Screening and the Welfare Trade-offs of Privacy

Anastasios Dosis ESSEC
Casey Rothschild Wellesley

July 2024

High-Level Motivating Questions

- Agents often have information which is
  1. Payoff relevant
  2. Known by them
  3. Potentially learnable by a Principal with whom they interact.
- We are interested, broadly, in the question of the welfare consequences of allowing such learning <u>when the principal can also screen agents</u>

High-Level Motivating Questions

1. When is information revelation efficient?
2. To what extent does it cause a tradeoff between efficiency and distributional goals?
   - E.g.: revealing consumer's information might improve efficiency...
   - ... but at the expense of consumer well-being
3. Do these answers change if agent's information is endogenous?
   - The result of an investment decision
   - Taken in advance, but with knowledge of the extent to which information will be revealed

Example 1 – Software Investment

- An organization plans to purchase a new software system from a monopoly provider.
- The monopoly can provide a range of system qualities.
- The organization's valuation of quality will depend on their organization readiness.
- To what extent should they let (or should laws prevent) the monopolist observe their readiness?
- (How) Does the answer change if readiness is endogenous via an ex-ante investment in preparation?

Example 2 – Patent Licences

- An inventor's new product relies on a previously patented idea.
- She is negotiating a licensing deal with the patent-holder.
- Licensing could be limited or expansive .
- She has information about the potential market for her new product.
- To what extent should they let (or should laws prevent) the patent holder from conducting research about their product and its market potential?
- (How) Does the answer change if the market potential depends on prior-to-negotiation product development and market research?

Example 3 – Firm-Specific Human Capital

- An employer will assign a set of tasks to their employee, in exchange for monetary compensation.
- The employee has (potentially) private information about the ease with which she can complete those tasks.
- To what extent should the employer learn (or be prevented from learning) the employee's skill?
- (How) Does this change if the skill depends on prior investments that the employee makes in firm-specific human capital acquisition?

Example 4 – Taxation by an Extractive State

- An extractive state (government) will raise revenue through a non-linear tax on the earnings of its residents.
- Residents have (potentially) private information about their skill in earning income (or disutility of unobservable effort at earning).
- To what extent is it desirable for the state to obtain information about resident's skills?
- (How) does that change if the skill depends on unobservable investments?

Meta-Preview of Results

1. Information revelation is:
   - Always good for the principal iff types are exogenous
   - Potentially (but not always) Pareto improving
2. With exogenous types, there is a general tradeoff within the class of Pareto optimal allocations:
   - More information is good for efficiency (and the principal)
   - But bad for agents

Meta-Preview of Results

1. Information revelation is:
   - Always good for the principal iff types are exogenous
   - Potentially (but not always) Pareto improving
2. With exogenous types, there is a general tradeoff within the class of Pareto optimal allocations:
   - More information is good for efficiency (and the principal)
   - But bad for agents
3. With endogenous types, there is not always such a tradeoff:
   - Limiting information can be Pareto improving
   - In fact, zero information provision can be the unique Pareto optimum

Meta-Preview of Results – Exogenous Types

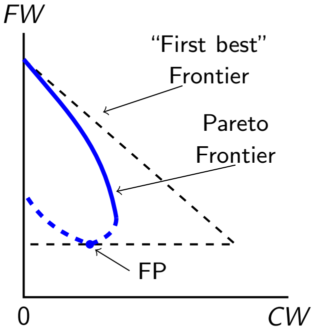

1. Full efficiency requires zero consumer welfare
2. Full privacy might or might not be Pareto optimal
3. Tradeoff between efficiency and consumer welfare

Our Contributions – Endogenous Types

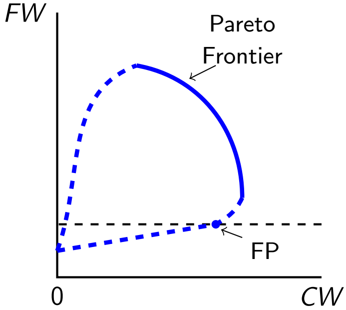

1. Information revelation may or may not be Pareto improving
2. Too much information revelation can be Pareto reducing
   - Intuition: anticipated information revelation can reduce investment incentives

Basic Model – Mussa and Rosen (1978)

- Monopolist sells a product of quality $q \in [0, \infty)$ to a consumer at price t.

$$\Pi(q, t) = t - C(q)$$

with $C(0) = C'(0) = 0,\ C'(q), C''(q) > 0,\ C'''(q) \geq 0$

- Types: $i = 1, ..., N$ (finite)
- Consumer of type i preferences:

$$U_i(q, t) = \theta_i q - t$$

- $\theta_1 < \theta_2 < ... < \theta_N$
- Prior of type i: $\lambda_i > 0$ for every i

Aside – Alternative Models

Basic analysis applies to related models with different interpretations.

- Firm asks for quality q from worker in exchange for wage t

$$\Pi(q,t) = g(q) - t$$

- Consumer of type i preferences:

$$U_i(q,t) = t - q/\theta_i$$

- Higher types are now higher skill types – who would efficiently work harder.
- Endogenizing skill $\leftrightarrow$ investing in firm-specific human capital.

Aside – Alternative Models

Basic analysis applies to related models with different interpretations.

- Firm asks for quality q from worker in exchange for wage t

$$\Pi(q, t) = g(q) - t$$

- Consumer of type i preferences:

$$U_i(q, t) = t - q/\theta_i$$

- Higher types are now higher skill types – who would efficiently work harder.
- Endogenizing skill $\leftrightarrow$ investing in firm-specific human capital.

Key: uninformed party has monopoly power

Basic Model: Information Structures

Information disclosure:

- Exogenous type distribution (prior) λ_i = fraction of type θ_i
- Signals: $s \in S$
- Conditional probabilities: $\pi_i(s)$ – type-conditional distribution over signals
- Implies (Bayesian) posteriors $\beta_i(s)$ = prob. of type i conditional on signal s
- Let $\Sigma = \{S, \{\pi_i\}_i\}$ denote a generic information structure

Basic Model: Information Structures

Information disclosure:

- Exogenous type distribution (prior) λ_i = fraction of type θ_i
- Signals: $s \in S$
- Conditional probabilities: $\pi_i(s)$ – type-conditional distribution over signals
- Implies (Bayesian) posteriors $\beta_i(s)$ = prob. of type i conditional on signal s
- Let $\Sigma = \{S, \{\pi_i\}_i\}$ denote a generic information structure

Special cases

1. $S = 1$: "full privacy" (posteriors equal priors)
2. $S = \{1, \cdots, N\}$, $\pi_i(s_j) = \delta_{ij}$ (0 if $i \neq j$, 1 if $i = j$): "full disclosure"

Basic Model:Timing

0. Planner sets Σ
1. Consumer's type i is realized
2. Signal s is determined via π_i
3. Firm observes s, but not i
4. Firm chooses a price schedule $t(q;s)$ with $t(0;s) = 0 \forall s$.

Solution concept: Perfect Bayesian Equilibrium for parts (1)-(5)

Basic Model:Timing

0. Planner sets Σ
1. Consumer's type i is realized
2. Signal s is determined via π_i
3. Firm observes s, but not i
4. Firm chooses a price schedule $t(q; s)$ with $t(0; s) = 0 \forall s$.

Solution concept: Perfect Bayesian Equilibrium for parts (1)-(5)

Question:

1. What Σs are (constrained) Pareto optimal?

Welfare

- Consumer's ex ante welfare:

$$CW(\Sigma) = \sum_i \lambda_i \sum_s \pi_i(s) U_i(q_i(s), t_i(s))$$

- Total ex ante welfare:

$$TW(\Sigma) \equiv CW(\Sigma) + FW(\Sigma) = \sum_i \lambda_i \sum_s \pi_i(s) \left(\theta_i q_i(s) - C(q_i(s))\right)$$

Definition

Σ is Pareto optimal if there does not exist Σ' with $CW(\Sigma) \leq CW(\Sigma')$ and $FW(\Sigma) \leq FW(\Sigma')$, and with at least one of the two inequalities strict.

Proposition 1 – The Leaky Bucket

Let TW^{max} be the expected consumer plus producer surplus associated with full information revelation.

Proposition (The Leaky Bucket with Exogenous Type Distributions)

1. *If* $CW(\Sigma) > 0$ *then* $TW(\Sigma) < TW^{max}$
2. *Suppose* Σ *and* Σ' *are both Pareto efficient. Then:*

$$CW(\Sigma') > CW(\Sigma) \Rightarrow TW(\Sigma') < TW(\Sigma)$$

In words:

1. Any allocation that gives consumers surplus is (first-best) inefficient
2. Along the (constrained) Pareto frontier, there is a tradeoff between consumer welfare and efficiency.

Proposition 1 – Illustrated

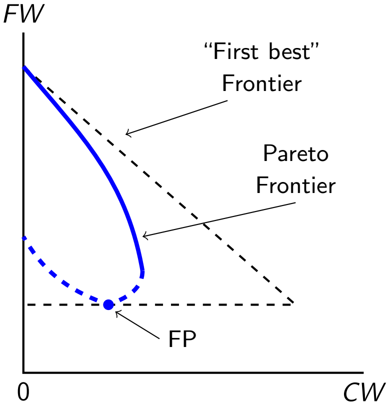

Proposition 1 – Proof (part 1)

- If $CW > 0$, then must be some signal s such that some type i gets positive surplus and $\pi_i(s) > 0$.
- Positive surplus implies some type $j < i$ with $\pi_j(s) > 0$.
- Solution to monopoly's s-conditional problem always distorts j's allocation.
  - Standard result follows from down-binding IC constraints:
    - Starting from no-distortion, adding a small 1 has only second-order efficiency costs, and hence second-order profit effects
    - But first-order eases incentive constraint, allowing greater rent extraction from higher types.
    - So all except the highest type sending s are down-distorted.
- Achieving $TW^{\max}$ requires no distortions
- So $CW(\Sigma) > 0$ implies $TW(\Sigma) < TW^{max}$

Proposition 1 – Proof (part 2)

- Set of implementable (CW, FW) pairs is convex:
  - Suppose (CW_1, FW_1) and (CW_2, FW_2) are feasible
  - Then exists $\Sigma_1 = \left\{\{s_n^1\}_{n\in N_1}, \left\{\pi_i^1\right\}\right\}$ and $\Sigma_2 = \left\{\{s_n^2\}_{n\in N_2}, \left\{\pi_i^1\right\}\right\}$
  - Define $S^3 = \{s_n^1,\}_{n\in N_1} \cup \{s_n^2\}_{n\in N_2}$ and
  - $\pi_i^3(s_n^1) = \alpha\pi_i^1(s_n^1)$, $\pi_i^3(s_n^2) = (1-\alpha)\pi_i^2(s_n^2)$.
  - Then conditional on signal s_n^k, expected CW and FW are the same under Σ_3 as under Σ_k.
  - So $(CW_3, FW_3) = \alpha(CW_1, FW_1) + (1-\alpha)(CW_2, FW_2)$.
- So upper boundary of the set of implementable (CW, FW) pairs is concave in (CW, FW) space.
- It includes $(0, TW^{\max})$.
- And any other point has strictly lower $CW + FW$

Characterizing Pareto Optimal Σ in the Two-Type Case

Results stated here for two-type case $N = 2$

- WLOG, can use revelation principle to restrict attention to incentive compatible menus $(q_1, t_1), (q_2, t_2)$
- Given any posteriors β_1, $\beta_2 = 1 - \beta_1$:

$$
\begin{aligned}
q_1(s) &= \max\left\{0, \xi\left(\theta_1 - \frac{\beta_2(s)}{\beta_1(s)}\Delta\theta\right)\right\} \\
t_1(s) &= \theta_1 q_1 \\
q_2(s) &= \xi(\theta_2) \\
t_2(s) &= \theta_2 q_2 - \Delta\theta q_1,
\end{aligned}
$$

Extra Notation

- Use $\xi = C'^{-1}$, $\Delta\theta = \theta_2 - \theta_1$
- Refer to $q_1 = 0$ case as "exclusion"

Characterizing Pareto Optimal Σ in the Two-Type Case

Proposition (Optimal Information Revelation with Two Types)

$\exists\ \underline{\pi} \geq 0$ such that an informational environment is Pareto optimal if and only if it is equivalent to a two-signal information structure $S = \{s^{priv}, s^2\}$ with $\pi_1(s^{priv}) = 1$ and $\pi_2(s^2) \in [\underline{\pi}, 1]$. Within this class, equilibrium consumer welfare is decreasing and both firm welfare and total (firm plus consumer) welfare are increasing in $\pi_2(s^2)$.

- In words: any Pareto optimum (except full revelation) involves full privacy for type 1 and partial privacy for type 2.
- Notes:
  - easy to construct examples with $\underline{\pi} > 0$, so full privacy for type 2 is Pareto inefficient.
  - ditto with $\underline{\pi} = 0$, so full privacy is Pareto efficient.

Proof

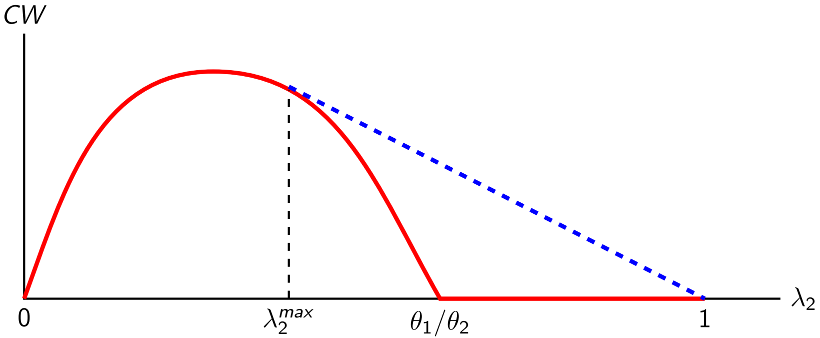

Proof

- 3 steps:
  1. Show that full privacy is not PO when $\lambda_2 > \lambda_2^{max}$ but is PO for $\lambda_2 \leq \lambda_2^{max}$.
  2. Show that if we have (more than two) signals, we can increase CW by having at most two signals.
  3. Show that CW is monotonically decreasing and FW is monotonically increasing in $\pi_2(s^2)$.

Bergemann, Brooks and Morris (AER, 2015; BBM)

- Single product sold by a monopolist to a consumer with a privately known valuation
- If "full privacy" is inefficient (i.e., involves exclusion), then:

1. Information revelation is efficiency improving
2. Information can be revealed in such a way that consumers reap all of the efficiency gains

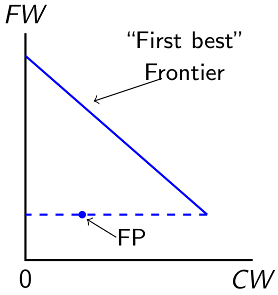

Hagphpanah and Siegel (AERi 2022, JPE 2023)

- Extend BBM to (finite!) multi-product monopolists

1. Information revelation is efficiency improving
2. Information revelation can generically Pareto improve (2023)
3. Whole BBM triangle can be achieved iff "full privacy" involves only "exclusion", not "screening" (2022)
4. At least some of the hypotenuse can be achieved

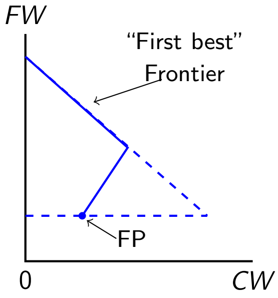

Extending to Endogenous Types

0. Planner sets Σ
1. Consumers choose effort e
2. Consumer's type i is realized probabilistically (to consumer) depending on e
3. Signal s is determined via π_i
4. Firm observe s, but not e or i.
5. Firm chooses a price schedule $t(q; s)$ with $t(0; s) = 0 \forall s$.

Extending to Endogenous Types

0. Planner sets Σ
1. Consumers choose effort e
2. Consumer's type i is realized probabilistically (to consumer) depending on e
3. Signal s is determined via π_i
4. Firm observe s, but not e or i.
5. Firm chooses a price schedule $t(q; s)$ with $t(0; s) = 0 \forall s$.

Additions:

1. "Prior" λ_i replaced by $\lambda_2 = e$, $\lambda_1 = 1 - e$ (WLOG)
   - ONLY two-type case here
2. Consumer preferences now: $\theta q - t - \psi(e)$, ψ increasing, strictly convex

Preliminary Observations – Consumers

Consumer welfare if firm has beliefs e, consumer chooses e':

$$CW(e'|e,\Sigma) = e'\Delta\theta \sum_s \pi_2(s) q_1(s|e,\Sigma) - \psi(e).$$

Equilibrium effort e^* must be optimal and beliefs correct in PBE, so:

$$\overbrace{\Delta\theta \sum \pi_2(s) q_1(s|\Sigma, e^*)}^{MB(e^*|\Sigma)} - \psi'(e^*) = 0.$$

- $\Rightarrow$ unique equilibrium $e^* \in (0,1)$ (as $MB \downarrow$ and $\psi' \uparrow$ in e)
- $\Rightarrow$ $e^*(\Sigma_2) \geq e^*(\Sigma_1) \Leftrightarrow MB(e_1^*|\Sigma_2) \geq MB(e_1^*|\Sigma_1) \Leftrightarrow CW(e_1^*|e_1^*,\Sigma_2) \geq CW(e_1^*|e_1^*,\Sigma_1)$

Preliminary Observations – Consumers

Consumer welfare if firm has beliefs e, consumer chooses e':

$$CW(e'|e,\Sigma) = e'\Delta\theta \sum_s \pi_2(s) q_1(s|e,\Sigma) - \psi(e).$$

Equilibrium effort e^* must be optimal and beliefs correct in PBE, so:

$$\overbrace{\Delta\theta \sum \pi_2(s) q_1(s|\Sigma,e^*)}^{MB(e^*|\Sigma)} - \psi'(e^*) = 0.$$

$\Rightarrow$ unique equilibrium $e^* \in (0,1)$ (as $MB \downarrow$ and $\psi' \uparrow$ in e)

$\Rightarrow$ $e^*(\Sigma_2) \geq e^*(\Sigma_1) \Leftrightarrow MB(e_1^*|\Sigma_2) \geq MB(e_1^*|\Sigma_1) \Leftrightarrow CW(e_1^*|e_1^*,\Sigma_2) \geq CW(e_1^*|e_1^*,\Sigma_1)$

Combine at equilibrium to get:

$$CW^*(\Sigma) = e^*(\Sigma)\psi'(e^*(\Sigma)) - \psi(e^*(\Sigma)),$$

$\Rightarrow$ $CW^*(\Sigma_2) \geq CW^*(\Sigma_1) \Leftrightarrow e^*(\Sigma_2) \geq e^*(\Sigma_1)$ (as $e\psi' - \psi \uparrow$ in e)

Preliminary Observations – Consumers

Consumer welfare if firm has beliefs e, consumer chooses e':

$$CW(e'|e,\Sigma) = e'\Delta\theta \sum_s \pi_2(s) q_1(s|e,\Sigma) - \psi(e).$$

Equilibrium effort e^* must be optimal and beliefs correct in PBE, so:

$$\overbrace{\Delta\theta \sum \pi_2(s) q_1(s|\Sigma, e^*)}^{MB(e^*|\Sigma)} - \psi'(e^*) = 0.$$

$\Rightarrow$ unique equilibrium $e^* \in (0,1)$ (as $MB \downarrow$ and $\psi' \uparrow$ in e)

$\Rightarrow$ $e^*(\Sigma_2) \geq e^*(\Sigma_1) \Leftrightarrow MB(e_1^*|\Sigma_2) \geq MB(e_1^*|\Sigma_1) \Leftrightarrow CW(e_1^*|e_1^*, \Sigma_2) \geq CW(e_1^*|e_1^*, \Sigma_1)$

Combine at equilibrium to get:

$$CW^*(\Sigma) = e^*(\Sigma)\psi'(e^*(\Sigma)) - \psi(e^*(\Sigma)),$$

$\Rightarrow$ $CW^*(\Sigma_2) \geq CW^*(\Sigma_1) \Leftrightarrow e^*(\Sigma_2) \geq e^*(\Sigma_1)$ (as $e\psi' - \psi \uparrow$ in e)

Combine: **consumers like a change in Σ with endogenous e IFF they like the change with exogenous e**

Preliminary Observations – Firms

Lemma

Consider any two information environments Σ_1 *and* Σ_2. *If* $e^*(\Sigma_2) \geq e^*(\Sigma_1)$ *and* $FW(e^*(\Sigma_1), \Sigma_2) \geq FW^*(\Sigma_1)$, *then* $FW^*(\Sigma_2) \geq FW^*(\Sigma_1)$, *and strictly so if the second inequality is strict.*

- Proof: follows from fact that $FW(e, \Sigma)$ is increasing in e.

Preliminary Observations – Firms

Lemma

Consider any two information environments Σ_1 and Σ_2. If $e^*(\Sigma_2) \geq e^*(\Sigma_1)$ and $FW(e^*(\Sigma_1), \Sigma_2) \geq FW^*(\Sigma_1)$, then $FW^*(\Sigma_2) \geq FW^*(\Sigma_1)$, and strictly so if the second inequality is strict.

- Proof: follows from fact that $FW(e, \Sigma)$ is increasing in e.

Corollary (Bootstraping)

Consider any informational environment Σ_1, with associated effort e_1^*. If Σ_1 is Pareto inefficient in the environment with exogenously given effort e_1^*, then it is also Pareto inefficient in the environment with endogenous effort.

Optimal Privacy with Two Types – Endogenous Skills

Proposition (Optimal Information Revelation with Two Types – Endogenous Effort)

$\exists\ \underline{\pi}^C \geq 0$ and $\underline{\pi}^F < 1$ such that an information structure is Pareto optimal if and only if it is equivalent to a two-signal information structure $S = \{s^{priv}, s^2\}$ with $\pi_1(s^{priv}) = 1$ and $\pi_2(s^2) \in \left[\underline{\pi}^C, \underline{\pi}^F\right]$. Within this class, equilibrium consumer welfare is decreasing and both firm welfare and total (firm plus consumer) welfare are increasing in $\pi_2(s^2)$.

- In words: any Pareto optimum ~~(except full revelation)~~ involves full privacy for type 1 and partial privacy for the type 2.
- Notes:
  1. Blue denotes differences from the exogenous-type case.
  2. Key difference: $\underline{\pi}^F < 1$, so full info revelation is necessarily inefficient.
  3. Can construct examples with $\underline{\pi}^C > 0$ – full privacy is Pareto inefficient.
  4. But also $\exists$ examples where $\underline{\pi}^C = \underline{\pi}^F = 0$: no info is unique P.O.
     - Start with any market; use the firm-optimum to construct a new one

Proof Logic

- "Only if" simply bootstraps exogenous case results:
  - If Σ doesn't take this form, it would be inefficient given the fixed $e^*(\Sigma)$
  - So, by Corollary 2, it can't be Pareto efficient with endogenous types
- But we don't have "if:" there could be more Pareto inefficient things
- Indeed, there are:
  - Full information revelation: $CW(e'|e, \Sigma^{FI}) = 0 \; \forall e, e'$ so: $e^*(\Sigma^{FI}) = CW^*(\Sigma^{FI}) = 0$ & $FW^*(\Sigma^{FI}) = \theta_1 \xi(\theta_1) - C(\theta_1)$
  - Full Privacy: $e^*(\Sigma^{Privacy})$, $CW^*(\Sigma^{Privacy}) > 0$, & $FW^*(\Sigma^{Privacy}) > \theta_L \xi(\theta_L) - C(\theta_L)$.
    - If firms believed $e = 0$, then they would offer $(q_1, t_1) = (\xi(\theta_1), \theta_1 \xi(\theta_1))$
    - So type 2 gets rent, so consumers choose $e > 0$, so $e^* \neq 0$
    - Offering $(\xi(\theta_1), \theta_1 \xi(\theta_1))$ to all (yielding $FW = FW(\Sigma^{FI})$) to all is feasible but suboptimal for $e > 0$.
  - So Σ^{FI} is strictly Pareto dominated by $\Sigma^{Privacy}$
  - By continuity, same for Σ close to full information – so $\underline{\pi}^F < 1$.

Comparative Statics

Proposition (Comparative Statics)

The consumer-optimal level of privacy $1 - \underline{\pi}^C$ *increases if any one of the following changes occurs:*

1. ψ' *increases (everywhere)*
2. C' *increases and* C''/C' *weakly increases (everywhere)*
   - *E.g., costs increase proportionally*
3. $\Delta\theta$ *decreases*

Intuition

- Privacy $\Leftrightarrow$ privacy of type 2 (by main proposition)
- Consumers face a tradeoff from more privacy:
  1. Benefit: more likely to earn information rents
  2. Cost: greater privacy increases distortions–hence reduced info rents–conditional on remaining private.

Intuition

- Privacy $\Leftrightarrow$ privacy of type 2 (by main proposition)
- Consumers face a tradeoff from more privacy:
  1. Benefit: more likely to earn information rents
  2. Cost: greater privacy increases distortions–hence reduced info rents–conditional on remaining private.
- Greater ψ' decreases e^* given any privacy level (by $\uparrow$ cost of effort).
  - That reduces distortions at any level of privacy
  - Decreasing the marginal cost of additional distortions
  - So consumers want more privacy

Intuition

- Privacy $\Leftrightarrow$ privacy of type 2 (by main proposition)
- Consumers face a tradeoff from more privacy:
  1. Benefit: more likely to earn information rents
  2. Cost: greater privacy increases distortions–hence reduced info rents–conditional on remaining private.
- Greater ψ' decreases e^* given any privacy level (by $\uparrow$ cost of effort).
  - That reduces distortions at any level of privacy
  - Decreasing the marginal cost of additional distortions
  - So consumers want more privacy
- Greater C' decreases e^*, distortions (by $\downarrow$ marginal benefit of effort)
  - Increasing the marginal cost of additional distortions
  - So consumers want more privacy

Intuition

- Privacy $\Leftrightarrow$ privacy of type 2 (by main proposition)
- Consumers face a tradeoff from more privacy:
  1. Benefit: more likely to earn information rents
  2. Cost: greater privacy increases distortions–hence reduced info rents–conditional on remaining private.
- Greater ψ' decreases e^* given any privacy level (by $\uparrow$ cost of effort).
  - That reduces distortions at any level of privacy
  - Decreasing the marginal cost of additional distortions
  - So consumers want more privacy
- Greater C' decreases e^*, distortions (by $\downarrow$ marginal benefit of effort)
  - Increasing the marginal cost of additional distortions
  - So consumers want more privacy
- Greater $\Delta\theta$ increases distortions given any privacy level.
  - Increasing the marginal cost of additional distortions
  - So consumers want less privacy

Interpretation/Take-Aways

Leaky Bucket + Privacy

1. Firm preferences for privacy are better aligned with efficiency goals
    - This is true quite generally – does not rely on two types
        - Leaky Bucket is general: within P.O. set, firms and "first best" efficiency goals are perfectly aligned
        - And providing additional information is always good for firms (they can ignore it).

Interpretation/Take-Aways

Leaky Bucket + Privacy

1. Firm preferences for privacy are better aligned with efficiency goals
    - This is true quite generally – does not rely on two types
        - Leaky Bucket is general: within P.O. set, firms and "first best" efficiency goals are perfectly aligned
        - And providing additional information is always good for firms (they can ignore it).

Endogenous vs Exogenous

1. Information revelation may or may not be Pareto improving (in both cases)
    - Consumers can sometimes capture some (never all) of the efficiency gains
2. With exogenous types, firms want full information, consumers strictly less
3. With endogenous types:
    - At least some privacy is mutually desirable
    - And maybe even full privacy.

Interpretation/Take-Aways

When types are endogenous

1. It can be better for firms to know less about their customers
   - Of course after the investment decision, they always want to know more
   - So really: valuable to commit/develop a reputation to not pry "too much"
2. In "firm-specific human capital" application:
   - Firms "should" commit to never learn that a worker is unskilled
   - Should (commit to) limit the amount they learn about skilled workers too
   - An argument for not prying into workers' time spent goofing off

Aside – Alternative Models

Same basic analysis applies to closely related models with different interpretations. Example:

- Firm asks for quality q from worker in exchange for wage t

$$\Pi(q, t) = g(q) - t$$

- Consumer of type i preferences:

$$U_i(q, t) = t - q/\theta_i$$

- Higher types are now higher skill types – who would efficiently work harder.
- Endogenizing skill $\leftrightarrow$ investments in firm-specific human capital.

Aside – Alternative Models

Same basic analysis applies to closely related models with different interpretations. Example:

- Firm asks for quality q from worker in exchange for wage t

$$\Pi(q, t) = g(q) - t$$

- Consumer of type i preferences:

$$U_i(q, t) = t - q/\theta_i$$

- Higher types are now higher skill types – who would efficiently work harder.
- Endogenizing skill $\leftrightarrow$ investments in firm-specific human capital.

1. Same analysis applies under appropriate concavity assumptions on g
2. Transform $g(q) \to y$, $q = g^{-1}(y) \equiv h(y)$, $t \to c$

A Tax Application

Same basic analysis applies to closely related models with different interpretations. Example:

- "Firm" asks for "quality" y from worker in exchange for "wage" c

$$\Pi(y, c) = y - c$$

- Consumer of type i preferences:

$$U_i(y, c) = c - h(y)/\theta_i$$

- A Mirrleesian (1971) optimal income tax problem
- y = pre-tax earning, c = post-tax income; Π = tax revenue, θ = skill
- Firm $\rightarrow$ extractive state
- Moreover: because preferences are quasilinear, Rawlsian state optimum yields <u>same</u> earnings profile

Tax Interpretation

- Suppose skill differences are due to effort (and luck)
- Then an extractive (or Rawlsian) state
  - Will benefit from a commitment to some privacy about skill
  - May want to commit to full privacy
- A novel normative justification for “unobservable skill” in Mirrlees (1971)?
  - Skill could be observable
  - But it (can be) better to commit not to observe it

Caveats/Work to Do

1. Characterization rely (for now) on: for two type models, with convex C'.
2. We are actively working on many-type extensions in the exogenous case;
3. The "bootstrap" trick from exogenous to endogenous is special:
   - The total benefit of effort e' is proportional to e':

$$e' \Delta\theta \sum_s \pi_2(s) q_1(s|e, \Sigma)$$

   - This implies two useful things:
     1. e^* is a sufficient stat. for consumer welfare: $CW^* = \psi'(e^*)e^* - \psi(e^*)$
     2. A $\Delta\Sigma$-induced increase in total benefit at fixed $e^* \Rightarrow MB \uparrow$.
   - But it relies on:
     1. L type gets zero surplus (OK – that's general)
     2. Total surplus from other types is proportional to e (not so much!)

Caveats/Work to Do

1. Characterization rely (for now) on: for two type models, with convex C'.
2. We are actively working on many-type extensions in the exogenous case;
3. The "bootstrap" trick from exogenous to endogenous is special:
   - The total benefit of effort e' is proportional to e':

$$e'\Delta\theta\sum_{s}\pi_2(s)q_1(s|e,\Sigma)$$

   - This implies two useful things:
     1. e^* is a sufficient stat. for consumer welfare: $CW^* = \psi'(e^*)e^* - \psi(e^*)$
     2. A $\Delta\Sigma$-induced increase in total benefit at fixed $e^* \Rightarrow MB\uparrow$.
   - But it relies on:
     1. L type gets zero surplus (OK – that's general)
     2. Total surplus from other types is proportional to e (not so much!)
4. Though it does work for:
   - $\lambda_1(e) = (1-e)$ (lowest type); $\lambda_i(e) = e\lambda_i$, $i > 1$
   - I.e., effort has a probability e of "success", and conditional on success there is a fixed probability of different levels of success.

Thanks

Feedback and ideas very welcome!

Thanks for having me!