Privacy Concerns, Voluntary Disclosure of Information, and Unraveling: An Experiment*

Volker Benndorf† Dorothea Kübler‡ Hans-Theo Normann§

January 2015

Abstract

We study the voluntary revelation of private information in a labor-market experiment where workers can reveal their productivity at a cost. While rational revelation improves a worker's payoff, it imposes a negative externality on others and may trigger further revelation. Such unraveling can be observed frequently in our data although less often than predicted. Equilibrium play is more likely when subjects are predicted to conceal their productivity than when they should reveal. This tendency of under-revelation, especially of low-productivity workers, is consistent with the level-k model. A loaded frame where the private information concerns the workers' health status leads to less revelation than a neutral frame.

JEL Classification numbers: C72, C90, C91

Keywords: information revelation, level-k reasoning, privacy.

* We are grateful to three referees, Ismael Martínez-Martínez, Scott Peppet, and audiences and seminar participants at EEA Gothenburg, ESA New York, Göttingen, Heidelberg, RES Manchester and SABE Granada for helpful comments. Thanks to Nina Bonge who programmed and helped conducting the experiments. Financial support through CRC 649 ("Economic Risk") is gratefully acknowledged.

† Duesseldorf Institute for Competition Economics (DICE), Tel: +49 211 8115058, +49 211 8115499; email: *benndorf@dice.hhu.de.*

‡ WZB Berlin and Technical University Berlin, Tel: +49 30 25491 440; email: *kuebler@wzb.eu.*

§ Duesseldorf Institute for Competition Economics (DICE), Tel: +49 211 8115297, +49 211 8115499; email: *normann@dice.hhu.de.*

1 Introduction

Privacy concerns and the treatment of personal data are at the center of current policy debates.[1] With the rise of digital data processing and the increased communication of information via the Internet, a wealth of personal data can be easily accumulated and distributed. As a result, enterprises and governmental institutions alike face new challenges of how to adequately handle the private data of their citizens and clients.

Situations where subjects *voluntarily* disclose private information are regarded as increasingly important. For example, prospective tenants or job applicants often voluntarily disclose verified personal information. Online services such as MyBackgroundCheck.com provide verified information on drug tests, criminal records and previous rental addresses to prospective landladies or employers.[2] New-generation passports and identity cards may contain biometric data; often these are optional. Finally, health or pregnancy tests are often voluntarily provided to existing or future employers.[3] A large part of the privacy debate regards it as highly relevant whether the revelation of sensitive information is voluntary.[4]

But even when the disclosure of personal information is voluntary, privacy issues can arise due to unraveling effects. Unraveling is a signaling process which works as follows: in a world where credible signals can easily be obtained and distributed, these signals will be used by those with the best medical records, credit scores, etc. This can force others to disclose similar information about themselves because not disclosing will be interpreted as a signal of low quality. In other words, unraveling is the result of signaling where revelation of high quality leads to expectations of low quality for those who do not disclose and, in turn, to more revelation. Thus, granting people the right to decide whether to disclose can be less of a voluntary choice than it seems at first sight. Or, as Posner (1998, p. 103) succinctly puts it "*As for privacy in general, it is difficult to see how a pooling equilibrium is avoided in which privacy is 'voluntarily' surrendered, making the legal protection of privacy futile.*"

The importance of the unraveling argument is also reflected in the legal debate. Peppet (2011) summarizes the legal perspective and argues that the voluntary disclosure of private information is crucial because of unraveling effects. The challenge to regulating voluntary

[1] To quantify this statement, we conducted a Google Books Ngram Viewer comparison of several keywords and compared them to the term "privacy concerns". We found that the use of the term "privacy concerns" in the English literature has been increasing steadily since the 1970s. This is in contrast to other topics like "nuclear threat" (in decline and nowadays occurring less frequently than "privacy concerns") or "racial discrimination" (more frequent than "privacy concerns", but also in decline).

[2] Connolly (2008) explicitly advises applicants in the job market to use such online services (pp. 59-60).

[3] Some of Apple Inc. suppliers screened their workers with health and pregnancy test (Apple Inc., 2012). See also New-York Times, January 26, 2012. Further examples, discussed in Peppet (2011), include car insurance policies or rental car contracts where drivers can voluntarily agree to have the car monitored with GPS-based systems.

[4] See Curtis (2006) for the debate in Australia, Acharya and Kasprzycki (2010) for Canada, Probst (2011) for Germany, Grijpink (2001) for the Netherlands.

disclosure is that there are always some agents in whose interest it is to disclose their information. Limits to inquiry that forbid an uninformed party from seeking information from an informed counterpart may not be sufficient as the informed party might feel that it is in her interest to disclose the information. A means to avoid unraveling may be to completely forbid the use of certain information, as for example in the 2008 US *Genetic Information Nondiscrimination Act* (GINA) which prohibits the use of genetic information by insurers.

We study the voluntary disclosure of information in a laboratory experiment with the help of a revelation game. In a labor market, workers can truthfully reveal their productivity at a positive cost. Costly and truthful revelation can be seen as a way to overcome the lemons problem (Akerlof, 1970). Rational revelation imposes an externality on others because it lowers the wage paid to other workers. Complete unraveling occurs when all workers reveal—except for the one with the lowest productivity who is identified by the fact that she does not reveal her productivity.

Our research questions are to what extent subjects reveal their productivity in an experiment, and whether these choices are in line with the equilibrium predictions. We further investigate how revelation choices depend on the productivity of the worker, the characteristics of the market, and the contextual framing of the choice.

We consider markets where revelation comes at a positive cost. Such costs may consist of the time and effort of the player involved or the payments to (legal, medical, etc.) specialists who conduct the certification. This does not only appear realistic in some cases, it also allows for a richer outcome space: not all players reveal their productivity in equilibrium. With zero revelation costs, revealing is always rational and there cannot be any mistaken revelation decisions. In contrast, if the costs of revelation are strictly positive, the share of workers who reveal depends on the revelation cost and the distribution of productivities. Depending on these parameters, there may be equilibria with complete unraveling, with partial unraveling, or with no revelation at all.

We implement three different experimental markets where either a high, a medium or a low degree of unraveling should occur according to the theory. We focus on the revelation decisions of the workers: employers are not represented by laboratory participants, so all results are driven by the behavior of the workers. Including employers in our experiments would come at the expense of adding another potential confound of unraveling (for example, social preferences between the worker and the employers).

Our results are as follows. The equilibrium predictions for the three markets capture the differences in observed aggregate revelation rates across these markets well. We observe a significant amount of unraveling. At the same time, we find that revelation rates are somewhat lower than in equilibrium in two of the three markets. Workers who are supposed to reveal their productivity in equilibrium fail to take the equilibrium choice significantly more often than workers who should conceal. We will argue that this pattern is consistent with behavioral models such as the level-k model of bounded rationality (see the literature survey below). We also find a statistically and economically significant framing effect: there is less revelation in the contextualized sessions. Thus, it appears that our labor-market frame where workers can provide employers with a health certificate triggers privacy concerns.

Taken together, our results confirm the concerns about voluntary revelation raised in the privacy debate. We observe robust revelation rates, suggesting this behavior is likely to occur in voluntary disclosure regimes in the field where incentives for revelation may be even stronger (see our Conclusion). Thus, unraveling effects should be considered in the context of privacy policies.

In the next section, we review the relevant literature. Section 3 introduces the revelation game and Section 4 the experimental implementation and the different treatments. Section 5 reports on the results. Section 6 investigates reasons for the behavioral patterns we observe and Section 7 concludes.

2 Related Literature

Following the introduction of the lemons problem by Akerlof (1970), the costly but truthful revelation of private information—the certificate solution to the lemons problem—was suggested by Viscusi (1978). Subsequently, it was shown by Grossman and Hart (1980), Grossman (1981), Milgrom (1981), Jovanovic (1982), and Milgrom and Roberts (1986) that taking no action (not acquiring a certificate) may reveal an agent's type when other agents have an incentive to disclose information. Specifically, they pointed out that complete unraveling will result when revelation costs are negligible.

More recently, Hermalin and Katz (2006) investigated the impact of privacy regimes on consumer and producer rents in markets with price discrimination, taking into account unraveling effects. They argue that markets may be ex post efficient due to unraveling. However, laws banning unraveling can improve welfare ex ante because the socially wasteful revelation costs can be avoided.

While information disclosure has received much attention in the theoretical literature, only Forsythe, Isaac, and Palfrey (1989) have studied unraveling in an experiment. They study a game where sellers have superior information about the good compared to the buyers and can decide whether to reveal this information. The game has multiple Nash equilibria. Full unraveling in the sense of sellers disclosing their private information about the good takes place in the unique sequential equilibrium, which experimental subjects learn to play in the course of several rounds of play. Our game differs from the one in Forsythe et al. (1989) in several aspects. As mentioned above, we introduce strictly positive revelation costs, and subjects in the experiment only take on the role of sellers (workers), not of buyers (employers). Also, our game has a unique Nash equilibrium with partial unraveling.

There is also an empirical literature on the topic based on field data. Jin (2005) reports evidence on incomplete unraveling among Health Maintenance Organizations which may disclose information on the quality of their services on a voluntary basis. As for mandatory disclosure Jin and Leslie (2003) find that the introduction of hygiene quality grade cards for restaurants increases the consumers' sensitivity for hygiene issues in restaurants. More recently, Lewis (2011) pointed out that a lack of (voluntarily provided) ex-post verifiable information on used cars (photos, text hinting at rust, scratches, etc.) has a negative influence on the selling price in internet auctions.

Our experiments have some bearing on the question of how people make choices regarding their personal data. To our knowledge, we are the first to study the unraveling of privacy experimentally.[5] Related to our framing treatment, there is a study on the framing effects of defaults used in electronic commerce for various privacy settings (Johnson, Bellman, and Lohse, 2002). Experiments have also been used to investigate decisions regarding personal data. When making purchasing decisions, consumers have been found to provide personal data freely, even when it is relatively easy and costless to avoid it (see Acquisti and Grossklags, 2005; Beresford, Kübler, and Preibusch, 2012). This behavior in combination with a strong concern for privacy protection voiced in surveys has been called the *privacy paradox*.

In an experiment on information acquisition and revelation, Schudy and Utikal (2012) investigate the impact of different data security schemes on information acquisition. Subjects can acquire the results of a binary test (for example, an HIV test). The data security regimes are perfect privacy (no one but the testee gets to know the test result), imperfect privacy (there is a 50% chance that the results of the test will be leaked to a player interacting with the testee), and automatic dissemination where the test results are automatically disclosed to both players. The authors find that almost all subjects take the test (that is, they acquire information) in perfect privacy and in imperfect privacy. The only treatment where incomplete information acquisition is observed is automatic dissemination, that is, the treatment where all test results are revealed to both players. We do not study information acquisition decisions, but our focus is on the externality imposed on others by revealing private information.

Our main behavioral model, the level-k model, was introduced by Stahl and Wilson (1995) and Nagel (1995). Its original application was to explain subjects' behavior in beauty-contest games (see, for example, Bosch-Domenech, Montalvo, Nagel, and Satorra, 2002; Kocher and Sutter, 2005; Brañas Garza, Garcia-Muñoz, and González, 2012). Further applications include private-value auctions (Crawford and Iriberri, 2007) or centipede games (Kawagoe and Takizawa, 2012; Ho and Su, 2013). Games such as the "11-20 money request game" have specifically been designed for the elicitation of k-levels (Arad and Ru-

[5] Signaling games are broadly related to the revelation game we study. The experimental literature on signaling includes early contributions like Miller and Plott (1985), Brandts and Holt (1992), Potters and van Winden (1996), and Cooper, Garvin, and Kagel (1997), and more recent papers like Kübler, Müller, and Normann (2008), Cooper and Kagel (2009) and de Haan, Offerman, and Sloof (2011).

binstein, 2012; Lindnera and Sutter, 2013). Extensions of the model have been developed by Camerer, Ho, and Chong (2004) and Goeree and Holt (2004).

3 The Revelation Game

Our design is based on a simple labor market which we call revelation game. There are $n \geq 2$ workers with $n \in \mathbb{N}$. Worker i has productivity θ_i. Let $\Theta = \{\theta_1, \theta_2, ..., \theta_n\}$ and assume w.l.o.g. that $\theta_1 \leq \theta_2 \leq ... \leq \theta_n$.

All n workers simultaneously choose between two actions, to *reveal* or to *conceal* their productivity. Revelation causes a cost of $c > 0$ and correctly reveals the worker's productivity. Let $I_i \in \{0, 1\}$ be a function indicating whether worker i has chosen to reveal her productivity, with $I_i = 1$ denoting revelation and $I_i = 0$ concealment.

Workers' payoffs are as follows. If worker i chooses to reveal, i earns her productivity minus the revelation cost. If not, she receives the average productivity of all workers who have chosen not to reveal. Formally, i's payoff is

$$\Pi_i = \begin{cases} \theta_i - c & \text{if } I_i = 1 \quad (\textit{reveal}) \\ \sum_{j=1}^{n}(1 - I_j)\theta_j / \sum_{j=1}^{n}(1 - I_j) & \text{if } I_i = 0 \quad (\textit{conceal}). \end{cases}$$

These payoffs can be thought to arise in a competitive labor market where two or more employers bid for workers and earn the workers' (expected) productivity. The employers in this labor market know the set Θ, that is, they know the n payoff functions of the n workers, but do not know which worker has which payoff function. Employers earn an expected payoff of zero. Since employers are not part of our game, this is a static game of complete information and the equilibrium concept is Nash equilibrium.[6]

Proposition 1. *In any pure strategy Nash equilibrium of the revelation game, we have* $I_n^* \geq I_{n-1}^* \geq ... \geq I_2^* \geq I_1^* = 0$.

The proof can be found in Appendix C. The proposition has two implications. First, there is a sorting effect in that $I_i < I_j$ for $i > j$ is impossible. The revelation decisions are monotonic in the productivity. Second, at least worker 1 (and possibly more workers) will

[6] If workers first take the revelation decisions and employers then bid a wage, this results in a dynamic game with incomplete information. This game with employers is analyzed in another paper (Benndorf, 2014). The same worker productivities are implemented for the three markets. Some of the predictions remain the same (in Markets A and B), but in Market C there is a unique equilibrium in pure strategies in the game without employers, but two pure and one mixed equilibria in the game with employers. As for differences in behavior, Benndorf (2014) shows that the revelation decisions of workers are similar in both games, the main differences being that high-productivity workers reveal slightly more frequently and that low-productivity workers reveal slightly less frequently compared to the same workers in the game without employers.

conceal in equilibrium. Here, the positive revelation cost in our model leads to interesting departures from the previous literature. For $c = 0$, our model suggests that all players (except for the worker with the lowest productivity, who is indifferent) reveal. For $c > 0$, the proposition allows for the pattern of equilibrium actions $I_1 = ... = 0 < I_m = ... = I_n = 1$, with $1 < m \leq n$. Accordingly, in our markets described below, the model predicts that several low-productivity workers will conceal.[7]

Furthermore, multiple equilibria can exist when $c > 0$. To characterize the conditions under which there is a unique equilibrium, the following definition is helpful:

Definition 1. *Let* $\bar{\theta}(s) = \frac{1}{s}\sum_{i=1}^{s}\theta_i$. *Further, define* $C = \left\{i|\theta_i - c \leq \bar{\theta}(i)\right\}$ *and* $R = \left\{i|\theta_i - c \geq \bar{\theta}(i)\right\}$.

In words, $\bar{\theta}(s)$ is the average of the productivities of all workers $1, 2, ..., s$. The set C contains all workers whose best response is to conceal given that all workers with lower (higher) productivity conceal (reveal). And R is the set of all workers whose best response is to reveal given that all workers with lower (higher) productivity conceal (reveal). When $c = 0$, $\theta_i - c \geq \bar{\theta}(i)$ holds for all $i \geq 1$, so $R = \{1, 2, 3, ..., n\}$.

Proposition 2. *The revelation game has a unique pure strategy equilibrium if and only if* $\max(C) < min(R)$.

See the Appendix for a proof. While the games we use in our experiment all have a unique pure strategy equilibrium, it is easy to construct cases with multiple equilibria (including mixed equilibria) with the help of the proposition.[8]

4 Experimental Design and Procedures

In each of our experimental markets, there are $n = 6$ workers. We design three different markets, A, B, and C, with different realizations of Θ. The cost of revelation, c, always equals 100 experimental currency units; it does not vary across workers, markets or treatments. The different productivities in each market are reported in Table 1. The entries in bold type indicate that the corresponding worker reveals her productivity in equilibrium.

The revelation costs in our experiment are relatively high. Such high costs may exist in the field with a complicated certification process requiring a number of costly tests

[7] Of course, when revelation costs are prohibitively high, the prediction is $I_1 = I_2 = ... = I_n = 0$.

[8] Suppose $n = 3$, $c = 100$, $\theta_1 = 200$, $\theta_2 = 402$ and $\theta_3 = 403$. We have $\bar{\theta}(1) = 200, \bar{\theta}(2) = 301$ and $\bar{\theta}(3) = 335$. Therefore, $C = \{1, 3\}$ and $R = \{2\}$. There are multiple equilibria: in one pure strategy Nash equilibrium, all workers conceal and we have $\theta_i - c < 335$ for $i = 1, 2, 3$; in a second pure strategy Nash equilibrium, workers two and three reveal and we have $\theta_1 - c < 200$, $\theta_{2,3} - c > \frac{200+\theta_{2,3}}{2}$; there is also a mixed equilibrium where worker one reveals with zero probability while the workers two and three reveal with the probabilities $\frac{64}{67}$ and $\frac{33}{34}$, respectively. Another possibility to generate mixed equilibria is to require $\bar{\theta}(j) = \theta_j - c$ for at least one worker j.

(say, fMRI scans or dangerous medical analyses like cardiac catheters) or elaborate legal procedures. When it comes to the digital dissemination of simple information in the field, revelation costs may be rather low.

The three markets are played on a rotating basis. In period 1 subjects play Market A, in period 2 they play Market B, and in period 3 Market C is played before they start all over again with Market A. Each market is played five times, totaling 15 periods altogether.

Productivity	Market A	Market B	Market C
θ_1	200	200	200
θ_2	210	**448**	280
θ_3	230	**510**	360
θ_4	260	**551**	**440**
θ_5	300	**582**	**520**
θ_6	**600**	**607**	**600**

Table 1: Workers' productivities in the three different markets. Entries in bold face indicate that the player reveals in equilibrium ($I_i = 1$).

At the beginning of the experiment, subjects were randomly allocated into groups of six, and they stayed in their group for the whole experiment (fixed matching). The productivities θ_i, expressed in experimental currency units, were randomly assigned to the workers in each period. The instructions emphasized that this allocation of productivities was without replacement such that each productivity value occurs exactly once in each group and in each period.[9]

The Nash equilibrium for the three markets is as follows. In Market A, only worker 6 reveals her productivity. That is, we have $I_6 = 1 > I_5 = ... = I_1 = 0$ in equilibrium. In Market B, all workers except for worker 1 reveal: $I_6 = ... = I_2 = 1 > I_1 = 0$, and in Market C, we have $I_6 = I_5 = I_4 = 1 > I_3 = I_2 = I_1 = 0$. The motivation for employing Markets A to C is that we need qualitatively different equilibrium outcomes to be able to infer whether there is too much or too little revelation. For example, Market B may show that subjects reveal too little, but given that almost all workers should reveal in equilibrium, we need to contrast this with Market A where only one of six workers reveals in Nash equilibrium.

We consider two treatments:

– The baseline treatment, called LOADED, is based on the revelation game described in the previous section with one peculiarity. It employs a loaded labor-market frame. Subjects are told that they are acting as *workers* in a *labor market*. Their productivity is referred to as their *health status*, and subjects are informed that they need to decide whether or not to *buy a health certificate*.

[9] The instructions are contained in our working paper available at: `http://hdl.handle.net/10419/104605`

– In our second treatment, NEUTRAL, we remove the labor market frame. The productivity is called *number* and the decision is whether to disclose this number (*yes*, *no*).

Our treatments are motivated as follows. The LOADED treatment is thought to resemble a real-world situation where the disclosure of sensitive personal data may play a role. We opted for a combination of a labor market and medical data, but there are other examples that could have been chosen (for example, landlords may require prospective tenants to disclose details on their lifestyle or banks might ask for information on a customer's personal situation). Of course, we do not know how these alternative frames affect behavior, but our intuition is that a loaded frame that explicitly asks for the disclosure of sensitive information should raise more privacy concerns among our subjects than a more neutrally framed treatment. Hence, the NEUTRAL treatment is implemented in order to control for the possibility of subjects' privacy concerns elicited by the framing. If the subjects care for privacy and if our frame elicits these privacy concerns, there should be more revelation in the NEUTRAL treatment compared to the baseline treatment with the loaded frame.

The feedback given to the participants at the end of a period was as follows. In all sessions, subjects were informed of their own profits and the market wage of that period. In 11 of the 23 groups of the LOADED treatment, we gave the subjects additional information about the choices of all six workers in the group. Our hypothesis was that the additional feedback would support learning. In Appendix A, we analyze the differences between the two feedback formats in detail. However, we do not find any effect of the additional feedback whatsoever and virtually all our results remain unchanged when we do not pool the data (see Appendix A). Therefore, we ignore the differences in feedback in the LOADED treatment and pool the data in the body of the paper.

	LOADED	NEUTRAL	Σ
number of participants	138	66	204
number of independent groups	23	11	34
number of sessions	6	3	9
groups per session	3-4	3-4	

Table 2: Treatments.

We note that several features of the experimental design suggest that there might be more unraveling in the field compared to our laboratory setting. First, the simultaneous move structure necessitates players to anticipate the decisions of the other market participants. In contrast, in a sequential setting unraveling occurs even if players are only myopically best responding to the choices of others. Second, our groups of six players are matched together for the entire experiment consisting of 15 rounds. Cooperation (in the sense of joint-payoff maximization) would induce them to conceal their productivity to save on the revelation costs, again making unraveling less likely. Field settings, by contrast, may often be of a one-shot nature. Third, in our experiment the productivity of the workers is not attained by merit, but assigned randomly. Revelation might increase if workers feel entitled to a higher wage because they have invested in their productivity (compare for example, Hoffman, McCabe, Shachat, and Smith, 1994; Konow, 1996). Finally, we point out that unraveling might also be less pronounced in the field because the information to be revealed actually sensitive. However, as mentioned above, we also implement the Loaded treatment which is thought to resemble such situations.

The experiments were conducted between July and September 2011 at the experimental lab at the Technical University Berlin, using the z-Tree software package by Fischbacher (2007) and Greiner's (2004) on-line recruitment system. In total, 204 subjects participated in the experiment, a session lasted around 60 minutes, and subjects earned between 8.19 € and 12.51 €. The average payment was 10.70 €. More details on the number of participants and the independent observations gathered in this experiment can be found in Table 2.

5 Results

When we report non-parametric tests, we conservatively count each group of six players as one independent observation, and we report two-sided p-values. For the regression analyses, we employ clustering at the group level.

5.1 Main Findings

Our first research question is to what extent subjects reveal their productivity. Figure 1 displays the relevant revelation rates per market, averaged across the six workers and all periods. We observe substantial revelation rates throughout. Conspicuously, the equilibrium predictions capture the differences in observed revelation rates between the three markets well. As in equilibrium, we observe more revelation in Market B than in Market C, and there is more revelation in Market C than in Market A.

Observation 1. *In all markets and all treatments, we observe substantial levels of revelation. The differences between markets are well organized by the predictions.*

Whereas prediction and behavior are rather consistent in Market A, too little revelation compared to the prediction is observed in Market B, especially for Loaded. Using a sign test, revelation rates are significantly below the prediction in Markets B and C of both Loaded and Neutral (all $p < 0.05$). Although statistically significant, the difference is small in Market C in the Neutral treatment.

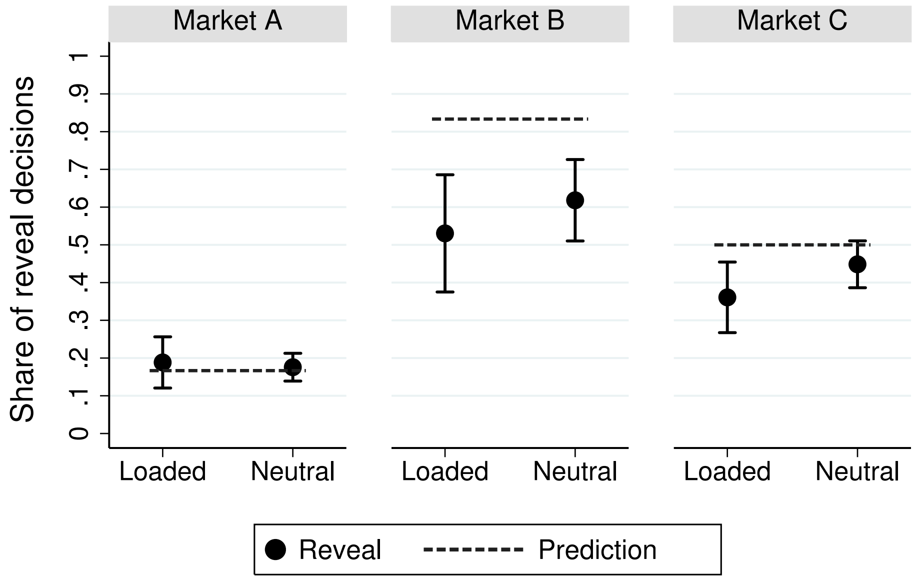

Figure 1: Average revelation rates (solid circle), standard deviations (calculated using group averages) and predictions (dashed line).

We next study the equilibrium consistency of choices at the individual level. Table 3 reports the results of probit regressions with *Consistency* as the dependent variable. *Consistency* indicates whether a subject behaves in line with the Nash prediction. The dummy variable *Reveal* indicates the equilibrium action of the corresponding decision (*Reveal* = 1). The dummy for the treatment with the loaded frame is *Loaded*, and *Period* captures possible time trends. We report results when the data from all markets are pooled (probit regression (1)) as well as probits for each market separately ((3), (4) and (6)). The regressions (2), (5) and (7) will be discussed in Section 6.

The regressions in Table 3 suggest three findings (results regarding the variable *Min_k* will be discussed below). The first result is that subjects who should reveal in equilibrium are less likely to behave according to this prediction compared to subjects who are predicted to conceal. We will analyze this further in the following paragraphs. Second, the loaded frame reduces the likelihood of observing the equilibrium action, with the exception of Market A (see Section 5.2). The third result is that there is a moderate positive and statistically significant time trend except in regression (6) (see Section 5.3).

The first and foremost observation in Table 3 is that *Reveal* is highly significant in regressions (1), (3) and (6). The marginal effect of *Reveal* on the consistency of the decision with the equilibrium prediction is -24.41% in regression (1). In the Loaded treatment, workers conceal when the equilibrium calls for revelation in 16%, 38%, and 33% of the corresponding decisions in Markets A, B and C, respectively. In contrast, the corresponding frequencies for workers who reveal when the equilibrium predicts concealment are only 6%, 10%, and 6%. This difference is highly significant using a non-parametric test (Wilcoxon signed-rank test, $p < 0.001$). In Neutral these figures read 11%, 26%, and

Consistency	All markets (1)	(2)	Market A (3)	Market B (4)	(5)	Market C (6)	(7)
Reveal	-1.098***	-0.423***	-0.577***	-1.108***	-0.143	-1.010***	-0.264*
	(0.105)	(0.107)	(0.162)	(0.175)	(0.222)	(0.134)	(0.147)
Loaded	-0.367***	-0.417***	-0.249	-0.402***	-0.485***	-0.400***	-0.455***
	(0.129)	(0.147)	(0.205)	(0.148)	(0.179)	(0.141)	(0.157)
Period	0.0639***	0.0732***	0.0863**	0.0722***	0.0893***	0.0433	0.0462
	(0.0175)	(0.0198)	(0.0382)	(0.0261)	(0.0316)	(0.0349)	(0.0388)
Min_k		-0.419***			-0.448***		-0.663***
		(0.0307)			(0.0408)		(0.0666)
Constant	1.693***	2.125***	1.581***	1.576***	2.044***	1.758***	2.455***
	(0.130)	(0.159)	(0.201)	(0.200)	(0.219)	(0.188)	(0.223)
Observations	3,060	3,060	1,020	1,020	1,020	1,020	1,020
Pseudo R2	0.136	0.225	0.0474	0.0694	0.206	0.122	0.202

Standard errors adjusted for 34 clusters (groups)
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3: Probit regression results on the equilibrium consistency of choices.

16% (if the prediction is to reveal) and 3%, 0%, and 5% (for workers with a prediction to conceal) and the difference is also significant (Wilcoxon signed-rank test, $p = 0.001$). It appears that subjects are biased against revealing their productivity, and therefore the equilibrium consistency is much higher for concealing.

Observation 2. *In both* Loaded *and* Neutral, *workers who are predicted to reveal in equilibrium violate the prediction significantly more often than workers who are predicted to conceal in equilibrium.*

We further break down the behavior of the workers for the different productivities. Table 4 summarizes the equilibrium prediction and actual play for all six workers separately in the three markets. Consider workers who are expected to reveal in equilibrium, denoted by the entry of 1 in the "Nash" column. The differences between actual revelation rates and predictions are large, ranging from 15% up to 68% in Loaded. The table shows that the inconsistencies with equilibrium play are negatively correlated with worker productivity for those workers who should reveal in equilibrium, that is, workers of lower productivity fail to reveal more often than workers of high productivity. (A sign test on the signs of Spearman's ρ calculated for each group yields $p \leq 0.001$ for both Loaded and Neutral). In contrast, inconsistencies with equilibrium play are not correlated with productivity when workers should conceal in equilibrium, and these inconsistencies are also generally small (analogous sign tests yield $p = 0.455$ for Loaded and $p = 0.508$ for Neutral). In Neutral, the under-revelation result is less pronounced, but again a correlation for out-of-equilibrium conceal decisions as well as no correlation for out-of-equilibrium reveal decisions can be found.

		Market A			Market B			Market C	
Worker	Nash	Loaded	Neutral	Nash	Loaded	Neutral	Nash	Loaded	Neutral
1	0	0.02	0.04	0	0.10	0.00	0	0.04	0.04
2	0	0.07	0.02	1	0.32	0.35	0	0.03	0.04
3	0	0.04	0.02	1	0.44	0.71	0	0.09	0.09
4	0	0.06	0.02	1	0.65	0.78	1	0.41	0.62
5	0	0.10	0.07	1	0.81	0.91	1	0.77	0.93
6	1	0.84	0.89	1	0.85	0.96	1	0.83	0.98

Table 4: Average revelation rates across markets and treatments. In the columns entitled "Nash", an entry of 0 denotes that the worker is predicted to conceal while 1 denotes that the worker is predicted to reveal.

Observation 3. *In both* Loaded *and* Neutral*, choices inconsistent with the equilibrium prediction are correlated with productivity when workers should reveal in equilibrium but not when they should conceal.*

From Table 4 it can also be taken that the discrepancy between revelation rates and predictions at the market level (see Figure 1) is related to the equilibrium share of reveal decisions. In Market A, only one in six workers is predicted to reveal and only minor discrepancies to the prediction occur. In Market B, it is the other way round: five in six workers are predicted to reveal and major discrepancies to the prediction occur. This reflects the observed tendency of under-revelation.

5.2 Framing Effect

Decision making differs between the Neutral and the Loaded treatment. As is apparent from Figure 1, there is more revelation in the Neutral treatment (Mann-Whitney U-Test, $p = 0.054$), at least in markets B and C.

When we look at equilibrium consistency, we find differences in three dimensions. First, averaging across all markets and all workers, we find that more decisions are in line with the equilibrium prediction in Neutral (87.8%) than with the loaded frame (79.8%). This difference is significant (Mann-Whitney U-Test, $p = 0.027$). Further evidence for this is delivered by our regression (1), Table 3. Here, the variable *Loaded* remains significant when we control for *Reveal* and *Period*. Its marginal effect is -7.54%.

Second, looking at markets separately, we find that the share of decisions that are in line with the prediction is higher in Neutral than in Loaded for Markets B and C. This is also captured by the regressions (4) and (6), Table 3 where the regressor *Loaded* is highly significant.

Third, the higher level of equilibrium play in neutral has a distinct pattern: there are more equilibrium revelation choices, but not more equilibrium conceal decisions in Neutral compared to Loaded. Averaging across the three markets in Loaded, we observe 34.2% conceal decisions of workers who should reveal in equilibrium and 6.2% reveal decisions for workers who should conceal in equilibrium. For NEUTRAL, the corresponding numbers are 20.8% and 3.6%. The decrease from 34.2% to 20.8% is significant (Mann-Whitney U-Test, $p = 0.024$), but the decrease from 6.2% to 3.6% is not (Mann-Whitney U-Test, $p = 0.356$).

Observation 4. *In* NEUTRAL, *subjects reveal their productivity more often than in* LOADED. *While there are significantly more choices consistent with equilibrium in* NEUTRAL, *this effect is quantitatively and statistically significant only for workers who should reveal in equilibrium.*

These findings suggest that the labor market frame in combination with the health certificate affects choices. It gives rise to preferences not restricted to the monetary incentives of the game. A share of subjects were more reluctant to disclose their productivity in the loaded treatment where private information concerned the subject's health status. Hence, the extent to which private information is revealed depends on the contextual frame, without any real privacy issues at stake in the experiment.

5.3 Learning

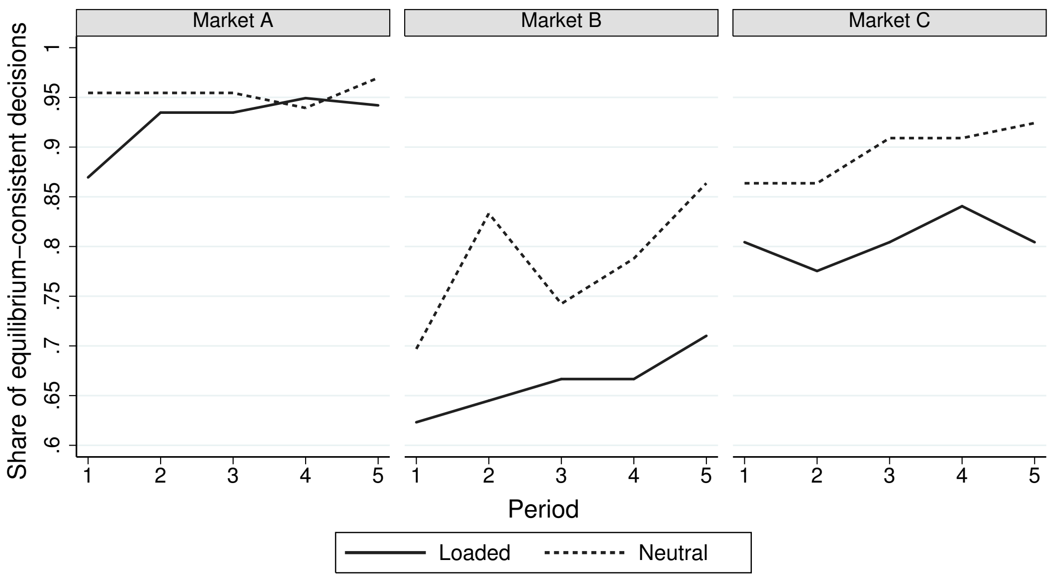

Figure 2: Fraction of decisions consistent with equilibrium across treatments and markets over time.

Do subjects learn to play the equilibrium over time? Figure 2 shows the frequency of choices consistent with equilibrium over the five periods that each market was played. There is a moderate increase in equilibrium decisions in all three markets and in both treatments. The significance of *Period* in Table 3 for all markets as well as for Market A and B separately confirm the learning effect. However, the learning effect is relatively small in Market A because decisions are already close to the equilibrium in the first period. In Markets B and C, the equilibrium is not reached even after five periods of play.

Observation 5. *The share of decisions consistent with equilibrium moderately increases in all markets of both treatments* LOADED *and* NEUTRAL*.*

6 Behavioral forces against revelation

Despite a relatively large congruence of the data with the theoretical predictions at the market level (Observation 1), we trust it is worthwhile to further investigate Observations 2 and 3 which suggest that there is a substantial tendency to under-reveal. We consider (i) level-k reasoning, (ii) quantal response equilibrium, and (iii) inequality aversion. Preferences for efficiency or surplus maximization may be relevant here, too. Revealing is socially wasteful and therefore subjects with a preference for maximizing total welfare may be disinclined to reveal. See Charness and Rabin (2002) and Engelmann (2012). While we discuss the impact of social preferences under the label inequality aversion, we note that a preference for total surplus may also have explanatory power.

Whereas the analysis is not aimed at conducting a horse race among behavioral models, we believe that level k or a limited depth of reasoning is a prime and parsimonious model for explaining our data. Accordingly, we focus on the level-k model here and relegate quantal response equilibrium and inequality aversion to the appendix. As none of the models can explain the observed framing effect, we focus on the treatment NEUTRAL while keeping in mind that in LOADED the under-revelation is even more pronounced.

To apply level-k reasoning to the revelation game, we assume that level-0 players randomize between their two actions (both equally likely), although the exact level-0 assumption qualitatively does not matter much for our game.[10] We then calculate the best replies for $k > 1$ where level-k' players (for $k' > 0$) believe that all other players reason at level $k = k' - 1$.

Figure 3 displays the required levels of reasoning for workers to pick their equilibrium action in the various markets for different productivities. The fewest iterations are required in Market A (k=1 throughout). The highest level-k requirement occurs in Market B for worker 2 who has to perform five steps of reasoning. Level-k reasoning yields Nash equilibrium choices for a finite number of steps.

Figure 3 also shows that taking the (equilibrium) decision to conceal merely requires a level of $k = 1$ throughout. By contrast, revealing requires up to $k = n - 1$ levels. As

[10] A different yet plausible assumption for level-0 types is that all workers conceal with probability one. In that case, it is straightforward to check that Markets A and B remain as in Figure 3, but in Market C also worker 5 reveals when $k = 1$ and worker 4 when $k = 2$. Less plausible in our view is the level-0 assumption that all workers reveal with probability one. If so, the prediction is the same as in the case where all workers conceal with probability one with all k-levels augmented by one. The logic is that when all players reveal, the $k = 1$ reply for all workers is to conceal with probability one.

this property is central to our research question, we prove this generally. It is impossible, for example, that a worker reveals when she is a level-1 type but conceals when she is level 2. The proof of the following proposition can be found in the appendix.

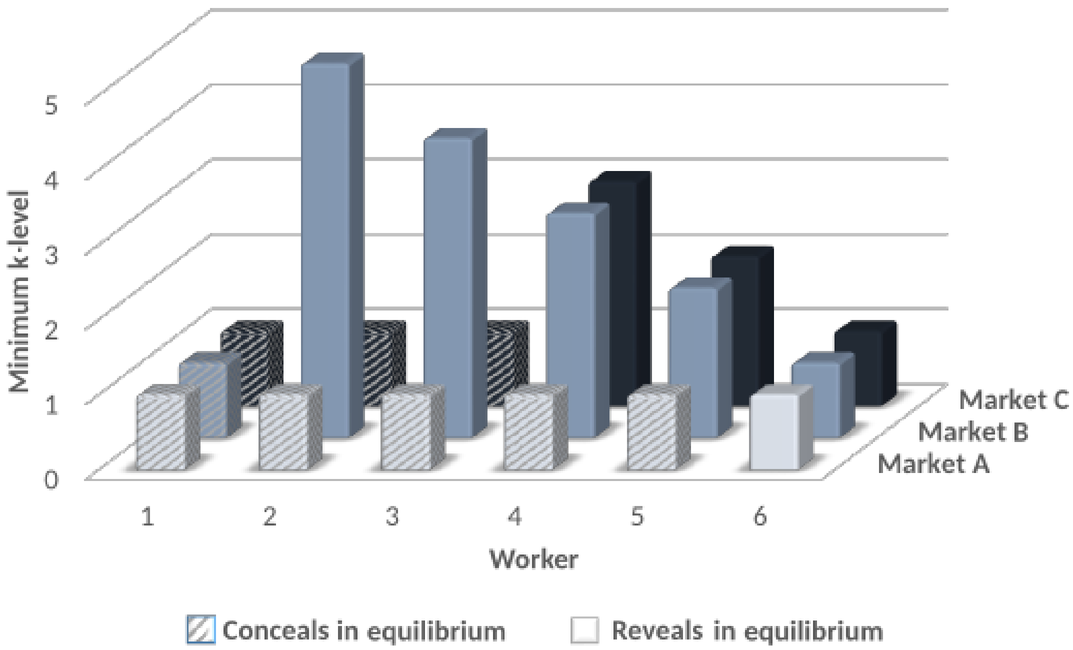

Figure 3: Minimum k-level required to play the Nash equilibrium action. Conceal decisions (shaded areas) require $k = 1$ throughout whereas revelation often demand $k > 1$.

Proposition 3. *Applying the level-k model to the revelation game, workers who conceal if they are level $k > 1$ also conceal if they are level $k = 1$.*

Proposition 3 and the level-k patterns in Figure 3 offer an explanation of Observations 2 and 3. Under the assumption that at least some players display limited depth of reasoning, Proposition 3 and Figure 3 imply (i) disproportionately more concealment in general, (ii) more consistency with the equilibrium prediction for workers who conceal than for those who reveal in equilibrium, and (iii) a positive correlation of equilibrium revelation decisions with productivities which does not hold for equilibrium conceal decisions.

An intriguing prediction the level-k model makes is that the frequency of conceal decisions should be equal to the frequency of reveal decisions by worker 6 because all of these decisions require at least reasoning at level $k = 1$ (see Figure 3). We find support for this hypothesis as we observe no significant differences in the fraction of equilibrium choices by those players (Wilcoxon signed-rank test, $p = 0.216$).[11]

[11] In LOADED, there is more under-revelation by worker 6 than over-revelation by those workers who conceal in equilibrium. This can be attributed to the framing effect causing lower revelation rates.

The relevance of level-k reasoning can also be taken from the regression analysis in Table 3. In regressions (2), (5) and (7), we additionally consider the cardinal variable *Min_k* defined as the minimum k-level required for an individual worker to choose her equilibrium action. (Since *Min_k* $= 1$ for all workers in Market A, we cannot perform the analysis for this market.) The explanatory variable *Min_k* is highly significant in all three regressions. Higher requirements on subjects' reasoning lower the likelihood for equilibrium play. *Min_k* also takes away some of the explanatory power of *Reveal*. Thus, beyond the bias subjects exhibit against revelation due to the framing effect, it is the cognitive requirement of revelation for low-productivity workers that explains our results. As low-productivity workers who are predicted to reveal need to anticipate other workers' behavior quite accurately, their decisions are more challenging compared to workers who should conceal and whose decisions are less dependent on others' choices. It follows that these low-productivity workers are more prone to making decisions that are inconsistent with the equilibrium.

7 Conclusion

We study experimental labor markets where workers can voluntarily reveal private data (their productivity) at a cost. We analyze whether and to what degree voluntary disclosure of private information may result in unraveling of privacy. The paper contains several arguments which, at first sight, appear to qualify Posner's (1998) view that privacy cannot be maintained. These arguments are as follows.

First, we show theoretically that if the cost of revelation is sufficiently high, the unraveling process may be mitigated. The equilibria of our different market games may entail unraveling and the complete disclosure of private information. However, unraveling can also be only partial in that, for instance, three or even five out of six workers are predicted to conceal and to maintain their privacy. Our experimental data documents that the equilibria are generally good predictors for the observed differences between markets.

Second, we identify behavioral forces that may further diminish workers' revelation behavior. Level-k reasoning predicts that players will conceal frequently, and only with fully rational players will there be complete unraveling. Out-of-equilibrium revelation, by contrast, will be rare as concealing in equilibrium does not require more than one step of reasoning (level $k = 1$). Our experimental data is nicely in line with the level-k predictions. We not only observe the tendency towards incomplete unraveling, our data also documents that this incompleteness is driven by exactly those agents who face an especially difficult decision according to the level-k model.

Third, we observe a framing effect which suggests that personal privacy concerns may also reduce the subjects' propensity to disclose private information. Our first treatment uses a contextualized labor-market frame where the wording suggests that the information to be disclosed is particularly sensitive; in a second treatment this loaded frame is removed.

Our data shows that subjects reveal significantly less frequently with the loaded frame compared to the neutral frame. We believe that this framing effect is driven by the subjects' privacy concerns that are switched on in the privacy-sensitive frame and switched off in the neutral frame.

Do these results and arguments refute Posner's (1998) suspicion that the protection of privacy is futile? By and large, we do not believe so. We observe a substantial degree of concealment in our lab experiment (roughly sixty percent overall where fifty percent is predicted). Yet, our results do not suggest that voluntary revelation is unimportant in the privacy debate. Unraveling is frequently observed. Furthermore, if the mechanism underlying unraveling is well understood, it can be avoided as evidenced by GINA legislation. Thus, the protection of privacy need not be futile in the end.

Finally, the level of revelation we observe may only constitute a lower bound since there are reasons why revelation in other settings may be stronger. First, decision making may be sequential rather than simultaneous. If so, unraveling may occur even for myopic players, weakening the level-k argument. Also, the repeated interaction within groups we employ makes the externality more salient and may limit unraveling. In a setting with one-shot interactions, there may be more unraveling. Third, players whose productivity is earned by merit (rather than being random) might feel entitled to a higher payoff and thus be more inclined to reveal. Further research will indicate whether these arguments have bite in explaining the unraveling of private information.

Bibliography

L. Acharya and T. Kasprzycki. Biometrics and government. Technical report, Parliamentary Information and Research Service, Library of Parliament, Ottawa, Canada, 2010.

A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security and Privacy*, 3:24–30, 2005.

G. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

Apple Inc. Apple supplier responsibility report. Technical report, Apple Inc., 2012.

A. Arad and A. Rubinstein. The 11-20 money request game: A level-k reasoning study. *American Economic Review*, 102(7):3561–73, 2012.

V. Benndorf. Voluntary disclosure of information and unraveling in the market for lemons: An experiment. Unpublished mimeo, 2014.

A. R. Beresford, D. Kübler, and S. Preibusch. Unwillingness to pay for privacy: A field experiment. *Economics Letters*, 117(1):25–27, 2012.

M. Blanco, D. Engelmann, and H. Normann. A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, 72(2):321–338, 2011.

A. Bosch-Domenech, J. Montalvo, R. Nagel, and A. Satorra. One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments. *American Economic Review*, 92(5): 1687–1701, 2002.

P. Brañas Garza, T. Garcia-Muñoz, and R. H. González. Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization*, 83(2):254–260, 2012.

J. Brandts and C. A. Holt. An experimental test of equilibrium dominance in signaling games. *American Economic Review*, 82(5):1350–1365, 1992.

C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898, 2004.

G. Charness and M. Rabin. Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3):817–869, 2002.

A. Connolly. *101 ways to score higher on your GMAT: What you need to know about the graduate management admission test explained simply*. Atlantic Publishing Company (FL), 2008.

D. J. Cooper and J. H. Kagel. The role of context and team play in cross-game learning. *Journal of the European Economic Association*, 7(5):1101–1139, 2009.

D. J. Cooper, S. Garvin, and J. H. Kagel. Signalling and adaptive learning in an entry limit pricing game. *RAND Journal of Economics*, 28(4):662–683, 1997.

V. P. Crawford and N. Iriberri. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770, 2007.

K. Curtis. Biometrics institute - privacy code. Technical report, Biometrics Institute Ltd, PO Box 567, Crows Nest NSW 1585, Australia, 2006.

T. de Haan, T. Offerman, and R. Sloof. Noisy signaling: theory and experiment. *Games and Economic Behavior*, 73(2):402–428, 2011.

D. Engelmann. How not to extend models of inequality aversion. *Journal of Economic Behavior & Organization*, 81(2):599–605, 2012.

E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868, 1999.

U. Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

R. Forsythe, M. R. Isaac, and T. R. Palfrey. Theories and tests of "blind bidding" in sealed-bid auctions. *RAND Journal of Economics*, 20(2):214–238, 1989.

J. Goeree and C. Holt. A model of noisy introspection. *Games and Economic Behavior*, 46(2):365–382, 2004.

B. Greiner. An online recruitment system for economic experiments. In K. Kremer and V. Macho, editors, *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63*, pages 79–93. Ges. fuer Wiss. Datenverarbeitung, Goettingen, 2004.

J. Grijpink. Biometrics and privacy. *Computer Law and Security Report*, 17(3):154–160, 2001.

S. J. Grossman. The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24(3):461–483, 1981.

S. J. Grossman and O. D. Hart. Disclosure laws and takeover bids. *Journal of Finance*, 35(2):323–334, 1980.

P. A. Haile, A. Hortaçsu, and G. Kosenok. On the empirical content of quantal response equilibrium. *American Economic Review*, 98(1):180–200, 2008.

B. Hermalin and M. Katz. Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics*, 4(3):209–239, 2006.

T.-H. Ho and X. Su. A dynamic level-k model in sequential games. *Management Science*, 59(2):452–469, 2013.

E. Hoffman, K. McCabe, K. Shachat, and V. Smith. Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7(3):346–380, 1994.

G. Z. Jin. Competition and disclosure incentives: An empirical study of HMOs. *RAND Journal of Economics*, 36(1):93–112, 2005.

G. Z. Jin and P. Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *Quarterly Journal of Economics*, 118(2):409–451, 2003.

E. J. Johnson, S. Bellman, and G. L. Lohse. Framing and privacy: Why opting in-opting out. *Marketing Letters*, 13(1):5–15, 2002.

B. Jovanovic. Truthful disclosure of information. *Bell Journal of Economics*, 13(1):36–44, 1982.

T. Kawagoe and H. Takizawa. Level-k analysis of experimental centipede games. *Journal Of Economic Behavior & Organization*, 82(2):548–566, 2012.

M. G. Kocher and M. Sutter. The decision maker matters: Individual versus group behaviour in experimental beauty-contest games. *Economic Journal*, 115(500):200–223, 2005.

J. Konow. A positive theory of economic fairness. *Journal of Economic Behavior & Organization*, 31(1):13–35, 1996.

D. Kübler, W. Müller, and H.-T. Normann. Job-market signaling and screening: An experimental comparison. *Games and Economic Behavior*, 64(1):219–236, 2008.

G. Lewis. Asymmetric information, adverse selection and online disclosure: The case of ebay motors. *American Economic Review*, 101(4):1535–46, September 2011.

F. Lindnera and M. Sutter. Level-k reasoning and time pressure in the 11–20 money request game. *Economics Letters*, 120(3):542–545, 2013.

R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.

P. Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12(2):380–391, 1981.

P. Milgrom and J. Roberts. Relying on the information of interested parties. *RAND Journal of Economics*, 17(1):18–32, 1986.

R. M. Miller and C. R. Plott. Product quality signaling in experimental markets. *Econometrica: Journal of the Econometric Society*, 53(4):837–872, 1985.

R. Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326, 1995.

S. R. Peppet. Unraveling privacy: The personal prospectus & the threat of a full disclosure future. *Northwestern University Law Review*, 105(3):1153–1203, 2011.

R. A. Posner. Privacy. In P. Newman, editor, *The new palgrave dictionary of economics and the law.* United Kingdom. Macmillan Reference, 1998.

J. Potters and F. van Winden. Comparative statics of a signaling game: An experimental study. *International Journal of Game Theory*, 25(3):329–353, 1996.

T. Probst. German privacy and data protection principles on biometrics for passports and public administration. Technical report, Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein, Kiel, Germany, 2011.

S. Schudy and V. Utikal. The influence of (im) perfect data privacy on the acquisition of personal health data. *Research Paper Series - Thurgau Institute of Economics and Department of Economics at the University of Konstanz*, 76:1–32, 2012.

D. Stahl and P. Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.

W. K. Viscusi. A note on "lemons" markets with quality certification. *Bell Journal of Economics*, 9(1):277–279, 1978.

Appendices

A Feedback and data pooling

As mentioned in the main text, we gave different feedback in the LOADED treatment. Generally, subjects were informed of their own profits and the market wage of that period. In our sub-treatment Base (12 groups) this was the only feedback subjects received. In 11 groups of the LOADED treatment, we gave the subjects additional information about the choices of all six workers in the group (sub-treatment Feed). Our hypothesis was that the additional feedback would support learning.

The data from both sub-treatments is hardly distinguishable. In Market A, 18.3% reveal in Base compared to 19.4% in Feed. The same pattern can be observed on the other markets. In Market B, we observe 52.8% and 53.3% of revelation in the sub-treatments without and with detailed feedback, respectively. In Market C, the revelation rate is 36.1% for both variants. There are no significant differences concerning subjects' revelation behavior on any market (Mann-Whitney U-Tests, $p = 0.819$, $p = 0.843$ and $p = 0.964$ for the markets A, B and C, respectively). Note that this result also holds true when comparing individual periods. The revelation rates in Base and Feed are not significantly different in any of the 15 periods.

Considering the behavior of the different workers, we find even less differences concerning the two feedback conditions. Table 5 reports the corresponding averages and p-values. Again, none of the differences between Base and Feed is significant.

	Market A			Market B			Market C		
Worker	Base	Feed	p-val.	Base	Feed	p-val.	Base	Feed	p-val.
1	0.017	0.018	1.000	0.133	0.073	0.603	0.050	0.036	1.000
2	0.067	0.073	0.857	0.333	0.309	1.000	0.033	0.036	1.000
3	0.033	0.055	0.727	0.400	0.491	0.536	0.117	0.055	0.314
4	0.067	0.055	0.964	0.650	0.655	0.929	0.383	0.436	0.770
5	0.100	0.091	0.947	0.833	0.782	0.641	0.767	0.764	0.948
6	0.817	0.873	0.632	0.817	0.891	0.316	0.817	0.836	0.538

Table 5: Average revelation and p-values (Mann-Whitney U-Tests) by workers in the two feedback conditions.

The feedback in our NEUTRAL treatment was like the LOADED sub-treatment Base and, accordingly, only LOADED Base and NEUTRAL can be directly compared. As a robustness check, we redid the entire empirical analysis without pooling the data and only using the LOADED Base data. It turns out that all of our results prevail, i.e. all coefficients are significantly different from zero and all tests with a significant result remain significant at least at the 5%-level. The marginally significant regressor *Reveal* in regression (7) of

Table 3 is no longer significant. In this sense, the pooling of our two LOADED treatments is without loss of generality.

B Further behavioral models

Quantal Response Equilibrium

Quantal Response Equilibrium (McKelvey and Palfrey, 1995) takes decision errors into account: workers do not always choose the best response with probability one but they choose better alternatives more frequently than others. Therefore, QRE allows for out-of-equilibrium choices to conceal and reveal one's productivity.

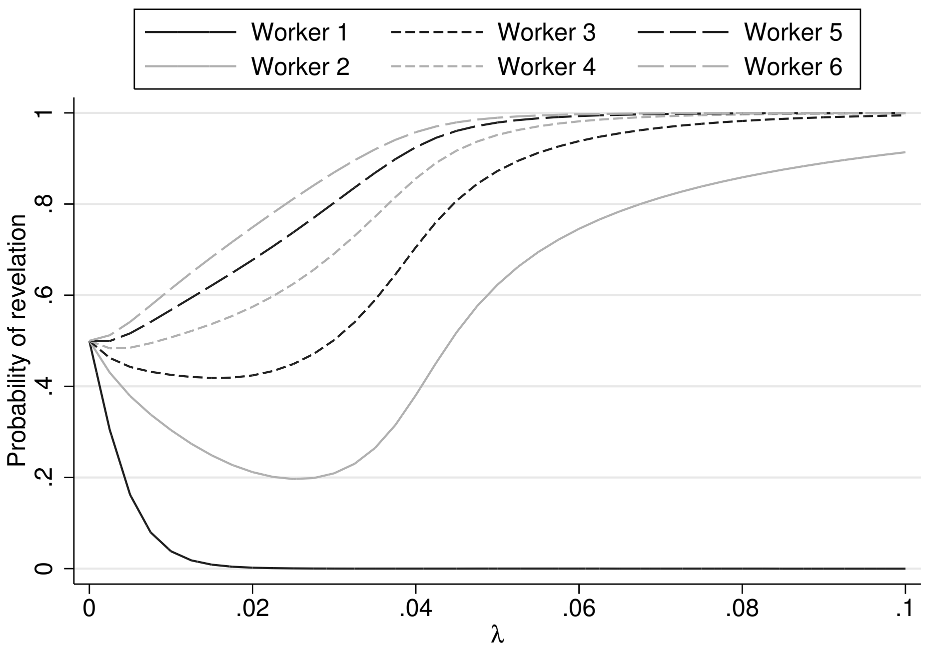

Figure 4: QRE predictions for Market B.

We employ the logit equilibrium variant of QRE. Worker i believes that the other workers will choose to reveal with certain probabilities and calculates her expected payoff from concealing based on this belief (the payoff from revealing is simply the productivity minus the revelation cost). Workers make better choices more frequently. In particular, choice probabilities are specified to be ratios of exponential functions where expected payoffs are multiplied with λ, the rationality parameter. This parameter captures deviations from the Nash equilibrium: if $\lambda = 0$, behavior is completely noisy and both choices are equally likely regardless of their expected payoff; as $\lambda \to \infty$, workers choose the best response with probability one. In the logit equilibrium, beliefs and choice probabilities are consistent.

As an example, Figure 4 displays the QRE predictions (revelation frequencies given the rationality parameter λ) for Market B. We note that QRE can explain the under-revelation result: for any $\lambda > 0$, the error probabilities are higher for workers 2 to 6 who reveal in equilibrium than for worker 1, suggesting too little revelation. The relationship between the probability of revealing and λ is even non-monotonic for some workers, so a higher λ can be associated with less revelation. The intuition is that the likelihood that worker 6 reveals must meet a certain threshold before other workers start preferring revelation over concealment. Thus, λ must be high enough. On the other hand, the higher the parameter λ, the more weight the best response gets such that the propensity to conceal can increase in λ as long as worker 6 does not reveal with a significantly high probability. Such a non-monotonicity cannot be observed for workers that should conceal in equilibrium in any of our markets. Qualitatively, Market C looks the same as Figure 4. In Market A, there are no non-monotonicities since all workers except worker 1 conceal in equilibrium.

We conduct a maximum-likelihood estimation of the QRE parameter λ for the treatment Neutral. Following Haile, Hortaçsu, and Kosenok (2008), the estimation is implemented jointly for our three markets such that there is only one free parameter in the model we estimate. We find an estimate of $\lambda = 0.035$ with a standard error of 0.0015, suggesting that λ significantly differs from zero.[12]

	Market A			Market B			Market C		
Worker	Nash	QRE	Neutral	Nash	QRE	Neutral	Nash	QRE	Neutral
1	0	<0.01	0.04	0	<0.01	0.00	0	<0.01	0.04
2	0	0.01	0.02	1	0.27	0.35	0	0.02	0.04
3	0	0.02	0.02	1	0.60	0.71	0	0.21	0.09
4	0	0.06	0.02	1	0.78	0.78	1	0.67	0.62
5	0	0.20	0.07	1	0.87	0.91	1	0.93	0.93
6	1	>0.99	0.89	1	0.92	0.96	1	0.99	0.98

Table 6: QRE estimates and data from Neutral

Table 6 summarizes the QRE prediction for the λ we estimated from the data and contrasts this prediction with our findings. Overall, QRE fits the data well. It correctly predicts the degree of under-revelation (of workers 2, 3, 4, 5, and 6 in Market B; and workers 4, 5, and 6 in Market C). The predicted and observed frequencies of revelation are remarkably similar. Also, the low revelation frequencies of low-productivity workers who should conceal in equilibrium are predicted rather well (workers 1, 2, and 3 in Market A; and workers 1 and 2 in Market C).

On the other hand, the QRE predictions for some of the workers who conceal in equilibrium do not perform as well. QRE predicts substantial rates of out-of-equilibrium

[12] Separate estimates for the three markets yield $\lambda_A = 0.026$, $\lambda_B = 0.039$ and $\lambda_C = 0.029$. The QRE estimate for the three markets in Loaded is $\lambda = 0.017$

revelation decisions with about 20% for worker 5 in Market A and for worker 3 in Market C. In both cases, less than 10% revelation is observed. Nevertheless, the QRE model is able to capture the overall patterns of behavior and can therefore account for the lower revelation rates that we observed compared to the equilibrium.

Inequality aversion

Does fairness prevent the full revelation of private information? In our game, when the highest productivity worker chooses to reveal, it increases her own payoff but imposes a negative externality on others.[13] Similarly, given worker n reveals, the same can hold true for worker $n-1$. Accordingly, inequality-averse subjects may be less inclined to reveal their productivity than the standard model of selfish payoff maximizers suggests. While such motives may play only a minor role in large markets like the labor market, they may be important in smaller groups (small teams or enterprises) and in our experimental groups of six.

We use the model of inequality aversion proposed by Fehr and Schmidt (1999) (henceforth F&S) where players are concerned not only about their own material payoff but also about the difference between their own payoff and other players' payoffs. As a consequence the player's utility is

$$U_i(x_i, x_j) = x_i - \frac{\alpha_i}{n-1} \sum_{j \neq i} \max[x_j - x_i, 0] - \frac{\beta_i}{n-1} \sum_{j \neq i} \max[x_i - x_j, 0]$$

where, x_i and x_j denote the monetary payoffs to players i and j, and α_i and β_i denote i's aversion toward disadvantageous inequality (envy) and advantageous inequality (greed), respectively. Standard preferences occur for $\alpha = \beta = 0$. Following F&S, we assume $0 \leq \beta_i < 1$.

There are two complications regarding the impact of inequality aversion. One issue is that the effect on inequality of (not) revealing one's productivity will often be ambiguous: a worker may find that concealing reduces the advantageous inequality with respect to less productive workers but it may also increase the payoff difference to the more productive workers provided they reveal. So this worker may stick with her (standard) equilibrium action even if she is inequality averse. Another complication is that there are multiple equilibria. It is not straightforward to show which of the $2^6 = 64$ possible outcomes can be an equilibrium for inequality-averse players and which cannot.

To tackle these issues we employ simulations based on a calibrated version of the model to identify the F&S equilibria of the revelation game. The model is calibrated

[13] As an example, consider Market A. The Nash equilibrium has only worker 6 revealing her productivity, and worker 6 earns 500 points in equilibrium whereas all others earn 240 points. If worker 6 did not reveal, everybody would earn 300 points. It follows that, for a sufficiently inequality-averse subject, concealing may yield a higher utility than revealing.

using the joint distribution of the α and β parameters observed in Blanco, Engelmann, and Normann (2011). For each subject, they derive an α_i from rejection behavior in the ultimatum game and a β_i from a modified dictator game. There are 61 subjects in this data set with 58 different α_i-β_i types.[14] Note that we need the *joint* distribution of the parameters, which is unavailable elsewhere. The computer simulations are implemented as follows: In each trial, the program randomly assigns an α_i-β_i parameter combination to each of the six workers (with replacement), where the 61 α_i-β_i types in the Blanco et al. (2011) data were equally likely. Given the realization of inequality parameters, the program then systematically checks which of the 64 possible outcomes turns out to be an equilibrium. Note that there can be multiple equilibria, which is also the reason why the percentages do not add up to 100%. For the each of the three markets A, B, and C, we ran 100,000 trials.

	Actions of workers							Market	
No.	I_1	I_2	I_3	I_4	I_5	I_6	A	B	C
1	0	1	1	1	1	1	–	90.5%	–
2	0	0	1	1	1	1	–	3.9%	–
3	0	0	0	1	1	1	–	7.2%	61.9%
4	0	0	0	0	1	1	–	9.5%	20.2%
5	0	0	0	0	0	1	80.3%	14.8%	17.0%
6	0	0	0	0	0	0	19.7%	55.6%	8.4%
7	0	0	0	0	1	0	–	–	5.2%

Table 7: Summary of F&S equilibria. Note: because of multiple equilibria, the figures do not always add up to one hundred percent.

The simulation results are summarized in Table 7 which can be read as follows. First, note that seven equilibria emerge out of the 64 possible outcomes where each is described in a separate row of the table. In equilibrium 1, only worker 1 chooses to conceal while all the other workers reveal. This strategy profile was an F&S equilibrium in 90.5% of the 100,000 simulations of Market B where it is also the standard Nash equilibrium. There appear to be no F&S parameters which support this outcome as an equilibrium for Market A or C.

Overall, inequality aversion is consistent with our results. The simulations show that there are equilibria with F&S preferences where fewer players reveal than in a standard Nash equilibrium, and there are no F&S equilibria where more players reveal than in a Nash equilibrium.

Having said that, there are some aspects of the simulations that show the limits of inequality aversion for rationalizing our data. Firstly, and perhaps surprisingly, the

[14] There are no significant differences between the distributions of α that Blanco et al. (2011) elicit and the one assumed in Fehr and Schmidt (1999). The β distributions differ, but they are still roughly comparable.

standard Nash equilibrium is very often also an equilibrium with F&S preferences. In all three markets, it is the most frequent equilibrium: 90.5% (Market B), 80.3% (Market A) and 61.9% (Market C) of the 100,000 random realizations of α_i-β_i parameter combinations. Relatedly, the F&S equilibria consistent with our data occur with rather low frequencies. In other words, there are only a few F&S parameter combinations that support these equilibria. For instance, in Market B, we observe especially little revelation by workers 2 and 3 which may be captured by the F&S equilibria numbered 2 and 3. However, these equilibria do not occur often in the simulations. Moreover, the calibrated F&S model predicts that there is an equilibrium in which no worker reveals (equilibrium 6), and this equilibrium occurs in more than 50% of the runs for Market B where we in fact observe a lot of revelation, in line with the equilibrium prediction.

Secondly, the simulations show that coordination problems may occur due to multiple equilibria. In Market B there is a unique equilibrium in 26.6% of the cases, two equilibria in 65.5%, and three equilibria in 7.9% of the cases. The corresponding values for Market C are 87.3% and 12.7% for one or two equilibria, respectively. It is unclear how inequality-averse players can resolve the coordination problems resulting from multiple equilibria. In Market A, we always found a unique equilibrium, but this market is not strongly supportive of inequality aversion either. The equilibrium for Market A can be found analytically: worker 6 will not reveal if and only if $\beta_i > 200/260 \approx 0.8$. This condition will be met for 20% in the data set of Blanco et al. (2011) (and our simulations indicate exactly the same frequency for the occurrence of this equilibrium). In NEUTRAL this equilibrium occurs, however, only with a frequency of 11%. While the loaded frame leads to results closer to the prediction, inequality aversion should not be driven by the frame.

Overall, we can explain the observed outcomes as equilibria when players have F&S preferences. However, there are a number of problems which limit the explanatory power of inequality aversion.

Discussion of behavioral models

Our goal has been to investigate behavioral models that might suggest how behavioral forces affect revelation. The three canonical behavioral models studied all suggest that players might be biased not to reveal. By contrast, we found hardly any support for a hypothesis suggesting that players reveal too much as compared to the equilibrium.

We believe that level-k rationality or a limited depth of reasoning is an intuitive model that explains our data surprisingly well. Level k implies that equilibrium concealment only requires $k = 1$ whereas revelation may require higher levels of reasoning. When higher levels of reasoning are increasingly rare among subjects, it follows for our setup that there is generally too little revelation; the lower the worker's productivity, the less likely the worker will reveal if the equilibrium calls for revelation; and there are virtually no equilibrium-inconsistent reveal decisions. This is what we see in the data.

In different manners level k and QRE take into account that the payoff from concealing will *ceteris paribus* become more attractive than the payoff from revealing, $\theta_i - c$, for workers with low θ. It requires a "high level of rationality" (high k or high λ) for unraveling to occur to the extent predicted in equilibrium. Thus, both models suggest that the unraveling process might be stuck after a few players, leading to less revelation.

Markets played by fully rational yet inequality-averse players may also unravel only incompletely. The negative externality imposed on others may make even high-productivity workers conceal. On the other hand, inequality aversion supports the Nash equilibrium with standard preferences. Another difficulty is that multiple equilibria occur, which reduces the predictive power of such preferences.

C Proofs

Proof of Proposition 1.

We first show that $I_1^* = 0$. By concealing, the lowest-productivity worker earns at least θ_1 (namely when all other workers reveal, otherwise more), but worker 1 earns $\theta_1 - c < \theta_1$ by revealing. Hence, concealing is strictly dominant for worker 1 and we have $I_1^* = 0$ in equilibrium.

Next, we prove that $I_i^* = 0 \wedge I_j^* = 1$ only if $\theta_i < \theta_j$ strictly. Consider an equilibrium outcome with $I_i^* = 0$ and $I_j^* = 1$ and denote $\theta' = \sum_{m \neq i,j}(1 - I_m)\theta_m$ and $I' = \sum_{m \neq i,j}(1 - I_m)$. Now $I_i^* = 0$ and $I_j^* = 1$ are best replies to action profile I' if and only if

$$\theta_i - c \leq \frac{\theta_i + \theta'}{1 + I'} \tag{1}$$

$$\theta_j - c \geq \frac{\theta_i + \theta_j + \theta'}{2 + I'} \tag{2}$$

where the inequality for player i follows from $I_i = 0$ and the inequality for j follows from $I_j = 1$. Solving both equations for θ', we obtain

$$-c + (\theta_i - c)I' \leq \theta' \leq \theta_j - \theta_i - 2c + (\theta_j - c)I' \tag{3}$$

and

$$0 \leq -c + (\theta_j - \theta_i)(1 + I') \tag{4}$$

which holds only if $\theta_i < \theta_j$ strictly. Since $I_i^* = 0 < 1 = I_j^*$ only if $\theta_i < \theta_j$, we cannot have $I_{i+1}^* < I_i^*$ and thus $I_n^* \geq I_{n-1}^* \geq ... I_2^* \geq I_1^*$ as claimed. □

Proof of Proposition 2.

We first show that, if $\min(R) > \max(C)$ as asserted in the proposition, we get a unique equilibrium. Assume that, say, $R = \{n, n-1, ..., m\}$ and $C = \{m-1, m-2, ..., 1\}$. Then the pure strategy action profile

$$1 = I_n^* = I_{n-1}^* = ... = I_m^* > I_{m-1}^* = ... = I_2^* = I_1^* = 0$$

is a Nash equilibrium by the definition of R and C.

Now consider another pure strategy equilibrium candidate where, from Proposition 1, we only need to consider outcomes where $I_n^* \geq I_{n-1}^* \geq ... \geq I_2^* \geq I_1^* = 0$. Assume first that more workers reveal in this equilibrium candidate than in the first equilibrium, that is, workers $m-1$ to $m-k$, $k \geq 1$, reveal in this alleged equilibrium (whereas they conceal in the first equilibrium):

$$1 = I_n^* = I_{n-1}^* = ... = I_m^* = I_{m-1}^* = ... = I_{m-k}^* > I_{m-k-1}^* = ... = I_2^* = I_1^* = 0$$

For this to be a Nash equilibrium, we necessarily need $\theta_{m-k} - c \geq \frac{1}{m-k} \sum_{j=1}^{m-k} \theta_j$. However, this requires that $m-k \in R$ which is a violation of $\min(R) > \max(C)$. Consider a different pure strategy equilibrium candidate where fewer workers reveal; say, workers m to $m+k$, $k \geq 0$ conceal (whereas they reveal in the first equilibrium). Here, we necessarily need $\theta_{m+k} - c \leq \frac{1}{m+k} \sum_{j=1}^{m+k} \theta_j$ for this outcome to be a Nash equilibrium. Hence, $m+k \in C$ which violates the assumption in the proposition. Hence, if $\min(R) > \max(C)$, the first Nash equilibrium is the unique pure-strategy equilibrium.

We now show the "only if" part of the proposition by proving that if $\min(R) > \max(C)$ is violated, multiple equilibria result. Let m be the highest worker in C and l be the lowest worker in R and assume the violation: $m > l$. First, note that $m \in C$ iff $\theta_m - c \leq \bar{\theta}(m)$ which, implies that concealment is a best-response for the workers $1, ..., m$ given that the remaining workers reveal. From $m = \max(C)$ it follows that $\theta_{m+1} - c > \frac{\theta_{m+1} + \sum_{j=1}^{m} \theta_j}{m+1}$ and, by the definition of m, the inequality will also hold for all workers $m+1, ..., n$. Hence, we have a Nash equilibrium where the workers $1, ..., m$ conceal and the workers $m+1, ..., n$ reveal. Second, note that $l = \min(R)$ implies $\theta_{l-1} - c < \frac{\theta_{l-1} + \sum_{j=1}^{m-2} \theta_j}{l-1}$, that is, concealment is a best-response for the workers $1, ..., l-1$ given that the remaining workers reveal. As for the remaining workers, $l \in R$ implies $\theta_l - c \geq \frac{\theta_l + \sum_{j=1}^{l-1} \theta_j}{l}$ and the same inequality will hold for all workers $l, ..., n$. Hence we have a second Nash equilibrium where the workers $1, ..., l-1$ conceal and the workers $l, ..., n$ reveal. Hence, if $\min(R) > \max(C)$ is violated, multiple equilibria occur and the proposition follows. $\square$

Proof of Proposition 3.

We prove the proposition by establishing a contradiction: suppose some worker conceals for $k = 2$ but reveals for $k = 1$. This yields a contradiction because, as we will show, the expected payoff from concealing is higher if $k = 1$ than if $k = 2$.

We first derive the best reply of a $k = 1$ player. Player i (when $k = 1$) believes that all other players randomize across both actions with a probability of 0.50. To calculate the payoff from concealing, player i needs to take into account all possible contingencies that may arise (no other player concealing, one of the $n - 1$ other players concealing and so on) which yields a complex combinatoric expression. Specifically, player i (when $k = 1$) will reveal if and only if

$$\theta_i - c \geq \frac{\theta_i \sum_{a=0}^{n-1} \frac{\binom{n-1}{a}}{a+1} + (\sum_{j \neq i} \theta_j)(\sum_{a=1}^{n-1} \frac{\binom{n-2}{a-1}}{a+1})}{2^{n-1}} \tag{5}$$

or

$$\theta_i - c \geq \frac{\theta_i \sum_{a=0}^{n-2} \frac{\binom{n-2}{a}}{a+1} + (\sum_{j \in I} \theta_j)(\sum_{a=1}^{n-1} \frac{\binom{n-2}{a-1}}{a+1})}{2^{n-1}} \tag{6}$$

where the numerator arises because all possibilities occur with equal probability.

Note that, if (6) is met for player i, it will also be met for all workers with $\theta_j \geq \theta_i$. This follows from the observation that the factor of θ_i on the RHS of (6) is strictly smaller than one. Hence (when $k = 1$), workers $\theta_1, ..., \theta_m$ will conceal and workers $\theta_{m+1}, ..., \theta_n$ will reveal for some $m \geq 1$, unless we have the trivial case where all workers conceal.

As a next step, we show that a necessary condition for worker i to reveal (when $k = 1$) is $\theta_i > \frac{\sum_{j=1}^{n} \theta_j}{n}$. To prove this, we evaluate RHS of (6) when $\theta_i = \frac{1}{n} \sum_{j=1}^{n} \theta_j$. Simple but tedious combinatorics show a rather intuitive result, namely that this expression is greater or equal than the average worker productivity if and only if $\theta_i \geq \frac{1}{n} \sum_{j=1}^{n} \theta_j$. Thus, (6) will be met only if worker i's productivity is above average.

We now establish the fact that the condition for a $k = 2$ worker to conceal is weaker than the condition for a $k = 1$ worker to conceal. When $k = 2$, worker i believes that all other players are level $k = 1$, and, accordingly, that $\{\theta_1, \theta_2, ...\theta_m\}$ will conceal with probability one. Player i will conceal (when $k = 2$) if and only if

$$\theta_i - c \leq \frac{\theta_i + \sum_{i=1}^{m} \theta_j}{m+1}. \tag{7}$$

Now, for worker i to reveal when $k = 1$ necessarily requires $\theta_i - c \geq \frac{1}{n} \sum_{j=1}^{n} \theta_j$ but to conceal when $k = 2$ requires (7). Putting these condition together, we obtain

$$\frac{\theta_i + \sum_{i=1}^{m} \theta_j}{m+1} \geq \theta_i - c > \frac{\sum_{i=1}^{n} \theta_j}{n}. \tag{8}$$

This, however, cannot hold: it cannot be that the average of the low-productivity workers

$1, ..., m$ plus worker i is larger than the average productivity of all workers because $\theta_i > \sum_j \theta_j/n$ for all $i > m$.

Since the condition for revealing as a $k = 1$ player contradicts the condition for concealing as a $k = 2$ player, it cannot be that player i reveals as a level $k = 1$ but conceals as level $k = 2$. Hence, if player i conceals for $k = 2$, she will do so with $k = 1$ steps of reasoning.

Finally and intuitively, similar arguments show that a worker will conceal if $k = 2$ if she conceals when $k = 3$ and so on for a higher k. With a higher k, high types will "drop out" by revealing, leading to even lower concealment wages. Hence, workers who conceal for some k' will not reveal when $k < k'$. $\square$