

Price of Privacy in the Keynesian Beauty Contest

Hadi Elzayn* Zachary Schutzman†

May 3, 2019

Abstract

The Keynesian Beauty Contest is a classical game in which strategic agents seek to both accurately guess the true state of the world as well as the average action of all agents. We study an augmentation of this game where agents are concerned about revealing their private information and additionally suffer a loss based on how well an observer can infer their private signals. We solve for an equilibrium of this augmented game and quantify the loss of social welfare as a result of agents acting to obscure their private information, which we call the ‘price of privacy’. We analyze two versions of this price: one from the perspective of the agents measuring their diminished ability to coordinate due to acting to obscure their information and another from the perspective of an aggregator whose statistical estimate of the true state of the world is of lower precision due to the agents adding random noise to their actions. We show that these quantities are high when agents care very strongly about protecting their personal information and low when the quality of the signals the agents receive is poor.

1 Introduction

In recent years, the mathematical study of privacy has become a major subject of inquiry. Much of the impetus for this work has been a series of data breaches and deanonymization of seemingly safe private information, perhaps most famously in the use of IMDb reviews to attack Netflix’s data set over a decade ago [19]. Models such as differential privacy alleviate this problem by providing formal guarantees about how much about any individual an adversary can learn from the release of some statistic computed from a dataset. However, such techniques are generally predicated on the presence of a trusted central agent which applies the differentially private mechanism to the data it collects from individuals. Alternatively, in the case of the *local* differential privacy model, the agents are typically instructed in how to privatize their data.

In this work, we endogenize a notion of privacy in the absence of a trusted party to coordinate the mechanism. In particular, we analyze a formalization of a classical game called the *Keynesian Beauty Contest*, which has been used to study strategic interaction involving information acquisition and the coordination of collective action. In particular, we show how the traditional formulation of this game neglects privacy concerns in its equilibrium predictions, and we then provide a framework for extending the game to incorporate a flexible notion of privacy into the utility of the agents. Using this, we can characterize a *price of privacy*, somewhat akin to quantities such as the price of anarchy, which measures the loss of social welfare in a population of agents who act selfishly to guarantee their own privacy.

Abstractly, we think about players in a game being perfectly rational Bayesian agents who observe some information, perform a utility-maximizing computation, and play an action. However, by observing some player i ’s chosen action, player j may be able to learn something about player i ’s private information. For this reason, if players fear that their public actions may reveal private information, they may be incentivized to deviate from the strategy which maximizes utility in the absence of privacy concerns. If all players share such concerns, or anticipate others harboring them, equilibrium behavior may be significantly different than standard predictions.

*Applied Mathematics and Computational Science, University of Pennsylvania. Email: hads@sas.upenn.edu.

†Department of Computer and Information Science, University of Pennsylvania. Email: ianzach@seas.upenn.edu.

One setting where there may be tension between privacy and coordination with other players and some underlying ground truth is intra-organizational information aggregation. Suppose that a firm wishes to poll a group of employees as to a particular decision, such as a prediction about the success of a particular product, an evaluation of a colleague, or a new procedure for evaluating and settling claims. In order to make the best decision possible, it is valuable to aggregate information, opinions, or signals received by individuals; however, individuals may be reluctant to fully share their opinion lest it be held against them or they prove to be ‘wrong’. In this situation, simple anonymization procedures may not be trustworthy or effective, and any such procedure requires trusting an internal mechanism to which the employee does not have access. If the respondent is concerned about other agents deanonymizing her survey, then it should be expected that she does not respond honestly.

1.1 Our results

We consider a formal model of the Keynesian Beauty Contest, a game in which each agent observes some information and submits an estimate about the true ‘state of the world’, then earns utility based on both how close her action is to the truth as well as how close it is to the *average action* over all agents. We describe the game, the information structure, and the strategy space and show the existence of a *symmetric linear Nash equilibrium*, where agents’ actions are a convex combination of the public and private signals they observe, extending the results in Morris and Shin [17].

We then turn to the privacy-augmented version of the game, where agents face the same utility function but also suffer a loss of utility based on the ability of other players to infer their private information. We show that this new game also has an equilibrium in strategies that can be written as strategies in the original game with added random noise. This leads to two different values which can be thought of as a ‘price of privacy’. In the first, we consider the perspective of the agents and quantify the total loss of the players’ utilities as a result of incorporating this concern for privacy. In the second, we think about an ‘untrusted aggregator’ who wants to compute some statistic using the agents’ private beliefs but cannot convince them to participate in a privacy-protecting mechanism. Here, we compute the decrease in the quality of the aggregation due to the players’ addition of noise to their previously optimal actions.

At a high level, we consider a setting in which there is no centralized planning mechanism, either to perform the differentially private computation or to instruct the agents who own the data to add a particular amount of noise to their information in order to perform a locally differentially private aggregation. Rather, we assume that agents are rational and derive some utility from both the aggregate-level computation being accurate as well as from the privacy gained by obscuring the information she releases. A major departure from the local differential privacy framework is in our treatment of outliers. In local differential privacy, an individual whose data is very different from the norm may have to add a considerable amount of noise before providing her information to the aggregator, since we need to worry that her unmodified data might shift the distributions of outcomes farther than we would like. In this work, we consider a different flavor of individual privacy, where an agent with outlying information does not necessarily care about being revealed as an outlier; she only cares how accurately an observer can guess her private information.

1.2 Related Work

1.2.1 Economics

Concretely, our work builds primarily on the results in Morris and Shin [17], which formalizes the modern version of a Keynesian Beauty Contest in order to study the tendency for individuals to over-weight public information. We similarly explore linear equilibria in a Keynesian Beauty Contest game with public and private signals; in fact, the structure of our model is substantially very similar, and we recover their results as a special case when there is no concern for privacy. Viewing our model as a generalization of theirs, we expand their result on the existence of such equilibria to a privacy-aware setting. The results in Hellwig and Veldkamp [11] build on those in [17], exploring a setting in which agents optimize a selection of information sources with different costs matching their qualities. While [11] is quite different mechanically and in spirit,

there is a certain sense in which our paper can be interpreted along similar lines: the equilibrium noisy action can be viewed as similar to choosing a private signal of differing variances, and the privacy cost can be thought of as a cost to more precise information. The Keynesian Beauty Contest is well-studied in the broader macroeconomics literature (see Nagel et al. [18] for a recent survey) and is frequently used to study settings where agents derive some value for correctness and for coordination, such as in financial markets or strategic voting. These settings are natural cases for the use of formal privacy methods, such as financial analysts wanting to conceal the model they use to predict asset prices or voters being concerned about their preferences being held against them. Gradwohl and Smorodinsky [10] study a setting in which strategic agents are concerned that their actions in a game will reveal sensitive private information, but their model and approach, which is based on the concept of *signaling games*, is considerably different from the setting we consider.

1.2.2 Differential Privacy

In the last decade, mathematical notions of privacy have been studied extensively in the computer science literature. The most influential is *differential privacy* (c.f. Dwork et al. [8]) which states roughly that a statistical algorithm is differentially private if the probability that the algorithm gives any particular output does not change by much when we modify one row of the database the algorithm is run on. Standard differential privacy is principally a *central*, rather than *individual* notion of privacy, as the algorithms operate under the assumption that the adversary does not have access to the raw data provided to the analyst. The concept was introduced in Kasiviswanathan et al. [14] and recent work has shown that this notion is a powerful framework for addressing privacy concerns (c.f. [2, 12, 13]). There is also a literature on *privacy in mechanism design*. The work in Chen et al. [6] and Nissim et al. [20] examines settings where agents participate in a data aggregation mechanism and earn utility which is increasing in the quality of the estimate but decreasing in the data leakage. The model in [6] is one of *truthful voting*. In this work, we study a game-theoretic setting where privacy is a concern, rather than designing a privacy-preserving mechanism. Our work differs from the formal study of privacy in that we use a definition of which does not (necessarily) satisfy the strictness of the various versions of *differential privacy*. Rather, we consider the perspective of strategic agents who are concerned with other players' ability to learn 'too much' about their own private information. In this way, our work is related to the issue of *response bias* on surveys, where respondents do not answer questions honestly in order to avoid revealing sensitive information. Work examining the use of randomization to alleviate this effect goes back several decades (e.g. Warner [22]).

Our extended game does bear a resemblance to the *local model* of differential privacy in that agents add noise to their actions, and one might intuitively map the case where agents add Gaussian noise to their actions to an instance of the local Gaussian mechanism (see, e.g. [8]). However, agents in our model are *not* using the Gaussian mechanism to achieve differential privacy, and such a guarantee cannot be recovered¹, as the fact that agents' data is unbounded means that this mechanism cannot give non-vacuous guarantees. To see this, recall that differential privacy requires that for any set of output of the mechanism must be approximately as likely when input is *neighboring*, which typically means differing in one record. In our case, this corresponds to having a different private signal. The problem is that for any choice of finite variance, adding Gaussian noise to private signals that are sufficiently distant will produce noisy actions that are far apart, and can thus be distinguished with high confidence.

More importantly, there is a significant conceptual difference between the models of differential privacy and our model. In differential privacy models, there is a strict separation of roles: there is an accuracy-concerned learner that wishes to perform and release the output of a query, and privacy-concerned agents that supply their data. The addition of noise either centrally or locally prevents any other party from learning the exact value of the private data with high confidence, and is assumed to be necessary to make agents willing to supply their data. In contrast, there is no external learner in our model; instead, agents attempt to coordinate their actions with other agents and the state of the world, but must add noise to their actions in order to prevent others from learning their private information. Noisier actions lead to greater

¹There are other, more complicated mechanisms that can achieve differential privacy with unbounded input data. We refer readers to [16] and [21] for more discussion.

variance around the state of the world and the average actions, and thus lower payoffs. Hence, the agents care about *both* privacy *and* accuracy, and this tradeoff determines how much noise they ultimately add.

2 Framework and Model

2.1 The Keynesian Beauty Contest

The Keynesian Beauty Contest dates back to John Maynard Keynes' 1936 *General Theory* [15] in which he formulates a game by analogy to a newspaper beauty contest. In this game, players select the 'most beautiful' entrants from an array of photographs printed in the newspaper, and those players who choose the *most popular* faces are eligible for a prize. The salient idea is that a rational contestant must consider her own opinions about which entrants are the most beautiful as well as her *beliefs* about the opinions of all other contestants, and perhaps her beliefs about the beliefs of the other players' opinions of other players, and her beliefs about the beliefs about the beliefs about the beliefs about the opinions of other players, and so on. Keynes originally proposed this as a model to explain the behavior of financial markets, where the value of an asset depends as much on its fundamental potential for returns as it does the *collective belief in its potential* for returns. Since its inception, this model has been used throughout the economics literature to describe the behavior of rational agents in strategic environments (c.f. [1, 4, 5, 9, 17, 18]), particularly in macroeconomic settings.

In this paper, we work with a common abstraction of the Keynesian Beauty Contest (see e.g. [11, 17]) in which there is some true state of the world s , and each agent submits a guess θ_i and earns utility equal to

$$u_i(\theta_i, \theta_{-i}) = -(1 - \alpha)(\theta_i - \bar{\theta})^2 - \alpha(\theta_i - s)^2,$$

where $\bar{\theta}$ is the average choice of θ_i over all players and θ_{-i} denotes the action of all other players other than i . We refer to the $(\theta_i - \bar{\theta})^2$ term as the *coordination component*, which rewards how close player i 's action is to that of the other players, and the $(\theta_i - s)^2$ term as the *guessing component*, which describes how close player i 's action is to being the correct guess for the true state of the world s . The parameter $\alpha \in [0, 1]$ is common to all agents and describes the relative weighting of the values of coordination and guessing the true state of the world correctly.

We let n denote the number of players. In the economics literature, it is standard to consider a continuum of agents indexed by the unit interval $[0, 1]$. We present our results from the perspective of the finite game and observe that the analogous result for the continuum case emerge by taking the limit as n grows to infinity. We state these as propositions immediately following the corresponding result in the finite game. For the sake of brevity, all omitted proofs as well as direct proofs of the results in the infinite game can be found in the appendix.

The fundamental difference between the two cases is in the definition of the average action $\bar{\theta}$. In the game with finitely many and infinitely many agents, these are

$$\bar{\theta} = \frac{1}{n} \sum_{j=1}^n \theta_j \quad \text{and} \quad \bar{\theta} = \int_0^1 \theta_j dj,$$

respectively. In the finite game, agent i 's action θ_i has a measurable effect on the average action whereas in the infinite game, her impact is infinitesimal. Therefore, when there are finitely many players, agent i must consider her own action's effect on $\bar{\theta}$; in the infinite case, she can treat $\bar{\theta}$ as simply the average action of all players other than herself. Intuitively, as n grows, any individual's ability to unilaterally affect $\bar{\theta}$, so we should expect that the results in the infinite game emerge as the limits of the results in the finite game as n grows to infinity.

2.2 Information Structure

We now describe the information structure of the game. Each player has a common *improper uniform* prior distribution over the value of s . There is a public signal y , drawn from a Gaussian distribution with mean s and variance σ_y^2 . Additionally, each player i observes a private signal x_i independently drawn from a Gaussian distribution with mean s and variance σ_x^2 . Each player observes y and her own x_i , but not any other player j 's private signal x_j . Each player i 's utility function $u_i(\theta_i, \theta_{-i})$ is identical as given above. The values of σ_y^2 , σ_x^2 , α , the utility functions, and the realization of y are public and common knowledge, and the form of the priors and structure of the game is also common knowledge. Furthermore, we assume that the agents are rational and Bayesian. That is, they seek to maximize their expected utility given their knowledge and their beliefs about other players.

We use the notion of an *improper uniform prior* belief to align with the economics literature; however, this concept is not necessarily standard across disciplines. One can think of a (proper) uniform prior with compact support as indicating that players believe any underlying parameter in the support is equally likely to be the truth. An ‘improper uniform prior’ expresses the same sentiment for an unbounded support. The term ‘improper’ arises because such an object is not a probability density function, but, when updated with the proper Gaussian signals in our setting, the posterior is a proper distribution. Berger et al. [3] provide a useful discussion of settings in which the improper uniform prior is appropriate.

For readers uncomfortable with such a notion, we present an alternative framing of the information structure which is equivalent to the above. The private signals x_i , the utility functions, and the values of α and σ_x^2 are as before. Rather than y representing a public signal, each agent instead has a prior belief about the true state of the world s , which is a Gaussian distribution with mean y and variance σ_y^2 . This prior is common to all agents and the parameters of the prior are common knowledge. This framing is equivalent to the previous one since an agent with an improper uniform prior, upon observing the realization of the public signal y , updates her posterior belief about s to be Gaussian with mean y and variance σ_y^2 , which is what she would have believed if she instead began with the common Gaussian prior.

We write \mathcal{I}_i for the information set of player i , which captures the structural information about the game as well as the values of x_i and y . Furthermore, we use $\mathbb{E}_i(\cdot|\mathcal{I}_i)$ to denote the *expectation of player i at \mathcal{I}_i* , which is player i 's belief about some quantity given that she knows everything at \mathcal{I}_i . Where it is clear from context, we drop the \mathcal{I}_i for notational clarity and simply write $\mathbb{E}_i(\cdot)$.

We now consider an equilibrium concept for this game.

Definition 1 Symmetric linear Nash equilibrium. *A symmetric linear Nash equilibrium of this game is an action θ_i for each player i which can be written as a linear combination of the private and public signals x_i and y such that no player can profitably deviate unilaterally and every player i chooses the same weight to put on x_i and y . That is, a Nash equilibrium in which $\theta_i = \kappa x_i + (1 - \kappa)y$ for all players and κ is identical for all players.*

The authors in [17] describe the *unique* symmetric linear Nash equilibrium of a slightly different version of the game which has the same first order condition; our proof follows a very similar structure.

Lemma 1 First order condition. *In equilibrium, an agent's optimal choice of θ_i^* must satisfy*

$$\theta_i^* = \frac{\alpha n^2 \mathbb{E}_i[s]}{\alpha(2n-1) + (n-1)^2} + \frac{(1-\alpha)(n-1)\mathbb{E}_i\left[\sum_{j \neq i} \theta_j\right]}{\alpha(2n-1) + (n-1)^2}$$

Proposition 1. *In the game with infinitely many agents, in equilibrium, an agent's optimal choice of θ_i^* must satisfy*

$$\theta_i^* = \alpha \mathbb{E}_i[s] + (1-\alpha)\mathbb{E}_i[\bar{\theta}]$$

These results say that the optimal action is a convex combination of agents' expectations about the state and the average actions; this follows directly from the structure of the payoff function. Now, if other agents'

strategies are convex combinations of signals, then the unique best response is also a convex combination of signals. We can match coefficients to solve for this symmetric linear Nash equilibrium.

Lemma 2 Nash equilibrium and Value of κ . *The symmetric linear Nash equilibrium is given by*

$$\theta_i = \kappa x_i + (1 - \kappa)y, \text{ where}$$

$$\kappa = \frac{\alpha n^2 \tau_x}{\alpha n^2 \tau_x + ((n - 1)^2 + \alpha(2n - 1)) \tau_y}, \tau_x = \frac{1}{\sigma_x^2}, \text{ and } \tau_y = \frac{1}{\sigma_y^2}$$

for all players i .

Proposition 2. *The symmetric linear Nash equilibrium in the game with infinitely many players is given by*

$$\theta_i = \kappa x_i + (1 - \kappa)y, \text{ where}$$

$$\kappa = \frac{\tau_x \alpha}{\tau_x \alpha + \tau_y}, \tau_x = \frac{1}{\sigma_x^2}, \text{ and } \tau_y = \frac{1}{\sigma_y^2}$$

for all players i . This recovers the analogous result in [17].

Notice that if $\alpha = 1$, the agent is only rewarded based on his closeness to the state, and θ_i^* then dictates that agent i play his best guess of the state. As α falls, the weight Agent i places on his private signal diminishes, until at $\alpha = 0$, the agent chooses θ_i^* to be equal to y , the public signal.

We can now write the expected utility to each agent in this equilibrium:

Lemma 3 Expected utility. *The expected utility of agents for playing $\theta_i = \kappa x_i + (1 - \kappa)y$, given s is:*

$$\mathbb{E}_i[u_i|s] = -\alpha(\kappa^2 \sigma_x^2 + (1 - \kappa)^2 \sigma_y^2) - \frac{(1 - \alpha)\kappa^2(n - 1)}{n} \sigma_x^2$$

Proposition 3. *The expected utility of agents for playing $\theta_i = \kappa x_i + (1 - \kappa)y$, conditional on s in the game with infinitely many players, is:*

$$\mathbb{E}_i[u_i|s] = -\alpha(1 - \kappa)^2 \sigma_y^2 - \kappa^2 \sigma_x^2$$

Using the value of κ from the infinite setting for simplicity² we can now examine how utility changes as the variances of the signals and number of agents do. We briefly analyze this next.

Corollary 1 Comparative statics.

$$\frac{\partial \mathbb{E}_i[u_i|s]}{\partial \sigma_x^2} = -\frac{(\alpha \sigma_y^2)^2}{(\alpha \sigma_y^2 + \sigma_x^2)^3} \left((2 - \alpha)\alpha^2 \sigma_x^2 \sigma_y^2 + \sigma_y^2 - \frac{n - 1}{n} (1 - \alpha)(\alpha \sigma_y^2 - \sigma_x^2) \right),$$

$$\frac{\partial \mathbb{E}_i[u_i|s]}{\partial \sigma_y^2} = -\frac{\alpha(\sigma_x^2)^2}{(\alpha \sigma_y^2 + \sigma_x^2)^3} \left(2\alpha^2 \sigma_y^2 - \alpha \sigma_y^2 + \sigma_x^2 + \frac{n - 1}{n} 2\alpha(1 - \alpha)\sigma_y^2 \right),$$

and

$$\frac{\partial \mathbb{E}_i[u_i|s]}{\partial n} = -\frac{(1 - \alpha)\alpha^2 \sigma_x^2 (\sigma_y^2)^2}{n^2 (\alpha \sigma_y^2 + \sigma_x^2)^2}.$$

²For a fixed n , the infinite version of κ differs from the finite version by a multiplicative factor of $1 - \frac{1}{n^2}$ while the utilities differ by a factor of $1 - \frac{1}{n}$. Therefore, the value of κ in the infinite case is very close to that in the finite case for modest values of n , and any countervailing effect of changing the parameters through κ will be dominated by the direct effects on the utility.

All else fixed, decreasing σ_x^2 unambiguously increases utility. The fractional term in $\frac{\partial \mathbb{E}_i[u_i|s]}{\partial \sigma_x^2}$ is always positive and the term in the parentheses is also always positive, so $\frac{\partial \mathbb{E}_i[u_i|s]}{\partial \sigma_x^2}$ is strictly decreasing as σ_x^2 increases. This aligns with the intuition that higher quality information means that agents will be able to both guess the true state of the world more precisely as well as coordinate better with one another. A similar intuition holds for the effect of decreasing σ_y^2 , since the term in the parentheses is always positive when $n \geq 2$. However, the *rate* at which utility increases when the signal variances decrease is not the same. Supposing $\sigma_x^2 = \sigma_y^2$, agents gain more value from decreasing σ_y^2 than σ_x^2 for two reasons. The first is the over-weighting of y . Since agents more heavily weight the public signal y compared to their private information x_i , a small improvement in the precision of y will help agents choose an action θ_i closer to s than an identical improvement in the precision of the x_i signals. The second reason can be thought of as a second-order effect of this over-weighting. Since agents over-weight it, improving the quality of the public signal y even further increases the weight agents will place on it. Therefore, not only will agents actions be closer (in expectation) to s , but they will be closer to *each other*, thus also improving the utility in the coordination component.

As n increases, utility decreases. There is an $\frac{n-1}{n}$ coefficient on the second term in the utility function, and as n grows, this grows towards 1. We can think of this term as measuring the amount of non-impact any one agent can have on the average action $\bar{\theta}$. When the number of agents is very few, agent i 's action θ_i can't be too far from the average since the construction of $\bar{\theta}$ will have a high weight on θ_i , and this weight decreases as n grows. For very large values of n , agent i has a small impact on the average action, so the risk of being far away is greater.

2.3 Privacy-Awareness

Now we come to the heart of the privacy issue. Suppose that at the end of the game, each player's action θ_i is made public. Then, upon seeing θ_i , player j can simply write

$$x_i = \frac{\theta_i - (1 - \kappa)y}{\kappa},$$

and since everything on the right hand side is known to player j , she can learn player i 's private signal x_i with perfect precision. This suggests that if the players care at all about preserving the privacy of their private signals, they ought not to play exactly the equilibrium prediction θ_i .

As written, privacy is not in the players' utility functions, and it is too much to ask for a prediction that incorporates privacy without actually incorporating privacy into the utility function. To overcome this, we let $\rho(\tilde{\theta}_i)$ denote, abstractly, a measure which describes how 'private' the (possibly randomized) action $\tilde{\theta}_i$ is. For example, ρ could represent the *maximum precision* (i.e. the reciprocal of the variance)³ to which some player j can infer the private signal x_i of player i after observing his action and given her information set \mathcal{I}_j and knowledge of the equilibrium strategies. Here, we consider both this precision measure of privacy as well as an *entropy* measure, where $\rho(\tilde{\theta}_i)$ is the (information-theoretic) entropy of this inference, rather than the variance.

It is clear that any equilibrium which prescribes a deterministic mapping from the available information to an action θ_i^* cannot offer any measure of privacy to the players. To correct for this, we enhance players' utility functions to incorporate the value they gain from obscuring their signals and extend the equilibrium concept to a *noisy* one, where players select not only a guess about the state, but also a *noise-generating distribution*. Formally, we extend the utility function to

$$v_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) = (1 - \beta)u_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) + \beta\rho(\tilde{\theta}_i),$$

where u_i is as before and $\beta \in [0, 1]$ denotes the agents' *relative value for obfuscation*.

We can observe that if $\beta \neq 0$, the equilibrium actions in the original game do not support an equilibrium in this game, since upon observing θ_i^* , any player j can exactly recover the private signal x_i , and player j 's

³Taken as a worst-case over all $j \neq i$.

‘distribution’ over the possible values of x_i is degenerate with variance or entropy equal to zero. Concretely, in the precision setting, $\beta\rho(\tilde{\theta}_i)$ is $-\infty$ and $v_i(\theta_i^*, \tilde{\theta}_{-i}) = -\infty$. Conversely, if player i were to choose $\kappa' = 0$ (and therefore $\tilde{\theta}_i = y$), her utility would be finite, thus demonstrating a profitable deviation.

Thus, any deterministic equilibrium cannot be supported if privacy is incorporated into the utility function, and some degree of randomization is necessary; however, the next section will show that the prediction of the deterministic optimum will serve as the core deterministic component to an optimal randomized strategy.

We remark here that if agents do not care about others learning their signal at all (i.e. $\beta = 0$), the utility function degenerates to that of the original game. If $\beta = 1$, then players *only* care about protecting their private signal, and a dominant strategy is to take the action which maximizes the obfuscation of the private signal.

3 The Extended Game

We now formalize this modified game and show that it still has a symmetric linear Nash equilibrium when we consider *randomized* strategies. Here, optimal actions are of the form $\tilde{\theta}_i = \theta_i + \eta_i$, where θ_i is the same linear aggregation of the public and private signals as in the original game and η_i is independent noise drawn from a distribution whose form depends on how we measure the obfuscation component $\rho(\tilde{\theta}_i)$. That is to say, there is a symmetric linear Nash equilibrium in which players’ optimal action is a noisy modification of their optimal action in the original game. Note that randomized strategies are *not* mixed strategies in the classical sense. Each player’s choice of a noise-generating distribution can be thought of as committing to a collection of parameters which describe that distribution. In this sense, our equilibrium concept is a ‘pure strategies’ one, since each player will commit to a single collection of parameters, not a distribution over such collections.

3.1 Defining the Extended Game

In the extended game, s , x_i , σ_x^2 , y , σ_y^2 , and $\bar{\theta}$ are as before. Each player chooses an action $\tilde{\theta}_i$, and receives utility

$$v_i(\tilde{\theta}_i) = (1 - \beta)u_i(\tilde{\theta}_i) + \beta\rho(\tilde{\theta}_i),$$

where $\beta \in [0, 1]$, and $u_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) = -(1 - \alpha)(\tilde{\theta}_i - \bar{\theta})^2 - \alpha(\tilde{\theta}_i - s)^2$ as before. We write $\bar{\theta}$ to represent the average (noisy) action of the players. After each player announces $\tilde{\theta}_i$, players also learn the function each player used to select $\tilde{\theta}_i$ as a function of \mathcal{I}_i . For example, if a player’s strategy is to play x_i plus some random noise, player j learns the distribution from which this noise was drawn, though (crucially) not the *realization* of this draw. We denote the *privacy loss* of each player after actions are revealed $\rho(\tilde{\theta}_i)$. The two possibilities we consider for ρ are both functions of the belief distribution that a Bayesian agent, given their information set, observation of $\tilde{\theta}_i$, and knowledge of equilibrium strategies and the randomization mechanism used by player i , assigns to x_i ; one is the *precision* of these beliefs, and the other is the *entropy*. Formally,

Definition 2 Precision privacy loss in equilibrium. *For a given equilibrium profile, Player i ’s precision privacy loss is the expected precision with which a Bayesian opponent, knowing the equilibrium profile and observing their own signal, the public signal, and Player i ’s action $\tilde{\theta}_i$, can estimate Player i ’s signal x_i . That is, if γ represents the belief distribution about x_i , we define:*

$$\rho_{prec}(\tilde{\theta}_i) = \frac{1}{\text{Var } \gamma(x_i | s, \bar{\theta}_i, \mathcal{I}_j, H)}$$

Definition 3 Entropy privacy loss in equilibrium. *For a given equilibrium profile, Player i ’s entropy privacy loss is the expected precision with which a Bayesian opponent, knowing the equilibrium profile and observing their own signal, the public signal, and Player i ’s action $\tilde{\theta}_i$, can estimate Player i ’s signal x_i . That is, if γ*

represents the belief distribution about x_i , we define:

$$\rho_{ent}(\tilde{\theta}_i) = - \int \gamma(x_i|s, \tilde{\theta}_i, \mathcal{I}_j, H) \log \gamma(x_i|s, \tilde{\theta}_i, \mathcal{I}_j, H) d\tilde{\theta}_i$$

Before we move on, it is worth considering kinds of distributions these choices consider to be ‘private’. For example, considering the following two belief distributions:

$$\gamma_1 : x_j = U([- \epsilon, \epsilon]) \quad \gamma_2 : \begin{cases} M & \text{with probability } \delta \\ -\epsilon & \text{otherwise} \end{cases}$$

By choosing M large and ϵ small, the appropriate choice of δ can be made to force γ_2 to have arbitrarily large variance (small precision), but extremely low entropy. On the other hand, a uniform distribution γ_1 has relatively high entropy as compared to its variance. If agents measure their privacy by precision, they will be indifferent between adding noise from a fairly narrow uniform distribution and a discrete distribution which almost always adds a very small amount of noise but rarely adds an enormous amount. However, if they measure their privacy with respect to entropy, the uniform option offers a much greater amount of privacy, and since (as we will show in Corollary 2) agents pay a price in their utility equal to the variance of the noise-generating distribution, for a fixed variance the higher entropy uniform distribution will be strictly preferred to the discrete one.

We aim to show the following theorem:

Theorem 1. *Consider the modified game as defined above. Suppose players’ actions $\tilde{\theta}_i$ and the (possibly randomized) mapping $(y, x_i) \mapsto \tilde{\theta}_i$ are revealed at the end of the game, and players’ loss of privacy $\rho(\tilde{\theta}_i)$ is measured as either the reciprocal of the variance or the entropy of a representative player j ’s posterior belief about x_i at the end of the game. Then, this game has a symmetric linear Nash equilibrium where each player chooses $\tilde{\theta}_i = \kappa x_i + (1 - \kappa)y + \eta_i$ where κ is as it is in the original game and η_i is a random variable drawn from a distribution whose form depends on whether we use reciprocal variance or entropy to measure privacy.*

Proof sketch. We prove this theorem over the course of several steps throughout the remainder of this section, which contains the formal statements of the following:

1. First, we show (Claim 1) that if players do indeed choose their action by adding noise to their privacy-unaware equilibrium action, then the distribution which generates the noise must have mean zero.
2. Using this, we prove (Lemma 4) that the utility function *separates* into the sum of three parts: the utility in the original game, a penalty paid in the variance of the noise distribution, and the privacy component.
3. Then, since the utility function separates additively, we can use the first-order conditions (Lemma 1) to solve for the optimal choice of distribution from which to draw the noise (Corollary 3), which depends on whether the loss of privacy is measured by reciprocal variance or negative entropy.
4. Finally, we demonstrate that this strategy profile supports a Nash equilibrium by arguing that there is no profitable deviation for any player in the game (Lemmas 6 and 7). We call this a *symmetric noisy linear Nash equilibrium*.

□

We first define a Noisy Strategy:

Definition 4 (Linear) *H*-noisy strategy. *An H-noisy strategy has the form*

$$\tilde{\theta}(x, y, \nu) = \theta_i(x, y) + \eta_i$$

where η_i is a random variable drawn from a distribution H and $\theta_i(x, y)$ is a deterministic function of signals x and y . We say that $\tilde{\theta}$ is a normal noisy strategy if H is a Gaussian distribution. If, moreover, $\theta_i(x, y)$ is linear and can be written as $\theta_i(x, y) = \kappa x + (1 - \kappa)y$, we say that $\tilde{\theta}$ is a linear noisy strategy.

Note that *any* randomized strategy that can be decomposed into a deterministic component and a random component can be described as a noisy strategy. This means that Lemma 2 applies whether or not the underlying deterministic component is linear.

Claim 1. *If there exists an equilibrium in noisy strategies, where player i chooses $\tilde{\theta}_i = \theta_i + \eta_i$ and each player's η_i is drawn independently from a distribution H_i , then there exists such an equilibrium strategy profile in which the mean of each H_i is zero.*

If the parameter α is greater than $\frac{1}{2}$, then *every* such equilibrium requires players to choose a mean-zero noise-generating distribution. To see this, suppose all players $j \neq i$ choose a mean $\mu > 0$. Then for a small enough value of ϵ , player i choosing a noise-generating distribution with mean $\mu - \epsilon$ is a profitable deviation, since $\alpha > \frac{1}{2}$ means that her gain from moving $\tilde{\theta}_i$ closer to s more than offsets the loss from moving further from $\bar{\theta}$. Going forward, we assume without loss of generality that any noisy equilibrium is one in which the added noise comes from a mean-zero distribution.

We now define our equilibrium concept, which is analogous to that of the original game.

Definition 5 Symmetric noisy linear Nash equilibrium. *A symmetric noisy linear Nash equilibrium of this game is a strategy profile $\tilde{\theta}$ where each player i chooses*

$$\tilde{\theta}_i = \theta_i + \eta_i$$

and $\theta_i = \kappa x_i + (1 - \kappa)y$, κ is the same for all agents, and each player's noise η_i is drawn independently from the same distribution H .

One important point is that $\rho(\tilde{\theta}_i)$ is a function of the *optimal Bayesian posterior* distribution of x_i given $\tilde{\theta}_i$ and y , so an agent must consider his choice of action both in relation to the change in coordination in addition to others' beliefs about him; however, these beliefs are only what an optimal Bayes estimator, knowing the equilibrium profile, could *infer*. In particular, agents are concerned only by the knowledge of agents that are *correct*. It is conceivable that agents could worry about opponents *incorrectly* inferring their signals, but that falls outside the scope of this model.

3.2 Solving the Extended Game

Having defined the privacy-extended game, we can solve for the parameter values which support a symmetric noisy linear Nash equilibrium. Much of the analysis follows either directly from or along similar arguments as in the original game. We proceed as follows. First, we demonstrate in Corollary 2 that the utility function separates additively into three components: the utility in the original game, a penalty in this utility due to the added noise, and the privacy term. We then derive the first order condition of a representative agent in Proposition 4 and use this to find the optimal parameter values in Corollary 4, depending on whether we consider the precision-based or entropy-based measure of obfuscation. Finally, in Lemmas 6 and 7, we show that these values support an equilibrium. This completes the proof of Theorem 1.

We begin with the separability of the utility function.

Lemma 4 Separability. *In the game with finitely many players, the players' utility functions in the privacy-aware game separate additively into the utility in the privacy unaware game, a penalty in v_i , and a privacy term as*

$$\begin{aligned} v_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) &= (1 - \beta) \left(-\alpha(\tilde{\theta}_i - s)^2 - (1 - \alpha)(\tilde{\theta}_i - \bar{\theta})^2 \right) + \beta \rho(\tilde{\theta}_i) \\ &= (1 - \beta) u_i(\theta_i, \tilde{\theta}_{-i}) + (1 - \beta) \left(\alpha + \left(1 - \frac{1}{n} \right)^2 (1 - \alpha) \right) v_i + \beta \rho(\tilde{\theta}_i), \end{aligned}$$

where v_i denotes the variance of the noise-generating distribution H_i of player i and u_i is the utility function in the privacy-unaware game.

Corollary 2. *Suppose that all agents play a noisy strategy $\tilde{\theta}_i = \theta_i + \eta_i$, with η_i being a random variable and $\mathbb{E}(\eta_i) = 0$. Then an agent's utility can be decomposed into*

$$\mathbb{E}_i[v_i(\tilde{\theta}, \tilde{\theta}_{-i})] = (1 - \beta)\mathbb{E}_i[u(\theta_i, \tilde{\theta}_{-i})] + (1 - \beta)\nu_i + \beta\rho(\tilde{\theta}_i)$$

where ν_i denotes the variance of the noise-generating distribution H_i of player i and u_i is the utility function in the privacy-unaware game.

This shows that agents can evaluate what their optimal action would be in the original game and then decide the optimal amount of noise to add, choosing their action as the realized value of the noise plus their optimal action.

We next derive the first order conditions for the privacy-extended game. By Lemma 4 and Corollary 2, the first order conditions can be disentangled into a first order condition on the deterministic component, which must be as in Proposition 1 and Lemma 1, and a separate first order condition on the variance of the random component. Because the added noise is mean-zero, expectations about average action and state are as before, and the κ in the optimal deterministic action is the same as before.

Lemma 5 Privacy-aware first order conditions. *In an equilibrium of the game with finitely many players where the optimal action is $\tilde{\theta}_i^* = \theta_i^* + \eta_i$, the optimal choice of θ_i^* and the variance ν^* for the noise-generating distribution H_i from which η_i is drawn must satisfy*

$$\theta_i^* = \frac{\alpha n^2 \mathbb{E}_i[s]}{\alpha(2n-1) + (n-1)^2} + \frac{(1-\alpha)(n-1)\mathbb{E}_i\left[\sum_{j \neq i} \theta_j\right]}{\alpha(2n-1) + (n-1)^2},$$

$$\frac{\partial \rho}{\partial \nu^*} = -\frac{-(1-\beta)\left(\alpha + \left(1 - \frac{1}{n}\right)^2(1-\alpha)\right)}{\beta}.$$

Proposition 4. *In the game with infinitely many players, in an equilibrium where the optimal action is $\tilde{\theta}_i^* = \theta_i^* + \eta_i$, the optimal choice of θ_i^* and the variance ν^* for the noise-generating distribution H_i from which η_i is drawn must satisfy*

$$\theta_i^* = \alpha \mathbb{E}_i[s] + (1-\alpha)\mathbb{E}_i[\tilde{\theta}] \quad \frac{\partial \rho}{\partial \nu^*} = -\frac{1-\beta}{\beta}.$$

Corollary 3 Finite game privacy parameters. *The optimal deterministic component θ_i is, as before,*

$$\theta_i^* = \kappa x_i + (1-\kappa)y$$

and the optimal choice of variance for the noise distribution is

$$\nu_{i,prec}^* = \sqrt{\frac{\beta}{1-\beta} \left(\alpha + (1-\alpha) \left(1 - \frac{1}{n}\right)^2 \right)} \quad \text{or} \quad \nu_{i,ent}^* = \frac{\beta}{1-\beta} \left(\alpha + (1-\alpha) \left(1 - \frac{1}{n}\right)^2 \right)$$

where $\nu_{i,prec}^*$ and $\nu_{i,ent}^*$ are the optimal variances under ρ being the precision and entropy privacy measures, respectively.

Corollary 4 Infinite game privacy parameters. *The optimal deterministic θ_i is, as before,*

$$\theta_i^* = \kappa x_i + (1-\kappa)y$$

and the optimal choice of variance for the noise distribution is

$$\nu_{i,prec}^* = \sqrt{\frac{\beta}{1-\beta}} \quad \nu_{i,ent}^* = \frac{\beta}{1-\beta}$$

where $\nu_{i,prec}^*$ and $\nu_{i,ent}^*$ are the optimal choices of variance when ρ is the precision-based or entropy-based privacy measure, respectively.

Before proving this, we state a fact about the entropy of Gaussian distributions.

Fact 1. *Among all distributions supported on the entire real line with a fixed mean μ and variance σ^2 , the Gaussian $\mathcal{N}(\mu, \sigma^2)$ achieves the maximum entropy, and its entropy is given by $\frac{1}{2} \log(2\pi e\sigma^2)$.*

Since agents pay a penalty in the variance of their noise distribution, agents who measure their privacy loss with entropy pay the same penalty for any distribution with a fixed variance, and their gain from privacy is maximized by picking the distribution with maximum entropy. Therefore, the choice of a mean-zero Gaussian with appropriate variance is the dominant strategy for such agents. A proof of this fact can be found in Chapter 9 of [7].

Proof of Corollaries 3 and 4. The fact that the first order condition on θ_i^* is the same as that of θ_i^* in the original game, and that agents add mean-zero noise (implying $\mathbb{E}_i[\tilde{\theta}] = \mathbb{E}_i[\theta]$) implies the optimal choice of θ_i^* . To find the optimal ν^* , we note that the reciprocal variance and entropy penalties (respectively) are

$$\rho_{prec}(\nu) = \frac{-1}{\nu} \quad \rho_{ent}(\nu) = \frac{1}{2} \log(2e\pi\nu).$$

Differentiating these with respect to ν and rearranging gives the results. □

We can observe that the optimal choice of ν_i^* is similar for both of these functions, even if their constructions are not. Each is increasing at a decreasing rate as ν grows, since their derivatives are of the form $\frac{1}{\nu^c}$ for $c \geq 1$. Since players gain a diminishing marginal benefit for adding additional noise as they increase ν but pay a penalty linear in ν , we should expect this trade-off to point to an equilibrium. This will hold for any penalty function whose derivative is of this type, although the interpretation of such a function may not be as natural of a property of a distribution as precision or entropy.

Using these, a symmetric linear noisy Nash equilibrium of this game follows:

Lemma 6 Equilibrium – precision loss. *There exists a symmetric linear noisy Nash equilibrium in which all agents use κ as defined in Lemma 2 and draw independent noise from a distribution H_i with variance ν^* , where H_i can be any distribution with mean zero and variance $\nu_{i,prec}^* = \sqrt{\frac{\beta}{1-\beta}}$. Here, agents may choose any distribution with these properties.*

Lemma 7 Equilibrium – entropy loss. *There exists a symmetric linear noisy Nash equilibrium in which all agents use κ as defined in Lemma 2 and draw independent noise from the Gaussian distribution $N(0, \nu_{ent}^*)$.*

Since the Gaussian is the maximum entropy distribution with fixed variance ν^ , all agents draw noise (independently) from the same distribution.*

This completes the proof of Theorem 1, since we have found a symmetric linear noisy action for each player which satisfies the first order conditions of the privacy aware game. We next analyze and discuss the price of privacy in the new game.

4 Price of Privacy

Informally, the price of privacy describes the loss in quality of some measure as a result of introducing privacy-awareness into the game. We describe two different, but related, quantities which represent this effect and discuss some settings where one may be interested in each.

4.1 The Agents' Cost

Our first notion of the price of privacy can be viewed as the utility the *agents* pay in the original game in order to express their value for obfuscation. Formally, recall that $u_i(\theta_i, \theta_{-i}) = -(1-\alpha)(\theta_i - \bar{\theta})^2 - \alpha(\theta_i - s)^2$ is the utility function in the unmodified game. In the symmetric linear Nash equilibrium, each player chooses the prescribed optimal action θ_i^* and earns utility $u_i(\theta_i^*, \theta_{-i}^*)$. Similarly, in the modified game, in the symmetric linear Nash equilibrium chooses the optimal noisy action $\tilde{\theta}_i^*$. We can now ask how much worse-off playing the actions $\tilde{\theta}^*$ in the original game makes a representative player as compared to playing the actions θ^* . Formally, we can quantify the price of privacy by taking the ratio of these utilities.

Definition 6 Agents' price of privacy. *The price of privacy given an information structure is*

$$\text{PoP}(\tau_x, \tau_y, \beta) = \frac{\mathbb{E}_i[u_i(\tilde{\theta}_i^*, \tilde{\theta}_{-i}^*)]}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]}$$

where the expected utility is with respect to the game with signal variances σ_x^2 and σ_y^2 .

Lemma 8 Agents' price of privacy – form. *The price of privacy in the game where agents play a linear strategy has the form*

$$\text{PoP}(\tau_x, \tau_y, \beta) = 1 + \frac{\nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]}$$

where the expected utility is with respect to the game with signal variances σ_x^2 and σ_y^2 .

Theorem 2 Price of privacy – value. *In the game with finitely many agents, the price of privacy is:*

$$\text{PoP}(\tau_x, \tau_y, \beta) = 1 + \frac{\nu_i^*}{\alpha(\kappa^2\sigma_x^2 + (1-\kappa)^2\sigma_y^2) + \frac{(1-\alpha)\kappa^2(n-1)}{n}\sigma_x^2},$$

for

$$\nu_{i,\text{prec}}^* = \sqrt{\frac{\beta}{1-\beta} \left(\alpha + (1-\alpha) \left(1 - \frac{1}{n}\right)^2 \right)} \quad \text{or} \quad \nu_{i,\text{ent}}^* = \frac{\beta}{1-\beta} \left(\alpha + (1-\alpha) \left(1 - \frac{1}{n}\right)^2 \right)$$

if we measure the privacy loss using precision or entropy, respectively, and

$$\kappa = \frac{\alpha n^2 \tau_x}{\alpha n^2 \tau_x + ((n-1)^2 + \alpha(2n-1)) \tau_y}.$$

Proposition 5. *In the game with infinitely many agents, the price of privacy is:*

$$\text{PoP}(\tau_x, \tau_y, \beta) = 1 + \frac{\nu_i^*}{(\alpha(1-\kappa)^2\sigma_y^2 + \kappa^2\sigma_x^2)},$$

where ν_i^* is $\sqrt{\frac{\beta}{1-\beta}}$ if privacy is measured by precision and $\frac{\beta}{1-\beta}$ if privacy is measured by entropy and recalling that $\kappa = \frac{\tau_x \alpha}{\tau_x \alpha + \tau_y}$.

Proof. This follows directly from the expected utility computation in Proposition 3 and the functional form of the price of privacy in Lemma 8. \square

We can observe several features of the price of privacy. First, it can be *arbitrarily large*, depending on the value of the parameters. If β is close to 1, then agents have a relatively high value for obfuscation and will add large amounts of noise to their actions in order to protect their private signals. As β becomes closer to zero, agents do not care too much about privacy and their actions become more and more similar to those in the original game, so the price of privacy is minimized when $\beta = 0$. Second, if we fix a value of β , then the price of privacy decreases as the variance of the public and private signals *increase*. This is because the value of β determines the noise that players will add to their signals, and therefore fixes the numerator of the fractional part of the price of privacy. The price of privacy is *decreasing* in any factor that improves the expected utility, such as decreasing the signal variances or n (as in Corollary 1). The effect is ambiguous when changing σ_y^2 , due to the over-weighting phenomenon.

This ratio measures the degree to which agents are worse-off in the coordination and guessing components by playing the privacy-aware equilibrium noisy actions as compared to the privacy-unaware equilibrium actions for a fixed α , σ_x^2 , and σ_y^2 and the realizations of x_i and y . This can be thought of as describing the increased risk of making the ‘wrong’ decision in an opinion-aggregation setting or of misvaluing an asset in a financial markets setting as a result of agents adding noise to their actions.

4.2 The Aggregator’s Cost

Suppose instead we are viewing this game from the position of a data aggregator. The aggregator can observe the actions of some finite number of agents n and takes the average of these to estimate the true state of the world s by taking a simple average. The aggregator does not know the realizations of any of the x_i or y , but she does know the signal variances and the value of κ ; she also knows that $\mathbb{E}[\theta_i] = \mathbb{E}[\tilde{\theta}_i] = s$ for all agents i , so the sample average will provide an unbiased estimate of s . If we define the aggregator’s ‘utility’ to be the variance of its sample about s , then we can make the following observation:

Lemma 9 Aggregator’s utility. *Consider an instance of the privacy-aware game where an aggregator observes the actions of n agents (either all of the agents in the finite case or some uniformly random sample in the finite or infinite case), the signal variances are σ_x^2 and σ_y^2 , and players choose to add mean-zero noise with variance ν_i^* . Then the utility of the aggregator, as measured by the variance of the sample average about the true state s is given by*

$$\begin{aligned} \mathcal{U}_{agg}(\sigma_x^2, \sigma_y^2, \nu_i^*, n) &= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \tilde{\theta}_i \right) - s \right)^2 \middle| s \right] \\ &= \frac{\kappa^2}{n} \sigma_x^2 + \frac{\nu_i^*}{n} + (1 - \kappa)^2 \sigma_y^2. \end{aligned}$$

We can find the aggregator’s utility in the privacy-unaware game by letting $\nu_i^* = 0$, which recovers a result in Morris and Shin [17].

We can again take the ratio of the aggregator’s utility in the privacy-aware game to that in the privacy-unaware game to quantify the extent to which privacy-awareness degrades the quality of the aggregator’s sample mean as an estimate of s .

Lemma 10 The Aggregator’s price of privacy. *The price of privacy for the aggregator in equilibrium is given by*

$$\text{PoP}_{agg}(\sigma_x^2, \sigma_y^2, n) = \frac{\mathbb{E}[\mathcal{U}_{agg}(\sigma_x^2, \sigma_y^2, \nu_i^*, n)]}{\mathbb{E}[\mathcal{U}_{agg}(\sigma_x^2, \sigma_y^2, 0, n)]} = 1 + \frac{\nu_i^*}{\kappa^2 \sigma_x^2 + n(1 - \kappa)^2 \sigma_y^2}$$

where ν_i^* is $\sqrt{\frac{\beta}{1-\beta}}$ if privacy is measured by precision and $\frac{\beta}{1-\beta}$ if privacy is measured by entropy and recalling that $\kappa = \frac{\tau_x \alpha}{\tau_x \alpha + \tau_y}$ in the game with infinitely many agents and $\kappa = \frac{\alpha n^2 \tau_x}{\alpha n^2 \tau_x + ((n-1)^2 + \alpha(2n-1)) \tau_y}$ in the game with finitely many agents.

Proof. This follows from plugging in the forms for the expected utilities from Lemma 9 and simplifying. \square

We can again observe that, for the same reason as the agents’ price of privacy, that this quantity can be arbitrarily large as β approaches one and that it is decreasing as the signal variances grow. We can also observe that this price decreases towards one as n grows, all else fixed. This is consistent with the intuition that an aggregator can offset the cost of agents adding more noise to their actions by making more observations.

The agents’ and aggregator’s prices of privacy are somewhat similar both in functional form and in interpretation and this results from their values almost aligning. The agents’ cost is measured by a weighted deviation from both the true state and the average while the aggregator’s is measured by a deviation from the true state alone. Any configuration of the game parameters which cause $\kappa = 1$ (i.e. agents ignore the public signal) causes them to be exactly equal. This again highlights the problem of over-weighting the public signal, and this issue is exacerbated for the aggregator, who, in effect, observes n ‘copies’ of y rolled into the actions θ_i and $\tilde{\theta}_i$. If the aggregator knows the realization of y , she can, similarly to the agents, use that information to find (an estimate of) the private signals x_i from the observations of θ_i or $\tilde{\theta}_i$. This motivates the same privacy concern as in the case where other agents were trying to infer agent i ’s private signal, this time in the context of a distrusted central aggregator.

5 Discussion and Future Directions

In this work, we used a game-theoretic model to examine a game with agents acting to obscure their private information. We began with the Keynesian Beauty Contest and modified the agents’ utility functions so as to endogenize a notion of privacy. Using this, we can quantify the social costs of this ‘selfish’ privacy protection as a ‘price of privacy’, both with respect to the agents as well as an (untrusted) central aggregator.

A clear next-step would be to repeat this analysis for other stylized models of strategic interaction, such as bargaining games or resource allocation problems, where individuals seek to balance expressing their preferences well enough to achieve a high payoff but choose to deviate slightly from this in order to obscure their true valuations.

Finally, the high-level motivation in this work is to examine how strategic agents behave when they are concerned with privacy, so another interesting direction would be to integrate this approach with existing formal notions of privacy. For example, examining what agents’ utility functions and action spaces must look like in order for there to exist an equilibrium which implements a locally differentially private mechanism could be a fruitful avenue of research.

Acknowledgments

We would like to graciously thank Annie Liang and Aaron Roth for their guidance in the research process and feedback on previous drafts of this project. We would additionally like to thank the anonymous reviewers for their helpful feedback.

References

- [1] Franklin Allen, Stephen Morris, and Hyun Song Shin. 2006. Beauty contests and iterated expectations in asset markets. *The Review of Financial Studies* 19, 3 (2006), 719–752.
- [2] Raef Bassily. 2018. Linear Queries Estimation with Local Differential Privacy. *arXiv preprint arXiv:1810.02810* (2018).
- [3] James O Berger, José M Bernardo, Dongchu Sun, et al. 2009. The formal definition of reference priors. *The Annals of Statistics* 37, 2 (2009), 905–938.

- [4] Antoni Bosch-Domenech, Jose G Montalvo, Rosemarie Nagel, and Albert Satorra. 2002. One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments. *American Economic Review* 92, 5 (2002), 1687–1701.
- [5] Giovanni Cespa and Xavier Vives. 2015. The Beauty Contest and Short-Term Trading. *The Journal of Finance* 70, 5 (2015), 2099–2154.
- [6] Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. 2016. Truthful mechanisms for agents that value privacy. *ACM Transactions on Economics and Computation (TEAC)* 4, 3 (2016), 13.
- [7] Thomas M Cover and Joy A Thomas. 2012. *Elements of information theory*. John Wiley & Sons.
- [8] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.
- [9] Pingyang Gao. 2008. Keynesian beauty contest, accounting disclosure, and market efficiency. *Journal of Accounting Research* 46, 4 (2008), 785–807.
- [10] Ronen Gradwohl and Rann Smorodinsky. 2017. Perception games and privacy. *Games and Economic Behavior* 104 (2017), 293–308.
- [11] Christian Hellwig and Laura Veldkamp. 2009. Knowing What Others Know: Coordination Motives in Information Acquisition. *The Review of Economic Studies* 76, 1 (2009), 223–251. <https://doi.org/10.1111/j.1467-937X.2008.00515.x>
- [12] Matthew Joseph, Janardhan Kulkarni, Jieming Mao, and Zhiwei Steven Wu. 2018. Locally Private Gaussian Estimation. *arXiv preprint arXiv:1811.08382* (2018).
- [13] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2016. Extremal Mechanisms for Local Differential Privacy. *J. Mach. Learn. Res.* 17, 1 (Jan. 2016), 492–542. <http://dl.acm.org/citation.cfm?id=2946645.2946662>
- [14] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2011. What can we learn privately? *SIAM J. Comput.* 40, 3 (2011), 793–826.
- [15] John Maynard Keynes. 1936. *The general theory of employment, interest, and money*. Macmillan.
- [16] Fang Liu. 2019. Generalized gaussian mechanism for differential privacy. *IEEE Transactions on Knowledge and Data Engineering* 31, 4 (2019), 747–756.
- [17] Stephen Morris and Hyun Song Shin. 2002. Social Value of Public Information. *American Economic Review* 92, 5 (December 2002), 1521–1534. <https://doi.org/10.1257/000282802762024610>
- [18] Rosemarie Nagel, Christoph Bühren, and Björn Frank. 2017. Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game. *Mathematical Social Sciences* 90 (2017), 191–207.
- [19] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*. IEEE, 111–125.
- [20] Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. 2012. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 774–789.
- [21] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. 2019. Collecting and Analyzing Multidimensional Data with Local Differential Privacy. In *Proceedings of IEEE ICDE*.
- [22] Stanley L. Warner. 1965. Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69. <http://www.jstor.org/stable/2283137>

A Omitted Proofs

We present here the statements and proofs of the results whose proofs were omitted from the main text. We also present from first-principles the proofs of the analogous results in the game with infinitely many players, which demonstrates that the proofs where we took the limit of the finite result as the number of players grows to infinity are indeed correct.

Omitted Proofs from Section 2

Lemma 1. *In equilibrium, an agent's optimal choice of θ_i^* must satisfy*

$$\theta_i^* = \frac{\alpha n^2 \mathbb{E}_i[s]}{\alpha(2n-1) + (n-1)^2} + \frac{(1-\alpha)(n-1) \mathbb{E}_i \left[\sum_{j \neq i} \theta_j \right]}{\alpha(2n-1) + (n-1)^2}$$

Proof. Fix the actions of all other players $j \neq i$ and consider the utility of player i , recalling that in the finite setting,

$$\bar{\theta} = \frac{1}{n} \theta_i^* + \frac{1}{n} \sum_{j \neq i} \theta_j.$$

We have

$$\begin{aligned} u_i(\theta_i^*, \theta_{-i}) &= -\alpha(\theta_i^* - s)^2 - (1-\alpha)(\theta_i^* - \bar{\theta})^2 \\ &= -\alpha(\theta_i^* - s)^2 - (1-\alpha) \left(\theta_i^* - \left(\frac{1}{n} \theta_i^* + \frac{1}{n} \sum_{j \neq i} \theta_j \right) \right)^2 \\ &= -\alpha(\theta_i^* - s)^2 - (1-\alpha) \left(\frac{n-1}{n} \theta_i^* - \left(\frac{1}{n} \sum_{j \neq i} \theta_j \right) \right)^2 \\ &= -\alpha(\theta_i^* - s)^2 - (1-\alpha) \frac{1}{n^2} \left((n-1)\theta_i^* - \left(\sum_{j \neq i} \theta_j \right) \right)^2 \end{aligned}$$

Player i 's first order condition will be to maximize this in expectation. Taking an expectation, we have

$$\begin{aligned} \mathbb{E}_i[u_i(\theta_i, \theta_{-i})] &= \mathbb{E}_i \left[-\alpha(\theta_i^* - s)^2 - (1-\alpha) \frac{1}{n^2} \left((n-1)\theta_i^* - \left(\sum_{j \neq i} \theta_j \right) \right)^2 \right] \\ &= -\alpha \mathbb{E}_i[(\theta_i^* - s)^2] - (1-\alpha) \frac{1}{n^2} \mathbb{E}_i \left[\left((n-1)\theta_i^* - \left(\sum_{j \neq i} \theta_j \right) \right)^2 \right] \\ &= -\alpha \mathbb{E}_i[(\theta_i^*)^2 - 2s\theta_i^* + s^2] \\ &\quad - (1-\alpha) \frac{1}{n^2} \mathbb{E}_i \left[(n-1)^2(\theta_i^*)^2 - 2(n-1)\theta_i^* \left(\sum_{j \neq i} \theta_j \right) + \left(\sum_{j \neq i} \theta_j \right)^2 \right] \end{aligned}$$

Differentiating with respect to agent i 's choice of θ_i gives us that, in equilibrium, θ_i^* must satisfy

$$\begin{aligned}\frac{\partial}{\partial \theta_i^*} = 0 &= -\alpha(2\theta_i^* - 2\mathbb{E}_i[s]) - (1-\alpha)\frac{1}{n^2} \left(2(n-1)^2\theta_i^* - 2(n-1)\mathbb{E}_i \left[\sum_{j \neq i} \theta_j \right] \right) \\ &= \alpha(\theta_i^* - \mathbb{E}_i[s]) + \frac{1-\alpha}{n^2} \left((n-1)^2\theta_i^* - (n-1)\mathbb{E}_i \left[\sum_{j \neq i} \theta_j \right] \right) \\ &= \theta_i^* \left(\alpha + \frac{(1-\alpha)(n-1)^2}{n^2} \right) - \left(\alpha\mathbb{E}_i[s] + \frac{(1-\alpha)(n-1)}{n^2} \mathbb{E}_i \left[\sum_{j \neq i} \theta_j \right] \right)\end{aligned}$$

Where we have used the fact that $\mathbb{E}_i[\theta_i^*] = \theta_i^*$, since it is agent i 's choice and not a random variable. Rearranging yields the claim. \square

Proposition 1. *In the game with infinitely many agents, in equilibrium, an agent's optimal choice of θ_i^* must satisfy*

$$\theta_i^* = \alpha\mathbb{E}_i[s] + (1-\alpha)\mathbb{E}_i[\bar{\theta}]$$

Proof. Since agent i chooses his action after observing signals x_i and y , they optimize their action given those signals. We can write:

$$\theta_i^* \in \arg \max \mathbb{E}_i[u_i(\theta_i)]$$

Expanding, we have:

$$\mathbb{E}_i[u_i(\theta_i, \theta_{-i})] = -\alpha(\mathbb{E}_i[\theta_i^2 - 2\theta_i s + s^2]) - (1-\alpha)(\mathbb{E}_i[\theta_i^2 - 2\theta_i \bar{\theta} + \bar{\theta}^2]).$$

We can differentiate with respect to θ_i^* and set the derivative equal to zero to find

$$0 = \frac{\partial}{\partial \theta_i} \mathbb{E}_i[u_i(\theta_i, \theta_{-i})] = -\alpha(2\theta_i - 2\mathbb{E}_i[s]) - (1-\alpha)[2\theta_i - \bar{\theta}].$$

Where we have used the fact that $\mathbb{E}_i[\theta_i^*] = \theta_i^*$, since it is agent i 's choice and not a random variable, as well as the fact that $\frac{\partial \bar{\theta}}{\partial \theta_i^*} = 0$ because there are infinitely many agents, so agent i 's action cannot affect the average of all of the actions.

Rearranging yields the claim. \square

Remark. *These results show that the optimal deterministic action is a linear combination of expectations about the state and average action; thus, the optimal deterministic action and the optimal linear action are one and the same. Because these propositions hold also for the first-order conditions in the extended game, this will be true in the extended game as well.*

Remark. *If we think of $\sum \theta_j$ as $(n-1)$ times the average action of other players and take the limit of this expression as n goes to infinity, we recover the result for the continuum case, as the coefficient on $\mathbb{E}_i[s]$ becomes α and the coefficient on the expected average action becomes $(1-\alpha)$.*

Lemma 2. *In the game with finitely many agents, the value of κ which supports a symmetric linear Nash equilibrium is*

$$\kappa = \frac{\alpha n^2 \tau_x}{\alpha n^2 \tau_x + ((n-1)^2 + \alpha(2n-1)) \tau_y}$$

Proof. Suppose that all agents play $\theta_i = \kappa_i + (1 - \kappa)y$, and consider a representative agent i .

Since agent i is Bayesian, she aggregates the public and private signals according to their precisions to compute

$$\mathbb{E}_i[s] = \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y}$$

and her belief about any other agent j 's private signal x_j is that

$$\mathbb{E}_i[x_j] = \mathbb{E}_i[s].$$

Supposing that each agent acts optimally and chooses θ_i^* according to Lemma 1, agent i will write the following. We let $\mathcal{C} = \frac{1}{\alpha(2n-1) + (n-1)^2}$ for clarity of notation.

$$\begin{aligned} \theta_i^* &= \frac{\alpha n^2 \mathbb{E}_i[s]}{\alpha(2n-1) + (n-1)^2} + \frac{(1-\alpha)(n-1) \mathbb{E}_i \left[\sum_{j \neq i} \theta_j \right]}{\alpha(2n-1) + (n-1)^2} \\ &= \mathcal{C} \left(\alpha n^2 \mathbb{E}_i[s] + (1-\alpha)(n-1) \mathbb{E}_i \left[\sum_{j \neq i} (\kappa x_j + (1-\kappa)y) \right] \right) \\ &= \mathcal{C} \left(\alpha n^2 \mathbb{E}_i[s] + (1-\alpha)(n-1) \left[\sum_{j \neq i} (\kappa \mathbb{E}_i[s] + (1-\kappa)y) \right] \right) \\ &= \mathcal{C} (\alpha n^2 \mathbb{E}_i[s] + (1-\alpha)(n-1)^2 (\kappa \mathbb{E}_i[s] + (1-\kappa)y)) \\ &= \mathcal{C} \left(\alpha n^2 \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y} + (1-\alpha)(n-1)^2 \left(\kappa \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y} + (1-\kappa)y \right) \right) \\ &= \mathcal{C} \left(\alpha n^2 \frac{\tau_x}{\tau_x + \tau_y} + (1-\alpha)(n-1)^2 \kappa \frac{\tau_x}{\tau_x + \tau_y} \right) x_i + \\ &\quad + \mathcal{C} \left(\alpha n^2 \frac{\tau_y}{\tau_x + \tau_y} + (1-\alpha)(n-1)^2 \left(\kappa \frac{\tau_y}{\tau_x + \tau_y} + (1-\kappa) \right) \right) y \end{aligned}$$

We can then equate κ and the coefficient on x_i or the coefficient on y to $(1 - \kappa)$. By construction, it's easy to see that the sum of these two coefficients will be one, so it suffices to perform the computation for κ .

$$\begin{aligned} \kappa &= \mathcal{C} \left(\alpha n^2 \frac{\tau_x}{\tau_x + \tau_y} + (1-\alpha)(n-1)^2 \kappa \frac{\tau_x}{\tau_x + \tau_y} \right) \\ &= \mathcal{C} \alpha n^2 \frac{\tau_x}{\tau_x + \tau_y} + \kappa \mathcal{C} (1-\alpha)(n-1)^2 \frac{\tau_x}{\tau_x + \tau_y} \\ &= \frac{\mathcal{C} \alpha n^2 \frac{\tau_x}{\tau_x + \tau_y}}{1 - \mathcal{C} (1-\alpha)(n-1)^2 \frac{\tau_x}{\tau_x + \tau_y}} \end{aligned}$$

Plugging the value for \mathcal{C} back in and simplifying, we get

$$\kappa = \frac{\alpha n^2 \tau_x}{\alpha n^2 \tau_x + ((n-1)^2 + \alpha(2n-1)) \tau_y}$$

□

Proposition 2. *The symmetric linear Nash equilibrium in the game with infinitely many players is given by*

$$\begin{aligned}\theta_i &= \kappa x_i + (1 - \kappa)y, \text{ where} \\ \kappa &= \frac{\tau_x \alpha}{\tau_x \alpha + \tau_y}, \quad \tau_x = \frac{1}{\sigma_x^2}, \quad \text{and } \tau_y = \frac{1}{\sigma_y^2} \\ &\text{for all players } i.\end{aligned}$$

Proof. Proposition 1 shows that agent i 's optimal action is a convex combination of his belief about the true state s and his belief about the average action of all players $\bar{\theta}$. $\mathbb{E}_i[s]$ is independent of the equilibrium profile, and given by

$$\mathbb{E}_i[s] = \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y},$$

which is the standard prior-free aggregation of independent Gaussian signals.

On the other hand, $\mathbb{E}_i[\bar{\theta}]$ does depend on the equilibrium profile. In a SLNE, all other players $j \neq i$ play

$$\theta_j = \kappa x_j + (1 - \kappa)y.$$

Since agent i is Bayesian:

$$\mathbb{E}_i[x_j] = \mathbb{E}_i[s]$$

and

$$\mathbb{E}_i[\theta_j] = \mathbb{E}_i[\kappa x_j + (1 - \kappa)y] = \kappa \mathbb{E}_i[x_j] + (1 - \kappa)y.$$

Notice that the expectation of agent i about the belief of any representative player $j \neq i$ is the same, since his information is symmetric. Thus, we write $\mathbb{E}_i[\theta_{-i}]$ to emphasize that this is the belief of player i about any other player. Thus:

$$\mathbb{E}_i[\bar{\theta}] = \mathbb{E}_i \int_0^1 \theta_j dj = \int_0^1 \mathbb{E}_i[\theta_{-i}] = \mathbb{E}_i[\theta_{-i}] = \kappa \mathbb{E}_i[x_{-i}] + (1 - \kappa)y$$

where we have used the Fubini-Tonelli theorem to exchange the integral with the expectation.

But now note that

$$\begin{aligned}\theta_i^* &= \alpha \mathbb{E}_i[s] + (1 - \alpha)[\kappa \mathbb{E}_i[s] + (1 - \kappa)y] \\ &= \alpha \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y} + (1 - \alpha) \left(\kappa \frac{\tau_x x_i + \tau_y y}{\tau_x + \tau_y} \right) + (1 - \alpha)(1 - \kappa)y\end{aligned}$$

We can rearrange this as:

$$\begin{aligned}\theta_i^* &= \left[\frac{\alpha \tau_x}{\tau_x + \tau_y} + (1 - \alpha) \frac{\kappa \tau_x}{\tau_x + \tau_y} \right] x_i + \left[\frac{\tau_y \alpha}{\tau_x + \tau_y} + (1 - \alpha) \frac{\kappa \tau_y}{\tau_x + \tau_y} + (1 - \alpha)(1 - \kappa) \right] y \\ &= \left[\frac{\tau_x (\alpha + (1 - \alpha)\kappa)}{\tau_x + \tau_y} \right] x_i + \left[(1 - \alpha)(1 - \kappa) + \frac{\alpha \tau_y + (1 - \alpha)\kappa \tau_y}{\tau_x + \tau_y} \right] y\end{aligned}$$

Matching coefficients and solving:

$$\left[(\alpha + (1 - \alpha)\kappa) \frac{\tau_x}{\tau_x + \tau_y} \right] = \kappa \implies \kappa = \frac{\alpha \tau_x}{\alpha \tau_x + \tau_y}$$

as desired. □

Remark. As we take the limit as n goes to infinity, we recover $\kappa = \frac{\alpha\tau_x}{\alpha\tau_x + \tau_y}$ as in the continuum setting and in Morris and Shin [17].

Lemma 3. The expected utility of agents for playing $\theta_i = \kappa x_i + (1 - \kappa)y$, given s is:

$$\mathbb{E}_i[u_i|s] = -\alpha(\kappa^2\sigma_x^2 + (1 - \kappa)^2\sigma_y^2) - \frac{(1 - \alpha)\kappa^2(n - 1)}{n} \sigma_x^2$$

Proof. We can write the expected utility of agents as

$$\mathbb{E}_i[u_i|s] = -\alpha\mathbb{E}_i[(\theta_i - s)^2|s] - (1 - \alpha)\mathbb{E}_i\left[(\theta_i - \bar{\theta})^2 \middle| s\right]$$

We can plug in the equilibrium strategies and recall the definition of $\bar{\theta}$ to write

$$\mathbb{E}_i[u_i|s] = -\alpha\mathbb{E}_i[(\theta_i - s)^2|s] - (1 - \alpha)\mathbb{E}_i\left[\left(\theta_i - \frac{1}{n}\sum_{j=1}^n\theta_j\right)^2 \middle| s\right]$$

We can write each $\theta_j = \kappa x_j + (1 - \kappa)y$ and since we are taking expectations conditional on knowing s , we can write the signals x_j and y as the state s plus mean-zero Gaussian noise with variance σ_x^2 and σ_y^2 , respectively. We write these as $s + \varepsilon_{x_j}$ and $s + \varepsilon_y$, so we can write $\theta_j = s + \kappa\varepsilon_{x_j} + (1 - \kappa)\varepsilon_y$.

$$\begin{aligned} \mathbb{E}_i[u_i|s] &= -\alpha\mathbb{E}_i[(s + \kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y - s)^2|s] \\ &\quad - (1 - \alpha)\mathbb{E}_i\left[\left(s + \kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y - \frac{1}{n}\sum_{j=1}^n s + \kappa\varepsilon_{x_j} + (1 - \kappa)\varepsilon_y\right)^2 \middle| s\right] \\ &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - (1 - \alpha)\mathbb{E}_i\left[\left(\kappa\varepsilon_{x_i} + -\frac{1}{n}\sum_{j=1}^n \kappa\varepsilon_{x_j}\right)^2\right] \\ &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - (1 - \alpha)\kappa^2\mathbb{E}_i\left[\left(\varepsilon_{x_i} + -\frac{1}{n}\sum_{j=1}^n \varepsilon_{x_j}\right)^2\right] \\ &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - (1 - \alpha)\kappa^2\mathbb{E}_i\left[\left(\frac{n-1}{n}\varepsilon_{x_i} + -\frac{1}{n}\sum_{j \neq i}^n \varepsilon_{x_j}\right)^2\right] \\ &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - \frac{(1 - \alpha)\kappa^2}{n^2}\mathbb{E}_i\left[\left((n-1)\varepsilon_{x_i} + -\sum_{j \neq i}^n \varepsilon_{x_j}\right)^2\right] \end{aligned}$$

Because all of the ε_{x_j} and ε_y are independent with mean zero, we can write

$$\begin{aligned} \mathbb{E}_i[u_i|s] &= -\alpha(\kappa^2\sigma_x^2 + (1 - \kappa)^2\sigma_y^2) - \frac{(1 - \alpha)\kappa^2}{n^2} ((n - 1)^2\sigma_x^2 + (n - 1)\sigma_x^2)^2 \\ &= -\alpha(\kappa^2\sigma_x^2 + (1 - \kappa)^2\sigma_y^2) - \frac{(1 - \alpha)\kappa^2(n - 1)}{n} \sigma_x^2 \end{aligned}$$

□

Proposition 3. *The expected utility of agents for playing $\theta_i = \kappa x_i + (1 - \kappa)y$, conditional on s , is:*

$$\mathbb{E}_i[u_i|s] = -\alpha(1 - \kappa)^2\sigma_y^2 - \kappa^2\sigma_x^2$$

Proof. We can write the expected utility of agents as

$$\mathbb{E}_i[u_i|s] = -\alpha\mathbb{E}_i[(\theta_i - s)^2|s] - (1 - \alpha)\mathbb{E}_i\left[(\theta_i - \bar{\theta})^2\middle|s\right]$$

We can plug in the equilibrium strategies and recall the definition of $\bar{\theta}$ to write

$$\mathbb{E}_i[u_i|s] = -\alpha\mathbb{E}_i[(\theta_i - s)^2|s] - (1 - \alpha)\mathbb{E}_i\left[\left(\theta_i - \int_0^1 \theta_j dj\right)^2\middle|s\right]$$

We can write each $\theta_j = \kappa x_j + (1 - \kappa)y$ and since we are taking expectations conditional on knowing s , we can write the signals x_j and y as the state s plus mean-zero Gaussian noise with variance σ_x^2 and σ_y^2 , respectively. We write these as $s + \varepsilon_{x_j}$ and $s + \varepsilon_y$, so we can write $\theta_j = s + \kappa\varepsilon_{x_j} + (1 - \kappa)\varepsilon_y$.

$$\begin{aligned}\mathbb{E}_i[u_i|s] &= -\alpha\mathbb{E}_i[(s + \kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y - s)^2|s] \\ &\quad - (1 - \alpha)\mathbb{E}_i\left[\left(s + \kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y - \int_0^1 (s + \kappa\varepsilon_{x_j} + (1 - \kappa)\varepsilon_y dj)\right)^2\middle|s\right] \\ &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - (1 - \alpha)\mathbb{E}_i\left[\left(\kappa\varepsilon_{x_i} + \int_0^1 (\kappa\varepsilon_{x_j} dj)\right)^2\right]\end{aligned}$$

Because $\mathbb{E}[\varepsilon_{x_j}] = 0$, $\int_0^1 (\kappa\varepsilon_{x_j} dj) = 0$. Furthermore, since all of the ε_{x_j} and ε_y are independent we have

$$\begin{aligned}\mathbb{E}_i[u_i|s] &= -\alpha\mathbb{E}_i[(\kappa\varepsilon_{x_i} + (1 - \kappa)\varepsilon_y)^2] - (1 - \alpha)\mathbb{E}_i\left[(\kappa\varepsilon_{x_i})^2\right] \\ &= -\alpha(\kappa^2\sigma_x^2 + (1 - \kappa)^2\sigma_y^2) - (1 - \alpha)\kappa^2\sigma_x^2 \\ &= -\alpha(1 - \kappa)^2\sigma_y^2 - \kappa^2\sigma_x^2\end{aligned}$$

□

Omitted Proofs from Section 3

Claim 1. *If there exists an equilibrium in noisy strategies, where player i chooses $\tilde{\theta}_i = \theta_i + \eta_i$ and each player's η_i is drawn independently from a distribution H_i , then in particular there exists such an equilibrium strategy profile in which the mean of each H_i is zero.*

Proof. Since the distributions which generate the η_i are revealed after each player announces $\tilde{\theta}_i$, choosing to draw noise from a distribution with non-zero mean cannot improve the privacy that player i achieves, since a representative player j can simply subtract this mean when constructing her posterior distribution over x_i . Additionally, choosing a mean other than zero makes the utility from the guessing component strictly worse. Finally, if we assume that all players other than i choose to add noise drawn from a distribution with the same mean, then in the coordination portion of the utility function, player i choosing a noise distribution mean other than the common one is dominated by choosing the common one. In particular, if everyone else chooses mean-zero noise, player i should as well.

□

Lemma 4. *In the game with finitely many players, the players' utility functions in the privacy-aware game separate additively into the utility in the privacy unaware game, a penalty in ν_i , and a privacy term as*

$$\begin{aligned} v_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) &= (1 - \beta)(-\alpha(\tilde{\theta}_i - s)^2 - (1 - \alpha)(\tilde{\theta}_i - \bar{\theta})^2) + \beta\rho(\tilde{\theta}_i) \\ &= (1 - \beta)u_i(\theta_i, \tilde{\theta}_{-i}) + (1 - \beta) \left(\alpha + \left(1 - \frac{1}{n}\right)^2 (1 - \alpha) \right) \nu_i + \beta\rho(\tilde{\theta}_i), \end{aligned}$$

where ν_i denotes the variance of the noise-generating distribution H_i of player i and u_i is the utility function in the privacy-unaware game.

Proof. Writing $\tilde{\theta}_i$ as $\theta_i + \eta_i$, i.e. a deterministic component plus random noise, we can decompose the various pieces of the utility function as follows. The first part is

$$-\alpha\mathbb{E}_i[(\theta_i + \eta_i - s)^2] = -\alpha\mathbb{E}_i[(\theta_i - s)^2 + \eta_i(\theta_i - s) + \eta_i^2] = -\alpha\mathbb{E}_i[(\theta_i - s)^2] - \alpha\nu_i,$$

where we have again used the independence of η_i to conclude that $E_i[\eta_i(\theta_i - s)] = 0$.

The second term is

$$-(1 - \alpha)\mathbb{E}_i \left[\left(\theta_i - \eta_i - \frac{1}{n} \sum_{j=1}^n \tilde{\theta}_j \right)^2 \right] = -(1 - \alpha)\mathbb{E}_i \left[\left(\theta_i \left(1 - \frac{1}{n}\right) + \eta_i \left(1 - \frac{1}{n}\right) - \frac{1}{n} \sum_{j \neq i} \tilde{\theta}_j \right)^2 \right]$$

Rewriting gives:

$$-(1 - \alpha)\mathbb{E}_i \left[\left(\theta_i \left(1 - \frac{1}{n}\right) - \frac{1}{n} \sum_{j \neq i} \tilde{\theta}_j \right)^2 \right] - (1 - \alpha) \left(1 - \frac{1}{n}\right) \nu_i - (1 - \alpha)\mathbb{E}_i \left[\eta_i \left(1 - \frac{1}{n}\right)^2 \frac{1}{n} \sum_{j \neq i} \tilde{\theta}_j \right]$$

After collecting terms, we see that it suffices to argue that the final term

$$(1 - \alpha)\mathbb{E}_i \left[\eta_i \left(1 - \frac{1}{n}\right)^2 \frac{1}{n} \sum_{j \neq i} \tilde{\theta}_j \right] = 0,$$

which follows from the fact that η_i is independent of all the other parameters of the game as well as the other η_j . □

Corollary 2. *The proof is nearly identical to that of Lemma 4. Suppose that all agents play a noisy strategy $\tilde{\theta}_i = \theta_i + \eta_i$, with η_i being a random variable and $\mathbb{E}[\eta_i] = 0$. Then an agent's utility can be decomposed into*

$$\mathbb{E}_i[v_i(\tilde{\theta}, \tilde{\theta}_{-i})] = (1 - \beta)\mathbb{E}_i[u(\theta_i, \tilde{\theta}_{-i})] + (1 - \beta)\nu_i + \beta\rho(\tilde{\theta}_i)$$

where u_i is the utility function in the original game and ν_i is the variance of the noise-generating distribution H_i .

Proof. By definition,

$$\mathbb{E}_i[v_i(\tilde{\theta}, \tilde{\theta}_{-i})] = (1 - \beta)\mathbb{E}_i[u_i(\tilde{\theta}_i, \tilde{\theta}_{-i})] + \beta\rho(\tilde{\theta}_i),$$

so if we show that $u_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) = u_i(\theta_i, \tilde{\theta}_{-i}) - \nu_i$ we will be done. We can write

$$\begin{aligned} \mathbb{E}_i[u_i(\tilde{\theta}_i, \tilde{\theta}_{-i})] &= -\alpha\mathbb{E}_i[(\tilde{\theta}_i - s)^2] - (1 - \alpha)\mathbb{E}_i[(\tilde{\theta}_i - \bar{\theta})^2] \\ &= -\alpha\mathbb{E}_i[(\theta_i + \eta_i - s)^2] - (1 - \alpha)\mathbb{E}_i[(\theta_i + \eta_i - \bar{\theta})^2] \\ &= -\alpha\mathbb{E}_i[(\theta_i - s + \eta_i)^2] - (1 - \alpha)\mathbb{E}_i[(\theta_i - \bar{\theta} + \eta_i)^2] \end{aligned}$$

Expanding these terms, we have

$$\begin{aligned}\mathbb{E}_i[u_i(\tilde{\theta}_i, \tilde{\theta}_{-i})] &= -\alpha \left(\mathbb{E}_i[(\theta_i - s)^2] + 2\mathbb{E}_i[\eta_i(\theta_i - s)] + \mathbb{E}_i[(\eta_i^2)] \right) \\ &\quad - (1 - \alpha) \left(\mathbb{E}_i[(\theta_i - \bar{\theta})^2] + 2\mathbb{E}_i[(\eta_i)(\theta_i - \bar{\theta})] + \mathbb{E}_i[\eta_i^2] \right)\end{aligned}$$

Now the first terms of each line sum to exactly $u_i(\theta_i, \tilde{\theta}_{-i})$. On the other hand, the sum of the last two terms is $-\mathbb{E}_i[\eta_i^2] = -\nu_i$. To complete the proof, we show that these middle two terms are, in fact, zero. To see this, notice that at \mathcal{I}_i , η_i is yet unrealized with $\mathbb{E}_i[\eta_i] = 0$, but is independent of s and θ_i and thus of $\bar{\theta}$. Hence,

$$\mathbb{E}_i[\eta_i(\theta_i - s)] = \mathbb{E}_i[\eta_i]\mathbb{E}_i[\theta_i - s] = 0,$$

Moreover, η_i is independent of each $\tilde{\theta}_{-i}$, and agent i 's action cannot unilaterally change $\bar{\theta}$, so

$$\mathbb{E}_i[\eta_i(\theta_i - \bar{\theta})] = \mathbb{E}_i[\eta_i(\theta_i - \bar{\theta}_{-i})] = \mathbb{E}_i[\eta_i]\mathbb{E}_i[\theta_i - \bar{\theta}_{-i}] = 0$$

□

Lemma 5. *In an equilibrium of the game with finitely many players where the optimal action is $\tilde{\theta}_i^* = \theta_i^* + \eta_i$, the optimal choice of θ_i^* and the variance ν^* for the noise-generating distribution H_i from which η_i is drawn must satisfy*

$$\theta_i^* = \frac{\alpha n^2 \mathbb{E}_i[s]}{\alpha(2n-1) + (n-1)^2} + \frac{(1-\alpha)(n-1)\mathbb{E}_i\left[\sum_{j \neq i} \theta_j\right]}{\alpha(2n-1) + (n-1)^2}$$

and

$$\frac{\partial \rho}{\partial \nu^*} = -\frac{-(1-\beta)\left(\alpha + \left(1 - \frac{1}{n}\right)^2(1-\alpha)\right)}{\beta}.$$

Proof. Using Lemma 4, we can decompose the expected utility of agent i as

$$\mathbb{E}_i[v_i(\tilde{\theta}_i, \tilde{\theta}_{-i})] = (1-\beta)\mathbb{E}_i[u_i(\theta_i, \tilde{\theta}_{-i})] - (1-\beta)\left(\alpha + \left(1 - \frac{1}{n}\right)^2(1-\alpha)\right)\nu_i + \beta\rho(\tilde{\theta}_i),$$

which is the sum of a piece that depends on θ_i and a piece that depends on ν_i .

The agent can therefore optimize each piece separately with her choice of θ_i and ν_i . Lemma 1 again gives the first order condition on θ_i^* - the optimal deterministic component is a convex combination of expectations about the state and the average action (which, under the assumption of a noisy symmetric linear equilibrium, contains added, but mean-zero, noise).

To find the first order condition on ν^* , we can write

$$0 = \frac{\partial v_i}{\partial \nu^*} = -(1-\beta)\left(\alpha + \left(1 - \frac{1}{n}\right)^2(1-\alpha)\right) + \beta\frac{\partial \rho}{\partial \nu^*}$$

and solve for $\frac{\partial \rho}{\partial \nu^*}$ to get the result.

□

Proposition 4. *In an equilibrium where the optimal action is $\tilde{\theta}_i^* = \theta_i^* + \eta_i$, the optimal choice of θ_i^* and the variance ν^* for the noise-generating distribution H_i from which η_i is drawn must satisfy*

$$\theta_i^* = \alpha \mathbb{E}_i[s] + (1 - \alpha) \mathbb{E}_i[\tilde{\theta}] \quad \frac{\partial \rho}{\partial \nu^*} = -\frac{1 - \beta}{\beta}.$$

Proof. Using Lemma 2, we can decompose the utility of agent i as

$$\mathbb{E}_i[v_i(\tilde{\theta}, \tilde{\theta}_{-i})] = (1 - \beta) \mathbb{E}_i[u(\theta_i, \tilde{\theta}_{-i})] - (1 - \beta) \nu_i - \beta \rho(\tilde{\theta}_i),$$

which is the sum of a piece that depends on θ_i and a piece that depends on ν_i .

The agent can therefore optimize each piece separately with her choice of θ_i and ν_i . Proposition 1 gives the first order condition on θ_i^* .

To find the first order condition on ν^* , we can write

$$0 = \frac{\partial v_i}{\partial \nu^*} = -(1 - \beta) - \beta \frac{\partial \rho}{\partial \nu^*}$$

and solve for $\frac{\partial \rho}{\partial \nu^*}$ to get the result. □

Omitted Proofs from Section 4

Lemma 8. *The price of privacy in the game where agents play a linear strategy has the form*

$$\text{PoP}(\tau_x, \tau_y, \beta) = 1 + \frac{\nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]}$$

where the expected utility is with respect to the game with signal variances σ_x^2 and σ_y^2 .

Proof. First note that

$$\mathbb{E}_i[u_i(\theta_i, \tilde{\theta}_{-i})] = \mathbb{E}_i[u_i(\theta_i, \theta_{-i})]$$

because the noise added to θ_i has a mean of zero. Now, the Corollary 2 lets us write

$$\text{PoP}(\tau_x, \tau_y, \beta) = \frac{\mathbb{E}_i[u_i(\tilde{\theta}_i^*, \tilde{\theta}_{-i}^*)]}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]} = \frac{\mathbb{E}_i[u_i(\tilde{\theta}_i^*, \tilde{\theta}_{-i}^*)] + \nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]}$$

where we have factored out all of the negative signs. Combining with the previous part, we have that

$$\text{PoP}(\tau_x, \tau_y, \beta) = \frac{\mathbb{E}_i[u_i(\tilde{\theta}_i^*, \tilde{\theta}_{-i}^*)] + \nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]} = \frac{\mathbb{E}_i[u_i(\tilde{\theta}_i^*, \theta_{-i}^*)] + \nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]} = 1 + \frac{\nu_i^*}{\mathbb{E}_i[u_i(\theta_i^*, \theta_{-i}^*)]},$$

as desired. □

Lemma 9. *Consider an instance of the privacy-aware game where an aggregator observes the actions of n agents (either all of the agents in the finite case or some uniformly random sample in the finite or infinite case), the signal variances are σ_x^2 and σ_y^2 , and players choose to add mean-zero noise with variance ν_i^* . Then the utility of the aggregator, as measured by the variance of the sample average about the true state s is given by*

$$\begin{aligned} \mathcal{U}_{agg}(\sigma_x^2, \sigma_y^2, \nu_i^*, n) &= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \tilde{\theta}_i \right) - s \right)^2 \middle| s \right] \\ &= \frac{\kappa^2}{n} \sigma_x^2 + \frac{\nu_i^*}{n} + (1 - \kappa)^2 \sigma_y^2 \end{aligned}$$

We can find the utility in the privacy-unaware game by letting $\nu_i^* = 0$. This recovers a result in Morris and Shin [17].

Proof.

$$\begin{aligned}
\mathcal{U}_{agg}(\sigma_x^2, \nu_i^*, \sigma^2, n) &= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \tilde{\theta}_i \right) - s \right)^2 \middle| s \right] \\
&= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \kappa x_i + (1 - \kappa)y + \eta_i \right) - s \right)^2 \middle| s \right] \\
&= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \kappa x_i + \eta_i \right) + (1 - \kappa)y - s \right)^2 \middle| s \right] \\
&= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \kappa(s + \varepsilon_{x_i}) + \eta_i \right) + (1 - \kappa)(s + \varepsilon_y) - s \right)^2 \middle| s \right] \\
&= \mathbb{E} \left[\left(\left(\frac{1}{n} \sum_{i=1}^n \kappa \varepsilon_{x_i} + \eta_i \right) + (1 - \kappa)\varepsilon_y \right)^2 \middle| s \right] \\
&= \mathbb{E} \left[\left(\frac{\kappa}{n} \sum_{i=1}^n \varepsilon_{x_i} + \frac{1}{n} \sum_{i=1}^n \eta_i + (1 - \kappa)\varepsilon_y \right)^2 \right] \\
&= \frac{\kappa^2}{n} \sigma_x^2 + \frac{\nu_i^*}{n} + (1 - \kappa)^2 \sigma_y^2
\end{aligned}$$

where we have decomposed x_i and y into s plus mean-zero Gaussian noise ε_{x_i} and ε_y . □