behavioral effects concentrated in the dense network.

This mechanism bridges field evidence on enforcement spillovers and laboratory evidence on tax compliance under uncertain audit risk. Field studies show that enforcement spills over through geography and shared tax preparers, but they cannot directly observe the information network or the beliefs that generate the spillover. Laboratory tax-compliance studies can control enforcement, but they usually announce the audit probability or study ambiguity without varying the reach of social evidence. Our design reconciles these findings: audit outcomes are noisy signals of a common enforcement environment, social reach determines how much evidence a taxpayer receives, and the perceived audit probability guides the next declaration. The relevant network effect is therefore not generic peer influence but social learning about hidden enforcement.

The policy counterfactuals sharpen this distinction. While keeping the audit probability hidden, making genuine audit activity more visible is most valuable when enforcement is strong; indiscriminate diffusion can instead erode compliance when enforcement is weak. Exact disclosure operates on a different margin. Because participants substantially overestimate audit risk, revealing the true rate improves belief accuracy but is predicted to reduce compliance, especially in the 5 percent environment. These exercises distinguish publicizing genuine enforcement activity from revealing its exact probability.

The policy implication therefore depends on the direction of the enforcement change.

If renewed investment strengthens audit enforcement, communicating the change and publicizing genuine audit activity can magnify deterrence. If budget cuts weaken enforcement, the same social-learning channel can spread evidence of lower audit risk and accelerate compliance erosion. Communication cannot replace enforcement capacity, but authorities should anticipate this multiplier and maintain the visibility of credible, targeted audit activity. In either direction, the evidence favors communicating genuine enforcement activity rather than an exact taxpayer-wide audit probability.

The experiment abstracts from several features of naturally occurring networks. It has two enforcement states, peer outcomes are accurate signals, and groups are temporary. Persistent relationships, selective communication, homophily, and misinformation may alter how enforcement information spreads outside the laboratory. Exact disclosure is a model-based counterfactual rather than a randomized treatment, so it does not capture how official announcements may affect credibility, salience, or trust, nor their broader consequences for fairness, accountability, or perceived legitimacy. These boundaries motivate field tests of how perceived enforcement changes in persistent networks and experiments that compare broad communication of enforcement activity with precise probability disclosure.

The disclosure counterfactual reveals a policy irony. Keeping audit probabilities opaque can preserve deterrence when taxpayers otherwise overestimate audit risk, but that same opacity makes social networks consequential. An audit affects not only the taxpayer who is audited but also the beliefs and declarations of those who learn about it through their social connections. Whether this network multiplier reinforces or erodes deterrence depends on the enforcement environment; its magnitude depends on how far the audit experience travels.

References

Allingham, M. G. and Sandmo, A. (1972), 'Income tax evasion: A theoretical analysis', *Journal of public economics* **1**(3-4), 323–338.

Alm, J. (2019), 'What motivates tax compliance?', *Journal of Economic Surveys* **33**(2), 353–388.

Alm, J., Bloomquist, K. M. and McKee, M. (2017), 'When you know your neighbour pays taxes: Information, peer effects and tax compliance', *Fiscal Studies* **38**(4), 587–613.

Alm, J., Jackson, B. R. and McKee, M. (1992), 'Institutional uncertainty and taxpayer compliance', *American Economic Review* **82**(4), 1018–1026.

Alm, J., Jackson, B. R. and McKee, M. (2009), 'Getting the word out: Enforcement information dissemination and compliance behavior', *Journal of Public Economics* **93**(3–4), 392–402.

Alm, J. and Kasper, M. (2021), Laboratory experiments, *in* B. van Rooij and D. D. Sokol, eds, 'The Cambridge Handbook of Compliance', Cambridge University Press, pp. 707–727.

Alm, J. and Malézieux, A. (2021), '40 years of tax evasion games: A meta-analysis', *Experimental Economics* **24**(3), 699–750.

Anderson, L. R. and Holt, C. A. (1997), 'Information cascades in the laboratory', *American Economic Review* **87**(5), 847–862.

Andreoni, J., Erard, B. and Feinstein, J. (1998), 'Tax compliance', *Journal of economic literature* **36**(2), 818–860.

Benjamin, D. J. (2019), Errors in probabilistic reasoning and judgment biases, *in* B. D. Bernheim, S. DellaVigna and D. Laibson, eds, 'Handbook of Behavioral Economics: Applications and Foundations 1', Vol. 2, Elsevier, pp. 69–186.

Bérgolo, M., Ceni, R., Cruces, G., Giaccobasso, M. and Perez-Truglia, R. (2018), 'Misperceptions about tax audits', *AEA Papers and Proceedings* **108**, 83–87.

Bérgolo, M., Ceni, R., Cruces, G., Giaccobasso, M. and Perez-Truglia, R. (2023), 'Tax audits as scarecrows: Evidence from a large-scale field experiment', *American Economic Journal: Economic Policy* **15**(1), 110–153.

Blackwell, C. (2010), A meta-analysis of incentive effects in tax compliance experiments, *in* J. Alm, J. Martinez-Vazquez and B. Torgler, eds, 'Developing Alternative Frameworks for Explaining Tax Compliance', Routledge, London, pp. 97–112.

Boning, W. C., Guyton, J., Hodge, R. and Slemrod, J. (2020), 'Heard it through the grapevine: The direct and network effects of a tax enforcement field experiment on firms', *Journal of Public Economics* **190**, 104261.

Bordalo, P., Conlon, J., Gennaioli, N., Kwon, S. and Shleifer, A. (2026), 'How people use statistics', *Review of Economic Studies* **93**(1), 250–285.

Chandrasekhar, A. G., Larreguy, H. and Xandri, J. P. (2020), 'Testing models of social learning on networks: Evidence from two experiments', *Econometrica* **88**(1), 1–32.

Chen, D. L., Schonger, M. and Wickens, C. (2016), 'oTree—an open-source platform for laboratory, online, and field experiments', *Journal of Behavioral and Experimental Finance* **9**, 88–97.

Chen, S. X. (2017), 'The effect of a fiscal squeeze on tax enforcement: Evidence from a natural experiment in china', *Journal of Public Economics* **147**, 62–76.

Choo, C. L., Fonseca, M. A. and Myles, G. D. (2016), 'Do students behave like real taxpayers in the lab? evidence from a real effort tax compliance experiment', *Journal of Economic Behavior & Organization* **124**, 102–114.

Congressional Budget Office (2020), 'Trends in the internal revenue service's funding and enforcement', https://www.cbo.gov/publication/56467. Accessed July 13, 2026.

Conlon, J. J., Mani, M., Rao, G., Ridley, M. W. and Schilbach, F. (2021), Learning in the household, NBER Working Paper 28844, National Bureau of Economic Research.

Conlon, J. J., Mani, M., Rao, G., Ridley, M. W. and Schilbach, F. (2022), Not learning from others, NBER Working Paper 30378, National Bureau of Economic Research. Revised January 2026.

De Filippis, R., Guarino, A., Jehiel, P. and Kitagawa, T. (2022), 'Non-bayesian updating in a social learning experiment', *Journal of Economic Theory* **199**, 105188.

Di Gioacchino, D. and Fichera, E. (2020), 'Tax evasion and tax morale: A social network analysis', *European Journal of Political Economy* **65**, 101922.

Drago, F., Mengel, F. and Traxler, C. (2020), 'Compliance behavior in networks: Evidence from a field experiment', *American Economic Journal: Applied Economics* **12**(2), 96–133.

Fortin, B., Lacroix, G. and Villeval, M. C. (2007), 'Tax evasion and social interactions', *Journal of Public Economics* **91**(11–12), 2089–2112.

Grether, D. M. (1980), 'Bayes rule as a descriptive model: The representativeness heuristic', *Quarterly Journal of Economics* **95**(3), 537–557.

Harrison, G. W., Martínez-Correa, J., Swarthout, J. T. and Ulm, E. R. (2017), 'Scoring rules for subjective probability distributions', *Journal of Economic Behavior & Organization* **134**, 430–448.

Hashimzade, N., Myles, G. D., Page, F. and Rablen, M. D. (2014), 'Social networks and occupational choice: The endogenous formation of attitudes and beliefs about tax compliance', *Journal of Economic Psychology* **40**, 134–146.

HM Revenue & Customs (2026), 'Tax compliance risk management: Risk assessment overview', https://www.gov.uk/hmrc-internal-manuals/tax-compliance-risk-management/tcrm4100. Updated June 9, 2026; accessed July 13, 2026.

HM Treasury (2025), 'Spending review 2025', https://www.gov.uk/government/publications/spending-review-2025-document/spending-review-2025-html. Accessed July 13, 2026.

Internal Revenue Service (2026), 'Irs audits', https://www.irs.gov/businesses/small-businesses-self-employed/irs-audits. Accessed July 13, 2026.

Kasper, M. and Rablen, M. D. (2023), 'Tax compliance after an audit: Higher or lower?', *Journal of Economic Behavior & Organization* **207**, 157–171.

Kastlunger, B., Kirchler, E., Mittone, L. and Pitters, J. (2009), 'Sequences of audits, tax compliance, and taxpaying strategies', *Journal of Economic Psychology* **30**(3), 405–418.

Khandelwal, V. (2024), 'Learning in networks with idiosyncratic agents', *Games and Economic Behavior* **144**, 225–249.

Korobow, A., Johnson, C. and Axtell, R. (2007), 'An agent–based model of tax compliance with social networks', *National Tax Journal* **60**(3), 589–610.

Krupka, E. L. and Weber, R. A. (2013), 'Identifying social norms using coordination games: Why does dictator game sharing vary?', *Journal of the European Economic Association* **11**(3), 495–524.

Lefebvre, M., Pestieau, P., Riedl, A. M. and Villeval, M. C. (2015), 'Tax evasion and social information: An experiment in belgium, france, and the netherlands', *International Tax and Public Finance* **22**(3), 401–425.

Mittone, L., Panebianco, F. and Santoro, A. (2017), 'The bomb-crater effect of tax audits: Beyond the misperception of chance', *Journal of Economic Psychology* **61**, 225–243.

Möbius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S. (2022), 'Managing self-confidence: Theory and experimental evidence', *Management Science* **68**(11), 7793–7817.

Möbius, M. and Rosenblat, T. (2014), 'Social learning in economics', *Annual Review of Economics* **6**(1), 827–847.

Reck, D., Slemrod, J. and Vattø, T. E. (2022), 'Public disclosure of tax information: Compliance tool or social network?', *Journal of Public Economics* **212**, 104708.

Shannon, C. E. (1948), 'A mathematical theory of communication', *Bell System Technical Journal* **27**(3–4), 379–423, 623–656.

Slemrod, J. (2019), 'Tax compliance and enforcement', *Journal of Economic Literature* **57**(4), 904–954.

Snow, A. and Warren, R. S. (2005), 'Ambiguity about audit probability, tax compliance, and taxpayer welfare', *Economic Inquiry* **43**(4), 865–871.

Snow, A. and Warren, R. S. (2007), 'Audit uncertainty, bayesian updating, and tax evasion', *Public Finance Review* **35**(5), 555–571.

Spicer, M. W. and Hero, R. E. (1985), 'Tax evasion and heuristics: A research note', *Journal of Public Economics* **26**(2), 263–267.

Weizsäcker, G. (2010), 'Do we follow others when we should? a simple test of rational expectations', *American Economic Review* **100**(5), 2340–2360.

A Additional Tables and Figures

This section collects the supporting figure and tables referenced in the main text. Definitions, samples, specifications, and inference are reported in the corresponding notes.

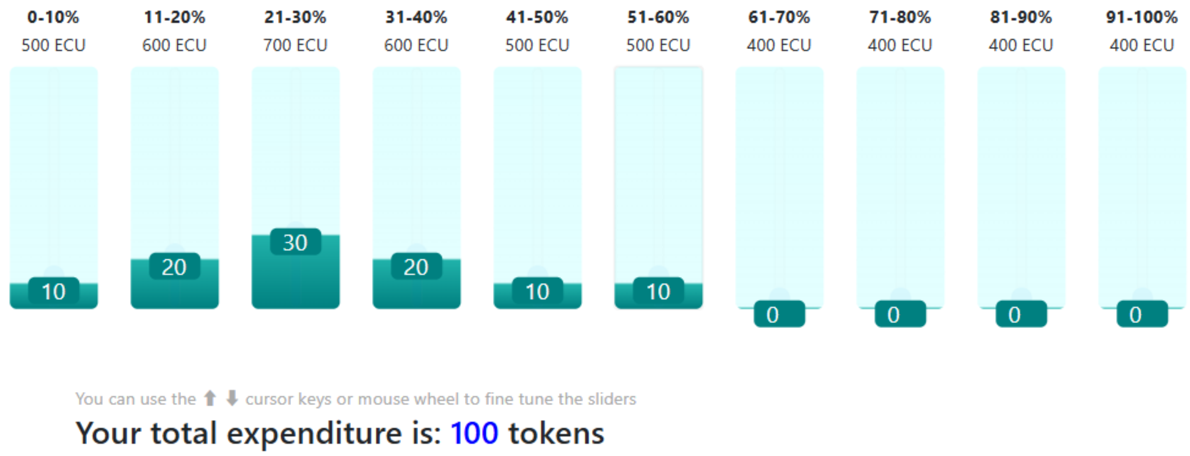

Figure A1: Audit-probability distribution elicitation interface

Notes: The figure provides an English translation of the participant-facing interface administered in Chinese. Participants allocated 100 tokens across ten audit-probability intervals. The filled sliders show an illustrative allocation and do not represent participant data. For each possible interval containing the true audit probability, the interface displays the corresponding belief-elicitation payoff in ECU.

Table A1: Participant characteristics and post-experimental questionnaire responses

	No peer information	Sparse network	Dense network	Joint p-value
Panel A: Predetermined characteristics				
Age (years)	20.230 (2.448)	19.886 (2.107)	20.287 (2.352)	.436
Female	0.638 (0.482)	0.663 (0.474)	0.707 (0.456)	.440
Monthly expenditure category (1–6)	3.755 (1.355)	3.739 (1.417)	3.771 (1.465)	.977
Expected career sector (1–4)	3.755 (0.752)	3.663 (0.890)	3.824 (0.683)	.139
Time with friends per day (hours)	3.954 (3.125)	4.060 (4.137)	3.872 (3.704)	.921
Communist Party member	0.204 (0.404)	0.179 (0.385)	0.271 (0.446)	.220
Bachelor's student	0.791 (0.408)	0.891 (0.312)	0.782 (0.414)	.053
Panel B: Preference and cognitive measures				
Risk willingness (1–7)	4.189 (1.559)	4.087 (1.580)	3.989 (1.414)	.358
Prosociality (1–7)	4.786 (1.567)	4.663 (1.671)	4.793 (1.542)	.359
Ambiguity aversion (1–5)	3.587 (1.066)	3.489 (1.136)	3.606 (1.163)	.466
Cognitive reflection score (0–3)	2.719 (0.580)	2.576 (0.697)	2.676 (0.599)	.025
Probability-reasoning score (0–2)	0.801 (0.578)	0.902 (0.628)	0.798 (0.622)	.038
Panel C: Post-experimental attitudes				
Acceptance of tax evasion (1–7)	3.281 (1.676)	3.174 (1.722)	3.266 (1.713)	.812
Perceived tax-system fairness (1–7)	4.383 (1.496)	4.293 (1.551)	4.426 (1.440)	.782
Perceived government-service fairness (1–7)	4.852 (1.341)	4.788 (1.388)	4.867 (1.395)	.830
Trust in government (1–7)	4.995 (1.524)	4.946 (1.571)	5.122 (1.399)	.502
Participants	196	184	188	
Sessions	9	8	9	

Notes: Entries are participant means with standard deviations in parentheses. Treatment was assigned at the session level. Joint p-values are from Wald tests of equality across the three information environments using session-clustered robust standard errors. All measures were collected in the questionnaire after the experimental decisions. Panels A and B contain participant characteristics and preference and cognitive measures. Panel C contains post-experimental attitudes and is not used as baseline balance evidence. Higher values indicate more of the named construct. Expected career sector and expenditure are categorical codes, so their means are compact descriptive summaries.

Table A2: Network-presence effects by experimental block

	Compliance		Perceived audit probability	
	(1) Initial block	(2) Post-change block	(3) Initial block	(4) Post-change block
Sparse network	0.096***	0.112***	0.044***	0.012
	(0.026)	(0.027)	(0.011)	(0.014)
Dense network	0.074***	0.146***	0.039***	0.016
	(0.023)	(0.025)	(0.009)	(0.013)
Constant	0.548*	0.821***	0.271**	0.349***
	(0.291)	(0.279)	(0.102)	(0.109)
Controls	Yes	Yes	Yes	Yes
Observations	568	568	568	568
Sessions (clusters)	26	26	26	26
R-squared	0.041	0.061	0.051	0.032
p-value: dense network = sparse network	0.459	0.314	0.650	0.639

Notes: Participant-round OLS using the signal-free first round of the indicated experimental block. No peer information is omitted. All columns use the full set of applicable controls from Table 2: the audit environment and experimental wave, together with age, gender, education level, Communist Party membership, expenditure category, and time spent with friends. The block indicator is constant within each column, and enforcement order is collinear with the audit environment. Perceived audit probability, $\widehat{p}_{it}$, is the midpoint-weighted mean of the elicited subjective distribution. The final row reports p-values from tests of equality between the dense- and sparse-network coefficients. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: Robustness of enforcement responsiveness to alternative clustering

	Dependent variable: enforcement responsiveness			
	Treatment factor model		Ordered model	
	Realized group-round (1)	Participant and group-round (2)	Realized group-round (3)	Participant and group-round (4)
Sparse network	0.021	0.021		
	(0.014)	(0.024)		
Dense network	0.081***	0.081***		
	(0.015)	(0.025)		
Expected peer signals			0.027***	0.027***
			(0.005)	(0.008)
Constant	0.070	0.070	0.066	0.066
	(0.088)	(0.127)	(0.088)	(0.127)
Controls	Yes	Yes	Yes	Yes
Observations	5,112	5,112	5,112	5,112
Group-round clusters (low/high)	1,278/1,278	1,278/1,278	1,278/1,278	1,278/1,278
Participants (clusters)	—	568	—	568
R-squared	0.039	0.039	0.039	0.039
p-value: dense network = sparse network	< 0.001	0.014		

Notes: Each observation pairs the same participant and round number across the initial and post-change blocks. Enforcement responsiveness is compliance in the 25 percent audit environment minus compliance in the 5 percent audit environment. Controls include indicators for enforcement order, experimental wave, and matched round, together with age, gender, education level, Communist Party membership, expenditure category, and time spent with friends. Each matched observation contains one realized group-round identifier from each enforcement environment. Columns (1) and (3) cluster standard errors over both component group-round identifiers; columns (2) and (4) additionally cluster by participant. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: Estimated information-network effects in early and later post-change rounds

	Rounds 2–5	Rounds 6–10	p-value: early = later
Panel A: After enforcement decreases to 5%			
Effects relative to no peer information			
Sparse network	0.003	0.042	0.288
	(0.035)	(0.052)	
Dense network	-0.133^{***}	-0.087	0.185
	(0.040)	(0.053)	
Pairwise comparison			
Difference: dense − sparse	-0.136^{***}	-0.128^{***}	0.865
	(0.034)	(0.044)	
Panel B: After enforcement increases to 25%			
Effects relative to no peer information			
Sparse network	0.034	0.014	0.534
	(0.035)	(0.026)	
Dense network	0.084^{***}	0.027	0.100
	(0.025)	(0.026)	
Pairwise comparison			
Difference: dense − sparse	0.049	0.013	0.221
	(0.032)	(0.021)	
Controls	Yes		
Observations	5,112		
Sessions (clusters)	26		

Notes: The table extends equation (10) by interacting the information conditions and direction of enforcement change with an indicator for early rounds 2–5. The first two columns report the estimated network effects for the indicated direction and window. The first two rows in each panel compare the sparse and dense networks with no peer information. The pairwise comparison reports the estimated difference between the two network effects. The final column reports p-values from tests that the early- and later-round effects are equal. Controls include round and experimental-wave indicators, age, gender, education, Communist Party membership, expenditure category, and time spent with friends. Session-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Baseline-adjusted learning by direction of enforcement change

	(1) Post-signal perceived-probability accuracy	(2) Post-signal true rate mass	(3) Post-signal reduction in belief entropy
Panel A: Enforcement decreases from 25% to 5%			
Treatment factor specification			
Sparse network	0.008	0.050**	0.027**
	(0.008)	(0.022)	(0.012)
Dense network	0.052***	0.157***	0.095***
	(0.012)	(0.024)	(0.011)
Ordered specification			
Expected peer signals	0.018***	0.053***	0.032***
	(0.004)	(0.008)	(0.004)
Panel B: Enforcement increases from 5% to 25%			
Treatment factor specification			
Sparse network	−0.001	0.030***	0.040***
	(0.006)	(0.011)	(0.010)
Dense network	0.017*	0.069***	0.045***
	(0.008)	(0.017)	(0.012)
Ordered specification			
Expected peer signals	0.005*	0.023***	0.016***
	(0.003)	(0.005)	(0.004)
Panel C: Equality across directions of enforcement change (p-values)			
Sparse network	0.385	0.411	0.431
Dense network	0.029	0.009	0.006
Expected peer signals	0.026	0.008	0.008
Baseline belief measure	Yes	Yes	Yes
Experimental-wave control	Yes	Yes	Yes
Observations	568	568	568
Sessions (clusters)	26	26	26

Notes: Participant-level ANCOVA estimates. The dependent variables are averages over post-signal rounds 2–10. Each specification conditions on the corresponding belief measure from the signal-free first round of the post-change block and controls for experimental wave. Panels A and B report direction-specific effects from pooled models that allow the treatment effects to differ by the direction of enforcement change. Panel C reports tests of equality between the corresponding effects in Panels A and B. The treatment factor and ordered estimates in each column come from separate regressions. No peer information is omitted from the treatment factor specification. The ordered specification replaces the treatment indicators with expected peer signals (0, 1.5, and 3). Higher values in columns (1)–(3) indicate greater accuracy, more probability mass on the true rate interval, and lower belief entropy (more concentrated beliefs), respectively. Standard errors clustered by session are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Belief–compliance relationship by information condition

	No peer	Sparse	Dense
Prior perceived audit probability, per 10 percentage points	0.049***	0.042***	0.036***
	(0.007)	(0.005)	(0.008)
p-value: equal slopes across conditions		0.506	
Prior belief entropy, per 10 percentage points	0.001	0.008	0.005
	(0.008)	(0.006)	(0.005)
p-value: equal slopes across conditions		0.796	

Notes: The dependent variable is the next declaration. All estimates come from one pooled model containing interactions between information condition and the preceding perceived audit probability and belief entropy. The model includes participant and session-round fixed effects and controls for the participant's preceding audit outcome, the number of observed peer audits, and the number of peer signals. Standard errors are clustered by session. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Perceived social norms across enforcement environments

	Full declaration appropriateness (1)	No declaration appropriateness (2)	Full-minus-none norm gradient (3)
Panel A: Average change from the 5% to the 25% environment			
Average change	0.054***	−0.075***	0.129***
	(0.011)	(0.013)	(0.021)
Panel B: Differences in the change relative to no peer information			
Sparse network	0.056*	−0.015	0.071
	(0.030)	(0.023)	(0.048)
Dense network	0.011	−0.034	0.045
	(0.022)	(0.033)	(0.044)
p-value: dense = sparse	0.152	0.606	0.662
Panel C: Ordered specification			
Expected peer signals	0.004	−0.011	0.015
	(0.007)	(0.011)	(0.015)
Controls		Yes	
Observations		568	
Sessions (clusters)		26	
R^2, treatment factor specification	0.070	0.054	0.068
R^2, ordered specification	0.065	0.054	0.067

Notes: Each observation is one participant's appropriateness score in the 25 percent audit environment minus the corresponding score in the 5 percent environment. Following the Krupka–Weber scoring convention, each appropriateness judgment is coded on an equally spaced index from −1 to 1; the full-minus-none gradient therefore ranges from −2 to 2. Panel A reports the average fitted change from the controlled treatment factor model. Panel B reports treatment coefficients, with no peer information omitted. Panel C replaces the treatment indicators with the expected number of peer signals (0, 1.5, and 3). Controls include indicators for enforcement order and collection wave and the predetermined participant characteristics reported in Appendix Table A1. Standard errors clustered by session are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: Robustness of the model-based exact disclosure counterfactual

	(1) Low enforcement (5% audit probability)	(2) High enforcement (25% audit probability)
Perceived probability and entropy	-0.111^{***}	-0.055^{**}
	(0.017)	(0.021)
Perceived probability only	-0.089^{***}	-0.029^{***}
	(0.006)	(0.002)
Lagged compliance	-0.071^{***}	-0.037^{*}
	(0.015)	(0.019)
Regime-specific slopes	-0.095^{***}	-0.024
	(0.020)	(0.024)

Notes: Entries are predicted changes in compliance when the prior perceived audit probability is replaced by the true rate and prior entropy is set to zero when included in the specification. The first row uses the belief–compliance model reported in Table 6. The three alternatives retain its sample of 10,224 participant-round observations, controls, participant fixed effects, session-by-round fixed effects, and session-clustered inference. They change only the belief specification by omitting entropy, adding the previous declaration, or allowing the perceived-probability and entropy coefficients to differ across enforcement environments. Predictions are evaluated for the no-peer-information environment. Standard errors clustered by session are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Responsiveness to Own and Peer Audit Outcomes

We estimate an updating regression as a reduced-form empirical counterpart to the signal weights in equation (4), following the sensitivity-comparison approach of Conlon et al. (2022).[12] This specification distinguishes responsiveness to each audit outcome from the cumulative effect of observing more outcomes. Own and peer audit outcomes are presented simultaneously and are generated by the same hidden audit probability. Let $a_{i,t-1}$ denote whether participant i was audited in the previous round and let $a_{j,t-1}$ denote whether observed peer j was audited in that round. For information condition g, we estimate

$$\Delta\widehat{p}_{it} = \beta_o^g(a_{i,t-1} - \widehat{p}_{i,t-1}) + \beta_s^g \sum_{j\in\mathcal{N}_{it}} (a_{j,t-1} - \widehat{p}_{i,t-1}) + \alpha_i + \eta_{st} + \varepsilon_{it}, \tag{11}$$

where α_i denotes participant fixed effects and η_{st} session-by-round fixed effects. The first forecast error captures surprise in the participant's own audit outcome. Because peer forecast errors enter as a sum, β_s^g measures responsiveness to each peer outcome. The coefficients are reduced-form updating sensitivities expressed in perceived probability units rather than structural likelihood weights. Table A9 reports the estimates.

Participants respond significantly to both own and peer audit outcomes. In the sparse network, the per-peer coefficient is .028, approximately 45% of the .062 own outcome coefficient. In the dense network, the corresponding estimates are .031 and .059, or approximately 53%. Participants respond significantly more to their own outcome than to each peer outcome in both network conditions. The per-peer estimates are statistically indistinguishable across the sparse and dense networks ($p = .475$), as are the own outcome estimates ($p = .765$). Their near equality across the two network conditions supports the interpretation that greater information reach operates by supplying more signals rather than changing the influence of each signal.

[12] Conlon et al. (2022) compare sensitivity to equally relevant information originating from oneself and from others and complement that comparison with a structural likelihood-weight model. We adapt their reduced-form comparison to our repeated belief reports. Related work estimates the weights placed on prior and signal information in log-odds formulations (Grether, 1980; Möbius et al., 2022).

Table A9: Responsiveness of perceived audit probabilities to own and peer audit outcomes

	No peer information (1)	Sparse network (2)	Dense network (3)
Own audit outcome	0.066***	0.062***	0.059***
	(0.011)	(0.006)	(0.006)
Peer audit outcome	—	0.028***	0.031***
		(0.003)	(0.004)
Participant fixed effects		Yes	
Session-by-round fixed effects		Yes	
Observations		10,224	
Participants		568	
Sessions		26	
Equality tests (p-values)			
Dense network = sparse network, own outcome		0.765	
Dense network = sparse network, peer outcome		0.475	
Own outcome = peer outcome, sparse network		< 0.001	
Own outcome = peer outcome, dense network		0.001	

Notes: The dependent variable is the round-to-round change in the perceived audit probability. Own-outcome surprise is the participant's preceding audit outcome minus the preceding perceived audit probability. Peer-outcome surprise is the sum, across observed peers, of the same forecast error; its coefficient therefore measures responsiveness to each peer outcome. The coefficients shown are linear combinations from a pooled regression allowing own- and peer-outcome sensitivities to differ across information conditions. Standard errors clustered by session are in parentheses. Stars test each coefficient against zero. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.